

IBM Tivoli Storage Manager for Databases
Version 7.1.6

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



IBM Tivoli Storage Manager for Databases
Version 7.1.6

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



Note:

Before you use this information and the product it supports, read the information in “Notices” on page 281.

Fifth edition (June 2016)

This edition applies to version 7, release 1, modification 6 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.

© **Copyright IBM Corporation 1997, 2016.**

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

Contents

Tables vii

About this publication ix

Who should read this publication ix

Publications x

Reading syntax diagrams x

| Data Protection for Microsoft SQL

| Server updates V7.1.6 xiii

Chapter 1. Getting started, planning information 1

Volume Shadow Copy Service framework 1

 Data protection in VSS environments 2

Data backup processing. 4

 Database backup types 4

 Data backup methods 6

 Policy management with Data Protection for SQL

 Server 7

Data restore processing 19

 VSS fast restore processing 19

 VSS instant restore processing 19

Data Protection for SQL Server with IBM SAN

Volume Controller and IBM Storwize V7000 20

 IBM System Storage requirements 21

AlwaysOn Availability Groups with Data Protection
for SQL Server in a Windows Failover Cluster

environment 21

 Availability database backup operations 22

 Availability database restore operations 23

 Enhanced data protection for Microsoft SQL

 AlwaysOn Availability Groups 24

Automated Tivoli Storage Manager server failover
for data recovery 26

Chapter 2. Installing, upgrading, and migrating. 29

Prerequisites 29

Installing and configuring Data Protection for SQL

Server 30

 Installing Data Protection for SQL Server 31

 Completing the installation configuration 31

 Verifying the configuration 33

 Customizing the configuration 33

Installing on a local system 34

Installing Data Protection for SQL Server on

Windows Server Core 35

Installing Data Protection for SQL Server in a cluster

environment 36

Implementing a silent (unattended) installation of

Data Protection for SQL Server 36

 Silently installing the Tivoli Storage Manager

 client 37

 Silently installing Data Protection for SQL Server
 with the setup program 37

 Silently installing with the Microsoft Installer
 program 39

 Options in silent installations 40

 Creating and testing a silent installation package
 on a DVD or a file server. 41

 Silent installation error messages 42

Upgrading Data Protection for SQL Server 42

Data Protection for SQL Server migration 43

Chapter 3. Configuring 45

Proxy node definitions for VSS backups 45

 Required node names for basic VSS operations 45

 Required node names for basic VSS offloaded
 backups 46

Specifying configuration parameters for Tivoli
Storage Manager 47

Specifying Data Protection for SQL Server node
name parameters 50

Specifying configuration and options files in
non-default locations 51

Setting user preferences 51

 Data Protection properties 52

Configuring VSS operations in an SQL Server

failover cluster environment. 62

Configuring Tivoli Storage Manager policy to set
automatic expiration and version control (VSS and

legacy backups) 63

 Setting automatic expiration for VSS backups 63

 Setting automatic expiration for legacy backups 64

Configuring Data Protection for SQL Server by
using the Tivoli Storage Manager Configuration

Wizard 65

Configuring a remote system with a TSM

Configuration 67

Manually configuring Data Protection for SQL

Server 68

 Configuring the computer that runs the SQL
 Server 68

 Configuring the Tivoli Storage Manager server 70

 Configuring the system that runs the offloaded
 backups 70

 Verifying the configuration of Data Protection for
 SQL Server 71

Manually configuring Data Protection for SQL

Server on Windows Server Core 73

 Creating a node on the Tivoli Storage Manager
 server 73

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

 environment 75

 Configuring the client in the Windows Server
 Core environment 75

Configuring Data Protection for SQL Server on Windows Server Core	77
Changing Data Protection for SQL Server configuration values on Windows Server Core.	79
Configuring Data Protection for SQL Server with Tivoli Storage Manager in a clustered environment	80
Configuring availability replicas to run scheduled data backups	81
Transitioning SQL Server backups from Tivoli Storage FlashCopy Manager to Tivoli Storage Manager	82
Configuring the Tivoli Storage Manager server	82
Configuring the workstation that runs the SQL Server	83
Transitioning standard SQL Server databases to the AlwaysOn node	85
Examples of IBM SAN Volume Controller and IBM Storwize V7000 configuration scenarios	86

Chapter 4. Protecting data 89

Prerequisites	89
Security requirements for backup and restore operations	89
VSS and legacy backup methods	89
Enhanced data protection for Microsoft SQL AlwaysOn Availability Groups	92
Starting Microsoft Management Console.	94
Starting the Data Protection for SQL Server command-line interface	95
Getting help for Data Protection for SQL Server commands.	95
Managing Data Protection for SQL Server installations remotely	95
Adding remote systems	96
Determining managed storage capacity	97
Managing backup, restore, and automation tasks in the Task Manager	97
Backing up SQL Server data	98
Enabling SQL Server backup compression	99
Verifying the integrity of legacy databases by using the checksum option	99
Creating VSS backups of SQL Server databases	100
Creating legacy backups of SQL Server databases.	102
Creating legacy backups of SQL Server groups or files	105
Backing up SQL Server databases on Windows Server Core	106
Deleting SQL Server backups	107
Deactivating legacy backups of SQL Server databases.	107
Mounting VSS snapshots to remote servers	108
Mounting SQL Server backups	109
Restoring SQL Server databases and files	109
Setting single-user mode for restore operations	109
Setting data restore options in Microsoft Management Console.	110
Restoring SQL Server data	112
Restoring an SQL Server database to an alternate instance	115
Restoring the master database	117

Restoring SQL Server databases with full-text catalogs and indexes	118
Restoring SQL Server databases from virtual machine snapshots	118
Restoring SQL Server databases on Windows Server Core	121
Restoring SQL Server file groups and files from legacy backups	122
Protecting SQL Server data in a Windows Server Core environment	123
Backing up SQL Server databases on Windows Server Core	124
Restoring SQL Server databases on Windows Server Core	124
Changing Data Protection for SQL Server configuration values on Windows Server Core	125
Viewing, printing, and saving reports	126

Chapter 5. Automating 127

Preparing to use Windows PowerShell cmdlets with Data Protection for SQL Server.	127
Cmdlets for Microsoft Management Console	128
Cmdlets for protecting Microsoft SQL Server data	129
Automating tasks	130
Tivoli Storage Manager task scheduler	131

Chapter 6. Troubleshooting 133

Diagnosing problems.	133
Error log files	133
Determining that the problem is a Data Protection for SQL Server issue or a general VSS issue	134
Resolving reproducible problems.	138
Troubleshooting VSS backup and restore operations	138
Troubleshooting VSS and SAN Volume Controller, Storwize V7000, or DS8000	142
Resolving problems with IBM Support	143
Viewing trace and log files	144
Gathering traces for the Data Protection client when using VSS technology	145
Gathering information about SQL Server with VSS before you call IBM.	145
Gathering files from SQL Server with VSS before calling IBM.	146
Viewing and modifying system information	147
Emailing files to IBM Support	148
Online IBM support	149

Chapter 7. Performance tuning 151

Buffering (Legacy only)	151
Data striping (Legacy only).	151
LAN-free data movement (Legacy and VSS)	152

Chapter 8. Reference information. 153

Command-line overview	153
Data Protection for SQL Server parameters by backup method.	154
Backup command	155

Backup syntax	156	VSS restore command-line considerations	226
Backup positional parameters	159	Restore syntax	227
Backup optional parameters	162	Restore positional parameters	230
Legacy backup examples	172	Restore optional parameters	232
VSS backup examples	174	Legacy restore output examples	247
Changetmpassword command	175	VSS restore output examples	254
Changetmpassword	175	Restorefiles command	257
Changetmpassword positional parameters	176	Restorefiles syntax.	258
Changetmpassword optional parameters	176	Restorefiles positional parameters	259
Changetmpassword output examples	179	Restorefiles optional parameters	259
Delete Backup command	179	Restorefiles examples.	263
Delete Backup syntax.	179	Set command	263
Delete Backup positional parameters	180	Set syntax	263
Delete Backup optional parameters	180	Set positional parameters	264
Delete Backup example	183	Set optional parameters	267
Help command.	183	Set output examples	267
Help syntax	184	Unmount Backup command	268
Help positional parameters.	184	Unmount Backup syntax	268
Help output examples	185	Unmount Backup positional parameter.	268
Inactivate command (Legacy only)	188	Unmount Backup optional parameters	269
Inactivate syntax	189		
Inactivate positional parameters	190	Appendix A. Frequently asked	
Inactivate optional parameters.	191	questions	273
Inactivate output examples.	194		
Mount Backup command	198	Appendix B. Accessibility features for	
Mount Backup syntax	198	the Tivoli Storage Manager product	
Unmount Backup positional parameter.	199	family.	279
Mount Backup optional parameters	199		
Query command	202	Notices	281
Query syntax	203		
Query positional parameters	205	Glossary	285
Query optional parameters	207		
Query output examples	213	Index	287
Query Managedcapacity command	224		
Query Policy command	225		
Restore command	225		

Tables

1. Data Protection for SQL Server backup types	5	13. Database restore options	110
2. Silent installation options	40	14. Database restore views	113
3. Silent installation features (base client only)	41	15. Database restore selection options.	113
4. Commands for creating a silent installation package.	41	16. Database backup views	115
5. Required node names for basic VSS operations	45	17. Database backup views	116
6. Required node names for basic VSS offloaded backups.	46	18. Node names used to set access.	120
7. Diagnostics modes and their usage.	55	19. Database restore views	123
8. Backup strategy characteristics	92	20. Cmdlets to protect Microsoft SQL Server data	129
9. Database backup views	101	21. VSSADMIN commands	136
10. Database backup views	102	22. Data Protection for SQL Server optional parameters	154
11. Database backup options.	103	23.	169
12. Database backup views	105	24.	210
		25.	242

About this publication

With Data Protection for Microsoft SQL Server software you can back up and restore Microsoft SQL Server databases to Tivoli® Storage Manager storage.

Data Protection for SQL Server provides a connection between an SQL Server and a Tivoli Storage Manager server. This connection allows SQL data to be protected and managed by Tivoli Storage Manager.

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

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

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

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.

To view IBM product documentation, see IBM Knowledge Center.

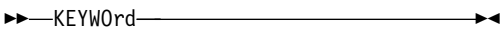
Reading syntax diagrams

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.


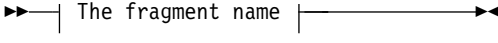
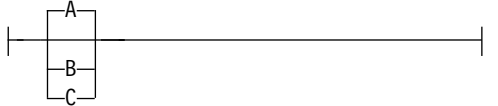
- The  symbol indicates the beginning of a syntax diagram.
- The  symbol at the end of a line indicates the syntax diagram continues on the next line.
- The  symbol at the beginning of a line indicates a syntax diagram continues from the previous line.
- The  symbol indicates the end of a syntax diagram.

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 																		
Symbols: Enter these symbols exactly as they appear in the syntax diagram.	<table><tr><td>*</td><td>Asterisk</td></tr><tr><td>{ }</td><td>Braces</td></tr><tr><td>:</td><td>Colon</td></tr><tr><td>,</td><td>Comma</td></tr><tr><td>=</td><td>Equal Sign</td></tr><tr><td>-</td><td>Hyphen</td></tr><tr><td>()</td><td>Parentheses</td></tr><tr><td>.</td><td>Period</td></tr><tr><td></td><td>Space</td></tr></table>	*	Asterisk	{ }	Braces	:	Colon	,	Comma	=	Equal Sign	-	Hyphen	()	Parentheses	.	Period		Space
*	Asterisk																		
{ }	Braces																		
:	Colon																		
,	Comma																		
=	Equal Sign																		
-	Hyphen																		
()	Parentheses																		
.	Period																		
	Space																		

Syntax Diagram Description	Example
Variables: Italicized lowercase items (<i>var_name</i>) denote variables. In this example, you can specify a <i>var_name</i> when you enter the KEYWORD command.	Variables
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 <i>repeat</i> 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 <i>must</i> specify one item. In this example, you <i>must</i> 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

Syntax Diagram Description	Example
<p>Repeatable Choices:</p> <p>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.</p> <p>In this example, you can choose any combination of A, B, or C.</p>	<p>Repeatable choices</p> 
<p>Syntax Fragments:</p> <p>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.</p>	<p>Syntax fragments</p>  <p>The fragment name:</p> 

Data Protection for Microsoft SQL Server updates V7.1.6

This document provides information about what's new or what has changed in Data Protection for Microsoft SQL Server version 7.1.6.

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

What's new

There are no new features.

What has changed

Data Protection for SQL Server installation program is renamed

The Data Protection for SQL Server installation program, **setup.exe**, is renamed to **spinstall.exe**. Use **spinstall.exe** to install Data Protection for SQL Server.

Tivoli Storage Manager products are delivered only electronically

The Tivoli Storage Manager family of products are now delivered as electronic releases. Physical media is no longer available as part of the standard product offering in Passport Advantage. The announcement letter also provides this information. Announcement letters are available at <http://www.ibm.com/common/ssi>.

IBM is developing products by using a continuous delivery model, which supports more frequent product updates. If you require physical media for legal or other reasons, contact your sales representative.

Instant restore operations using Microsoft VSS ResyncLuns API is updated

In Microsoft Windows Server 2012 or later, you can use Tivoli Storage FlashCopy Manager to run instant restore operations on all storage systems. With Microsoft updates for Microsoft Windows Server 2012, restrictions to running instant restore operations using the Microsoft VSS ResyncLuns API no longer exist. You can also now run an instant restore to a shared cluster disk or file system using the Microsoft VSS ResyncLuns API.

Tip:

You must have a Tivoli Storage FlashCopy Manager licence to avail of this feature.

Chapter 1. Getting started, planning information

With Tivoli Storage Manager for Databases: Data Protection for Microsoft SQL Server software, you can back up and restore Microsoft SQL Server databases to Tivoli Storage Manager storage or local shadow volumes. A *local shadow volume* contains data that is stored on shadow volumes, which are local to a disk storage system.

Data Protection for SQL Server provides a connection between an SQL Server and a Tivoli Storage Manager, which allows SQL Server data to be protected and managed by Tivoli Storage Manager. Data Protection for SQL Server protects SQL Server data and improves the availability of SQL Server databases. You can continue to run primary applications on your database servers while data is backed up and restored.

You can use a command-line interface or graphical user interface (GUI) to back up and restore SQL Server databases. For more information about backing up and restoring SQL Server databases, see your SQL Server documentation.

Microsoft supports the Microsoft Legacy application programming interface (API) for streaming backup and restore operations. Microsoft also supports the use of Volume Shadow Copy Service (VSS) technology for backup and restore operations.

Data Protection for SQL Server uses the Tivoli Storage Manager API to communicate with the Tivoli Storage Manager, and the SQL Server API to communicate with SQL Server.

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

Volume Shadow Copy Service framework

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

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

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

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

configure the VSS hardware and software provider as required.

Data protection in VSS environments

Windows

The characteristics of Volume Shadow Copy Service (VSS) backup and restore operations can affect management tasks, for example, the backup types that you can run, the backup granularity, and the backup storage location options.

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

VSS backup characteristics

Backups can be stored on local shadow volumes, a Tivoli Storage Manager server, or at both locations. You can define different policy settings for each backup location.

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

VSS backup requirements

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

Follow these guidelines when you plan for VSS backups:

- When you perform VSS operations, ensure that at least 200 MB of free disk space is on your Windows System Drive. This space is used to store the metadata files for Data Protection for SQL Server.
- Schedule and run legacy backups as part of your strategy.
- Use basic disks, which are initialized for basic storage. A basic disk consists of basic volumes, such as primary partitions, extended partitions, and logical drives.
- If you plan to keep VSS snapshot backups only on local shadow volumes, know how to implement the configuration options of your VSS hardware provider. For example, if your VSS hardware provider supports a full-copy snapshot versus a copy-on-write snapshot mechanism, full-copy type implementations have greater disk storage requirements. However, full-copy type implementations do not rely on the original volume to restore the data and are less risky. Copy-on-write implementations require less disk storage but rely on the original volume to restore the data.
- If you run parallel VSS backups, stagger the start time of the backups by at least 10 minutes. This interval ensures that the snapshot operations do not overlap.
- If you run parallel VSS backups, configure the parallel instance backups so that snapshots of the same volumes are not created.
- If you run parallel VSS backups, ensure that parallel backups do not create a snapshot of the same LUN.
- Do not place multiple volumes on the same LUN. Configure a single volume, single partition, and single LUN as one-to-one.

VSS restore characteristics

In a VSS restore operation, VSS backups SQL database files and log files) that are on Tivoli Storage Manager server storage are restored to their original location on the SQL Server.

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

- You can restore only SQL Server VSS backups to the same SQL Server instance.
- You can restore full and copy-only full backup types.
- You cannot run differential, individual filegroups, individual files, and set backup operations because VSS cannot restore that data.
- Data is restored at the database level.
- You can restore one or more databases from a VSS snapshot backup on Tivoli Storage Manager server storage.
- You can run restore operations in a Microsoft Windows Failover Clustering or Veritas Cluster Server (VCS) environment.
- By using the **/relocatedir** option, you can restore a VSS backup directly from Tivoli Storage Manager server storage to an alternate location.
- You cannot use parallel VSS fast restore or VSS instant restore operations with Microsoft Windows Server 2008 or later versions.

VSS restore requirements

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

If you complete VSS snapshot backups with the backup destination parameter set to TSM, *restore* processing also refers to an image-level restore from the Tivoli Storage Manager server.

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

VSS instant restore

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

- If you do not want to overwrite the source volumes, ensure that you set the **Instant Restore** option to **No** in Microsoft Management Console (MMC).
- VSS instant restore processing requires that the local disk is not accessed by other applications, for example, Windows Explorer.
- When you run a VSS instant restore operation, verify that there is no other data on the volumes that are being restored.
- Before you start a VSS instant restore operation, ensure that any previous background copies that contain the volumes that are being restored are completed. XIV, SAN Volume Controller, or Storwize V7000 with space-efficient target volumes do not need to be completed.

VSS fast restore

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

When you complete a VSS restore operation from local shadow volumes, the bytes that transfer are displayed as 0 because no data (0) is restored from the Tivoli Storage Manager server.

Because of an 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 is taken.

VSS operations in IBM N-series and NetApp environments

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

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

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

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

In addition, Write Anywhere File Layout (WAFL) reserves space, that is, blocks equal to two times the specified size of the LUN to be used. This space reservation ensures that write operations are allowed for virtual disks. The following example shows how to calculate the size of the volumes:

```
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
```

Data backup processing

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

Database backup types

With Data Protection for SQL Server, you can use the common interface in the Volume Shadow Copy Service (VSS) framework to create database backups.

VSS backups are at the volume and file-level. Legacy backups are a stream of bytes that Data Protection for SQL Server stores on the Tivoli Storage Manager server.

You can back up Data Protection for SQL Server data by using the following methods:

Table 1. Data Protection for SQL Server backup types

Data Protection for SQL Server	
Full database backup (Legacy and VSS)	With this method, Data Protection for SQL Server backs up an SQL Server database and the portion of the transaction log that is necessary to provide a consistent database state. With this backup type, the copy includes enough information from any associated transaction log to create a backup that is consistent with itself. The portion of the log that is included contains only the transactions that occur from the beginning of the backup until its completion.
Copy-only full backup (Legacy and VSS)	With this method, Data Protection for SQL Server creates data backups that do not affect existing backup and restore processes and can be retained in the longer term. For example, you can use this type to back up a log before an online file restore operation. In this example, the copy-only full backup is used once. After the backup is restored, it is deleted.
Differential backup (only Legacy)	<p>With this method, Data Protection for SQL Server backs up only the data pages in an SQL Server database instance that changed after the last full backup. A portion of the transaction log is also backed up.</p> <p>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 application. For example, if you run a full SQL Server-to-disk backup, and run a differential backup, the differential backup is associated with the SQL Server disk backup.</p> <p>You cannot use differential backup for databases on the secondary replica in Microsoft SQL Server 2012.</p>
Log backup (only Legacy)	<p>With this method, Data Protection for SQL Server backs up only the contents of an SQL Server database transaction log since the last successful log backup. This type of backup is preceded by a full backup or an equivalent type of backup.</p> <p>Log backups normally follow full backups. The portion of the log that is 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 follows 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.</p>

Table 1. Data Protection for SQL Server backup types (continued)

Data Protection for SQL Server	
File backup (only Legacy)	With this method, 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 conflicts if you must back 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 backup is not required.
Group backup (only Legacy)	<p>With this method, Data Protection for SQL Server backs up only the contents of a specified SQL Server file group. You can back up the set of database tables and indexes within a specific group of files.</p> <p>The group is specified as part of the setup within SQL Server when you define the database files. If no group is specified and all the database files are part of the primary group, you cannot partially back up or partially restore the database by using the group.</p>
Set backup (only Legacy)	With this method, Data Protection for SQL Server backs up the contents of specified SQL Server file groups and files as a unit.

Data backup methods

Windows

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

For SQL Server, you can also run legacy backups that create a copy of all or part of an SQL Server database or logs on Tivoli Storage Manager storage media.

VSS data backups

You can store VSS backups on local VSS shadow volumes, or, when integrated with Tivoli Storage Manager, in Tivoli Storage Manager server storage.

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

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

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

- For data that is backed up only to 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.
- For data that is backed up to VSS shadow volumes and 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.

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

Offloaded VSS backups

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

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

SQL Server legacy backups

With Data Protection for SQL Server, you can run legacy backups and store the backup on Tivoli Storage Manager server.

Legacy backups are unlike VSS backups because volume and file-level data are not backed up with this method.

Policy management with Data Protection for SQL Server

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

Although 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. In addition, verify that enough available storage space is assigned to the volumes to accommodate your backup operations. The shadow copy volume that is the storage destination of a snapshot must have sufficient space for the snapshot.

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

Backup expiration 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 when the expiration date is past or the maximum number of backup versions that must be retained is reached.

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

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

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

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

How policy affects backup management on Data Protection for SQL Server

A 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 and Data Protection for SQL Server as a *node*.

Data that is backed up to Tivoli Storage Manager storage from the Data Protection for SQL Server node is stored and managed according to settings that you specify in the Tivoli Storage Manager server policy.

The Tivoli Storage Manager policy manages the VSS backups that are placed in Tivoli Storage Manager server storage pools. The Tivoli Storage Manager server manages VSS backups.

If you use IBM Tivoli Storage Manager for Copy Services and upgrade 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.

Tivoli Storage Manager requires that sufficient storage space is available to create shadow volumes 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 temporarily.

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

For example, if **verexists**=3, then at least three sets of target volumes must be allocated for the backup to complete successfully. If only two sets of target volumes are allocated, the third and subsequent backup attempts fail. If more sets of target volumes exist than the number specified by the **verexists** parameter, these sets are ignored by the 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 in the GUI, the **VSSPOLICY** parameter is set in the `tdpsql.cfg` file.

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

You must manage the policy for local backups to reconcile the local backup repository with the information that is stored on the Tivoli Storage Manager server. For example, if target volume LUNs that are used for a local backup are removed from the storage system, the information that represents the backup on the Tivoli Storage Manager server must be reconciled. Similarly, if a 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 system. The local backup manager is released so that these LUNs can be used for future backup operations. Tivoli Storage Manager automatically detects when these situations occur and completes the reconciliation.

Recommended Tivoli Storage Manager policy settings

When you use Data Protection for Microsoft SQL Server, verify that the policy contains the recommended settings.

Domain

A policy domain contains policy sets, management classes, and copy groups. Create a policy domain on the Tivoli Storage Manager server to be used exclusively for Data Protection for SQL Server backups.

Policy sets

Policy sets contain management classes (which contain copy groups) that determine the rules by which Data Protection for SQL Server backups are performed and managed.

Define the policy set to the policy domain to which Data Protection for SQL Server backups belong. The policy set must be activated and only one policy set can be active in the policy domain.

Management class

A management class is a policy object that users can bind to each file to specify how the file is managed.

Define a management class for backups on local shadow volumes, and a management class for backups on Tivoli Storage Manager server storage. Different management classes provide the opportunity for specialized policies for each storage destination. For example, you can maintain six versions of local VSS backups of a given storage group (*verexists=6*) while you maintain only two versions of the same storage group on Tivoli Storage Manager server storage (*verexists=2*). In addition, you can create a separate management class for copy backup types for use in long term storage. Such policies can maximize storage resources and provide more control over your storage strategy.

Important: Because VSS backup processing requires sufficient storage space to create shadow volumes, ensure that you specify *verexists=N+1* to keep *n* backups on local shadow volumes.

You can have five active backups of the same storage group because legacy backups on Tivoli Storage Manager server storage, VSS backups on Tivoli Storage Manager server storage (COPY and FULL), and VSS backups on local shadow volumes (COPY and FULL) all have different Tivoli Storage Manager server naming and therefore, can each have their own management class. Ensure that you plan your backup strategy before you define management classes.

Metadata considerations (Legacy only)

The management classes for Data Protection for SQL Server metadata are to be identical to the corresponding management classes for database data. The exception is that the metadata management classes do not allow migration to removable media. If any Data Protection for SQL Server metadata is on removable media, queries might require media mounts, and backups or restores might require more media mounts.

Data objects and their associated meta objects are to have the same version limits and retention values. However, because meta objects might be restored as a result of a Data Protection for SQL Server query, consider storing meta objects in a disk-only storage pool. In doing so, a media mount is not necessary to resolve the query. To achieve this scenario, implement these steps:

1. Define a separate management class with a Copy Destination pointing to a disk pool that does not have any removable media in its hierarchy.
2. Bind all meta objects to that management class by using an **include** statement in the Data Protection for SQL Server options file.

Alternatively, you can choose to use the same management class (and storage pools) for both meta and data objects if you rarely need the meta objects, or need them only immediately preceding a restore when a volume mount is required anyway. In many cases, you can also obtain the meta object information from SQL Server

as recorded in its *msdb* database. For more information about binding backup objects to specific management classes, see “Configuring Tivoli Storage Manager policy to set automatic expiration and version control (VSS and legacy backups)” on page 63.

Copy group

A copy group controls how backup versions are generated, located, and expired.

Define the copy group as a backup copy group and not as an archive copy group. Because Data Protection for SQL Server stores all objects as backup objects on Tivoli Storage Manager in backup storage pools, an archive copy group is not required, although an archive copy group can exist. The following backup copy group parameters significantly influence your backup policy:

VERExists

Determines the maximum number of SQL Server database backup versions to retain for databases that exist on the Data Protection for SQL Server client system.

VERDeleted

Determines the maximum number of SQL Server database backup versions to retain for databases that were deleted from the Data Protection for SQL Server client system after being backed up by Tivoli Storage Manager.

RETEExtra

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

RETOnly

Determines the number of days to retain the last SQL Server database backup version of a database that was deleted from the Data Protection for SQL Server client system. Because of version limit, incremental backups are not expired. Only one version of an incremental backup object exists at a time because incremental backups are always uniquely named. However, all legacy backup objects for an SQL Server storage group are inactivated when a new full backup of that SQL Server storage group is completed. VSS backup objects remain active. The retention period that is set in the **RETOnly** parameter controls the expiration of incremental backup objects.

When you set the value of the **RETOnly** parameter for incremental backups, the value must be, at a minimum, as long as the value that is set for the full backup objects to which the incremental backups are associated. You can use the same management class for incremental backups and the full backup objects (that are retained the longest) to ensure that an adequate value is used. However, when a new legacy full backup of that SQL Server storage group is completed, all legacy backup objects for an SQL Server storage group are inactivated. In this scenario VSS backup objects remain active.

MODE, SERialization, FREQuency

Because these parameters settings do not apply to Data Protection for SQL Server, accept the default values.

When you plan a backup strategy, as a best practice, consult with the Tivoli Storage Manager server administrator about recommended parameter settings.

Storage pool

A storage pool is a named set of storage volumes and the destination that is used by the Tivoli Storage Manager server to store data.

A single restore operation can require a full backup, a differential backup, and multiple incremental backups. Use collocation if backups are stored on removable media. Specify collocation by file space (**define stgpool COLlocate=Filespace**) if you plan to restore multiple storage groups in parallel.

Tip: As a best practice, use collocation because data for any one storage group is stored within one Tivoli Storage Manager server file space.

Creating a local backup policy

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

Before you begin

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

Procedure

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

What to do next

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

Specifying policy binding statements

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

About this task

A default policy binds any backups that are not explicitly bound to a named policy. Policy binding is available in environments with or without a Tivoli Storage Manager server.

Procedure

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

The way that you set policy is different for VSS and legacy backups:

- VSS backups: Specify the VSSPOLICY statement in the Data Protection for SQL Server configuration file. By default, the configuration filename is `tdpsql.cfg`.
- Legacy backups: Specify the INCLUDE and EXCLUDE statements in the Data Protection for SQL Server options file. By default, the options filename is `dsm.opt`.

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

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

Binding backups to a policy

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

About this task

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

Procedure

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

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

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

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

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

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

VSSPOLICY statements for backup types

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

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

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

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

Sample VSSPOLICY statements

The following example shows the syntax of a VSSPOLICY statement:

```
VSSPOLICY * * COPY TSM VSS_FULL_TSM_MC
```

Legacy examples

Legacy backups use INCLUDE and EXCLUDE statements in the Data Protection for SQL Server options file. The examples in the following tables group statements that are intended to be used together. For example:

```
...\full*  
...\full*\*
```

and

```

\...\file\fl*\*
\...\file\fl*
\...\fl\file*

```

Object matches for backuptype	Specification
Example for all objects	\...*
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 (<i>fl</i>) and group (<i>gl</i>)	\...\gl\group* \...\fl\file*
Example for INCLUDE statements with file (<i>fl</i>) and group (<i>gl</i>)	\...\group\gl** \...\group\gl* \...\gl\group* \...\file\fl** \...\file\fl* \...\fl\file*
Example for EXCLUDE statements with group or file object names beginning with <i>g</i> or <i>f</i>	\...\g*\group* \...\f*\file*
Example for INCLUDE statements with group or file object names beginning with <i>g</i> or <i>f</i>	\...\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 <i>Db1</i>	\...\Db1\...*
Example for all objects with database name <i>Db1</i> beginning with <i>Db</i>	\...\Db*\...*
Ambiguous	\...*\...*

backuptype object with database matches	Specification
Example for EXCLUDE statements using full, diff, copyfull objects with database name <i>Db1</i>	\...\Db1\full* \...\Db1\copyfull* \...\Db1\diff*
Example for EXCLUDE and INCLUDE statements using full, diff, copyfull objects with database name <i>Db1</i>	\...\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 <i>Db1</i>	\...\Db1\...\log* \...\Db1\...\group*
Example for EXCLUDE and INCLUDE statements using log, group, file, set objects with database name <i>Db1</i>	\...\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 (<i>g1</i> , <i>f1</i>) with database name <i>Db1</i>	\...\Db1\g1\group* \...\Db1\f1\file*
Example for EXCLUDE and INCLUDE statements using all group or file object names (<i>g1</i> , <i>f1</i>) with database name <i>Db1</i>	\...\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 <i>g</i> or <i>f</i> with database name <i>Db1</i>	\...\Db1\g*\group* \...\Db1\f*\file*
Example for EXCLUDE and INCLUDE statements using all group or file object names beginning with <i>g</i> or <i>f</i> with database name <i>Db1</i>	\...\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**

backuptype object with database matches	Specification
Example for EXCLUDE statements using \...\group* (there is no equivalent for INCLUDE statements)	\... * * \group*
Example for EXCLUDE statements using \...\g1\group* (there is no equivalent for INCLUDE statements)	\... * \g1 \group*
Ambiguous	\... * \... \log*
Nothing (typeInfo missing)	\... \Db1 \set*

Examples in the following table reflect these guidelines:

- If you use only EXCLUDE statements with only \meta\, all objects (including data) are excluded.
- If you use 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 (<i>g1</i>)	\... \meta \... \g1 \group*
Example for all data group object names (<i>g1</i>)	\... \data \... \group \g1* \... \data \... \group \g1* *
Example for all meta group object names beginning with <i>g</i>	\... \meta \... \g* \group*
Example for all data group object names beginning with <i>g</i>	\... \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 <i>Db1</i>	\... \meta \... \Db1 \... * \... \data \... \Db1 \... *
Example for all meta objects with database name <i>Db1</i>	\... \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 <i>Db1</i>	\... \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 (<i>g1</i>) with database name <i>Db1</i>	\... \meta \... \Db1 \g1 \group*

Meta and data object with database matches	Specification
Example for all data group object names (<i>g1</i>) with database name <i>Db1</i>	<code>\...\data\...\Db1\group\g1*</code>
Example for all meta object names beginning with <i>g</i> with database name <i>Db1</i>	<code>\...\meta\...\Db1\g*\group*</code>
Example for all data group object names beginning with <i>g</i> with database name <i>Db1</i>	<code>\...\data\...\Db1\group\g*</code>
Same as <code>\...\meta\...\Db1\...\group*</code> (No equivalent for data objects)	<code>\...\meta\...\Db1*\group*</code>
Same as <code>\...\meta\...\full*</code> (No equivalent for data objects)	<code>\...\meta\...*\full*</code>
Same as <code>\...\meta\...\group*</code> (No equivalent for data objects)	<code>\...\meta\...**\group*</code>
Same as <code>\...\meta\...\g1\group*</code> (No equivalent for data objects)	<code>\...\meta\...*\g1\group*</code>
Ambiguous	<code>\...\meta\...*\...\log*</code> <code>\...\data\...*\...\log*</code>
Nothing (qualifiers missing)	<code>\...\meta*\...\data*</code>
Server matches	Specification
Example for all objects from all servers beginning with <i>SQL</i>	<code>SQL*\...*</code>
Example for all objects from all server instances with host <i>SQL2012</i>	<code>SQL2012\...*</code>
Example for all objects from server <i>SQL2012\INST1</i>	<code>SQL2012\INST1\...*</code>
Example for all objects from all servers beginning with <i>SQL2012\INST</i>	<code>SQL2012\INST*\...*</code>
Same as <code>SQL2012\...*</code>	<code>SQL2012*\...*</code>
Example for all meta or data objects from server <i>SQL2012\INST1</i>	<code>SQL2012\INST1\meta\...*</code> <code>SQL2012\INST1\data\...*</code>
Example for all meta or data objects from all named server instances with host <i>SQL2012</i>	<code>SQL2012*\meta\...*</code>
Example for all meta or data objects from all server instances with host <i>SQL2012</i>	<code>SQL2012\...\meta\...*</code>
Example for all objects from server default instance (if no instance name matches <i>??ta</i>)	<code>SQL2012\meta\...*</code> <code>SQL2012\data\...*</code>

Data restore processing

Windows

Data Protection for SQL Server can use the Microsoft Volume Shadow Copy Service (VSS) framework to complete fast and instant restores of database backups. For SQL Server, you can run legacy restore operations from Tivoli Storage Manager server.

In a VSS restore operation, you restore one or more databases from a VSS backup on Tivoli Storage Manager server storage to the original location on the SQL Server.

VSS fast restore processing

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

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

- Restore SQL Server full and copy-only full backup types. You cannot restore differential, individual file groups, individual files, and set backups.
- Restore data at the database level. For custom application and file system data, however, you can only restore data at the file system level. File overwriting occurs even if the file is marked read-only.
- Restore only SQL Server VSS backups to the same SQL Server instance.
- Restore SQL Server VSS backups to an alternate location by using the **/relocatedir** option.
- Restore a VSS backup to an alternate database.
- Restore one or more databases or file systems from a VSS snapshot backup on local shadow volumes that are managed by Tivoli Storage FlashCopy Manager.
- For SQL Server, custom application, and file system data, restore the data in Microsoft Windows failover clustering environments.

VSS instant restore processing

A VSS instant restore operation restores data by using a hardware-assisted restore method. A FlashCopy operation is an example of a hardware-assisted restore method.

You can run a VSS instant restore operation only when all of the data from the storage group or database to restore is on storage systems 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 operation is completed.

The data that is to be restored must be on a storage system that is valid for VSS instant restore operations. If data is not on an XIV, SAN Volume Controller, or Storwize V7000 systems with space-efficient target volumes, you must ensure that background copies that use the volumes are restored.

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

- IBM System Storage DS8000® series requires IBM System Storage Support for Microsoft Volume Shadow Copy Service software.

- SAN Volume Controller requires IBM System Storage Support for Microsoft Volume Shadow Copy Service software.
- Storwize V7000 requires IBM System Storage Support for Microsoft Volume Shadow Copy Service software.
- XIV has separate VSS Provider software.
- Backups can be restored only to the same storage system from which they are backed up.

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

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

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

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

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

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

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

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

- To determine how much storage space is required for each local backup, the backup LUNs require the same amount of storage space as the original LUNs. For example, if you have a 100 GB database on a 200 GB LUN, you need a 200 GB LUN for each backup version.
- Do not use a combination of persistent and nonpersistent VSS snapshots.
- Do not mix COPY and NOCOPY FlashCopy relationships from the same source volume or volumes.
- Enable the autoexpand option for the space-efficient target volumes to avoid out-of-space conditions.
- Allocate enough space for space-efficient target volumes to hold 120 % of the data that is expected to change on the source volume in the time between snapshots. For example, if a database changes at a rate of 20 % per day, VSS backups complete every six hours, and a steady rate of change throughout the

day is assumed. The expected change rate between snapshots is 5 % of the source volume (20/4). Therefore, the allocated space for the space-efficient target volumes is to be 1.2 times 5 % equal to 6 % of the source volume size. If the rate of change is not consistent throughout the day, allocate enough space to the target volumes to accommodate the highest expected change rate for the period between snapshots.

You can use VSS instant restore operations with

- 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 occurs in the correct order.

IBM System Storage requirements

If you use IBM System Storage DS8000 series, SAN Volume Controller, or Storwize V7000 storage systems, be aware of database, log, file, and LUN settings.

Follow these guidelines when you plan for IBM System Storage:

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

AlwaysOn Availability Groups with Data Protection for SQL Server in a Windows Failover Cluster environment

The AlwaysOn Availability Groups (AAG) function provides high availability and disaster recovery at the SQL database level. A SQL AlwaysOn failover cluster instance provides high availability and disaster recovery at the SQL Server level.

Data Protection for SQL Server protects availability databases in an AAG and AlwaysOn Failover cluster instances. An AAG can contain a set of primary databases and multiple 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 Node manages backups of availability databases. This node is a shared node that allows data backups and restores of availability databases from any replica. Data Protection for SQL Server treats a backup as originating on a single SQL Server regardless of which node of the cluster is backed up.

You can use AAGs with SQL AlwaysOn failover cluster instances to complete the following tasks:

- In an AAG, you can deploy a group of single or clustered server instances, each holding a copy of all databases
- As many as nine online copies of a database (one primary and up to eight secondary copies) can be in an AAG
- You can use synchronous and asynchronous replication
- You can use log shipping
- You can use automatic and manual failover modes
- Databases within an AAG fail over as a group

Follow these guidelines when you set up AAGs in a Microsoft Windows failover cluster environment or in a Veritas cluster server cluster environment:

- Install Data Protection for SQL Server on each cluster node and configure each node identically. Specify identical configurations in the Data Protection for SQL Server options file.
- Ensure that each availability replica of an availability group is on a different node in the same Windows failover cluster environment.
- Use the configuration wizard to register an AlwaysOn Node on the Tivoli Storage Manager server.
- To access a clustered SQL Server, identify the virtual server name and specify that name in Data Protection for SQL Server.
- If you use the Tivoli Storage Manager scheduler to automate data backups, install the scheduler service on each node of the cluster to enable failover support.
- You cannot restore a VSS backup to an alternate instance. Restore VSS backups on the same SQL Server instance where the snapshot is taken.

Tip: VSS and legacy full backups of availability databases on secondary replicas are copy-only. The copy-only option is not automatically used with log backups because you can run log backups that truncate logs on secondary replicas.

Availability database backup operations

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

An AlwaysOn Availability Group (AAG) 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 a replica for one or more availability groups. For example, node1 might be a replica for availability group 1 and another availability group.

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 complete the backup.

You can run the following types of VSS backup operations:

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

You can run the following type of legacy backup operations:

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

Restriction: Microsoft does not support legacy full backups on secondary replicas. Data Protection for SQL Server, however, does permit you to run a full backup of a secondary replica based on Tivoli Storage Manager policy.

When you run a full legacy backup of a secondary replica, the underlying implementation of Data Protection for SQL Server is to back up the data as copyfull. However, Data Protection for SQL Server detects the intended full backup operation and applies the Tivoli Storage Manager policy that is associated with the full backup type.

Microsoft Management Console (MMC) and CLI views honor the Tivoli Storage Manager policy that applies to the backup type and in this instance, show the backup type as full. For more information, see Active Secondaries: Backup on Secondary Replicas (AlwaysOn Availability Groups)(<https://msdn.microsoft.com/en-us/library/hh245119.aspx>).

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

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

Availability database restore operations

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

Certain restrictions apply when you restore 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 the database is restored, you must verify that the data on all replicas is transactionally consistent.

To verify that the data is transactionally consistent, 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, manually add the database to the availability group.

VSS restore

You cannot restore a VSS backup to an alternative SQL Server instance. Restore VSS backups to the same SQL Server instance where the snapshot is taken.

AlwaysOn availability databases

For AlwaysOn availability databases, you must set up Data Protection for SQL Server to use an AlwaysOn node name. By default, the AlwaysOn node name is set to the cluster node name for the Availability Group in SQL Server 2012, and later versions.

Enhanced data protection for Microsoft SQL AlwaysOn Availability Groups

You can run VSS (full) and legacy (full, differential, file/set/group, and log) backup operations on a primary replica. You can run copy-only VSS and legacy backup operations, and normal log backups on a secondary replica. You cannot run a differential backup on a secondary replica.

For backups on a secondary replica, the replica must be in the synchronized or synchronizing state. You can have multiple AlwaysOn Availability Groups (AAGs) in a SQL Server cluster. You can also have a mix of standard databases and AAGs on a SQL Server cluster.

When you back up data, you can distribute the backup workload for scalability and isolate backup activity to a dedicated backup node. When you isolate backup activity, it minimizes the effect on production databases.

Given that replicas are copies of the same database, avoid redundant backups of the same databases. Apply retention policies to unique databases.

As a best practice, allow backups from any node in the availability group and enable restore operations from any node in the availability group.

Best practices for backing up data in an AAG

When you use Tivoli Storage FlashCopy Manager for SQL Server to manage AAG backups, consider the following backup options:

Backup priority

Specified per database in an AAG, the backup priority option defines the order in which replicas are used to back up a database in an AAG.

Preferred replica

Specified at an AAG level, the preferred replica option defines whether primary or secondary replicas can be used for backup operations.

- Prefer secondary replica: Scheduled backups occur on a secondary replica, if available. If the secondary replica is not available, you can use the primary replica.
- Secondary only: Scheduled backups can occur only on a secondary replica.
- Primary: Scheduled backups can occur only on the primary replica.
- Any replica: Scheduled backups can occur on any replica.

/USEALWAYSONNode parameter

A parameter option on the **backup** command that provides a common namespace for all backups. Each node authenticates separately with Tivoli

Storage Manager. Backed up data is stored in the AlwaysOnNode namespace by using the Asnode option.

/ALWAYSONPriority parameter

A parameter option on the **backup** command that specifies that a local availability database is backed up only if it has the highest backup priority among the availability replicas that are working properly. This parameter applies only to scheduled backups.

Typical data protection deployments in AAG environments

You can back up data in an AAG in the following ways:

- Distribute a legacy backup across AAG replicas
- Distribute a VSS backup across AAG replicas

Approach®: Legacy backups are distributed across AAG replicas

When you configure your environment to distribute a legacy backup across AAG replicas, follow these steps:

1. Set the preferred replica to **Prefer secondary replica**.
2. Install Tivoli Storage FlashCopy Manager for SQL Server on all replicas that are eligible to run a backup.
3. Create a command script to run a .CMD file with a **backup** command similar to the following example:

```
tdpsqlc backup db1,db2,db3 full /alwaysonpriority
```
4. Associate each Tivoli Storage FlashCopy Manager for SQL Server node with the defined schedule.
5. Run backups on the SQL node according to defined priorities for each database.

Scenario: VSS backups are distributed across AAG replicas

When you configure your environment to distribute a VSS backup across AAG replicas, follow these steps:

1. Set the preferred replica to **Prefer secondary replica**.
2. Install Tivoli Storage FlashCopy Manager for SQL Server on all replicas that are eligible to run a backup.
3. Create a command script to run a .CMD file with a separate **backup** command per database similar to the following sample

```
tdpsqlc backup db1 full /alwaysonpriority /backupmethod=VSS  
backupdest=TSM  
tdpsqlc backup db2 full /alwaysonpriority /backupmethod=VSS  
backupdest=TSM  
tdpsqlc backup db3 full /alwaysonpriority /backupmethod=VSS  
backupdest=TSM
```
4. Associate each Tivoli Storage FlashCopy Manager for SQL Server node with the defined schedule.
5. Run backups on the SQL node according to defined priorities for each database.

Automated Tivoli Storage Manager server failover for data recovery

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

The Tivoli Storage Manager server that Data Protection for 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
- VSS requestor 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 and later versions)

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 is replicated to the secondary server. When the primary server is online again, the backup-archive client automatically fails back to the primary server the next time the backup-archive client connects to the server.

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

- Data Protection for SQL Server must be at least at V7.1 level or later.
- The primary server, secondary server, and backup-archive client must be at least at V7.1 level or later.
- The primary and secondary servers must be set up for node replication.
- The following nodes must be configured for replication with the replstate=enabled option in each node definition on the server:
 - Data Protection node
 - VSS requestor node
 - Remote DSM agent node for offloaded backups
 - 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 restore data in failover mode and the replication status is not current, the recovered data might not be usable. You must wait until the primary server comes back online before you can restore the data.

Chapter 2. Installing, upgrading, and migrating

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

Prerequisites

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

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

For the latest requirements, review the Hardware and Software Requirements technote that is associated with the level of your Data Protection for SQL Server program. This technote is available at this web page: TSM for Databases - All Requirement Documents (<http://www.ibm.com/support/docview.wss?uid=swg21218747>). Follow the link to the requirements technote for your specific release or update level.

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.

Installation process might require a reboot

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

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

Virtualization environment resources

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

For more information about virtualization environments that can be used with Data Protection for SQL Server, see this web page: [IBM Tivoli Storage Manager \(TSM\) guest support for Virtual Machines and Virtualization](#).

Installing and configuring Data Protection for SQL Server

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

Before you begin

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

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

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

About this task

Data Protection for 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 information 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.

Installing Data Protection for SQL Server

Procedure

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

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

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

`package_name.exe`

where `package_name` is like this example:

`7.1.4.0-TIV-TSMSQL-Win.exe`

4. Follow the installation instructions that are displayed on the screen.
5. Click **Finish**.
6. 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.

Completing the installation configuration

Procedure

1. To start Microsoft Management Console (MMC), click **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.
2. If the Tivoli Storage Manager configuration wizard does not start automatically, click **Manage > Configuration > Wizards** in the navigation tree, select the wizard, and click **Start** in the Actions pane.
3. Complete the following pages of the wizard:

Data Protection Selection

Select **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 Server 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 with SQL Server 2012 or later versions, 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.

Verifying the configuration

Procedure

1. Verify that VSS is working correctly.

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

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

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

Snapshot Volume Selection

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

VSS Snapshot Tests

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

Completion

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

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

```
query tdp
query tsm
query sql
```
 - d. 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.
 - e. When verification is complete, you can use Data Protection for SQL Server to back up and restore SQL Server data.
 - f. Back up and restore a set of test data.

Customizing the configuration

Procedure

After you successfully configure and verify Data Protection for SQL Server, customize your settings by defining your policy settings and scheduled operations to meet your business requirements.

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.

Installing on a local system

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

Before you begin

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

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

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

About this task

Data Protection for 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 information 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

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. Download the appropriate package file from one of the following websites, or you can access the files from the product DVD:
 - For a first time installation or a new release go to Passport Advantage at: IBM Passport Advantage. Passport Advantage is the only website from which you can download a licensed package file.
 - For a maintenance fix, go to this FTP site and to the directory that contains the maintenance fix version that you require: Index of Data Protection for Microsoft SQL Server patch files (<ftp://public.dhe.ibm.com/storage/tivoli-storage-management/patches/tivoli-data-protection/sql/>).
 - c. If you download the package from one of the download sites, complete the following steps:

- Verify that you have enough space to store the installation files when they are extracted from the product package.
- Change to the directory where you placed the executable file.

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

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

`package_name.exe`

where `package_name` is like this example:

`7.1.4.0-TIV-TSMSQL-Win.exe`

- Follow the installation instructions that are displayed on the screen.
If you are configuring Data Protection for SQL Server with SQL Server 2012 or later versions, 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.
 - If prompted, restart your system before the installation is completed.
 - Click **Finish**. Microsoft Management Console (MMC) is shared among Data Protection for SQL Server, and Tivoli Storage FlashCopy Manager. If one of these products is installed in a location other than the default location, the setup wizard defaults to the existing installation directory. Use the same directory when you install any of these products on the same computer. The default base directory is `c:\program files\tivoli`.
- 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.

Installing Data Protection for SQL Server on Windows Server Core

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

Before you begin

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

Procedure

- Log on as an administrator.
- Install Data Protection for SQL Server by using the setup wizard.
 - 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**.
 - Follow the installation instructions that are displayed on the screen.
 - Click **Finish**. If prompted, restart your system.

What to do next

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

Installing Data Protection for SQL Server in a cluster environment

You can install Data Protection for SQL Server in a Windows failover cluster environment, and protect clustered SQL Server 2008 databases and later versions.

Before you begin

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

Procedure

1. Install Data Protection for SQL Server on all nodes of your cluster where you intend to perform backups and restore operations.
2. If you use 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. Install Data Protection for SQL Server on a local system if required.

Related tasks:

“Installing on a local system” on page 34

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

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 a silent installation.

About this task

A 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 distribute 25 SQL Server installations 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 in a shared directory on a file server for distribution across the different sites.

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

Silently installing the Tivoli Storage Manager client

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

Ensure that the latest, supported version of the Tivoli Storage Manager backup-archive client is available on your computer. You must be an administrator to complete this installation.

About this task

You use the Windows Installer program (**msiexec.exe**) to install the Tivoli Storage Manager client.

Procedure

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 /! *v "C:\downloads\logs\ba_logs.log"
```

Silently installing Data Protection for SQL Server 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 and later versions, you can also use the setup program to silently install Data Protection for SQL Server on Windows Server Core.

Before you begin

- If you want to install Data Protection for SQL Server on Windows Server Core, first install the Tivoli Storage Manager client on the same computer as Data Protection for SQL Server.
- 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.

About this task

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\7140\enu\setup.exe
- (64-bit) x:\fcm\x64\sql\7140\enu\setup.exe

Procedure

1. Issue the following command to silently install both Data Protection for SQL Server components to the default installation directory:

```
x:\fcm\aaa\sql\7140\enu\setup.exe /s /v/qn
```

where x: is your DVD drive and aaa is either x64 or x86. The following command is an example.

```
x:\fcm\x64\mmc\4140\enu\setup.exe /s /v/qn
x:\fcm\x64\sql\7140\enu\setup.exe /s /v/qn
```

where x: is your DVD drive.

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

```
x:\fcm\x64\mmc\4140\enu\setup.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\7140\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\"
```

3. Review these guidelines as you complete the 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.

Batch files usage in silent installations

Windows

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

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

```
@echo off
rem =====
rem sample silent install script
rem
call x:\fcm\x64\mmc\4140\enu\setup.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\7140\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 =====
```

Silently installing with the Microsoft Installer program

You can use the Microsoft Installer (MSI) program, `msiexec.exe`, to implement a silent installation of Data Protection for SQL Server. If you are protecting Microsoft SQL Server 2012 and later versions, you can also use the MSI program to silently install Data Protection for SQL Server on Windows Server Core.

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 include a number of prerequisites that is installed by Data Protection for SQL Server. When you use `msiexec.exe`, you must install all prerequisites manually.

Before you install and configure Data Protection for SQL Server, verify that you satisfy the hardware and software requirements. For more information, see the topic that describes the minimum hardware and software requirements.

About this task

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

Procedure

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

```
msiexec /i"x:\fcm\aaa\mmc\4140\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. To install Data Protection for SQL Server, issue each of these **`msiexec`** commands on a single line.

```
msiexec /i"x:\fcm\aaa\sql\7140\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

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.

Options in silent installations

The following options can be applied to both silent installation methods, the setup program and the Microsoft Installer (MSI) program.

Review the list of silent installation options in the following tables:

Table 2. 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.
/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: <ul style="list-style-type: none"> • 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. <p>Force Always prompts user to restart after silent installation.</p> <p>Suppress Suppresses prompt to restart after silent installation.</p> <p>ReallySuppress Suppresses all restarts and prompts to restart after silent installation.</p>

Table 2. Silent installation options (continued)

Option	Description
rebootyesno	Specifies whether to restart the system after silent installation. Specify <i>Yes</i> to restart the system after silent installation. Specify <i>No</i> not to restart the system after silent installation.

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 for Data Protection for SQL Server to install if those prerequisite components are not already installed on 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 tables list the silent installation features (case-sensitive) that apply to the base client only.

Table 3. Silent installation features (base client only)

Feature	Description
Client	Data Protection for SQL Server code

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

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

Before you begin

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

About this task

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

Procedure

1. Issue the following commands to create the package:

Table 4. Commands for creating a silent installation package

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

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

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

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

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

Silent installation error messages

The **setup.exe** program can produce error messages if it cannot start properly.

Upgrading Data Protection for SQL Server

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

Procedure

1. Download the updates.
2. To install the updates, run setupfcm.exe.
3. To start Microsoft Management Console (MMC), click **Start > All Programs > Tivoli Storage Manager > Data Protection for Microsoft SQL Server > DP for SQL Management Console**. When you start MMC after you install the updates, the configuration wizard automatically starts. The configuration wizard guides you through the process of provisioning and installing the remaining files.
Depending on the software licenses that are found on the system, the configuration process varies. The wizard provides instructions to guide you through the process.
4. If the configuration wizard does not start automatically, click **IBM Tivoli Storage Manager** in the navigation tree, and click **Configuration**. Then, double-click **Wizards**.

Data Protection for SQL Server migration

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

After you upgrade from an older version of Data Protection for SQL Server to a newer version, you can use VSS data restore operations to restore 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 need to retain the active backup, you must delete the active backup by issuing the **delete backup** command. The existing backup copies can be restored.

Chapter 3. Configuring

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

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.

Proxy node definitions for VSS backups

Because Data Protection for SQL Server use the Tivoli Storage Manager backup-archive client to implement VSS backup operations, you must use node names specifically for VSS operations. In addition, you must use 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 node names. By default, this proxy relationship is defined when you run the configuration wizard. Follow the guidelines in this topic to manually complete the configuration.

The proxy relationship allows node names to process operations on behalf of another node name. When you register these nodes to the 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.

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

- *Target node*: A node name that controls data backup and restore operations and also owns the data on the 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.

To define these nodes, issue 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 5. 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)

Table 5. Required node names for basic VSS operations (continued)

Proxy node type	Node name	Where to specify
Agent node	The Local DSMAGENT Node name that must match the VSS requestor node name	Use the localdsmagentnode parameter in the Data Protection for SQL Server configuration file (tdpsql.cfg)

Target node

The target node name specifies the name of the node that owns the Data Protection for SQL Server backup data on the Tivoli Storage Manager server. 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

The agent node name specifies the name of the node operating for the target node. 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 Microsoft Management Console (MMC) by selecting **VSS backup**. From here, you can update the Local DSMAGENT Node name. Otherwise, use the **tdpsqlc set** command to specify this parameter.

Tip: The agent node and target node are in a proxy relationship, and are used for VSS operations on the SQL server that is being protected.

Required node names for basic VSS offloaded backups

VSS offloaded backups require specific node name settings.

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

Table 6. 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)
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 - Local DSMAGENT 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). To specify this parameter with the Properties window of Microsoft Management Console (MMC), select VSS backup. From the Properties window, you can update the Local DSMAGENT Node name. Otherwise, use the **tdpsqlc set** command to specify this parameter.

Agent node - Remote DSMAGENT Node

This node name is a separate system that must also have the backup-archive client 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 Local DSMAGENT Node and is specified with the **localdsmagentnode** parameter in the Data Protection for SQL Server configuration file (`tdpsql.cfg` by default). To specify this parameter with the Properties window of MMC, VSS backup. From the Properties window, you can update the Local 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 shadow copies that can be transported. You cannot specify a node name if you are using the default VSS system provider.

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 configuration parameters for Tivoli Storage Manager

After Data Protection for SQL Server is registered to Tivoli Storage Manager, you must configure the node name, password, the communications method, and the appropriate parameters to connect to the Tivoli Storage Manager server.

About this task

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

Procedure

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

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

COMMethod

Specify the communication protocol to use between the Data Protection for SQL Server node and the Tivoli Storage Manager server. Depending

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

- For all backups, specify the **COMMMethod** option in the Data Protection for SQL Server options file.
- 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.

For all backups, specify the **COMMMethod** option in the Data Protection for SQL Server options file. For all 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, also specify the **COMMMethod** option in the backup-archive client options file that is used as the remote DSMAGENT node.

NODename

Specify the Tivoli Storage Manager node name that Tivoli Storage Manager uses to identify the system that runs Data Protection for SQL Server.

PASSWORDAccess

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

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

COMPRESSION

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

- The network adapter has a data overload
- Communications between Data Protection for SQL Server and Tivoli Storage Manager server are over a low-bandwidth connection
- Heavy network traffic exists

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

- The computer that runs Data Protection for SQL Server has a processor overload; the added processor usage might cause issues for other applications that include the server. You can monitor processor and network resource usage with the Performance Monitor program that is included with Windows.
- You are not constrained by network bandwidth; you can achieve the best performance by leaving the compression no option unchanged and enabling hardware compaction on the tape drive, which also reduces storage requirements.
- For 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, specify the **COMPRESSION** option in the backup-archive client options file that is used as the remote DSMAGENT node

DEDUPLICATION

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

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

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

ENABLECLIENTENCRYPTKEY

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

Restriction: You can back up encrypted VSS databases only to the Tivoli Storage Manager server. You cannot back up encrypted data to a Data Protection for SQL Server.

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

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

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

If you make changes in the backup-archive client options file, ensure that you restart the Tivoli Storage Manager Client Acceptor Daemon (CAD) service for the SQL Server.

ENABLELANFree

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

What to do next

You can create more Data Protection for SQL Server options files to point to another Tivoli Storage Manager server. You can create more than one options file, where each file contains different parameters to use with a single Tivoli Storage Manager server.

Related tasks:

“Enabling SQL Server backup compression” on page 99

Specifying Data Protection for SQL Server node name parameters

You must register the system where Data Protection for SQL Server is installed to the Tivoli Storage Manager server with a node name.

About this task

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

The node name owns and manages all Data Protection for SQL Server data that is backed up to the Tivoli Storage Manager server.

If you run 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. The node name must match the SQL virtual server name.

Procedure

1. Specify the node name with the **nodename** option in the `dsm.opt` options file. By default, the `dsm.opt` options file is in the Data Protection for SQL Server installation directory.
2. To run VSS operations, register node names for more systems if required.
3. Configure the following Tivoli Storage Manager parameters when you register your Data Protection for SQL Server node name to the Tivoli Storage Manager server:
 - **BACKDElete** Specify that the Data Protection for SQL Server node can delete its own backup files from the Tivoli Storage Manager server. You must specify the yes value for this parameter.
 - **MAXNUMMP** Specify the maximum number of mount points that a client node is allowed to use on the Tivoli Storage Manager server during a backup operation. If you use SQL data-stripping with data that is sent directly to a tape pool, set this parameter to a number that is 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 operations when removable media such as tapes are used, or if migration occurs during the backup or restore operation. If other backups or restores operations occur at the same time, the value of this parameter must be large enough to allow for all of the required mount points. If the storage pool for the backup operation has Active Data or Backup Stgpools that are written to simultaneously, the **MAXNUMMP** parameter must also include these mount points.
 - **TXNGroupmax** Specify 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 12 or greater.
 - **COMPRESSIon** (Legacy only) Specify 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` in the backup-archive client options file (`dsm.opt`) in the Data Protection for SQL Server directory.

Specifying configuration and options files in non-default locations

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

Before you begin

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

About this task

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

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

For example:

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

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

```
/SQLCONFigfile=filename # SQL configuration file  
/SQLOPTfile=filename # SQL OPT file  
/SQLINSTancenames=Instance1,Instance2,... # SQL instances to show in the MMC
```

Procedure

Start MMC with the parameters by using `flashcopymanager.exe`, as shown in the following example.

```
flashcopymanager.exe /FSCONFigfile=newcfg.xml /SQLCONFigfile=altsql.cfg  
/SQLINSTancenames=mysql1,mysql2
```

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

Setting user preferences

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

Before you begin

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

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

Procedure

To configure user preferences, complete the following steps:

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

What to do next

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

Data Protection properties

Use property pages to customize your configuration preferences.

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

Server Information

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

The fields that display depends on whether the product is configured for a stand-alone snapshot environment or for a Tivoli Storage Manager environment.

Tip: References to the stand-alone snapshot environment are specific to Tivoli Storage FlashCopy Manager.

Node name

Specifies the name that is used to identify the client node for stand-alone backup operations or backup operations to Tivoli Storage Manager server.

TSM API version

Specifies the version of the Tivoli Storage Manager application programming interface (API).

Server name

For backups to Tivoli Storage Manager, specifies the name of the Tivoli Storage Manager server that you are connected to.

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

Server Network Host name

Specifies the network host name for the Tivoli Storage Manager server.

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

Server type

For backups to Tivoli Storage Manager, specifies the type of operating system of the Tivoli Storage Manager server.

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

Server version

Specifies the version of the Tivoli Storage Manager server.

Compression mode

Specifies whether compression is used during backup operations to the Tivoli Storage Manager server. The possible values are Yes, No, and Client Determined.

Domain name

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

For SQL systems, the domain name, policy set, and management class are listed for the Data Protection or AlwaysOn node.

Active Policy Set

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

Default Management Class

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

Server Password

Use the Server Password property page to change the password for the Data Protection node that you use to access the Tivoli Storage Manager server. This property page applies only to Tivoli Storage Manager configurations.

The following fields are displayed in the property page:

Old password

Type the Tivoli Storage Manager password that you want to change.

New password

Type a new password. Follow the Tivoli Storage Manager server password policy rules.

Confirm new password

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

Policy Management

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

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

The following fields are displayed in the property page:

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

Number of Snapshots to keep

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

This parameter does not apply to incremental backup versions of SQL Server data. Incremental backups do not participate in expirations because of version limit because there is never more than one version of an

incremental backup object. There is only one version of an incremental backup object because incremental backups are always uniquely named.

Days to keep a Snapshot

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

VSS Policy Binding

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

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

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

Managed Capacity

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

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

Diagnostics

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

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

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

Diagnostic modes

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

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

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

Table 7. Diagnostics modes and their usage

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

Normal

Click **Normal** to collect trace and log files for legacy operations.

Complete

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

Custom

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

Enable snapin tracing

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

Set Default Trace Flags

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

Enable Data Protection tracing

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

Enable DSM Agent tracing

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

Enable API tracing

Select this box to enable tracing for the Tivoli Storage Manager API. Click **Review** to view the trace file. Add or update trace flags in the field.

Event log entries

Click **Event log entries** to view Windows application event log for SQL Server, Exchange Server, or File System workloads. For SQL workloads, you can also click **SQL Server Log file** to view the SQL Server log.

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

Screenshot

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

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

Tracing details for each component

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

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

TraceFm.trc
TraceUx.trc

Data Protection

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

TraceFileSql.trc

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

TraceFileAgent.trc

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

TraceFileSqlAPI.trc

SQL Login

Use this property page to set preferences for logging on to the Microsoft SQL Server.

Use Windows Authentication

Select this option to use a trusted connection and allow Microsoft Windows to authenticate the logon.

Use SQL Server Authentication

Select this option to use SQL user ID security. With this type of security, you must enter the logon ID and the password to log on to the Microsoft SQL Server.

User name

Specifies the SQL Server user ID.

Password

Specifies the password to log on to the Microsoft SQL Server.

General (SQL)

Use this property page to specify general preferences for the **SQL Server** workload. This property page applies if the product is configured to back up data to stand-alone storage or Tivoli Storage Manager.

SQL Server

Specify the unique name that identifies the SQL Server instance.

From Server

Specify the SQL Server backups that you want to use for the restore. By default, this field displays the same name for the **SQL Server**.

Wait for tape mounts for backup or restore

Select this box when you want Data Protection for Microsoft SQL Server to wait for tape media to be mounted for backup and restore operations. This setting is applicable when the Tivoli Storage Manager server is configured to store the backup data on tape media. With backup data on removable media, during backup and restore operations, a wait period occurs during storage volume mounts. If a wait occurs, this setting specifies whether Data Protection for Microsoft SQL Server waits for the media mount or stop the current operation. By default, this option is not selected.

Use VSS backups as the default backup method

Select this box to set VSS backups as the default backup method. Ensure that the **Local DSMAGENT Node name** field is specified in the VSS Backup property page. Backups can be restored only by using VSS.

Compress backup by using SQL Server compression

Select this box to enable SQL Server compression during legacy backup operations. This check box is available only if you are running Microsoft SQL Server 2008 or later versions.

Compute SQL Server checksum for legacy backup

When selected, this option is written to the Data Protection for SQL Server preferences file (`tdpsql.cfg`), and can be applied to all legacy backups. If you clear the check box, you ensure that the integrity check does not apply to any legacy database backup.

Estimate % change for differential backup

Specify the value for the estimated change to database pages for differential backups. This estimate is used by Data Protection for Microsoft SQL Server to determine whether enough storage space is available for the backup. The default value is 20. This value becomes the default value for all differential backups.

This field applies only to Data Protection for Microsoft SQL Server legacy backups.

Logging

Use the Logging property page to specify activity log preferences.

Log File Name

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

Enable pruning

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

Number of days to keep old entries

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

Prune now

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

Regional

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

Regional and Language options

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

Language

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

Date Format

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

Time Format

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

Number Format

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

Match MMC Language

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

VSS Options

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

Default Backup Destination

Select the default storage location for your backups.

Tip: You must have the Tivoli Storage FlashCopy Manager license to use the Tivoli Storage Manager software. If you have only the Data Protection license, only the Tivoli Storage Manager option is enabled. You can select from the following storage locations:

TSM The backup is only stored on Tivoli Storage Manager server storage. For SQL Server, Tivoli Storage Manager server is the default backup destination.

Local The backup is only stored on local disk.

Both The backup is stored on both Tivoli Storage Manager server storage and local disk.

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

Local DSMAGENT Node name

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

Remote DSMAGENT Node name

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

Import VSS snapshots only when needed

By default, local persistent VSS snapshots are automatically imported to the Windows system where the snapshots are created. If you select this option, the VSS snapshots are imported to the local host when needed to perform FlashCopy Manager operations.

To automatically import local persistent snapshots to the Windows system where the snapshots are created, the check box should be cleared.

If you intend to keep many backup versions (for example, more than 100 backup versions), or, if there is a limit to the number of LUNs that your server can use (for example, in virtual machine environments), select this option.

When you select this option, if the VSS hardware provider does not support transportable snapshots, or, if no hardware provider is available, the backup is completed, but the VSS snapshot is imported and is not transportable.

If you are running in a VMware environment and want to use VMware vMotion, ensure the LUNs are properly zoned to the ESX hosts and select this option. By choosing to import VSS snapshots when the snapshots are needed, the import process maps the VSS snapshot to the ESX host where the Windows virtual machine is running.

Custom Settings

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

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

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

Automatic and manual refresh modes differ in the following ways:

- In automatic refresh mode, a view automatically refreshes the first time that you select it. If there are thousands or millions of objects on the server, the refresh can take a long time to complete.
- In manual refresh mode, the view is not automatically refreshed. A name filter is available on the **Refresh Options** toolbar that you can use to narrow down the number of objects selected. After you enter a name pattern, you can click **Refresh**. By using manual refresh mode and limiting your query by using filters, you can reduce the amount of information that is returned from the server. Reducing the amount of information that is returned from the server can improve query and restore performance.

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

AlwaysOn Node

All availability databases in an availability group are backed up under this AlwaysOn node.

When you configure Data Protection for SQL Server with SQL Server 2012 and later versions, the AlwaysOn node name is a required parameter. The AlwaysOn node name can be changed at any time. To change the parameter, use the configuration wizard. From the AlwaysOn Node property page, view the parameter that is set.

AlwaysOn node name

The AlwaysOn node name that is used to back up availability databases is displayed. The databases that are not in an availability group are backed up under the standard Data Protection for SQL Server node name unless you select the **Use AlwaysOn node name for all databases** check box.

You cannot change the node name from this property page. To change the node name, use the configuration wizard.

Use AlwaysOn node name for all databases

Select this check box to specify that the AlwaysOn node is the default node for backing up all availability and standard databases. This option can be used to change your database backups from the standard Data Protection for SQL Server node to an AlwaysOn node. By selecting this check box, you can back up all standard and availability databases under a single node to help you manage your database backups more easily.

AlwaysOn Preferences

Use this property page to configure the backup preference settings for scheduled backups of availability groups and availability replicas in an SQL Server environment.

The settings are intended for scheduling backups of availability groups and availability replicas, and not for interactive backups with Microsoft Management Console (MMC).

The following settings are available in this property page:

Availability group

Select an AlwaysOn Availability Group for which you want to set schedule backup preferences.

Primary replica

Displays the primary replica for the selected availability group.

Preferred replica

Backup commands are scheduled on all nodes in an availability group. When the schedules run, the backup occurs only on the preferred replica. Other replicas receive, at run time, a warning message. The backup is skipped. You can make the following selections:

- 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.

Availability replicas

For each availability replica in this 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.

Availability databases

Displays the availability databases that are in an availability replica. The synchronization state is also displayed.

Performance

Use this property page to set preferences that affect performance for legacy backups.

DP Buffers

Specifies a number in the range 2 - 8 that specifies the number of communication data buffers that Data Protection for SQL Server uses when it transfers data to the Tivoli Storage Manager server. Each buffer has the size that is specified by the **DP Buffer Size** option. This option applies only to legacy backups.

DP Buffer size

Specifies the size of the buffers that are used by Data Protection for SQL Server to transfer data to the Tivoli Storage Manager server. This option applies only to legacy backups. Specify a value in the range 64 - 8192.

SQL Buffers

Specifies the number of communication data buffers that Data Protection for SQL Server uses when it transfers data between the SQL Server and Data Protection for SQL Server. Each buffer has the size that is specified in the **SQL Buffer Size** option. This option applies to legacy backups only. Specify a value from 0 to 999.

SQL Buffer size

Specifies the size of the buffers that are used by Data Protection for SQL Server to transfer data from the SQL Server to Data Protection for SQL Server. This option applies only to legacy backups. Specify a value in the range 64 - 4096.

Stripes

Specifies the number of data stripes to use in a legacy backup or legacy restore operation. Specify a value in the range 1 - 64. The default value is 1. This option applies to legacy backup and restore operations only.

When using a multiple stripes number for legacy backups, and setting the **Verify Only** parameter to **Yes** to restore the legacy backup, the number of stripes for legacy restore should be equal or greater than the number of stripes for the legacy backup.

Configuring VSS operations in an SQL Server failover cluster environment

You can configure Data Protection for SQL Server to protect VSS operations in a clustered SQL Server environment.

About this task

The following procedure applies if you are configuring Data Protection for SQL Server with Tivoli Storage Manager in a clustered SQL Server environment. The **VSSALTSTAGINGDIR** path must point to an existing directory on a shared disk or cluster shared volume that all cluster nodes can access.

Procedure

- For each node in your cluster environment, specify the **VSSALTSTAGINGDIR** directory path only in the VSS Requestor options file, `\baclient\dsm.opt`.

```
# Tivoli Storage Manager backup-archive dsm.opt file
NODename variable
PASSWORDAccess generate
TCPServeraddress variable.domain.com
TCPPort variable
CLUSTERnode no
CLUSTERDISKOnly no
VSSALTSTAGINGDIR C:\ClusterStorage\Volume1\vss_staging
```

Specify the **VSSALTSTAGINGDIR** parameter only if the following is true of your cluster environment:

—

- Tivoli Storage Manager completes the VSS operations.
- VSS backups are stored on local shadow volumes.
- At least 200 megabytes of free disk space is available on the drive that the **VSSALTSTAGINGDIR** parameter specifies. This disk space contains the metadata files for Data Protection for SQL Server.
- In each of the `dsm.opt` files that are used for local DSMAGENT and remote DSMAGENT components, ensure that the options for the **CLUSTERnode** and **CLUSTERDISKOnly** parameters are set to no.

```
CLUSTERNODE no
CLUSTERDISKONLY no
```

- In the Data Protection for SQL Server options file, `tdpsql\dsm.opt`, ensure that the option for the **CLUSTERnode** parameter is set to yes.

```
# Data Protection for SQL dsm.opt file
NODename variable
PASSWORDAccess generate
TCPServeraddress variable.domain.com
TCPPort variable
HTTPport variable
CLUSTERnode yes
```

Configuring Tivoli Storage Manager policy to set automatic expiration and version control (VSS and legacy backups)

With Data Protection for SQL Server, you can use Tivoli Storage Manager policy to set automatic expiration and version control.

About this task

You can set automatic policy for data backups by editing the Data Protection for SQL Server options file, or by specifying the policy statements in Microsoft Management Console (MMC). If you edit the options file, use INCLUDE and EXCLUDE statements to define the files that you want to automatically process, and to assign specific management classes to files by using object naming conventions.

Setting automatic expiration for VSS backups

Procedure

1. Complete one of the following steps:
 - Specify the VSSPOLICY statement in your Data Protection for SQL Server configuration file
 - Specify the VSSPOLICY statement in MMC (**Properties > VSS Policy Binding**).

You cannot specify VSSPOLICY statements by issuing the **tdpsqlc set** command.

2. Specify the following information in the VSSPOLICY statement:

VSSPOLICY *srvname dbname backuptype backupdest mgmtcls*

where:

srvname

Specify the name of the SQL Server or wildcard character (*)

dbname

Specify the name of the database or wildcard character (*)

backuptype

Specify the backup type as FULL or as a wildcard character (*). When you specify a wildcard character for *backuptype*, a FULL backup type is completed because you can run only full backup types.

backupdest

Specify the backup destination to be the Tivoli Storage Manager server, a local serve, or a wildcard character (*)

mgmtcls

Specify the management class name for the backup type.

Setting automatic expiration for legacy backups

About this task

Metadata is stored as a data object on the Tivoli Storage Manager server. You can migrate metadata to removable media if that action complies with Tivoli Storage Manager policy. A Data Protection for SQL Server backup object name comprises a series of qualifiers separated by \.

Procedure

1. Ensure that metadata is available to query without causing a volume mount.
2. Adhere to the INCLUDE and EXCLUDE syntax for object naming as follows:

```
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 | copyfull | diff | log | group | file | set

The following are examples of individual *objectNameSpecifications* in INCLUDE and EXCLUDE statements:

SqlServerNames:

SQL2008, SQL2012, SQL2014

InstanceNames:

INST1, INST2

DatabaseNames:

Db1, Db2, Db3

GroupNames:

g1, g2, g3

LogicalFileNames:

f1, f2, f3

Configuring Data Protection for SQL Server by using the Tivoli Storage Manager Configuration Wizard

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 to use VSS operations, the Tivoli Storage Manager backup-archive client, serving as the VSS Requestor, must also be installed and configured.

Before you begin

When you are remotely configuring Data Protection for SQL Server, you must first install Tivoli Storage FlashCopy Manager for Windows on the Data Protection node server, as shown in the example that is used in the procedure. You must then run the Tivoli Storage Manager Configuration Wizard, on at least one occasion, on the Data Protection node server.

Procedure

1. Start Microsoft Management Console (MMC) by clicking **Start > All Programs > Tivoli Storage Manager > Data Protection for Microsoft SQL Server > DP for SQL Management Console**.
2. From the start page, click **Configuration**. Alternatively, from the navigation tree, 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, which exist on the same system, to use for the applications that you want to protect.
 - In the **VSS Requestor** field, enter the node name that communicates with the VSS Service to access the SQL Server data. This node name is the Tivoli Storage Manager VSS requestor 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 node name is used to store the Data Protection for SQL Server backups.
 - If you are configuring Data Protection for SQL Server for SQL Server 2012 and later versions, enter a node name in the **AlwaysOn Node** field. This node name 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 to distinguish the type of backup that runs. For example, if your host name is *MALTA*, you can name the VSS requestor node name *MALTA*, and you can create a Data Protection node name that is called *MALTA_SQL*. For an SQL Server 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 your Tivoli Storage Manager administrator.
 - Specify whether to have the wizard 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 and 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 fields and perform the following actions:

- 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, the VSS Requestor is configured in addition to the applications that you selected. The client and agent services are also registered and configured, and a schedule to support historical managed capacity is defined.
 - f. Review the results of the configuration process. Click **Show Details** to view a list of individual configuration results.
5. Click **Finish** in the Completion page.
 6. Optional: For a VSS configuration, verify that the **Run VSS diagnostics when this wizard exits** option is selected. When this option is selected, a diagnostic process tests the VSS snapshots on your system after you complete the wizard.
- Attention:** If the configuration is for space-efficient target volumes for 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

Verify the configuration.

Related tasks:

“Verifying the configuration of Data Protection for SQL Server” on page 71

Configuring a remote system with a TSM Configuration

You can use a configuration wizard to configure a remote system to communicate with a Tivoli Storage Manager server.

Before you begin

Your system must run Microsoft Windows 2008 or later versions, PowerShell 3.0 or later, and Data Protection for SQL Server. On Windows 2012 and later versions, PowerShell version 4.0 is installed by default. For information about downloading, installing, and enabling Windows PowerShell, see this web page: Microsoft Windows Management Framework 3.0 Downloads (<http://www.microsoft.com/en-us/download/details.aspx?id=34595>)

Procedure

To configure a remote system with a TSM Configuration, complete the following steps:

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

For system with an AlwaysOn Availability Group, the corresponding AlwaysOn node must be detected.

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

14. Click **Finish** to complete the wizard.

What to do next

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

1. In the navigation tree, for the remote system, expand **Protect and Recover** and click on the application that is configured.
2. Open the Properties and click **Server Information**. Verify that the correct information is displayed.
3. Query the components and verify that a successful backup can be completed.

Manually configuring Data Protection for SQL Server

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

Configuring the computer that runs the SQL Server

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

Before you begin

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

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

Procedure

1. Specify your Data Protection for SQL Server node name and communication method in the `dsm.opt` file that is located by default in the Data Protection for SQL Server installation directory.
2. Using the **set** command, specify your Data Protection for SQL Server preferences (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 with SQL Server 2012 and later versions, 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 and later versions: 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 change 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 about installing the Tivoli Storage Manager backup-archive client for Windows, see Installing the Tivoli Storage Manager backup-archive clients (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.6/client/c_inst_baunix.html).

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) Install Tivoli Storage FlashCopy Manager if you want to manage local persistent VSS snapshots, which are created for VSS backups to LOCAL, VSS Instant Restores, and you want to run offloaded backups.
10. (VSS only) Install and configure a VSS provider. Consult the VSS provider documentation for information about configuration of that software. You do not need to install and configure any components if you use the default Windows VSS System Provider.
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 store all copies of the VSS backups as designated by your policies. Provisioning storage space to manage VSS snapshots depends on the VSS provider that you use. Consult the VSS Provider documentation for more details.

Related tasks:

“Specifying configuration parameters for Tivoli Storage Manager” on page 47

Configuring the Tivoli Storage Manager server

Perform these steps on the Tivoli Storage Manager server

Before you begin

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

Procedure

1. Define the policy domains, policy sets, management classes, copy groups, and storage pools. Choose what is necessary to meet your Data Protection for 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 by issuing 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 VSS requestor 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 VSS requestor node name and password for the system. For example, the agent node is the Remote DSMAGENT Node. *BAOFF* is used here (and in Step 5) to differentiate between this Remote DSMAGENT Node and the Local DSMAGENT Node (Step 3). You can replace *BAOFF* with the node name of your backup-archive client, and remove the *BAOFF* from the **grant proxynode** command.
5. (VSS only) Define the proxy node relationship (for the target node and agent nodes) by issuing the Tivoli Storage Manager **grant proxynode** command. For example:

```
grant proxynode target=alwayson node name agent=BAnodename
```

What to do next

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

Configuring the system that runs the offloaded backups

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

Procedure

1. Configure the 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 these services are not already installed. If a client CAD Service is already installed and configured, you can use an existing one. Use the backup-archive client Setup Wizard to guide you through the CAD installation process by selecting **Utilities > Setup Wizard > Help me configure the TSM Web Client**.
 3. Install and configure a VSS provider if you do not use the default system VSS provider. Consult the VSS provider documentation for information about the configuration of that software.

Verifying the configuration of Data Protection for SQL Server

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

Verifying the server configuration from the integrated command line

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

Procedure

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

```
query tdp
query tsm
query sql
```
4. With the cursor in the command-line panel, press Enter to run the commands to verify your configuration.

Verifying the SQL Server is ready to start VSS operations

Complete the following tests to verify that your SQL Server is ready to perform VSS operations.

Before you begin

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

About this task

The following list identifies the diskshadow tests to complete before any Tivoli Storage Manager components are installed.

Procedure

1. Test non-persistent shadow copy creation and deletion. Run diskshadow in a command window and issue 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 on drive F and drive G represent the SQL Server database and log volumes. Repeat this test four times, and verify that the Windows System and Application Event Log contains no errors.

2. Test persistent shadow copy creation and deletion. Run diskshadow in a command window and issue 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 on drive F and drive G represent the SQL Server database and log volumes. Repeat this test four times, verify that the Windows System and Application Event Log contains no errors.

3. Test non-persistent transportable shadow copy creation and deletion. Run diskshadow in a command window and issue the following commands:

```
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
```

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

What to do next

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

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

Common errors returned from VSS backup operations

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

The following two errors are commonly returned:

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

This message is displayed when the Tivoli Storage Manager backup-archive client CAD is either not running or is not configured properly.

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

This message is displayed when the Tivoli Storage Manager server is not configured correctly for the proxy nodes.

Manually configuring Data Protection for SQL Server on Windows Server Core

You can manually configure Data Protection for SQL Server to protect your SQL Server 2012 and later versions on Windows Server Core.

Before you begin

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.

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

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, 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\dsmadm
```

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

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
TCPSEVERADDRESS 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 that 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.

TCPSEVERADDRESS 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
TCPSeveraddress gijoe
TCPPort 1500
```

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 VSS requestor node. The default service name is "TSM Remote Client Agent". The value for the **/partnertype** 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
/partnertype:"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 on Windows Server Core

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

Before you begin

Restriction: You cannot use the following special characters in SQL Server database names on Data Protection for SQL Server:

- Question mark (?)
- Multibyte character (,)
- Multibyte character (^)
- Asterisk (*) character
- Colon (:) character cannot be used with Data Protection for SQL Server version 7.1.0 or earlier versions
- Backslash character(\) cannot be used with Data Protection for SQL Server version 7.1.0 or earlier versions

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
TCPServeraddress  tsm_server
TCPPort           1500
TCPWindowSize     63
TCPBuffSize       32
```

Where **nodename** is the Data Protection for SQL Server node name, and **TCPServeraddress** 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
```

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   [Yes|No]
ENABLEREPlacementchars [Yes|No]
LOGFile           tdpsql.log
```

Descriptions of the key options in the `tdpsql.cfg` file follows:

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 backups, 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 that is used to back up availability databases in an AlwaysOn Availability Group.

USEALWAYSONnode

Specify *Yes* 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 *No* to back up standard databases to the Data Protection for SQL Server node. Availability databases are always backed up with the AlwaysOn node.

ENABLEREPlacementchars

Specify *Yes* to enable Data Protection for SQL Server to process backslash (\) or colon (:) characters in a database name, and back up the database to Tivoli Storage Manager.

Specify *No* to prevent database backups to Tivoli Storage Manager if a user-defined string is substituted for a backslash (\) or colon (:) character in the database name.

Restriction: The **ENABLEREPlacementchars** parameter applies only to Data Protection for SQL Server version 7.1.1 and later versions. The maximum length of the database name is 128 characters.

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`.

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

To configure preferences for Data Protection for SQL Server, use the **set** command at 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`.

Procedure

At the command prompt, enter the following command:

```
tdpsqlc set parameter=value [/configfile=filename]
```

where *parameter* is the parameter or option for which you want to change the value, 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.

Examples:

Task Set the preferred SQL Server in the `tdpsql.cfg` file.

```
Command: tdpsqlc set sqlserver=your_SQL_instance  
/configfile=tdpsql.cfg
```

```
Command: tdpsqlc set fromsqlserver=your_SQL_instance  
/configfile=tdpsql.cfg
```

Task Change the name of the activity log file to tdpsql.log.

```
Command: tdpsqlc set logfile=tdpsql.log
```

Configuring Data Protection for SQL Server with Tivoli Storage Manager in a clustered environment

You can configure Data Protection for SQL Server to communicate with Tivoli Storage Manager, and protect SQL Server workloads in a clustered environment.

About this task

By default, the VSS Requestor dsm.opt file is in the Tivoli Storage Manager backup-archive client installation directory, c:\Program Files\tivoli\tsm\baclient\dsm.opt.

The location of application-specific dsm.opt files depends on the application that is being protected, for example:

Microsoft SQL Server

```
c:\Program Files\tivoli\tsm\TDPSQL
```

Procedure

1. Start Microsoft Management Console (MMC) 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, go 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. In the configuration wizard, select the application that you want to protect, for example, the SQL Server.
5. On the Requirements Check page, locate the main panel. For the **Cluster check**, click **Warnings**.
6. In the Issue Resolution dialog, specify the path to the VSSALTSTAGINGDIR directory that all cluster nodes can access. The VSSALTSTAGINGDIR path must point to an existing directory on a shared disk, for example, X:\vss_staging, or to a cluster shared volume, for example, C:\ClusterStorage\Volume1\vss_staging.
7. Click **OK** and complete the configuration wizard. In a clustered environment, the **Import VSS snapshots only when needed** configuration option is selected by default.
8. After you complete the configuration wizard, verify the VSSALTSTAGINGDIR path in the backup-archive client options file, baclient\dsm.opt. If the VSSALTSTAGINGDIR path is not correct, edit the file to specify the correct path. For example:

```
# Tivoli Storage Manager backup-archive dsm.opt file  
NODename variable  
PASSWORDAccess generate  
TCPServeraddress variable.domain.com
```



```

TCPPort variable
CLUSTERnode no
CLUSTERDISKOnly no
VSSALTSTAGINGDIR C:\ClusterStorage\Volume1\vss_staging

# Data Protection for SQL dsm.opt file
NODename variable
PASSWORDAccess generate
TCPServeraddress variable.domain.com
TCPPort variable
HTTPport variable
CLUSTERnode yes

```

For a Tivoli Storage Manager Server configuration, ensure that the VSSALTSTAGINGDIR directory path is specified only in the backup-archive client options file, baclient\dsm.opt.

For a stand-alone configuration, ensure that the VSSALTSTAGINGDIR directory path is specified in the backup-archive client options file, baclient\dsm.opt, and in the Data Protection for SQL Server options file, tdpsql\dsm.opt.

9. Repeat these steps on the other nodes in the cluster.

Configuring availability replicas to run scheduled data backups

When an availability database is replicated across multiple availability replicas in an availability group, a configuration option is available to enable you to select a single replica on which to run a backup operation instead of backing up all replicas.

About this task

Microsoft SQL Server 2012 and later versions provide a set of configuration options that you can use to specify whether scheduled backups are run on the primary or secondary availability replica. You can use the Data Protection for SQL Server GUI to set these options.

The configuration option can also be used to offload the backup from a primary replica to a secondary replica for load balancing. When databases fail over, backups must continue to run from other replicas to ensure that high availability is maintained.

Procedure

1. Start Microsoft Management Console (MMC).
2. In the Management section of the window, click **Protect Data** next to the SQL Server workload.
3. On the Actions pane, click **Properties**.
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 your 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 that is listed in the Availability replicas list box, specify whether it is a candidate to run 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.

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 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.

Before you begin

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

About this task

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

Configuring the Tivoli Storage Manager server

Procedure

1. Select or create the policy definitions that are used for each type of backup you plan to use. You can provide the administrator with the existing local-defined policy settings in your 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 VSS requestor node name and password for the workstation where the SQL Server is installed. For example:
`register node BAnodename BApasword`

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`

Configuring the workstation that runs 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 Server 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 are displayed.
 - b. Start 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 these conditions are true:

- You specified policy bindings during the use of Tivoli Storage FlashCopy Manager.
 - You are updating the policy bindings to reflect changes in your policy specifications for your Tivoli Storage Manager server usage.
3. In the Tivoli Storage Manager backup-archive client installation directory, make a copy of the VSS requestor options file named `dsm.opt`. Use the Windows **copy** command. For example:
`C:\Program Files\Tivoli\TSM\baclient>copy dsm.opt dsm_local.opt`
 4. In all of the files named `dsm.opt`, modify the `TCPSERVERADDRESS` line. Replace `FLASHCOPYMANAGER` with the IP address of the Tivoli Storage Manager server. For example:
`TCPServeraddress 9.52.170.67`

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

5. To access the Tivoli Storage FlashCopy Manager stand-alone environment during the transition period, open a Windows command prompt and change the directory to the Tivoli Storage Manager backup-archive client installation directory. This path is the default:
`C:\Program Files\Tivoli\TSM\baclient`

Create an alternative Windows service for the Tivoli Storage Manager Client Acceptor service by using the **dsmcutil** command. For example:

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

For more information about using the **dsmcutil** command, see the dsmcutil command (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.2/com.ibm.itsm.client.doc/c_cfg_dsmcutil_usewin.html).

6. Create an alternate Windows service for the Tivoli Storage Manager remote agent service. For example:


```
dsmcutil install cad /name:tsmcad4remote
/node:my_backup-archive_client_node
/password:my_TSM_server_password
/optfile:"C:\Program Files\Tivoli\TSM\baclient\dsm_remote.opt"
/httpport:1583
```
7. Edit the dsm_local.opt file in the Data Protection for 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 action through the Windows Services GUI or by using the **dsmcutil** command:


```
dsmcutil stop /name:"TSM Remote Client Agent"
dsmcutil stop /name:"TSM Client Acceptor"
dsmcutil start /name:"TSM Client Acceptor"
```
10. As backups start occurring and are managed in the Tivoli Storage Manager server environment, phase out the remaining backups that are created in the Tivoli Storage FlashCopy Manager stand-alone environment. You can choose between two ways of achieving the phase-out:
 - a. In the Tivoli Storage FlashCopy Manager stand-alone environment, define a time-based policy that automatically causes the old backups to expire and delete. For example, if you want to expire each backup after it is 30 days old, update the time-based policy by using the command:


```
tdpsqlc update policy mypolicy /daysretain=30
/tsmoptfile=dsm_local.opt
/configfile=tdpsql_local.cfg
```

You can also change these parameters by using the Local Policy Management dialog that is accessed from the **Utilities** menu of the Data Protection for SQL Server Backup/Restore GUI. Information about how to start the GUI is in the section that describes how to access the Tivoli Storage FlashCopy Manager stand-alone environment.

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

```
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 to run weekly to cause the expiration of out-of-date backups.

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

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

```
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:

```
tdpsqlc /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 access to the Tivoli Storage FlashCopy Manager stand-alone environment, you can remove the alternate services. To remove the services, use the Windows Services GUI or the **dsmcutil** command:

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

Transitioning standard SQL Server databases to the AlwaysOn node

You can back up standard SQL Server 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

The AlwaysOn node name is required when you configure Data Protection for SQL Server with SQL Server 2012 and later versions. 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.

Procedure

1. To back up your standard SQL Server databases to the file space for the AlwaysOn node, specify the **/USEALWAYSONnode** parameter with the **backup** command as follows.

For example, issue the following command

```
TDPSQLC Backup *|dbname[,dbname,...] Full /USEALWAYSONnode
```
2. To back up all databases, specify the wildcard character (*) or specify a list of database names that are separated by commas.

For example:

TDPSQLC Backup standard_db01,standard_db02 Full /USEALWAYSOnnode

3. To regularly back up standard SQL Server databases to the file space for the AlwaysOn node, specify the `/USEALWAYSOnnode` parameter with the `set` command.

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

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

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

SAN Volume Controller and Storwize V7000 settings

Create 14 space-efficient target volumes for each source volume to be protected. Enable the autoexpand option for the space-efficient target volumes. Add the space-efficient target volumes to the VSS_FREE pool.

VSS Provider settings

Set the background copy rate to 0.

Data Protection for SQL Server settings

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

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

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

SAN Volume Controller and Storwize V7000 settings

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

VSS Provider settings

Set the background copy rate to 0.

Data Protection for SQL Server settings

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

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

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

SAN Volume Controller and Storwize V7000 settings

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

VSS Provider settings

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

Data Protection for SQL Server settings

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

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

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

SAN Volume Controller and Storwize V7000 settings

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

VSS Provider settings

Use the default background copy rate of 50.

Data Protection for SQL Server settings

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

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

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

SAN Volume Controller and Storwize V7000 settings

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

VSS Provider settings

Use the default background copy rate of 0.

Data Protection for SQL Server settings

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

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

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

SAN Volume Controller and Storwize V7000 settings

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

VSS Provider settings

Use the default background copy rate of 50.

Data Protection for SQL Server settings

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

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

Chapter 4. Protecting data

You can back up and restore your Microsoft SQL Server data by using Microsoft Management Console (MMC) or the command-line interface.

About this task

If required, you can manage your installations remotely.

Tip: For information about protecting Microsoft SQL Server data in VMware environments, see the chapter that explains in-guest application protection in the *Data Protection for VMware User's guide*.

Prerequisites

With Data Protection for SQL Server, you can back up and restore SQL Server data and protect your SQL Server environment.

Security requirements for backup and restore operations

Data Protection for SQL Server requires certain settings to process backup and restore operations in a secure environment.

To install Data Protection for SQL Server, you must have Windows administrator authority. You must register Data Protection for SQL Server to the Tivoli Storage Manager server and you must use the appropriate node name and password when it connects to the Tivoli Storage Manager server.

You can specify SQL Server logon information in one of the following ways:

- Accept the default sa account and system administrator password. Ensure that you secure your sa login account with a 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.

You must add the SQL logon user name or Windows user name to the SQL Server SYSADMIN fixed server role before Data Protection for SQL Server can use those credentials.

VSS and legacy backup methods

Depending on your SQL Server environment, you can run full backups only, copy-only full backups, full plus log backups, full plus differential backups, or file and group backups. Your backup strategy might also be to back up data to Tivoli Storage Manager or local shadow volumes.

Full backup method (Legacy and VSS)

Use the full backup method for system databases such as *master*, *model*, and *msdb* because of their typical small size. A full backup can take a long time to run. However, the restore process is the most efficient because only the most recent full backup is restored.

Copy-only full backup method (Legacy and VSS)

Use the copy-only full backup method to periodically create copy-only full backups for long-term retention without affecting existing backup schedules or retention policies that you use for disaster recovery. Copy-only full backups do not affect the transaction logs or the sequence of backups, such as differential backups or full backups.

Full backup plus log backup method (Legacy and VSS)

Use the full backup plus log backup method when the normal backup schedule or network capacity cannot support a full backup.

To minimize the effect on the backup schedule and network traffic during peak times, you can run a periodic full backup, followed by a series of log backups. For example, you can schedule full backups on the weekend and incremental backups during the week. You can run full backups during low usage times and when increased network traffic can be tolerated.

Restriction: If you run multiple full backups, the SQL database log can become full. Subsequent backups might fail as a result. If necessary, use basic SQL Server tools to truncate the log of your SQL databases.

Full backup plus differential backup method (Legacy and VSS)

Use the full backup plus differential backup method if your backup schedule and network capacity can facilitate backing up all transaction logs that accumulate between full backup operations. This strategy requires that only one differential backup plus the last full backup be transferred to complete a restore operation. However, the same amount of data must be transferred in the differential image, as in the series of log backup operations.

Although you can run only VSS full backups, you can apply legacy differential backups to the VSS full backup.

Full backup plus differential plus log backup method (Legacy and VSS)

Use the full backup plus differential plus log backup method to reduce the number of transactions that must be restored and applied. Restore operations are faster as a result.

If, for example, you complete a full legacy or VSS backup weekly, and a differential backup nightly, and a log backup every four hours, the restore processing would include the full backup, a differential backup, and at most five log backups. However, if you only complete a full plus log backup scheme on the same cycle, the restore processing would include a full backup plus up to 41 log backups (six days multiplied by six log backups per day plus up to five backups on the day the full backup is completed).

Although you can run only VSS full backups, you can apply legacy log backups and legacy differential backups to the VSS full backup.

File or group backup method (only Legacy)

Use the file or group backup method when it is impractical to back up an entire database because of the size of the data, or associated time and performance issues.

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

File or group options can save backup and restore processing time when certain tables or indexes have more updates than others and must be backed up more often. It is time-effective to contain such data in their own file group or files, and to back up only those items.

Except for logical log files, you can back up your transaction logs after you back up a data file or file group.

Tivoli Storage Manager backups versus local shadow volumes backups

When you create a policy for your backups, you must choose whether to back up data to Tivoli Storage Manager storage versus VSS disks. Data backups to Tivoli Storage Manager typically takes longer to process than backups to local shadow volumes.

Backing up SQL Server data to Tivoli Storage Manager is necessary when long-term storage is required. For example, saving SQL Server data on tape for archival purposes requires 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.

Restriction: If you run legacy log backups to a Tivoli Storage Manager server, the SQL database log files can be truncated.

Local shadow volumes

When you back up data to local shadow volumes, ensure that sufficient local storage space is assigned to the local shadow volumes. Create different sets of policies for backups to both local shadow volumes and to Tivoli Storage Manager server storage. If you use a VSS provider other than the Windows VSS System Provider, follow the backup recommendations of the VSS provider.

Environment and storage resources also impact how many backup versions you can maintain on local shadow volumes for VSS fast restore and VSS instant restore operations, and on Tivoli Storage Manager server for VSS restore operations.

Combined VSS and legacy backups

After you restore a full VSS backup, you can use VSS backups and legacy backups together as a highly-effective backup solution for Data Protection for SQL Server data.

Table 8. Backup strategy characteristics

Strategy characteristics	Legacy backup only	Legacy backup plus VSS backup
Available backup types	<ul style="list-style-type: none"> • FULL 1+ per week • DIFF 1+ per day • LOG 1+ per day 	<ul style="list-style-type: none"> • 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	<ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • FULL • Database level restore granularity • Point-in-time recovery² • Roll-forward recovery² Legacy: <ul style="list-style-type: none"> • 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, the backups that are applied to a full VSS backup must be legacy log backups or legacy differential backups.

Enhanced data protection for Microsoft SQL AlwaysOn Availability Groups

You can run VSS (full) and legacy (full, differential, file/set/group, and log) backup operations on a primary replica. You can run copy-only VSS and legacy backup operations, and normal log backups on a secondary replica. You cannot run a differential backup on a secondary replica.

For backups on a secondary replica, the replica must be in the synchronized or synchronizing state. You can have multiple AlwaysOn Availability Groups (AAGs) in a SQL Server cluster. You can also have a mix of standard databases and AAGs on a SQL Server cluster.

When you back up data, you can distribute the backup workload for scalability and isolate backup activity to a dedicated backup node. When you isolate backup activity, it minimizes the effect on production databases.

Given that replicas are copies of the same database, avoid redundant backups of the same databases. Apply retention policies to unique databases.

As a best practice, allow backups from any node in the availability group and enable restore operations from any node in the availability group.

Best practices for backing up data in an AAG

When you use Tivoli Storage FlashCopy Manager for SQL Server to manage AAG backups, consider the following backup options:

Backup priority

Specified per database in an AAG, the backup priority option defines the order in which replicas are used to back up a database in an AAG.

Preferred replica

Specified at an AAG level, the preferred replica option defines whether primary or secondary replicas can be used for backup operations.

- Prefer secondary replica: Scheduled backups occur on a secondary replica, if available. If the secondary replica is not available, you can use the primary replica.
- Secondary only: Scheduled backups can occur only on a secondary replica.
- Primary: Scheduled backups can occur only on the primary replica.
- Any replica: Scheduled backups can occur on any replica.

/USEALWAYSONNode parameter

A parameter option on the **backup** command that provides a common namespace for all backups. Each node authenticates separately with Tivoli Storage Manager. Backed up data is stored in the AlwaysOnNode namespace by using the Asnode option.

/ALWAYSONPriority parameter

A parameter option on the **backup** command that specifies that a local availability database is backed up only if it has the highest backup priority among the availability replicas that are working properly. This parameter applies only to scheduled backups.

Typical data protection deployments in AAG environments

You can back up data in an AAG in the following ways:

- Distribute a legacy backup across AAG replicas
- Distribute a VSS backup across AAG replicas

Scenario: Legacy backups are distributed across AAG replicas

When you configure your environment to distribute a legacy backup across AAG replicas, follow these steps:

1. Set the preferred replica to **Prefer secondary replica**.
2. Install Tivoli Storage FlashCopy Manager for SQL Server on all replicas that are eligible to run a backup.
3. Create a command script to run a .CMD file with a **backup** command similar to the following example:

```
tdpsqlc backup db1,db2,db3 full /alwaysonpriority
```
4. Associate each Tivoli Storage FlashCopy Manager for SQL Server node with the defined schedule.
5. Run backups on the SQL node according to defined priorities for each database.

Scenario: VSS backups are distributed across AAG replicas

When you configure your environment to distribute a VSS backup across AAG replicas, follow these steps:

1. Set the preferred replica to **Prefer secondary replica**.
2. Install Tivoli Storage FlashCopy Manager for SQL Server on all replicas that are eligible to run a backup.
3. Create a command script to run a .CMD file with a separate **backup** command per database similar to the following sample

```
tdpsqlc backup db1 full /alwaysonpriority /backupmethod=VSS
backupdest=TSM
tdpsqlc backup db2 full /alwaysonpriority /backupmethod=VSS
backupdest=TSM
tdpsqlc backup db3 full /alwaysonpriority /backupmethod=VSS
backupdest=TSM
```
4. Associate each Tivoli Storage FlashCopy Manager for SQL Server node with the defined schedule.
5. Run backups on the SQL node according to defined priorities for each database.

Starting Microsoft Management Console

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

Before you begin

If you try to use Data Protection for SQL Server before you complete the configuration process, the software does not function correctly.

About this task

Data Protection for SQL Server software is displayed in MMC as a plug-in. MMC uses a navigation tree to organize the computer data that is registered. Each computer icon that is followed by the word *Dashboard* represents a physical computer.

When you register a computer, information about the computer is collected and stored. Password information is encrypted and stored separately. The computers that are registered are tracked with a globally unique identifier (GUID). The GUID is used when you back up and restore data.

You can create groups of computers. These groups consolidate information when you view the dashboard, prepare reports, and run group commands. By default, the computers in a group are selected when you complete tasks for the group, but you can select more computers in the tree to include in an operation.

Procedure

To start MMC, click **Start > All Programs > Data Protection for Microsoft SQL Server > DP for SQL Management Console**.

Starting the Data Protection for SQL Server command-line interface

You can start the Data Protection for SQL Server command-line interface by using a Windows command prompt with administrative privileges. Alternatively, you can start the command-line interface from Microsoft Management Console (MMC).

Procedure

1. Start MMC.
2. In the navigation tree, select the computer node where you want to run the commands.
3. Expand the **Protect and Recover Data** node.
4. In the navigation tree, select an 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. From the drop-down list, change **PowerShell** to **Command Line**.

Getting help for Data Protection for SQL Server commands

By issuing the **help** command at the command prompt, you can display a complete list of Data Protection for SQL Server commands, and associated parameters.

Procedure

At the command prompt, issue **tdpsqlc ?command_name** Where *command_name* is the name of the Data Protection for SQL Server command.

For example:

```
tdpsqlc ? restore full
```

Managing Data Protection for SQL Server installations remotely

From a single Data Protection for SQL Server installation you can manage all of the Data Protection for SQL Server installations in your enterprise.

Before you begin

Your system must run Microsoft Windows 2008 or later versions, PowerShell 3.0 or later, and Data Protection for SQL Server. On Windows 2012 and later versions, PowerShell version 4.0 is installed by default. For information about downloading, installing, and enabling Windows PowerShell, see this web page: Microsoft Windows Management Framework 3.0 Downloads (<http://www.microsoft.com/en-us/download/details.aspx?id=34595>)

Procedure

1. Enable remote management for Data Protection for SQL Server installations by issuing the following command:

```
Enable-PSRemoting -force
```
2. Add the Data Protection for SQL Server servers to the trusted hosts list by issuing the following command:

```
Set-Item WSMAN:\localhost\Client\TrustedHosts -Value remote_server_name -Force
```
3. After you make configuration changes, restart the winrm service by issuing the following command.

Adding remote systems

You can add remote systems in MMC.

Procedure

1. From Microsoft Management Console (MMC), in the Actions pane, click **Manage Computers**.
2. Verify that the local system is listed in both the Tree Nodes and Computers panes.
3. From the Tree Nodes pane, click the add icon. The icon is green and resembles the symbol for addition.
4. Type the name and description for the new tree node.
5. From the Computers pane, click the add icon. The computers that you add are associated with the tree node that you are creating. If you add only one computer, the tree node type can be either **Dashboard** or **Group**. If you add more than one computer, the tree node type is **Group**. If you add only one computer, from the Tree Nodes pane, you can toggle between the **Dashboard** and **Group** types.
6. Type the system name and a description. For systems that are not in the domain, provide the fully qualified address. Alternatively, to select a system that is based on another system in the domain or to read a list of computers from a file, on the Computers pane, click **Import**. Clicking **Import** displays a dialog called Add Computers. From the Add Computers dialog, there are two tabs: Active Directory and Import. To complete the Add Computers dialog window entries, complete the following steps:
 - a. For the Active Directory tab, complete these fields

Domain

The current domain is displayed. The domain cannot be changed.

Location

The organizational unit that is used to search for computers. The default value is displayed.

Name By default, the wildcard character (*) is displayed. You can leave the default value or enter a specific name.

Account

The current account is displayed. If you want to use a different account, click **Search** to search the domain for other computers. The Search capability is enabled only when the **Location** and **Name** fields have values.

- b. For the Import tab, browse to find a comma-separated values (.CSV) file that contains computer entries. After you find a .CSV file and click **Import**, the contents of the .CSV file are read as entries and are added to the list. The following .CSV file is an example of a valid .CSV file for the import activity:

```
NewNode1,Group1,CurrentUser,Test node 1
NewNode2
NewNode3,,Description of NewNode3
NewNode4,Group2,CurrentUser,Test node 4
```

The first column (the node name) is required. The other data columns are optional. The list is processed by position. For the group, if a group does not exist, the group is created.

7. From the Computers pane, click **Test Connection**. The test status is reported in the Message column of the Computers table.
8. Click **OK** to close the Manage Computers window.
9. Verify that the tree node is listed in the navigation tree. The remote node does not have all of the functionality available for local systems. For example, entries for learning, online support, and favorite links are not displayed.
For tree node type **Dashboard**, the main window displays the Protect, Recover, and Automate tabs. For tree node type **Group**, the main window displays the Group Dashboard, Group Reports, and Group Commands tabs.
10. After you add systems, you can remove (delete) the systems. You can also select the system to edit the properties, including tree node type, that you entered when you added the system. If you want to change the order of the systems that are displayed in the navigation tree, use the GUI controls from the Manage Computers window.

Determining managed storage capacity

You can track the capacity of managed storage assets. This information can be useful when you are calculating storage requirements for license renewal.

About this task

Typically, the capacity that is used by server data differs from the capacity of the volume that contains that data. For example, a set of databases might require a capacity of 1 GB and be on a 10 GB volume. When a snapshot of the volume is created, the Data Protection for SQL Server managed capacity measurement is 10 GB.

Procedure

1. From Microsoft Management Console (MMC), select an SQL Server instance.
2. On the **Protect** tab, click **Properties** in the Action pane.
3. Select **Managed Capacity** from the list of available property pages. The managed capacity is calculated and displayed.
4. To view a list of the volumes that contain backups and their respective managed capacities, click **Show Details**.

Managing backup, restore, and automation tasks in the Task Manager

The Task Manager provides you with a centralized view of Microsoft Management Console (MMC), from which you can view, stop, remove, or manage backup, restore, and automation tasks.

Procedure

1. Start MMC.
2. On the Welcome page of MMC, click the **Protect Data** or **Recover Data** task that is appropriate for your data requirements.
3. In the **Action** pane, click **Show Activity**. The Task Manager pane opens under the results pane.
4. Choose your preferred view for the current task:

Option	Description
Task List (default)	<p>Click this view to see the following information:</p> <p>Name State Result Progress Start Time Duration Messages</p> <p>Use the Task List view to complete these tasks:</p> <ul style="list-style-type: none"> • Click Up and Down to modify the order in which incomplete operations are processed. 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. 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	<p>Click this view to see the information that is available in the Task List in detailed format. Click Mode: Navigate and use the arrows to view operation details. Summary and error information is also available when applicable.</p>

Backing up SQL Server data

Windows

By using Microsoft Volume Shadow Copy Service (VSS), you can back up SQL Server data and mount the backup if required. You can also create legacy backups of SQL Server databases, groups, and files.

Enabling SQL Server backup compression

With Data Protection for SQL Server, you can compress SQL Server backups. Compression reduces traffic and storage requirements. SQL Server backup compression is generally faster and more effective than using Tivoli Storage Manager compression.

About this task

- You can run SQL Server backup compression only with legacy backups on SQL Server.
- For SQL Server, you can run backup compression only on Enterprise Edition.
- For SQL Server 2008 R2 and later versions, you can run backup compression on Standard, Enterprise, and Datacenter editions. Beginning with SQL Server 2008, any edition can restore a compressed backup.

Procedure

Use either Microsoft Management Console (MMC) or the command line to enable SQL Server backup compression:

- From the General properties page in MMC, specify the SQL Server native backup compression option. 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 issue this command:

```
SQLCOMPRESSION Yes | No
```

The default value is No.

Verifying the integrity of legacy databases by using the checksum option

With Data Protection for SQL Server, you can verify the integrity of legacy database backups by setting a checksum option.

About this task

A *checksum* is a value that is calculated and written in the data page header of the database data file. When a data file is read again, the checksum value is recalculated. Checksum processing validates the values in a file or configuration for unexpected changes. Values are verified between the current state and the baseline state.

Restriction: Checksum integrity checking is available only with legacy backups on SQL Server.

Procedure

1. Open the General Properties window in Microsoft Management Console (MMC).
2. Select **Compute SQL Server checksum for legacy backup**.

If you select this option, all legacy backups are checked by default. You can override this setting to set integrity checking for a particular backup. For example, if you bypassed integrity checking on all backups, you can set integrating checking on a particular legacy backup by selecting the **SQL**

Checksum backup option on the **Protect** tab for the SQL Server instance. You can also issue the SQLCHECKSUM option with the **backup** command on the command line to temporarily enable or disable the checksum option.

Results

When you select the **Compute SQL Server checksum for legacy backup** check box, the setting is written to the Data Protection for SQL Server preferences file, `tdpsql.cfg`, and is applied to all legacy backup operations. If you clear the check box, integrity checking does not apply to any legacy database backup.

Related tasks:

“Creating legacy backups of SQL Server databases” on page 102

Creating VSS backups of SQL Server databases

You can back up standard SQL Server databases or availability databases by using Microsoft Volume Shadow Copy Service (VSS).

Before you begin

To manage local VSS backups or to run offloaded backups to Tivoli Storage Manager server storage, ensure that Tivoli Storage FlashCopy Manager is configured in your environment.

If you use VSS to back up data to a Tivoli Storage Manager server, Tivoli Storage FlashCopy Manager is not required.

About this task

On SQL Server 2012 and later versions, you can back up availability databases in an AlwaysOn Availability Group (AAG) regardless of which availability replica is used for the backup operation.

Restriction: When you complete a full backup of a secondary replica in an AAG, only a copyfull backup of that database is created.

To back up availability databases, ensure that is configured to use an AlwaysOn node. Additionally, 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 Server 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

1. Start Microsoft Management Console (MMC).
2. If you plan to use offloaded backups, and your environment is configured for use with a Tivoli Storage Manager server, specify a value in the **Remote DSMAGENT Node name** field.
 - a. Select the **SQL Server** instance in the navigation tree, and click **Properties** in the Actions pane.

- b. Select the VSS Backup property page. If the **Remote DSMAGENT Node name** is blank, enter a node name.

An offloaded backup uses another system (specified with the **Remote DSMAGENT Node name** parameter) to move SQL Server data to Tivoli Storage Manager server storage. Offloaded backups can reduce the load on network, I/O, and processor resources during backup processing.

3. On the **Protect** tab of an SQL Server instance, select an option for viewing databases.

Table 9. Database backup views

Task	Action
View a list of SQL Server databases that are available for a backup operation	Click View: Databases .
View a list of SQL Server 2012 and later version availability databases that are available for a backup operation	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. Toggle the Standard Databases / Availability Databases button for the respective database views.

Refine the list of available databases in the results pane by entering a keyword in the **Search** field. Then, select the databases to back up.

4. Verify the backup options. If the backup options are not displayed, click **Show Backup Options**. If you want to use offloaded backups, select **Yes** in the **Offload** field.
5. In the Actions pane, click **Backup Method** and select **VSS**.
6. On the Actions pane, select **TSM** for the **Backup Destination**. The only option that is available to you is **TSM** because 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 runs the command that is required to complete the task.
8. To start the backup operation, in the Actions pane, take one of the following actions:
 - a. Click **Full Backup**. Alternatively, right-click a database and select the backup action that you require from the menu.
 - b. Click **Copy-Only Full Backup**. A copy-only full backup is independent of the sequence of SQL Server backups, and is not used as a base for a differential backup. A differential backup is not associated with the copy-full backup, but is associated with the previous full backup that completed. You might use a copy-only full backup as a special purpose backup that does not affect backup and restore operations, and retain such a backup for longer than conventional backups.
9. Review the status of the backup operation by clicking **Task List** in the results pane. Click **Task Details** to view detailed status information.

Results

During backup processing, Data Protection for SQL Server bypasses database snapshots and databases that are in offline, mirroring, and restoring states.

What to do next

To determine which databases backups are bypassed during backup processing, review the `tdpsql.log` in the directory where Data Protection for SQL Server is installed.

Creating legacy backups of SQL Server databases

You can create a legacy backup of your standard SQL Server databases by using Microsoft Management Console (MMC). You can also use the legacy method to back up availability databases with SQL Server 2012 and later versions.

Before you begin

- To run a legacy backup, ensure that the Data Protection for SQL Server license file is installed.
- On SQL Server 2012 and later versions, you can also back up availability databases in an AlwaysOn Availability Group (AAG) regardless of which availability replica is used for the backup operation. To back up availability databases, ensure that is configured to use an AlwaysOn node. Additionally, specify the AlwaysOn node in the **AlwaysOn Node** field in the TSM Node Names page of the Tivoli Storage Manager Configuration Wizard.

About this task

When you change log and set backups to a new management class, the new management class setting takes effect only for new backups. Existing backups are not rebound to the new management class. Therefore, schedule a new full backup after you change the management class for either log or set backups.

Procedure

1. Start MMC.
2. Select the **SQL Server** instance in the tree view.
3. On the **Protect** tab of an SQL Server instance, select an option for viewing databases.

Table 10. Database backup views

Task	Action
View a list of SQL Server databases that are available for a backup operation	Click View: Databases .
View a list of SQL Server 2012 and later version availability databases that are available for a backup operation	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. Toggle the Standard Databases / Availability Databases button for the respective database views.

Refine 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 displayed, click **Show Backup Options**.

Table 11. Database backup options

Option	Action
Data Stripes	<p>Use this option to specify the number of data stripes to use in a backup or restore operation.</p> <p>The <i>numstripes</i> variable can be in the range 1 - 64. The default value is 1.</p> <p>When you use a multiple stripes number for legacy backups, and set the Verify Only parameter to Yes to restore the legacy backup, the number of stripes for legacy restore must be equal or greater than the number of stripes for the legacy backup.</p>
Estimated Database % Change	<p>Use this option to specify the estimated percentage of the database that changed since its last full database backup. The default value is 20.</p> <p>This estimate is useful 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.</p>
Estimated Log % Change	<p>Use this option to specify the estimated percentage of an SQL Server database that changed due to non-logged operations since the last log backup. The default value is 0.</p>
Truncate Logs	<p>Use this option to specify whether to dispose of entries that you no longer need in the SQL Server database transaction log after you back up the log. The default value is Yes.</p> <p>In general, you do not want to truncate the log when you rebuild a corrupted database. This option enables the server to back up the transaction log but does not affect the data. All transaction log entries are written 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.</p>

Table 11. Database backup options (continued)

Option	Action
Back Up Tail-Log	<p>Use this option to store log records that are not backed up.</p> <p>By storing these records, also known as the <i>tail of the log</i>, the log chain is kept intact. Before you can recover an SQL Server database to the last 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.</p>
SQL Server Checksum	<p>Use this option to verify the integrity of a legacy database backup. Integrity checking is a process that validates the values in a file or configuration for unexpected changes. Values are verified between the current state and the baseline state.</p> <p>In the Performance Properties window of MMC, you can enable or disable the checksum option for all your legacy databases at once. You can override the global setting, and temporarily enable or disable the checksum option for a database backup, by setting this SQL Checksum option to Yes or No.</p>

5. In the Actions pane, click **Backup Method** and select **Legacy**.
6. In the Actions pane, select **TSM** for the **Backup Destination**. The only option that is available to you is **TSM** because 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 runs the command that is required to complete the task.
8. To start the backup operation, in the Actions pane, take one of the following actions:
 - **Full Backup**
 - **Copy-Only Full Backup**
 - **Differential Backup to TSM**
 - **Log Backup to TSM**
9. Review the status of the backup operation by clicking **Task List** in the results pane. Click **Task Details** to view detailed status information.

What to do next

- To determine which databases backups are bypassed during backup processing, review the `tdpsql.log` in the directory where Data Protection for SQL Server is installed. Data Protection for SQL Server bypasses database snapshots and databases that are in offline, mirroring, and restoring states.

- To determine whether the checksum option is applied to a legacy database backup, issue the **tdpsqlc query tsm *** command on the command-line interface, or the equivalent **Get-DpSqlBackup** cmdlet.

Related tasks:

“Verifying the integrity of legacy databases by using the checksum option” on page 99

Creating legacy backups of SQL Server groups or files

You can create a legacy backup of SQL Server groups or files in a normal SQL Server database by using Microsoft Management Console (MMC). You can also use the legacy method to back up groups of files in availability databases with SQL Server 2012 and later versions.

Before you begin

On SQL Server 2012 and later versions, you can back up availability databases in an AlwaysOn Availability Group regardless of which availability replica is used for the backup operation. To back up availability databases, ensure that Data Protection for SQL Server is configured to use an AlwaysOn node. Additionally, specify the AlwaysOn node in the **AlwaysOn Node** field in the TSM Node Names page of the Tivoli Storage Manager Configuration Wizard.

Procedure

1. Start MMC.
2. Select the **SQL Server** instance in the tree view.
3. On the **Protect** tab of an SQL Server instance, select an option for viewing databases.

Table 12. Database backup views

Task	Action
View a list of SQL Server databases that are available for a backup operation	Click View: Databases .
View a list of SQL Server 2012 and later version availability databases that are available for a backup operation	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. Toggle the Standard Databases / Availability Databases button for the respective database views.

Refine 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 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 be in the range 1 - 64. The default value is 1.

When using a multiple stripes number for legacy backups, and setting the **Verify Only** parameter to **Yes** to restore the legacy backup, the number of stripes for legacy restore should be equal or greater than the number of stripes for the legacy backup.

5. Optional: Choose a mode for the current task:
 - **Run Interactively:** Click this item to run the current task interactively. This selection is the default.
 - **Run Scheduled:** Click this item to convert the current action into a scheduled task. When you select this item, the schedule wizard runs the command that is required to complete the task.
6. To start the backup operation, in the Actions pane, take one of the following actions:
 - **Group Backup to TSM:** This option backs up the contents of the specified file group.
 - **File Backup to TSM:** This option backs up the contents of the specified logical data file.
 - **Set Backup to TSM:** This option backs up the contents of the specified groups and files.
7. After you complete a Group, File, or Set backup operation, back up the transaction logs. You cannot back up a logical log file.

What to do next

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

Backing up SQL Server databases on Windows Server Core

To back up Microsoft SQL Server 2012 and later version databases, use the **backup** command.

About this task

Use the following procedure to back up SQL Server databases to the Tivoli Storage Manager server, or to take local VSS snapshots.

Procedure

1. To back up all or part of an SQL Server database on Windows Server Core, enter the following command at the command prompt:
`tdpsqlc backup database_name backup_type [other_options]`

where *database_name* specifies the name of the database, and *backup_type* specifies the type of backup such as a full backup. You can specify other options, such as the back up method. For example, to create a full legacy backup of SQL Server databases DB_01 and DB_02, enter the following command:

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

For example, to create a full legacy backup of all databases on the SQL Server, enter the following command:

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

2. To back up a file group, enter the following command at the command prompt:
`tdpsqlc backup database_name file_group`

where *database_name* specifies the name of the database, and *file_group* specifies the file group in the database. For example, 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
```

Deleting SQL Server backups

You can remove an SQL Server backup that you created with the VSS backup method. Complete this task only if necessary.

Before you begin

Typically, backups are deleted automatically based on user-defined policy management settings. This procedure is necessary only if you need to delete backups that are outside the scope of Data Protection for SQL Server policy definitions.

If you delete a remotely-mounted backup, the snapshots and the relationship between the source and target volumes on the storage device are also deleted. However, the target volume that is imported and mounted might continue to exist. In addition, the target volume might not be available to the server where the remote mount occurred. The operations to the target volume depend on the VSS hardware provider and the storage device implementation.

After the maximum number of remotely-mounted backup versions or the maximum number of days to retain a backup is exceeded, the associated backup is expired and deleted.

Procedure

1. Start Microsoft Management Console (MMC).
2. Click **Recover Data > SQL** in the Management window.
3. On the **Recover** tab for the SQL Server instance, select **View: Database Restore**. In the results pane, browse to and select one or more database backups to delete. The corresponding node type, for example, DP or AlwaysOn, must also be selected.
4. In the Actions pane, click **Delete backup**. When a backup is deleted, two tasks display in the task window to show you that the deletion is in progress, and that the view is being refreshed.

Deactivating legacy backups of SQL Server databases

Tivoli Storage Manager deactivates an SQL Server database backup as a part of Tivoli Storage Manager policy management. Data backups are typically deactivated when an SQL Server database is deleted from the SQL Server as part of the scheduled backup processing.

Before you begin

The SQL Server database that you want to deactivate must be a legacy backup. You cannot use this procedure to deactivate VSS backups. The **Delete** action is available in the Actions pane when you select a VSS backup from the **Recover** view.

About this task

For legacy backups, you can deactivate any or all of the following backup object types: full, differential, copyfull, log, file, group, or set. You can also deactivate any object or object type that is older than a specified number of days.

When you deactivate database backups, any existing backups on Tivoli Storage Manager server are subject to be deleted, as specified by the **verdeleted** setting.

When automatic processing is insufficient, the **inactivate** function explicitly deactivates one or more active data backups on the Tivoli Storage Manager server.

Procedure

1. Under the **Protect and Recover Data** node in the tree view, select the SQL Server.
2. Open the **Recover** view to see the status of the backup. Active backups are displayed.
3. Select the database backup that you want to deactivate, and in the Actions pane, click **Inactivate**.
4. To view the results, take one of the following actions:
 - To display the database that you made inactive, click **All Backups** on the toolbar.
 - To display only active database backups, click **Active Backups** on the toolbar.

Mounting VSS snapshots to remote servers

You can use the command-line interface to mount VSS snapshots to remote servers that other users can access.

About this task

The following procedure is specific to Tivoli Storage FlashCopy Manager and assumes that you have at least three servers in your environment: Server A, Server B, and Server C. The backup that is created on Server A is mounted remotely to Server B.

To mount a backup remotely, the hardware provider must allow transportable snapshots. In addition, you must enable the **Import VSS snapshots only when needed** configuration option and remote Windows PowerShell.

When a backup is mounted remotely and the backup is deleted, the state of the mount point varies. The state of the mount point depends on the VSS hardware provider and storage device that is used. When a backup is mounted remotely, the backup can be deleted. When a local persistent VSS snapshot is created, a source and target volume relationship is created. The local persistent VSS snapshot is created on your storage device. In this case, when a remote mount operation occurs, the target volume is imported and mounted to the server that sends the request for the remote operation.

Procedure

1. On Server A, use the command-line interface to complete a local backup query. The query shows that the backup is mounted on Server B.
2. On Server C, use the command-line interface to complete a remote backup query of Server A. The query shows that the backup is mounted on Server B.
3. When you enter a **mount** or **query** command with the **/remotecomputer** option, enable command-line interface tracing and enable tracing on the agent. Enable tracing on both the local and remote systems by appending **/tracefile=filename.trc /traceflag=service** to the command.

Mounting SQL Server backups

To see a copy of SQL Server data from a specific point in time, mount a snapshot backup.

About this task

A copy of data from a specific time is also known as a point-in-time consistent copy or online snapshot.

Restriction: You cannot use Microsoft Management Console (MMC) to mount a backup to a different server. To mount a VSS snapshot to a remote server, enter the **mount backup** command at the command line.

Procedure

1. Start MMC.
2. Click **Recover Data** in the welcome page of MMC.
3. In the Actions pane on the **Recover** tab, click **Mount Backup**.
4. Either type the path to the empty NTFS or ReFS folder where you want to mount the backup or browse to find the path. Click **OK**. On the **Recover** tab, the backup that you mounted is displayed.
5. In the Actions pane, select the **Explore** and **Unmount Backup** tasks for the backup that you mounted.

Restoring SQL Server databases and files

You can restore legacy backups of SQL Server databases, file groups, and files by using Microsoft Management Console (MMC).

Setting single-user mode for restore operations

You might have to start an SQL Server instance in single-user mode during certain restore operations. For example, you might use single-user mode when you are restoring a damaged master database or a system database, or when you are changing server configuration options.

Before you begin

Restriction:

- You cannot restore SQL databases that are in use. By placing SQL databases to be restored in single-user mode, you can avoid system attempts to restore those databases.
- Microsoft Management Console (MMC) cannot connect to a SQL Server instance that is started in single-user mode. If you want to use MMC when the SQL Server instance is in single-user mode, you must use the command-line interface, `tdpsqlc.exe`, to restore the master database.

Procedure

1. To determine which users are using the databases, use the SQL stored procedure, `SP_WHO`.
2. To force users off the SQL database and set the SQL Server to single-user mode, issue this `TRANSACT-SQL` command.

```
ALTER DATABASE DBNAME SET SINGLE_USER  
WITH ROLLBACK AFTER N SECONDS
```

3. To start the SQL Server in single-user mode, use the `-m SQL SERVER` startup option.
4. To return the database to multiple-user mode, issue this TRANSACT-SQL command.

```
ALTER DATABASE DBNAME SET MULTI_USER
```

Setting data restore options in Microsoft Management Console

To optimize the data restore process for your environment, modify the default options that are available in Microsoft Management Console (MMC).

Procedure

1. On the Recover tab, select **Database Restore**.
2. Click **Show Restore Options** to modify the default restore options as follows: You can modify the default values for the restore options as described in the following table:

Table 13. Database restore options

Option	Action
Auto Select	<p>For this option, specify a value of Yes (default) to quickly select the backup objects to restore. With automatic selection, when you select the most recent backup to restore, all associated backups are automatically selected, up to the previous full backup. This option affects backups in the following ways:</p> <ul style="list-style-type: none"> • When you click a differential backup, the associated full backup is also selected. • When you click a log backup, the associated full backup and all associated earlier differential or log backups are also selected.
Performance	
Stripes	<p>For this option, 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.</p> <p>The value that you enter must correspond to the value that you set for SQL Server buffers.</p> <p>Restriction: This restore option is available only with legacy backups.</p> <p>When you use a multiple stripes number for legacy backups, and set the Verify Only parameter to Yes to restore the legacy backup, the number of stripes for legacy restore must be equal to or greater than the number of stripes for the legacy backup.</p>
Restore Behavior	
Database Owner Only	<p>To mark a database for owner use only, set this value to Yes. The default value is No, which specifies not to mark the database for owner use.</p> <p>Restriction: This restore option is only available with legacy backups.</p>

Table 13. Database restore options (continued)

Option	Action
Replace	<p>To replace a database during a restore operation, set this value to Yes. The default value is No, which specifies not to replace databases.</p> <p>Restriction: This restore option is available only with legacy backups.</p>
Recovery	<p>Use this option to restore data to an SQL Server database that is not on a standby SQL Server. The default value is Yes.</p> <ul style="list-style-type: none"> • Select Yes when you run a sequence of restore operations to an SQL Server database and the current restore operation is the final one in the sequence, or when it is the only restore operation. • Select No when you run a sequence of restore operations to an SQL Server database and the current restore operation is not the final one in the sequence. Select No for all restore operations in the sequence except for the final one.
Stand By Undo File Name	<p>For this option, specify a value of Yes to change the target SQL Server database to a standby SQL Server database. The default value is No.</p> <p>This option is available for full, differential, and log backup types. When you specify this option for a database, it applies to all backup objects for that database. Similarly, when you remove this option for a backup object, the option is removed for all backup objects.</p>
Verify Only	<p>Before you restore a legacy database backup, set this option to Yes to verify that the database backup is complete and can be read. The default value is No.</p> <p>Restriction: This restore option is available only for legacy database backups.</p> <p>When you use a multiple stripes number for legacy backups, and set the Verify Only parameter to Yes to restore the legacy backup, the number of stripes for legacy restore must be equal to or greater than the number of stripes for the legacy backup.</p>
Source Server	
From SQL Server	<p>Use this option to specify the name of the SQL Server that the backup is created from.</p> <p>To specify the name of a virtual environment SQL Server, change IncludeTsmVM to Yes to view Virtual Environment backup SQL Server databases in the Databases view. The backup method is listed as TSMVM to distinguish these databases from the other databases that are listed.</p>
Tape	

Table 13. Database restore options (continued)

Option	Action
Wait for Tape Mounts for Restore	Use this option to 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. The default value is Yes .
Wait for Tape Mounts for File Information	When you query Tivoli Storage Manager for file information, use this option to specify whether Data Protection for SQL Server waits for the Tivoli Storage Manager server to mount removable media. The default value is Yes . Restriction: This restore option is available only with legacy backups.
VSS	
Instant Restore	For this option, specify a value of Yes to use volume-level snapshot restore (instant restore) for local VSS backups if the backup exists on SAN-attached volumes. Specify a value of No to disable instant restore, which bypasses volume-level copy and uses file-level copy (fast restore) to restore the files from a local VSS backup. The default value is Yes , which uses volume-level snapshot restore if it is available. This option is available for VSS operations only. If you use instant restore for SAN Volume Controller earlier than version 5.1 or DS8000, ensure that any previous background copies that involve the volumes that are being restored are completed before you initiate the instant restore. In an instant restore operation, files on the destination file system are overwritten. Incremental and differential backups are automatically converted to file-level restores. An instant restore operation requires that the drive or volume where the mailbox database is located must be available. Any other process or application must not have access to the drive or volume.

Related tasks:

“Troubleshooting VSS backup and restore operations” on page 138

Restoring SQL Server data

You can restore SQL Server databases or parts of databases only from full, differential, and log backups. You can also restore availability databases with SQL Server 2012 and later versions.

About this task

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

- The Regional settings, which are defined in the Regional property page, must match the date format that is defined for the Microsoft SQL Server.

- You can use VSS to run backup operations of type full. You can apply legacy differential and legacy log backups after a full VSS backup is restored.
 - When Virtual Environment restore operations are configured from the Tivoli Storage Manager server, you can restore and view these databases from the Recover tab.
 - You can also restore availability databases that you backed up with the AlwaysOn node with SQL Server 2012 and later versions. Backups of availability databases can be restored to any availability replica in an availability group.
 - You can restore a legacy database backup that is verified as valid and complete with the **Verify Only** option in Microsoft Management Console (MMC), or with the **/VERIFYOnly** option of the **restore** command on the command line.

Procedure

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

Table 14. Database restore views

Task	Action
View a list of SQL Server databases that are available for a restore operation	Click View: Databases .
View a list of SQL Server database backup files that are available for a restore operation	Click View: Files .
View a list of SQL Server 2012 and later version availability databases that are available for a restore operation	Click DP Node Backups to show AlwaysOn node backups. Toggle the DP Node Backups / AlwaysOn Node Backups button for the respective database views.

4. On the **Recover** tab of an SQL Server instance, select an option for viewing databases. In the Results pane, browse to the databases that are available to restore. The following options are available:

Table 15. Database restore selection options

Option	Action
Search	Enter a keyword in the Search field to refine and filter the list of databases.

Table 15. Database restore selection options (continued)

Option	Action
Filter	<p>Use the filter options to refine and filter the list of databases.</p> <ol style="list-style-type: none"> 1. Click Show Filter Options and Add Row. 2. In the Column Name field, click the down arrow and select an item to filter. 3. In the Operator field, select an operator. 4. In the Value field, specify a filter value. 5. If you want to filter on more items, click Add Row. 6. Click Apply Filter.
Backups	<p>Select the database to restore. You can click Active Backups to show only active backups, or click All Backups to show both active and inactive backups.</p>
Refresh	<p>Click Refresh to update the view with your changes.</p>

If you applied a filter, the objects on the server that match the filter or search criteria are listed on 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.

5. Verify the options for the restore operation. If the restore options are not displayed, click **Show Restore Options**.
6. Optional: Choose a mode for the current task:
 - **Run Interactively**: Click this item to run the current task interactively. This selection is the default.
 - **Run Scheduled**: Click this item to convert the current action into a scheduled task. When you select this item, the schedule wizard runs the command that is required to complete the task.
7. To start the restore operation, in the Actions pane, take one of the following actions:
 - Click **Restore**.
 - Click **Restore VerifyOnly**. The **Restore VerifyOnly** task is available only if all the selected database backups are legacy backups.
8. To view the status of the restore operation, click **Task List** in the results pane. Click **Task Details** to view detailed status information.

Related tasks:

"Troubleshooting VSS backup and restore operations" on page 138

"Setting data restore options in Microsoft Management Console" on page 110

Restoring an SQL Server database to an alternate instance

By using Microsoft Management Console (MMC), you can restore an SQL Server database backup to an alternate SQL Server instance or database. You can also restore availability databases to an alternate location on any availability replica in an availability group.

Before you begin

Install Data Protection for SQL Server on both systems. Unlike legacy backups, you cannot restore VSS backups to an SQL Server that has a different name.

About this task

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

You can select only one database at a time when you restore a database to an alternate location.

Procedure

1. Copy the Data Protection for SQL Server options file (dsm.opt) from the source system to the target system.

Source system

The system from which the original backup (to be restored) is created.

Target system

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

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 MMC.
3. On the **Recover** tab for the SQL Server instance, specify the type of SQL Server data to restore.

Table 16. Database backup views

Task	Action
View a list of SQL Server databases that are available for a restore operation	Click All Backups .
View a list of SQL Server 2012 and later version availability databases that are available for a restore operation	Click DP Node Backups to show AlwaysOn node backups. Toggle the DP Node Backups / AlwaysOn Node Backups button for the respective database views.

4. Verify the options for the restore operation. If the restore options are not displayed, click **Show Restore Options**.
 - a. Ensure that **Wait for Tape Mounts for Restore** is set to **Yes**.
 - b. Ensure that **Wait for Tape Mounts for File Information** is set to **Yes**.
 - c. If the database to be restored is 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 **Yes** to use Instant Restore. Click **No** to disable Instant Restore if you want to use Fast Restore.

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

5. To start the backup operation, in the Actions pane, take one of the following actions:
 - a. Click **Restore to Alternate Location**.
 - b. Click **Restore VerifyOnly to Alternate Location**. The **Restore VerifyOnly to Alternate Location** task is available only if all the selected database backups are legacy backups.
6. In the **Restore Into** section of the Alternate Location Restore Settings window, click **Restore to new database**, and specify a target SQL Server instance name and target database name to restore a backup object to. VSS backups cannot be restored into an SQL Server that has a different name.

Tip: Any type of **Restore Into** processing automatically disables the VSS instant restore function.

7. In the **Relocate** section of the window, filter the restore processing operations.

Table 17. Database backup views

Task	Action
Specify new destination locations in which to restore backed up SQL Server databases, logs, and FILESTREAM files (SQL Server 2008 or later versions)	Click Restore all files into one directory .
Restore the log files into a location that is different from where the SQL Server database and other related files are restored	Select Relocate logs into and specify a new path in the text entry field.
Restore FILESTREAM files (SQL Server 2008 or later versions) into a location that is different from where the SQL database and logs are restored	Select Relocate other files into , and specify a new path in the text entry field.
Restore one or more individual SQL Server database, log, and FILESTREAM files	Click Relocate files individually , and click Browse to open a folder selection window. Select a folder or create 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 a partial restore operation in MMC. You must use the command-line interface to complete a partial restore operation that requires these parameters.

8. Click **Restore** to close the Alternate Location Restore Settings window and begin the restore.
9. To view the status of the restore operation, click **Task List** in the lower half of the results pane. Click **Task Details** to view detailed status information.

Restoring the master database

A damaged master database can prevent the SQL Server from starting and cause other errors. To protect your data if the master database is damaged, you must routinely complete a full database backup of the master database (msdb).

Before you begin

- Set single-user mode for restore operations.
- Always 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 you use any of these statements:
 - 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

About this task

If the master database is damaged while a server instance is running, fix the damaged database by restoring a recent full master database backup. If a server instance cannot start because the master database is damaged, the master database must be rebuilt. When you rebuild a master database, all system databases revert to their original state.

Restriction: Microsoft Management Console (MMC) cannot connect to an SQL Server instance that is started in single-user mode. When the SQL Server instance is in single-user mode, you must use the command-line interface, `tdpsqlc.exe`, to restore the master database.

Procedure

1. Click **Start > All Programs > Tivoli Storage Manager > Data Protection for Microsoft SQL Server > SQL Client - Command Line**.
2. Start the SQL Server in single-user mode.
3. Use Data Protection for SQL Server to restore the master database. When the master database finishes the restoration process, the SQL Server shuts down and an error message is displayed. The message indicates that the connection to the SQL Server is lost. This loss of connection is expected.
4. Restart the database engine to restore SQL Server to the typical multiuser mode.
5. Run the SQL Server setup program to rebuild the master database. When you rebuild the master database, use the same character set and sort order as the master database backup that is to be restored.

6. Manually reapply any changes to the master database that occurred after the date of the database backup that is used to complete the restore operation.
7. Restore the msdb database. During the process of rebuilding the master database, the SQL Server setup program drops, and then re-creates, the msdb database. Therefore, you must restore the msdb database with the master database.

Results

After the master database is restored, you can use MMC to back up and restore individual databases that are operating in single-user mode.

Related tasks:

“Setting single-user mode for restore operations” on page 109

“Troubleshooting VSS offline restore of a master database” on page 139

Restoring SQL Server databases with full-text catalogs and indexes

You can protect SQL Server 2005 and 2008 data with full-text catalog files.

About this task

When protecting SQL Server 2005 and the full-text index is part of a full-text catalog, note that the full-text catalog has a physical path. In this scenario, the full-text catalog is treated as a database file.

When you back up an SQL Server 2008 database and later data, a full-text catalog is either a logical or virtual object that contains a group of full-text indexes. This full-text catalog does not have a physical path. When you restore a database with SQL Server 2008 and later full-text catalog files, no data is explicitly stored. The file is automatically backed up and restored as part of the filegroup.

Procedure

- To restore a database with the SQL Server 2005 physical full-text catalog file from the command-line interface, use the **/RELocate** and **/T0** parameters. For example:

```
Restore DATABASE full /relocate=database,sysft_docindex,database_log
/T0={database_dir}\database.mdf,{database_dir}\docindex,
{database_log_dir}\database_log.ldf
```
- To restore a database with the SQL Server 2005 physical full-text catalog file from the GUI, use the **Relocate files individually** option. From the command-line interface, use **/relocate** and **/T0** instead of **/RELOCATEDir**.

Restoring SQL Server databases from virtual machine snapshots

You can restore SQL Server databases from virtual machine snapshots when Tivoli Storage Manager for Virtual Environments, Version 7.1 and later is used to back up the data. Before you restore the SQL Server database, you must verify that the data is backed up correctly.

Before you begin

- Before you restore SQL Server databases from virtual machine snapshots, verify that the data is backed up according to the following procedure.

- You can restore the SQL Server data to basic disks with MBR-style partitions. Because of an SQL Server limitation, however, 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 is created.

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

Verifying the SQL Server backup 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 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 scheduled to occur. This notification allows the application to truncate transaction logs and commit transactions so that 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 an SQL Server. Specifying this parameter allows the SQL Server administrator to manually manage the SQL Server logs. The logs can be preserved as required and be used to restore SQL Server transactions to a specific checkpoint, after the virtual machine is restored. When this option is specified, the SQL Server 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
```

Tivoli Storage Manager does not back up the SQL Server log files. The SQL Server administrator must back up the log files so that those logs 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 operations. 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.
5. After you back up data, verify that the virtual machine backup contains the necessary database metadata.
 - a. Enter the following Tivoli Storage Manager Backup-Archive Client command on the data mover:


```
dsmc query vm <vmname> -detail
```
 - b. In the command output, ensure that the Application(s) protected: value specifies (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
```

Restoring the SQL Server database Procedure

1. Log on to the system where you want to restore the SQL Server database. The Data Protection for VMware Recovery Agent license and Data Protection for SQL Server must be installed on the system where you restore the data.
2. When Data Protection for 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. Set access to the virtual machine that is backed up in a data center node as shown in the following table:

Table 18. 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 Server	Node for Data Protection for SQL Server	Agent (data owner)
VSS_NODE	In guest virtual machine running Microsoft SQL Server	Node for Data Protection for SQL Server DSMAGENT	Agent (data worker)

4. Ensure that the Tivoli Storage Manager administrator provides access to the virtual machine that is 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. Enabling access is required for Data Protection for SQL Server to restore the data that is owned by the DC_NODE.

Running the commands from the DC_NODE is best practice 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.

5. The **set access** command cannot be run if the ASNODE option is used. To issue the **set access** command, use an option file that does not contain ASNODE.

- a. Copy dsm.opt and dsm.setaccess.opt files.
- b. If you run the **set access** command from a node with ASNODE in the options file, 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, repeat these steps.

6. From the Protect and Recover Data section in MMC, select an **SQL Server**.
7. On the **Recover** tab for that SQL Server, select **View: Databases** to see a list of SQL Server database backups that are available to restore. SQL Server databases that are backed up with Tivoli Storage Manager for Virtual Environments software are listed with the VMVSS backup method.
8. Verify the restore options. If the restore options are not displayed, click **Show Restore Options**.
9. 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.
10. On the Actions pane, click **Restore** to begin the restore operation.
11. To view the status of the restore operation, click **Task List** in the lower half of the results pane. Click **Task Details** to view detailed status information.

Restoring SQL Server databases on Windows Server Core

To restore Microsoft SQL Server 2012 and later version databases, use the **restore** command.

About this task

Use the following procedure to recover all or part of one or more SQL Server databases.

Procedure

To restore all or part of an SQL Server database on Windows Server Core, enter the following command at the command prompt:

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

where `database_name` specifies the name of the database, and `backup_type` specifies the type of backup such as a full backup. You can specify other options, such as the file group. For example, to create a full database restore of databases DB_01 and DB_02, and to 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 group=DB_01_group1
```

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=*
```

Restoring SQL Server file groups and files from legacy backups

By using Microsoft Management Console (MMC), you can restore SQL Server file groups and files from legacy backups. You can also restore availability databases to any availability replica in an availability group.

Before you begin

- You can restore databases or parts of databases from **group**, **file**, **set**, **log**, and **full legacy** backups.
- Microsoft SQL Server requires that all files in the Primary filegroup is restored before or with a user-defined filegroup. To bring the database back to a usable state, you must perform a log restore after the user-defined filegroup is restored.
-

You can also restore availability databases that you backed up with the AlwaysOn node with SQL Server 2012 and later versions. Backups of availability databases can be restored to any availability replica in an availability group. For more information, consider the guidelines and restrictions in the topic that describes the restore of availability databases.

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 **No**.
- 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 **Yes**.
- 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 **Yes**.

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 Availability Group in SQL Server 2012, and later versions.

Procedure

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

Table 19. Database restore views

Task	Action
View a list of SQL Server database backup files that are available for a restore operation	Click View: Files .
View a list of SQL Server 2012 and later version availability databases that are available for a restore operation	Click DP Node Backups to show AlwaysOn node backups. Toggle the DP Node Backups / AlwaysOn Node Backups button for the respective database views.

4. Verify the restore options. If the restore options are not displayed, click **Show Restore Options**.
5. Optional: Choose a mode for the current task:
 - **Run Interactively**: Click this item to run the current task interactively. This selection is the default.
 - **Run Scheduled**: Click this item to convert the current action into a scheduled task. When you select this item, the schedule wizard begins, complete with the command that is required to complete the task.
6. On the Actions pane, click **RestoreFile** or **RestoreGroup** to begin the restore operation.
7. To view the status of the restore operation, click **Task List** in the lower half of the results pane. Click **Task Details** to view detailed status information.

Protecting SQL Server data in a Windows Server Core environment

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

About this task

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

In such a minimal environment, only the command-line interface is available for Data Protection for SQL Server on Windows Server Core unless you use the Data Protection for SQL Server remote management support. Additionally, if you use Windows Installer (MSI) to install Data Protection for SQL Server, you can use only the unattended mode.

You can use the **backup** and **restore** commands to protect databases that are stored on Microsoft SQL Server 2012 or later versions.

Backing up SQL Server databases on Windows Server Core

To back up Microsoft SQL Server 2012 and later version databases, use the **backup** command.

About this task

Use the following procedure to back up SQL Server databases to the Tivoli Storage Manager server, or to take local VSS snapshots.

Procedure

1. To back up all or part of an SQL database on Windows Server Core, enter the following command at the command prompt:

```
tdpsqlc backup database_name backup_type [other_options]
```

where *database_name* specifies the name of the database, and *backup_type* specifies the type of backup such as a full backup. You can specify other options, such as the back up method. For example, to create 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
```

For example, to create a full legacy backup of all databases on the SQL Server, enter the following command:

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

2. To back up a file group, enter the following command at the command prompt:

```
tdpsqlc backup database_name file_group
```

where *database_name* specifies the name of the database, and *file_group* specifies the file group in the database. For example, 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 Server databases on Windows Server Core

To restore Microsoft SQL Server 2012 and later version databases, use the **restore** command.

About this task

Use the following procedure to recover all or part of one or more SQL databases.

Procedure

To restore all or part of an SQL database on Windows Server Core, enter the following command at the command prompt:

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

where *database_name* specifies the name of the database, and *backup_type* specifies the type of backup such as a full backup. You can specify other options, such as the file group. For example, to create a full database restore of databases

DB_01 and DB_02, and to 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 group=DB_01_group1
```

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

To configure preferences for Data Protection for SQL Server, use the **set** command at 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`.

Procedure

At the command prompt, enter the following command:

```
tdpsqlc set parameter=value [/configfile=filename]
```

where *parameter* is the Data Protection for SQL Server parameter or option for which you want to change the value, 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.

Examples:

Task Set the preferred SQL Server in the `tdpsql.cfg` file.

```
Command: tdpsqlc set sqlserver=your_SQL_instance  
/configfile=tdpsql.cfg
```

```
Command: tdpsqlc set fromsqlserver=your_SQL_instance  
/configfile=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
```

Viewing, printing, and saving reports

You can access reports on recent activity and historical managed capacity. You can determine which licenses and software are installed.

Procedure

1. Select **Reporting** in the **Manage** section. A list of available reports is displayed. Each report provides a summary of the report contents.
2. Select a report from the list. The selected report displays.
3. To print or save the current report, click the appropriate icon at the top of the report.

Chapter 5. Automating

With Data Protection for SQL Server *automation* capability, you can run commands from the command line, create scripts, schedule tasks, and use Microsoft Management Console (MMC) to start tasks. The tasks that you can automate are based on the scripts and schedules that you create.

About this task

Data Protection for SQL Server supports you automating tasks from the command-line interface or Microsoft Windows PowerShell command prompt (Version 3.0 or later). You can also use the **Automate** tab in MMC.

Preparing to use Windows PowerShell cmdlets with Data Protection for SQL Server

Data Protection for SQL Server includes a set of Windows PowerShell cmdlets to help you manage Data Protection for SQL Server components in your environment.

About this task

You can issue cmdlets that are provided with Data Protection for SQL Server in Windows environments.

Data Protection for SQL Server cmdlets help support a seamless management environment and greatly improve remote management and automation capabilities. You can aggregate cmdlets together to form commands and use the large volume of existing cmdlets from other vendors.

Before you use the cmdlets, complete the following steps.

Procedure

1. Log on to the system as an administrator.
2. From a Windows PowerShell command prompt, issue the following command:
`set-executionpolicy remotesigned`
3. Import the Windows PowerShell modules from the FlashCopyManager folder:
 - FmModuleSQL.dll
 - FmModuleMMC.dll

From the Windows PowerShell command prompt, use your administrator credentials to import modules as follows:

- a. Browse to the FlashCopyManager 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 that 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 or later documentation. For more information about Windows PowerShell cmdlets, consistent naming patterns, parameters, arguments, and syntax, see this web page as a starting point: Microsoft TechNet: Getting Started with Windows PowerShell (<http://technet.microsoft.com/en-us/library/hh857337.aspx>).

Cmdlets for Microsoft Management Console

The following list identifies the cmdlets that you can use when interacting with Microsoft Management Console (MMC).

- **Clear-FcmMmcManagedCapacityHistory**
- **Clear-FcmMmcScheduledActivityHistory**
- **Disable-FcmMmcSchedule**
- **Enable-FcmMmcSchedule**
- **Get-FcmMmcActivity**
- **Get-FcmMmcComputerInformation**
- **Get-FcmMmcManagedCapacityHistory**
- **Get-FcmMmcReport**
- **Get-FcmMmcSchedule**
- **Get-FcmMmcScheduledActivity**
- **New-FcmMmcSchedule**
- **Remove-FcmMmcSchedule**
- **Set-FcmMmcSchedule**
- **Start-FcmMmcSchedule**

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

```
Get-Help New-FcmMmcSchedule
```

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

```
get-help New-FcmMmcSchedule -examples
```

For more information, enter:

```
get-help New-FcmMmcSchedule -detailed
```

For technical information, enter:

```
get-help New-FcmMmcSchedule -full
```

For online product information, enter:

```
get-help New-FcmMmcSchedule -online
```

For information about a specific parameter, enter:

```
help New-FcmMmcSchedule -Parameter backupdestination
```


To display the help in a separate window, include the **-showwindow** parameter with the **help** command.

Cmdlets for protecting Microsoft SQL Server data

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

Table 20. 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 Server 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.
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 an SQL Server database, or deactivate 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
```

For online product information, 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.

Automating tasks

You can automate a workload by entering Windows PowerShell cmdlets or command-line interface commands in the integrated command line.

About this task

You use the Automate view to work with commands. You can create, save, store, and schedule commands to run at the scheduled time.

Procedure

1. To open the Automate view, select a workload that you want to work with and click **Automate**. An integrated command line is available in the task window from which you can enter PowerShell cmdlets or command-line interface commands.

2. Change **PowerShell** to **Command Line**.

3. To run a command, type a command in the details pane and click the **Execute** icon. You can issue the commands with or without specifying `tdpsqlc`.

For example, for each selected workload instance, you can enter a single command or multiple commands, such as:

```
q tsm  
q sql
```

You can also run a saved task by clicking the **Open** icon, selecting the command file, and clicking the **Execute** icon. The output is displayed in the main window.

4. Click the **Save** icon and follow the prompts to save a command for future use.
5. To schedule a command, click the **Schedule this command** icon to open the scheduling wizard. Follow the prompts in the wizard to create a schedule for the command. The output of the command is displayed in the results pane.
6. (Optional) Save or send the command output to an email address.

What to do next

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

1. Start MMC and select the **SQL Server** instance in the navigation tree.
2. Click the tab for the task you want to do (**Protect** or **Recover**).
3. Automate the command by using one of the following methods:

Result pane

Select the item for your task in the result pane, and select **Run Scheduled** in the toolbar menu. Click the appropriate task in the **Action** pane. When the 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.

Tivoli Storage Manager task scheduler

Review these guidelines when you define a Tivoli Storage Manager schedule.

- If you use the Tivoli Storage Manager-prompted scheduling mode, ensure that the Data Protection for SQL Server option file specifies the `tcpclientaddress` and `tcpclientport` options. If you want to run more than one scheduler service, use the same `tcpclientaddress`. However, you must use different values for `tcpclientport` in addition to the different node names. As an example, you might want to run more than one scheduler service when you schedule Data Protection for SQL Server and the regular Windows backup client.

You can use server-prompted scheduling only when TCP/IP communication is used. By default, Data Protection for SQL Server uses the client polling schedule mode.

- If you make any changes that affect the scheduler in the Data Protection for SQL Server options file, restart the scheduler to activate the changes. For example, the Tivoli Storage Manager address, the schedule mode, or the client TCP address or port can affect the scheduler. To restart the scheduler, issue the following commands:

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

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

```
d:\Program Files\Tivoli\TSM\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 a log file with statistics about the backed up database objects when the `/logfile` parameter is specified during the **tdpsqlc** command. Outputs from the scheduled commands are sent to the scheduler log file (`dsmsched.log`). After scheduled work is completed, check the log to verify that the work is completed successfully.

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

```
Scheduled event eventname completed successfully
```

This result indicates that Tivoli Storage Manager successfully issued the scheduled command that is associated with the *eventname*. No attempt is made to determine whether the command succeeded or failed. To assess the success or failure of the command, evaluate the return code from the scheduled command in the scheduler log. The scheduler log entry for the command return code is prefaced with the following text:

```
Finished command. Return code is:
```

If any scheduled backups fail, the scheduler script exits with the same error code as the failed backup command. A non-zero error code means that the backup failed.

- If `passwordaccess generate` is not specified in the `dsm.opt` file, then the Tivoli Storage Manager password must be specified on the **tdpsqlc** command. To specify the password, use the **/tsmpassword** parameter in the command file that is 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=mars1 /tsmpassword=newpassword
```

Chapter 6. Troubleshooting

Data Protection for SQL Server supports you in protecting Microsoft SQL Server databases. Data Protection for SQL Server uses the Microsoft Server Managed Objects (SMO) application programming interface (API), and the Microsoft Virtual Shadow Copy Service (VSS).

About this task

If an error condition occurs during a Data Protection for SQL Server event, you typically start with a symptom, or set of symptoms, and trace the root cause. Problem determination, however, is not the same as problem solving. During the process of problem determination, you might obtain sufficient information to enable you to solve the problem. In some cases, you cannot solve a problem even after you determine its cause. For example, a performance problem might be caused by a limitation of your hardware.

Diagnosing problems

One of the most difficult challenges of troubleshooting in a client-server environment is determining which component is the origin of the problem. VSS diagnostic wizards are available to help you test VSS snapshots on your system. You can determine whether the source of the problem is a general VSS issue, or a Tivoli Storage Manager issue, or a Data Protection for SQL client issue.

Error log files

If an error condition occurs during a Data Protection for SQL Server event, you can view several log files to help diagnose the problem.

For example, you can confirm that Data Protection for SQL Server failed over by searching entries about the secondary server in the following log files:

- `Tivoli\tsm\TDPSQL\dsierror.log`
- `Tivoli\tsm\baclient\dsmerror.log`
- Data Protection for SQL Server logs information about backup, restore, and delete commands to the Tivoli Event Console.
- Data Protection for SQL Server logs information, by default, to the `tdpsql.log` file in the directory where Data Protection for SQL Server is installed. This file indicates the date and time of a backup, the data that is 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 Server is installed. No backup statistics are contained in this log. The `dsierror.log` file cannot be marked as read-only.
- Data Protection for SQL Server logs information to the SQL Server error log. The 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.

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

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

Scheduled event *eventname* completed successfully

This entry 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 System and Application Event Log.
- For VSS operations, view the dsmerror.log file in the backup-archive client installation directory.

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

The Data Protection client interacts closely with the backup-archive client (DSMAGENT). The client completes all of the Virtual Shadow Copy Service (VSS) operations. You can test the connectivity between the Data Protection client and the Tivoli Storage Manager DSMAGENT. You can determine whether the source of the problem is the Microsoft VSS service or a problem within the Tivoli Storage Manager code.

About this task

- The vssadmin and diskshadow tools are applications that can run backups that use the Microsoft SQL Server VSS APIs.

vssadmin

A utility that is installed with your operating system, and can show current volume shadow copy backups and all installed shadow copy writers and providers in the command window.

diskshadow

The diskshadow tool is available with Windows Server 2008 and later versions.

With these tools, you can determine the following items:

- Verify VSS provider configurations
- Rule out any possible VSS problems before you run the Tivoli Storage Manager VSS functions
- That you might have a VSS configuration problem or a real hardware problem if an operation does not work with diskshadow or vssadmin
- That you might have a Tivoli Storage Manager problem if an operation works with diskshadow or vssadmin but not with the Tivoli Storage Manager
- For VSS operations, you can re-create the problem with the Microsoft diskshadow tool. If you are able to re-create the problem with the diskshadow tool, the source of the problem is likely to be within the VSS provider or the SQL Server.

Procedure

1. Test the connectivity between the Data Protection client and the Tivoli Storage Manager DSMAgent.
 - a. Select the SQL Server workload that you want to work with and click the **Automate** tab to open the **Automate** view.
 - b. To verify that your installation and configuration is correct, issue the **TDPSQLC QUERY SQL** command in the lower details pane and click **Execute** (or **Enter**). Alternatively, issue the **TDPSQLC QUERY SQL** command on the computer where the SQL Server is installed. The results are displayed in the pane.

The **TDPSQLC QUERY SQL** command returns information about the following items:

- SQL Server status
- Circular logging
- VSS components

The following example shows a sample of the output that is 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 1.0
(C) Copyright IBM Corporation 1997, 2014. All rights reserved.

Connecting to SQL Server, please wait...

SQL Server Information
-----

SQL Server Name      ..... RHINO
SQL Server Version   ..... 12.0.2000 (SQL Server 2014)

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

Writer Name          : SqlServerWriter
Local DSMAgent Node   : RHINO
Remote DSMAgent Node  :
Writer Status        : Online
Selectable Components : 6

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

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. To determine whether the problem is a Microsoft VSS service issue or a problem within the Tivoli Storage Manager code, use the **vssadmin** and **diskshadow** tools to re-create the error as follows:
 - a. Issue **VSSADMIN** and **diskshadow** commands, as shown in the following table:

Table 21. VSSADMIN commands

Option	Run this command:
To list the VSS writer:	VSSADMIN LIST WRITERS
To list the VSS providers:	VSSADMIN LIST PROVIDERS
To list the shadow copies that are created by using the Microsoft Software Shadow Copy provider:	VSSADMIN LIST SHADOWS
To list the shadow copies that are created by using a VSS hardware provider starting with Windows Server 2008:	diskshadow diskshadow > list shadows all

- b. Before you install Data Protection for SQL Server, test the core VSS function. Complete the following diskshadow testing before you install any Tivoli Storage Manager components:

- Test non-persistent shadow copy creation and deletion by issuing the following **diskshadow** commands:

```
diskshadow>set verbose on
diskshadow>begin backup
diskshadow>add volume f: (database volume)
diskshadow>add volume g: (log volume)
diskshadow>create
diskshadow>end backup
diskshadow>list shadows all
diskshadow>delete shadows all
diskshadow>list shadows all
```

Volumes *f:* and *g:* represent the SQL Server database and log volumes. Repeat issuing the **diskshadow** commands four times and verify that the Windows System and Application Event Log file contains no errors.

- Test persistent shadow copy creation and deletion by issuing the following **diskshadow** commands:

```
diskshadow>set context persistent
diskshadow>set verbose on
diskshadow>begin backup
diskshadow>add volume f: (database volume)
diskshadow>add volume g: (log volume)
diskshadow>create
diskshadow>end backup
diskshadow>list shadows all (this might take a few minutes)
diskshadow>delete shadows all
diskshadow>list shadows all
```

Volumes *f:* and *g:* represent the SQL Server database and log volumes. Repeat issuing the **diskshadow** commands four times and verify that the Windows System and Application Event Log file contains no errors.

- Test persistent transportable shadow copy creation and deletion by issuing the following **diskshadow** commands:

```
diskshadow>set context persistent
diskshadow>set option transportable
diskshadow>add volume f: (database volume)
diskshadow>add volume g: (log volume)
diskshadow>set metadata c:\metadata\SQLmeta.cab
(the path where you want the metadata stored)
diskshadow>create
```

You must copy the `sqlmeta.cab` file from the source server to the offload server. After you copy the file, issue the following commands:


```

diskshadow>load metadata newpath/sqlmeta.cab
diskshadow>import
diskshadow>list shadows all (this might take a few minutes)
diskshadow>delete shadows all

```

Volumes *f*: and *g*: represent the SQL Server database and log volumes.

Repeat issuing the **diskshadow** commands four times and verify that the Windows System and Application Event Log file contains no errors.

3. Test nonpersistent shadow copy creation and deletion.

- a. Issue the following **diskshadow** commands. In this example, the volumes *k*: and *l*: represent the SQL Server database and log volumes.

```

diskshadow>set context volatile
diskshadow>set verbose on
diskshadow>begin backup
diskshadow>add volume k: (database volume)
diskshadow>add volume l: (log volume)
diskshadow>create
diskshadow>end backup
diskshadow>list shadows all (this might take a few minutes)
diskshadow>exit
diskshadow
diskshadow>list shadows all (no shadow copies)

```

- b. Verify that the Windows System and Application Event Log file contains no errors.

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

- a. Issue the following **diskshadow** commands. In this example, the volumes *k*: and *l*: represent the SQL Server database and log volumes.

```

diskshadow>set context volatile
diskshadow>set option transportable
diskshadow>add volume k: (database volume)
diskshadow>add volume l: (log volume)
diskshadow>set metadata c:\metadata\SQLmeta.cab
(the path where you want the metadata stored)
diskshadow>create

```

- b. Repeat issuing the **diskshadow** commands four times and verify that the Windows System and Application Event Log file contains no errors.

- c. Copy the sqlmeta.cab file from the source server to the offload server. Then, issue the following commands:

```

diskshadow>load metadata newpath/sqlmeta.cab
diskshadow>import
diskshadow>list shadows all (this might take a few minutes)
diskshadow>delete shadows all

```

If any of these tests fail repeatedly, you have a hardware configuration problem or a VSS problem. Consult your hardware documentation for known problems or search Microsoft Knowledge Database for information.

If all tests pass, continue to Step 5.

5. Re-create your specific problem by using diskshadow. If you can re-create your problem, only through a series of steps (for example: a backup fails only when you perform two consecutive local backups), try to perform the same tests by using diskshadow.

- Simulate SQL Server VSS backups to Local by running a diskshadow persistent snapshot.
- Simulate SQL Server VSS backups to the Tivoli Storage Manager by running a diskshadow nonpersistent snapshot.

- Simulate SQL Server VSS backups to Local and to the Tivoli Storage Manager by running a diskshadow persistent snapshot.
- Simulate offloaded SQL Server VSS backups to the Tivoli Storage Manager by running a diskshadow nonpersistent, transportable snapshot.

See the diskshadow documentation for the specific commands that are required to complete backups. If you can re-create the problem, you are most likely experiencing a general VSS issue. See the Microsoft Knowledge Database for information. If your operation passes successfully with diskshadow, you are most likely experiencing a Tivoli Storage Manager or Data Protection for SQL client problem.

6. For legacy backup operations, try to re-create the problem by using the Backup or Restore utility in the SQL Server administrator program. If the following error message is displayed, the SQL Server encountered an unexpected situation: AC05350E An unknown SQL API error has occurred. Microsoft assistance might be needed if the problem continues.

Data Protection for SQL Server error messages occasionally contain an HRESULT code. For more information about the problem, use the HRESULT code to search Microsoft documentation and the Microsoft Knowledge Base.

Resolving reproducible problems

When a component fails to operate as designed, try to reproduce the problem and capture information about the current operating environment at the time of the error. You can troubleshoot VSS backup and restore operations, VSS and SAN Volume Controller, Storwize V7000, or DS8000 problems.

Troubleshooting VSS backup and restore operations

If you encounter a problem during VSS backup and restore processing, attempt to reproduce the problem in your environment.

Before you begin

If a VSS backup fails, verify that sufficient disk space is available to store the snapshot.

Procedure

1. Try the operation that failed again.
2. If the problem still exists, close other applications, especially those applications that interact with SQL Server, for example, antivirus applications. Try the operation that failed again.
3. If the problem persists, look for information in the event logs: `tdpsql.log` and `dsmerror.log`. You can also review the messages in the Windows System and Application Event Log. Log entries might exist to help you identify the VSS event that triggers the issue.
4. If you do not find a resolution to the problem in the log files, complete the following steps:
 - a. Shut down the SQL Server or the computer.
 - b. Restart the SQL Server or the computer.
 - c. Run the operation that failed.

Troubleshooting issues with tail-log backups

A database restore operation might fail if transaction log records in the *tail of the log* are not backed up.

About this task

During the restore operation, you might see the following error 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)
```

To resolve the error, complete the tail-log backup.

Procedure

1. On the Protect tab of the SQL Server instance, click **Show Backup Options** and set the Back Up Tail-log option to **Yes**.
2. On the Actions pane, select **Log Backup to TSM**.

Troubleshooting VSS offline restore of a master database

Microsoft SQL Server only supports offline VSS restores of the master database. Data Protection for SQL Server does not support offline restore operations. Therefore, you cannot use Data Protection for SQL Server to restore the master database.

Procedure

1. Ensure that the SQL Server is online.
2. Restore the master database to a new database in Microsoft Management Console (MMC), or at the command line. For example: Enter the **tdpsqlc** command with the **/recovery=no** option.
3. After the restore operation is complete, verify that all data files are restored successfully.
4. Stop the SQL Server instance, and rename all data files of the master database.
5. Copy all data files from the new master_restore database to the location of the master database. Verify that all data files are copied.
6. Start the SQL Server instance and verify that the master database is restored successfully.

Related tasks:

“Restoring the master database” on page 117

Troubleshooting SQL Server failover cluster limitations for VSS operations

Certain limitations apply when you complete VSS operations in an SQL Server failover cluster environment.

Procedure

- Ensure that all servers within the cluster use the same levels of Tivoli Storage Manager, Windows, and other applicable software.
- Install the Tivoli Storage Manager client acceptor daemon (CAD) on each cluster node so that it can continue operations if servers fail over. Ensure that the CAD service name is named the same on all cluster nodes so that it can be started by a generic cluster service.

Tip: The local DSMAGENT client node must be a separate node from your normal backup-archive client because the CAD service must be a non-cluster option. When you use the remote DSMAGENT client node, you do not have to register a separate node for each server within the cluster.

- Use the Microsoft **vssadmin** and **diskshadow** commands to verify the environment.
- Configure a Data Protection for SQL Server configuration file for each node in the cluster. These configuration files are almost identical, except that the **localdsmagentnode** parameter points to the corresponding local DSMAGENT on each node.
- If you run scheduled VSS operations in an SQL Server failover cluster environment, complete the following steps.
 - Install the Tivoli Storage Manager scheduler as a Windows service on all cluster nodes.
 - If the command file is on a local drive, ensure that the file remains consistent on all cluster nodes. Optionally, you can create the command file on a shared drive. Make sure that the **objects** parameter, which is specified with the **define schedule** command on the Tivoli Storage Manager server, points to this command file.

What to do next

To resolve timeout issues, schedule VSS backups so that enough time elapses between backups, or increase the copy rate of the IBM SAN Volume Controller or IBM Storwize V7000 background copy.

Troubleshooting VSS limitations with IBM SAN Volume Controller and IBM Storwize V7000

When you run a Data Protection for SQL Server VSS backup (non-offloaded) to a Tivoli Storage Manager server, the IBM SAN Volume Controller or IBM Storwize V7000 LUNs can sometimes remain mapped to the Windows host even though the backup is complete.

Procedure

Use a backup destination other than Tivoli Storage Manager server (BOTH or LOCAL).

Results

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. Typically, the SQL Server data is on IBM SAN Volume Controller or IBM Storwize V7000 disks. However, the second backup is waiting for the IBM SAN Volume Controller or IBM Storwize V7000 background copy of the first backup to complete before proceeding. IBM SAN Volume Controller or IBM Storwize V7000 does not allow two background copies of the same volume to occur at the same time. You might not know that the second backup is waiting for the first background copy to complete.

You might also see timeout errors if the previous IBM SAN Volume Controller or IBM Storwize V7000 background copy takes too long.

What to do next

To resolve timeout issues, schedule VSS backups so that enough time elapses between backups, or increase the copy rate of the IBM SAN Volume Controller or IBM Storwize V7000 background copy.

Troubleshooting VSS limitations with IBM N-series and NetApp FAS series

If you plan to run VSS backups with backup destination set to LOCAL, understand the limitations in the VSS Provider for NetApp FAS series and IBM N-series, and in SnapDrive 4.2 and earlier versions, that affect the way in which you can run your VSS backup operations. You must configure your environment correctly to avoid snapshot deletions, backup failure, and out-of-space conditions on the production volumes.

Before you begin

- Ensure that a NAS file server LUN that is used by SQL Server databases is fully dedicated to the database. SQL Server databases cannot share LUNs.
- Ensure that a NAS filer LUN that is used by SQL Server databases is the only LUN on the filer volume. For example, if the SQL Server uses four LUNs, four corresponding filer volumes must exist, where each volume contains one LUN.
- If the NetApp volume type is Traditional, ensure that VSS backups with backup destination set to LOCAL are bound to a management class that has `verExists=1`. This setting is not required if flexible volumes are used.
- Ensure that VSS backups with backup destination set to LOCAL are either a full or copy backup. You cannot mix local backups of type FULL and COPY.
- Ensure that VSS backups with backup destination set to TSM are a full or copy backup. There are no restrictions on Tivoli Storage Manager backups.
- When you run VSS backups, ensure that previous backups finish completely before you start a new backup. To avoid issues on SQL Server, the VSS service, and, the NAS filer, avoid backup overlaps.

About this task

The following backup procedure is an example that shows how to optimally run VSS backups by using both Tivoli Storage Manager and local backup destinations. The following assumptions apply to this example backup procedure:

- Stated configuration requirements are in place.
- Daily VSS full backups to a local destination occurs every four hours - 12 a.m., 4 a.m., 8 a.m., 12 p.m., 4 p.m., 8 p.m.
- The VSS backup to Tivoli Storage Manager takes one hour to complete.
- The VSS backup to a local destination takes five minutes to complete.

Procedure

Set your daily VSS full back schedule to run in one of the following ways:

- Run daily VSS full backups to a local destination every four hours - 12 a.m., 4 a.m., 8 a.m., 12 p.m., 4 p.m., 8 p.m.
- Run daily VSS full backups to Tivoli Storage Manager storage by one of the following two methods:
 - Set **backupdestination** to BOTH to run at 12 a.m. Because this setting runs a 12 a.m. backup to a local destination, do not separately schedule a 12 a.m. backup to a local destination.
 - Set full offloaded-backup to run at 1 a.m. No VSS local backup is available to restore VSS backups between 1 a.m. and 4 a.m., when the next VSS backup to a local destination occurs.
- Set weekly VSS full backups to run to Tivoli Storage Manager, as an offloaded backup, at 5 a.m.

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

If you experience VSS and SAN Volume Controller, Storwize V7000, or DS8000 problems, use these troubleshooting tips to help you discount some common configuration issues.

Procedure

1. Verify connectivity to the CIMOM (Common Information Model Object Manager) as follows:
 - a. Refer to your SAN Volume Controller, Storwize V7000, or DS8000 documentation.
 - b. Run the **IBMVCFG LIST** command. The default location is %Program Files%\IBM\Hardware Provider for VSS-VDS.
 - c. Issue the **IBMVCFG SHOWCFG** command to view the provider configuration information.
 - d. Check that the CIMOM is properly configured. Run `verifyconfig.bat -u username -p password` on the Master Console.
 - e. Check the username and password. If the problem is with the truststore, follow the procedure in the documentation to generate a new truststore.

2. Verify CIMOM operational issues as follows:

- a. If your backup or restore operation fails, check the IBMVSS.log file.

If the backup or restore failure is from a CIMOM failure, the log displays output similar to the following example:

```
Wed Jan 13 17:34:34.793 - Calling AttachReplicas
Wed Jan 13 17:34:35.702 - AttachReplicas: 909ms
Wed Jan 13 17:34:35.702 - returnValue: 34561
Wed Jan 13 17:34:35.718 - AttachReplicas returned: 34561
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 the backup or restore operation is successful.

- b. To determine why a backup or restore operation failed, look at the log files. The files are generated by the CLI or graphical user interface (GUI), depending on how you run your operation. The log files might provide more information about the failure.
3. If the failure seems to be for a different reason than a CIMOM failure, verify your host configuration. Run the latest support levels of the software for SAN Volume Controller, Storwize V7000, or DS8000.
4. If you are unable to resolve these problems, provide the following information to IBM Support:
 - Information that is listed in the Tivoli Storage Manager diagnostic information section
 - HBA type, firmware, and driver levels
 - SDD version
 - SAN Volume Controller microcode version (if applicable)
 - DS8000 microcode version (if applicable)
 - Storwize V7000 microcode version (if applicable)
 - SAN Volume Controller or Storwize V7000 Master Console version (if applicable)
 - For DS8000, the CIM Agent version (if applicable)
 - IBMVSS.log
 - IBMVDS.log
 - Application Event Log
 - System Event Log
 - CIMOM logs if the problem seems to be related to CIMOM. 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\svconconsole\support, and the default location for DS8000 is C:\Program Files\IBM\cimagent.

Resolving problems with IBM Support

Contact IBM Support for further assistance if you have a problem that you are unable to solve by applying maintenance fixes, reproducing the issue, or using the information in previous topics. IBM Support might request to see some or all of the trace and log files while investigating a problem that you report.

About this task

Also, you might be asked to set a trace the Data Protection client when using VSS technology, and then collect the log. IBM Support uses the information that is captured in the log file to trace a problem to its source or to determine why an error occurred.

Viewing trace and log files

Data Protection for SQL Server uses several components. Each component is in its own directory along with its respective troubleshooting files. By using the Trace and Log Files view, you can easily view these files in a central location.

About this task

The following files are examples of the files that you can view, including default log and trace files:

Examples of Data Protection for SQL Server default log and trace files:

- Installation directory: C:\Program Files\Tivoli\TSM\TDPSql
- dserror.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

Procedure

1. When you encounter a problem in Microsoft Management Console (MMC), create trace files by using the Diagnostics property page.
 - a. Click **Properties > Diagnostics**, and click **Begin**.
 - b. Close the property page and reproduce the problem.
 - c. Open the Diagnostics property page and click **Stop**. Clicking the **Diagnostics** button is the preferred method for gathering information to send to your service representative. This method gathers all the information that is needed. Even if a problem occurs only on the command-line interface, command, you can always gather information by using the Automate tab. The log files are displayed in the Trace and Log Files view.
2. Click the trace or log file that you want to view. The contents of the file are displayed in the results pane.

Gathering traces for the Data Protection client when using VSS technology

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

About this task

To diagnose Data Protection for SQL Server VSS operational problems, gather these traces:

- Data Protection for SQL Server trace
- Tivoli Storage Manager API trace
- DSMAGENT trace

Procedure

1. To gather the Data Protection for SQL Server trace, open Microsoft Management Console (MMC) and go to the diagnostics property page to turn tracing on. Tracing is turned 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
2. Gather the Tivoli Storage Manager API trace as follows: Enable tracing with the DP/SQL dsm.opt file and the **"TRACEFILE"** and **"TRACEFLAGS"** keywords. The following text is an example of the entry in the DP/SQL dsm.opt file:

```
TRACEFILE APITRACE.TXT
TRACEFLAG SERVICE
```
3. Gather the DSMAGENT trace as follows: Enable tracing with the dsmagent (baclient) dsm.opt file and the **"TRACEFILE"** and **"TRACEFLAGS"** keywords. The following text is an example of the entry in the dsmagent (baclient) dsm.opt file:

```
TRACEFILE AGTTRACE.TXT
TRACEFLAG SERVICE PID TID ENTER ALL_VSS SBRM RESTORE
```

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

4. Enable the Volume ShadowCopy service debug trace features in Windows.

Gathering information about SQL Server with VSS before you call IBM

The Data Protection client depends on the operating system and the SQL Server application. Collecting all the necessary information about the environment can significantly assist IBM Support in determining the source of problem.

Procedure

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

- The exact level of the Windows operating system, including all service packs and hotfixes that are applied.

- The exact level of the SQL Server, including all service packs and hotfixes that are applied.
- The exact level of Data Protection for SQL Server with Volume Shadow Copy Service (VSS) backup and 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 Server **TDPSQLC QUERY SQL** command.
- The device type (and connectivity path) of the SQL Server databases and logs.
- (SAN only) The specific hardware that is being used. For example: HBA, driver levels, microcode levels, SAN Volume Controller or Storwize V7000 levels, DS8000 hardware details.
- Permissions and the name of the user ID that is 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. For information about how to enable tracing and collect the trace log files, see the documentation of the specific VSS hardware provider.
- (SAN only) The IBM CIM agent level for DS8000, SAN Volume Controller, or Storwize V7000.
- A list of other applications that are running on the system.
- A list of the steps that are needed to re-create the problem (if the problem can be re-created).
- If the problem cannot be re-created, list the steps that caused the problem.
- Is Data Protection for SQL Server running in a Microsoft Failover Clustering environment?
- Does the problem occur on other SQL Servers?

Gathering files from SQL Server with VSS before calling IBM

You can collect several log files and other data for Data Protection for SQL Server server diagnosis.

About this task

Microsoft Management Console (MMC) automatically collects information in a package file, which you can send to IBM Support. To collect this information manually, refer to the following file list.

Procedure

1. Gather as many of the following files as possible before you contact IBM Support.
 - The contents of the C:\adsm.sys\vss_staging directory and subdirectories. Or gather the appropriate directories if you are using the VSSALTSTAGINGDIR option.

- The Data Protection for SQL Server configuration file. The default configuration file is `tdpsql.cfg`.
 - The Data Protection for SQL Server Tivoli Storage Manager 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 captures or command-line output of failures or problems.
2. Gather the following Tivoli Storage Manager log files, which can indicate the date and time of a backup, the data that is backed up, and any error messages or completion codes that might help to determine your problem:
 - The Data Protection for SQL Server 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 `dsmcrash.dmp` and DSMAGENT crash log file, if requested. The default location is `C:\Program Files\Tivoli\TSM\baclient\dsmcrash.log`.
 3. Gather the following VSS provider log files, if applicable:
 - System Provider - (Windows System and Application Event Log)
 - IBM System Storage SAN Volume Controller, IBM Storwize V7000, or IBM System Storage DS8000 series - `%Program Files%\IBM\Hardware Provider for VSS\IBMVss.log`
 - NetApp - `%Program Files%\SnapDrive*.log`
 - XIV - zip up all of the files in the `C:\Windows\Temp\xProvDotNet` directory
 4. Use the Data Protection for SQL Server console to list the events that originate by Data Protection for SQL Server. Select **Dashboard - ServerName > Diagnostics > System Information** and double-click the `dpevents.ps1` script in the PowerShell section of the **System Information** page.

Viewing and modifying system information

You can view and edit scripts that provide information about system components including, for example, Windows-related services for Data Protection for SQL Server, Windows System and Application Event Log entries, and Volume Shadow Copy Service (VSS) information.

About this task

The System Information view is extensible. You can take advantage of this flexibility to add and share customized scripts.

Procedure

1. Open the System Information view as follows:
 - a. In the welcome page, click **Diagnostics**.

- b. On the Results pane, double-click **System Information**. 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 scripts.
2. Add, update, or delete your scripts, as follows:

Action	Steps
Add your own scripts	<ol style="list-style-type: none"> 1. Click New in the Actions pane. 2. If you want to copy your scripts directly to the ProgramFiles\Tivoli\FlashCopyManager\Scripts directory, make sure that your scripts follow these extension requirements: <ul style="list-style-type: none"> • PowerShell scripts: <i>filename.ps1</i> • Windows Management Instrumentation (WMI) scripts: <i>filename.wmi</i> • scripts: <i>filename.tsm</i> Tivoli Storage FlashCopy Manager uses the file type extension to determine how to run the script.
View or edit an existing script	<ol style="list-style-type: none"> 1. From the list of script files in the results pane, select the name of a script that you want to view or edit. Tip: The name of the script is displayed in the Actions pane. Click the name of the script in the Actions pane to reveal or hide a list of actions to process. 2. To open the script file for viewing or editing, click Command Editor in the Actions pane. 3. View or edit the script. 4. Click OK to save your changes, or click Cancel to exit the System Information Command Editor without saving any changes.
Delete a script	<ol style="list-style-type: none"> 1. From the list of script files in the results pane, select the name of a script that you want to delete. Tip: The name of the script is displayed in the Actions pane. Click the name of the script in the Actions pane to reveal or hide a list of actions to process. 2. Click Delete in the Actions pane.

Emailing files to IBM Support

You can send diagnostic information to IBM Support.

Before you begin

About this task

The Email Support files feature collects all detected configuration, option, system information, trace, and log files. It also collects information about services, operating systems, and application versions. These files are compressed and then attached in an email.

Procedure

1. Start Microsoft Management Console (MMC).
2. Click **Diagnostics** in the results pane of the welcome page.
3. Click the **E-Mail Support files** icon in the Action pane.
4. Enter the required information in the various fields and click **Done**. The information is sent to the designated support personnel and the dialog closes.

Results

Files are collected, compressed, and stored in the `flashcopymanager\problemdetermination` folder. The files are deleted and replaced each time that you email the support files. If the Email feature is not configured, or is blocked by a firewall, or if the files are large, use another method to transfer them. You can copy the files directly from the `flashcopymanager\problemdetermination` folder and transfer the files to another site by using another method such as FTP.

Online IBM support

Multiple online support resources are available for you to use.

The following list identifies where you can find information online:

- Tivoli Storage Manager wiki (<https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Tivoli%20Storage%20Manager>).
- Storage Management community on Service Management Connect (<https://www.ibm.com/developerworks/servicemanagement/sm/index.html>).
- Tivoli Storage Manager for Databases (<http://www.ibm.com/software/products/en/tivostormanafordata>). Enter the search term to narrow the search criteria for your support requirements. Examples of search terms that you might use include an authorized program analysis report (APAR) number, release level, or operating system.

Chapter 7. Performance tuning

Many factors can affect the backup and restore performance of Data Protection for SQL Server.

Some of these factors, such as hardware configuration, network type, and capacity, are not within the scope of Data Protection for SQL Server. Some options that are related to Data Protection for SQL Server can be tuned for optimum performance. The following issues affect performance.

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. As the value increases, the client can start more sessions. The range for the option is from 1 to 10.

Legacy backups are a stream of bytes that Data Protection for SQL Server stores on the Tivoli Storage Manager server.

VSS backups differ because these backups operate at the volume and file-level. In a situation where an 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.

Buffering (Legacy only)

Data Protection for SQL Server is a multi-threaded application that uses asynchronous execution threads to transfer data between the SQL Server and Tivoli Storage Manager server.

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 read data from an SQL Server while another sends 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 Server 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 Server database, should show an improvement over the unstriped case. The improvement is approximately proportional to the number of data stripes that are used, given the constraints of the devices and the network that is 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 Microsoft Management Console (MMC), 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 Server buffers, be certain that the number of SQL Server 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 MMC to modify the configuration file.

LAN-free data movement (Legacy and VSS)

Windows

Running Data Protection for SQL Server in a LAN-free environment means that data can be directly sent to storage devices.

When you implement a LAN-free environment, data bypasses potential network congestion. However, you must be properly equipped to operate in a LAN-free environment. For more information about setting up a LAN-free environment, see LAN-free client-data backup: Scenario (http://www.ibm.com/support/knowledgecenter/SSSQZW_7.1.1/com.ibm.itsm.sta.doc/c_scenario_lanfree.html).

In addition to specific LAN-free requirements, you must specify the following options.

- Use **lanfreetcpserveraddress** to specify the TCP/IP address for the Tivoli Storage Manager storage agent.
- 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 more information about these options, see Installing and configuring the client (https://www.ibm.com/support/knowledgecenter/SSSQZW_7.1.1/com.ibm.itsm.sta.doc/t_extlib_inst_client.html).

Chapter 8. Reference information

Reference information includes the backup and restore commands that you can issue at the command-line interface as an alternative to using Microsoft Management Console (MMC).

Command-line overview

The name of the Data Protection for SQL Server command-line interface is `tdpsqlc.exe`. This program is in the directory where Data Protection for SQL Server is installed.

Command-line parameter characteristics

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 example, 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
```

Data Protection for SQL Server uses the following command line syntax:

```
tdpsqlc command positional parameter 0 or more optional parameters
```

The command-line parameters have the following characteristics:


- Positional parameters do not include a leading slash (/) or dash (-).
- Optional parameters can display in any order after the required parameters.
- Optional parameters begin with a forward slash (/) or a dash (-).
- Minimum abbreviations for keywords are indicated in uppercase text.
- All SQL Server names of databases or parts of databases are case-sensitive.
- Some keyword parameters require a value.
- For those keyword parameters that require a value, the value is separated from the keyword with an equal sign (=).
- If a parameter requires more than one value after the equal sign, the values are separated with commas.
- Each parameter is separated from the others by using spaces.
- If a parameter value includes spaces, the value must be enclosed in double quotation marks.
- A positional parameter can display only once per command invocation.
 - 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 need to be contiguous. For example: /files=a /groups=y /files=b /groups=z

Where repeatable syntax exists, multiple values are separated with commas as shown in the following syntax diagram:



To select all instances of database names or file names on the server, follow the command with the wildcard asterisk (*).

Data Protection for SQL Server parameters by backup method

You can set different optional parameters for your VSS and legacy data backups.

The following table classifies the optional parameters that apply to VSS and legacy backups.

Table 22. Data Protection for SQL Server optional parameters

Optional Parameters	Legacy	VSS
/ACtive	Yes	Yes
/ADJUSTKBtsmestimate	Yes	No
/ADJUSTPERcenttsmestimate	Yes	No
/ALI	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

Table 22. Data Protection for SQL Server optional parameters (continued)

Optional Parameters	Legacy	VSS
/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
/REPlace	Yes	No
/SQLAUTHentication	Yes	Yes
/SQLBUFFers	Yes	No
/SQLBUFFERSize	Yes	No
/SQLCHECKSum	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

With Data Protection for SQL Server, you can back up and restore Microsoft SQL Server databases to Tivoli Storage Manager server storage by using the command-line interface or GUI.

Backup command

Use the **backup** command to back up all or part of one or more SQL Server 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, consider the following guidelines:

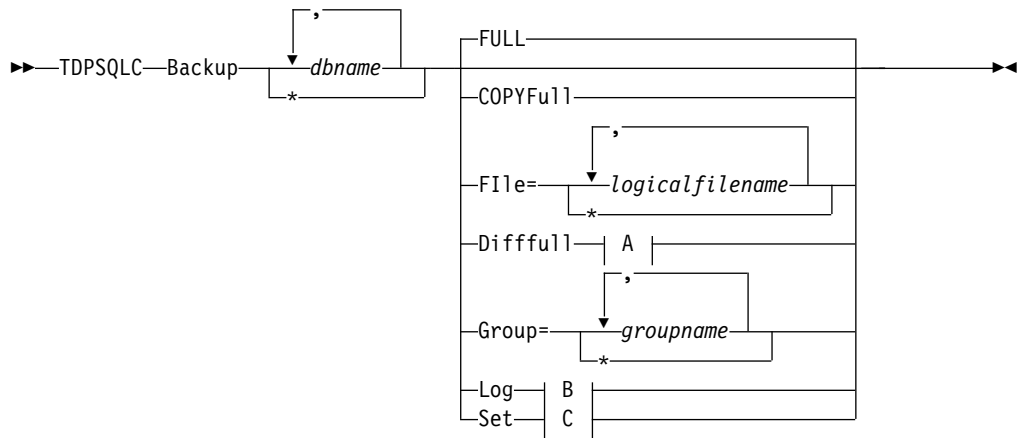
- Simple recovery model databases are automatically excluded from log backups.

- 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 that the server is started.
- Although full and differential backups include a part of the transaction log, that part is the only part that is required to make a restore operation consistent. The partial transaction log is not a log backup and does not truncate the log.
- The user id that is used by Data Protection for SQL Server 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 Server databases before you back up the databases.

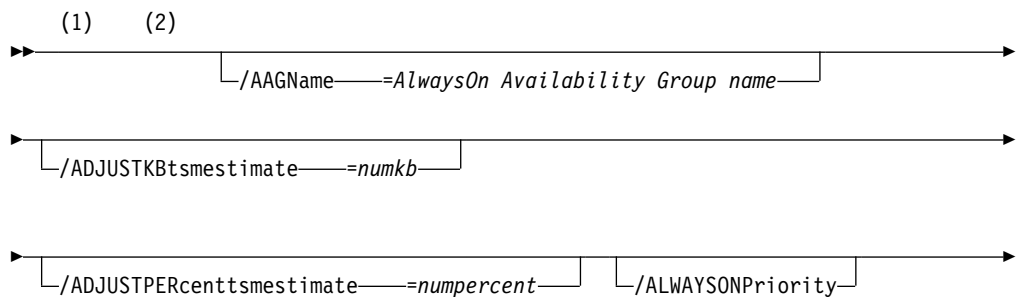
Backup syntax

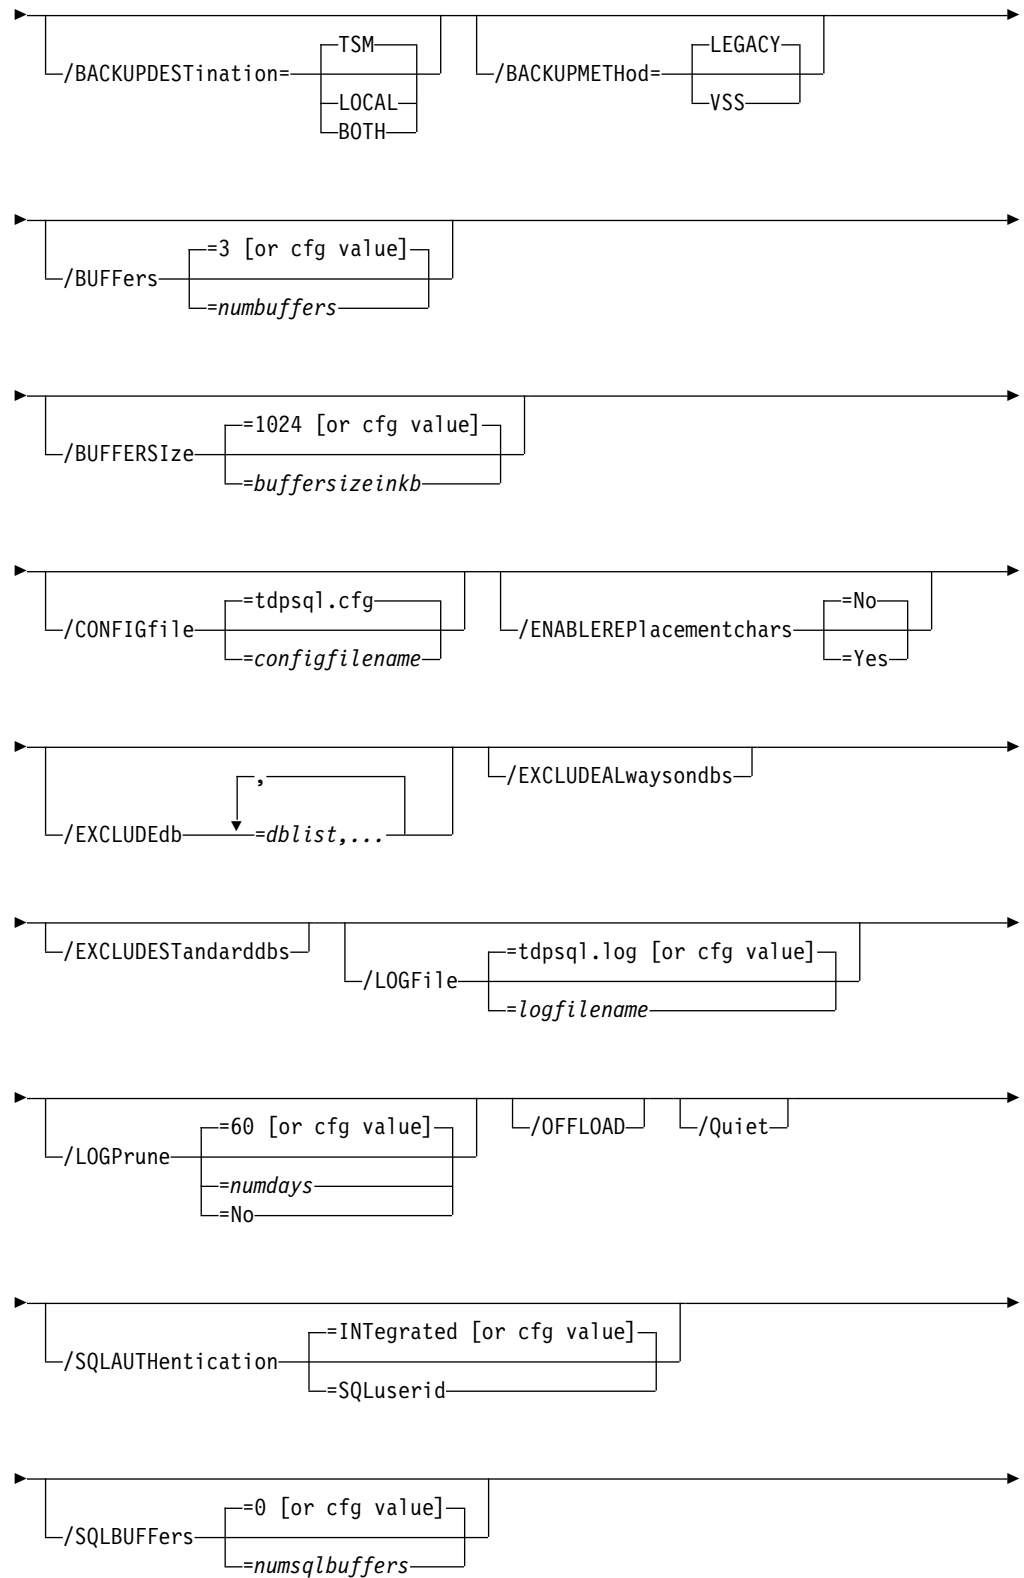
Use the **backup** command syntax diagrams as a reference to view available options and truncation requirements.

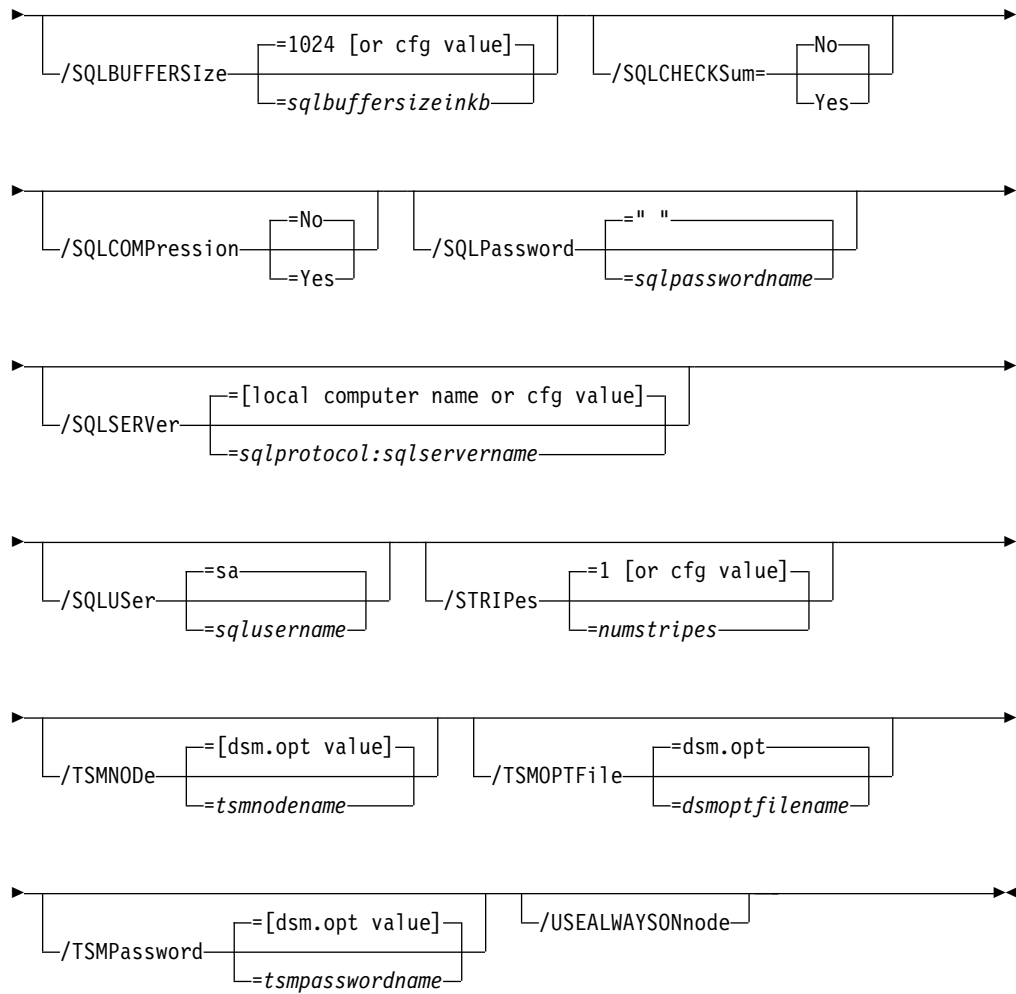
TDPSQLC backup




Backup optional parameters









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. You cannot issue the **/BACKUPMETHod=** parameter 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 data SQL Server logical file that you specify. You can use this option when it is not practical to back up an entire SQL Server 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 data files that you want to back up or restore to.

Considerations:

- For each SQL Server database that you back up, back up the corresponding transaction logs. The Data Protection for SQL Server log file, `tdpsql.log`, indicates the date and time of a database backup, the data that is backed up, and any error messages or completion codes.
- You can specify this parameter many times per command invocation.
- A new backup object deactivates any active backup object of the same name in the same SQL Server 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 that is 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 quotation marks.
- The *logicalfilename* variable is case-sensitive.

- You cannot specify the **/recovery** parameter with **restore file** operations.
- An 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 complete the unit backup. An SQL Server error message indicates which filegroups are affected. You must complete a full database backup or a set backup of at least the affected filegroups for the file backup to succeed.

FULL A **FULL** legacy database backup contains all of the contents of an SQL Server database and enough of the database transaction log to restore the database consistently. A **FULL** VSS database backup contains all of the contents of an SQL Server database (database files, log files, and full-text index files).

Each SQL Server database that is backed up by using the legacy backup method creates a separate backup object on the Tivoli Storage Manager server. A new full database backup object deactivates all prior legacy active backup objects for the same SQL Server database. This deactivation includes any active full backup object and any active file, group, set, differential, and log backup objects.

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 an SQL Server database that changed since the latest full backup and enough of the SQL Server database 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 does 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 Server database table or indexes are 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 Server 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 Server database's filegroups and that volume needs to be restored. You need restore only that filegroup for that SQL Server database.

The *groupname* variable specifies the names of the SQL Server database filegroups you want to back up.

Considerations:

- You can specify this parameter many times per command invocation.
- A new group backup object deactivates any active group backup object of the same name in the same SQL Server 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 quotation marks.
- The *groupname* variable is case-sensitive.
- You must follow group backups with transaction log backups for all SQL Server databases you back up.
- You cannot complete group backups for the following SQL Server databases:
 - Databases with the SQL Server attribute TRUNCATE LOG ON CHECKPOINT.
 - Databases that use the SIMPLE recovery model.
- You cannot specify the **/recovery** parameter with **restore group** operations.
- An 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 complete the unit backup. An SQL Server error message indicates which filegroups are affected. You must complete 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 Server 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 Server 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 option is the default.
- By using the **/truncate=no** parameter, you might be able to back up 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 deactivate 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 deactivate all active backup objects for an SQL Server database. Therefore, you can deactivate log backup objects by using the **inactivate** command if full database backups are not completed frequently or at all.

- You cannot complete log backups for the following SQL Server databases:
 - Databases with the SQL Server attribute TRUNCATE LOG ON CHECKPOINT.
 - Databases that use 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 Server 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, specifically regarding the Create Index operation.
- Each SQL Server database backed up creates a separate backup object on the Tivoli Storage Manager server. All of the files and filegroups that are backed up as part of a set backup for the same SQL Server database are contained in a single backup object. Group and file backups 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 deactivate 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 deactivate all active backup objects for an SQL Server database. Therefore, if full database backups are not performed or not performed completed, you can deactivate set backup objects by using the **inactivate** command.
- You must follow set backups with transaction log backups for all SQL Server databases you back up.
- The **file**, **group**, **log**, and **set** parameters can take a list of values (repeatable syntax) and might 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.

/ADJUSTPERcentsmestimate=numpercent

The **/adjustpercentsmestimate** 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 and later versions. 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 option 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, you cannot issue the **/BACKUPDESTination** parameter 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 backup is the

SQL Server 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, you cannot issue the **/BACKUPMETHod** parameter 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=bufferizeinkb

The **/buffersize** parameter specifies the size of each Data Protection for SQL Server buffer specified by the **/buffers** parameter. The *bufferizeinkb* 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 *bufferizeinkb*, 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*.

/ENABLEREPlacementchars=No|Yes

The **/ENABLEREPlacementchars** parameter enables SQL Server databases that have backslash (\) or colon (:) characters in the database name to be backed up. The maximum length of the database name is 128 characters. This parameter applies only to Data Protection for SQL Server version 7.1.1 and later versions.

You can specify the following values:

- Yes** Enable Data Protection for SQL Server to process backslash (\) or colon (:) characters in a database name, and back up the database to Tivoli Storage Manager. This value is the default.
- No** Prevent database backups to Tivoli Storage Manager if a user-defined string is substituted for a backslash (\) or colon (:) character in the database name.

/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 log file that you specify 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

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, a certain number of days of data are saved. By default, 60 days of log entries are saved. The option **No** can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify **no**, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, 60, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

/MOUNTWait=Yes | No

The **/mountwait** parameter is used to specify whether Data Protection for SQL Server waits for removable media to mount, such as tapes or CDs, or stops the current operation. This parameter is not valid for all backup types; the parameter 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. You cannot specify this parameter 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 Server 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 Server determines how many buffers should be used.

Considerations:

- The default value is the value specified by the SQL Server 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 Server 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.

/SQLCHECKSum=No | Yes

The **/SQLCHECKSum** parameter is used to verify the integrity of a legacy database backup. Integrity checking is a process that validates the values in a file or configuration for unexpected changes. Values are verified between the current state and the baseline state.

You can specify the following values:

No Do not enable integrity checking for a legacy database backup. This value is the default.

Yes Enable integrity checking for a legacy database backup.

In the Performance Properties window of Microsoft Management Console, you can enable or disable the checksum option for all your legacy databases at once. You can override the global setting, and temporarily enable or disable the checksum option for a database backup, by setting this **SQLCHECKSum** parameter value to **Yes** or **No**.

/SQLCOMPression=No | Yes

The **/SQLCOMPression** parameter specifies whether SQL compression is applied. If you do not specify **/SQLCOMPression**, the **No** default value is used. This parameter is only applicable on systems running SQL Server 2008 and later versions. For SQL Server 2008, you can run backup compression only on the Enterprise Edition. For SQL Server 2008 R2, you can run backup compression on Standard, Enterprise, and Datacenter editions.

/SQLPassword=sqlpasswordname

This parameter specifies the SQL Server 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 Server 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, port

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. With this variable, you can specify an *sqlservername* and, in the case of the *tcp* protocol, you can also define a port. You can choose from the following protocols:

Table 23.

Protocol Name	Description	Example Usage
lpc	Shared Memory	/sqlserver=lpc: <servername>\<instancename>
np	Named Pipes	/sqlserver=np:<servername>\pipe Note: You can optionally specify a specific named pipe. For example, /sqlserver=np: \\hostname\pipe\pipe name By default, the pipe name is <i>sql\query</i> . If you connect to a named instance, the pipe name is typically in the following format: \\<servername>\pipe\ MSSQL\$<instancename>\SQL\ query
tcp	Transmission Control	/sqlserver=tcp:<servername>[\<instancename>],<port> Note: For tcp protocols only, you have the option of defining a <i>port</i> . If you do not define a port, the default port value is 1443.
via Note: SQL Server only supports this protocol through SQL Server 2008 R2.	Virtual Interface Architecture	/sqlserver=via:<virtualservername>[\instancename]

If you do not specify a protocol, 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) That is, 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?
<i>local-computername</i>	default	no	no
<i>local-computername\instancename</i>	named	no	yes
<i>virtualservername</i>	default	yes	yes
<i>virtualservername\instancename</i>	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 Server user id for this password must both be configured for SQL Server authentication.
- The SQL Server 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 *dsmerv.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 more information about the Tivoli Storage Manager options file, see Creating and modifying the client system-options file (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.4/client/t_cfg_crtmodoptunix.html).
- 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 more information about the Tivoli Storage Manager options file, see Creating and modifying the client system-options file (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.4/client/t_cfg_crtmodoptunix.html).

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	<i>generate</i>	yes	<i>/tsmpassword</i> ignored

/tsmpassword	PASSWORDACCESS in Tivoli Storage Manager options file	Password already stored in registry?	Result
specified	<i>generate</i>	no	<i>/tsmpassword</i> used and stored
specified	<i>prompt</i>	—	<i>/tsmpassword</i> used
not specified	<i>prompt</i>	—	user is prompted

/USEALWAYSOnode

Specify this parameter to back up standard databases on SQL Server 2012 and later versions 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 and later version availability databases are backed up to the AlwaysOn node.

Legacy backup examples

The following examples are provided to show how you can issue the **backup** command with various parameters and options.

Full backup examples

If you want to complete a full backup, the following examples are provided to help you to 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 /EXCLUDEAlwaysondb
```
- To complete a log backup of all availability databases, enter the following command:

```
tdpsqlc backup * log /EXCLUDEStandarddb
```
- 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 diff full /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 * diff full /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 /EXCLUDEAlwaysondb
```

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 an SQL Server 2012 AlwaysOn Availability Group environment, and later versions, enter the following command:

```
tdpsqlc backup filestreamdb CopyFull /backupdestination=TSM  
/backupmethod=legacy
```

VSS backup examples

The following examples are provided to show you how you can issue the **backup** command 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 you 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 an SQL Server 2012 availability database and later versions (*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

To change the Tivoli Storage Manager password that is used by Data Protection for SQL Server, use the **changetsmpassword** command. The password is used to log on to the Tivoli Storage Manager server.


Changetsmpassword

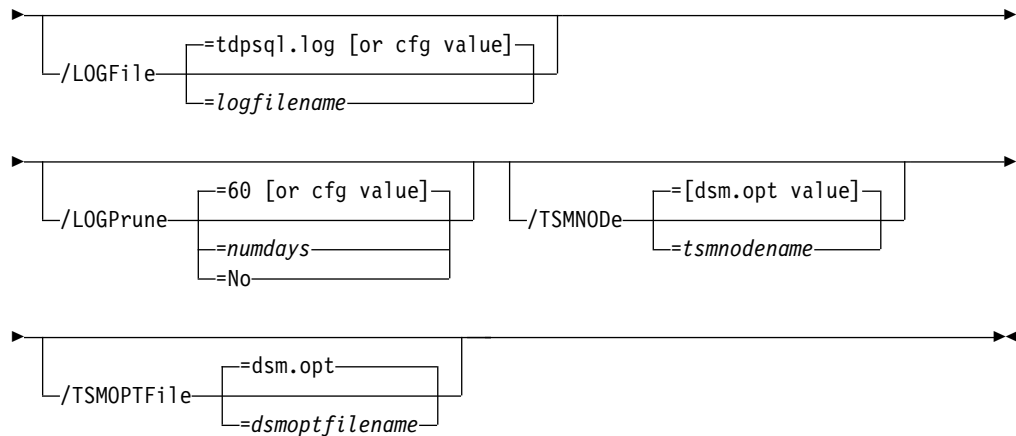
Use the **changetsmpassword** command syntax diagrams as a reference to view available options and truncation requirements.

TDPSQLC CHANGETSMPassword command



Optional Parameters





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 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*.

- If you specify **/configfile** but not *configfilename*, the default value *tdpsql.cfg* is used.

/LOGPrune=numdays | No

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, a certain number of days of data are saved. By default, 60 days of log entries are saved. The option **No** can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify **no**, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, 60, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

/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.

/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 more information about the Tivoli Storage Manager options file, see *Creating and modifying the client system-options file* (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.4/client/t_cfg_crtmodoptunix.html).
- 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 more information about the Tivoli Storage Manager options file, see *Creating and modifying the client system-options file* (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.4/client/t_cfg_crtmodoptunix.html).

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*.

Related reference:

“Set positional parameters” on page 264

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 changetsm smp sqlv2old sqlv2new sqlv2new
```

Output:

```
IBM Tivoli Storage Manager for Databases:  
Data Protection for Microsoft SQL Server  
Version 7, Release 1, Level 3.0  
(C) Copyright IBM Corporation 1998, 2015.  
All rights reserved.  
  
AC00260I Password successfully changed.
```

Delete Backup command

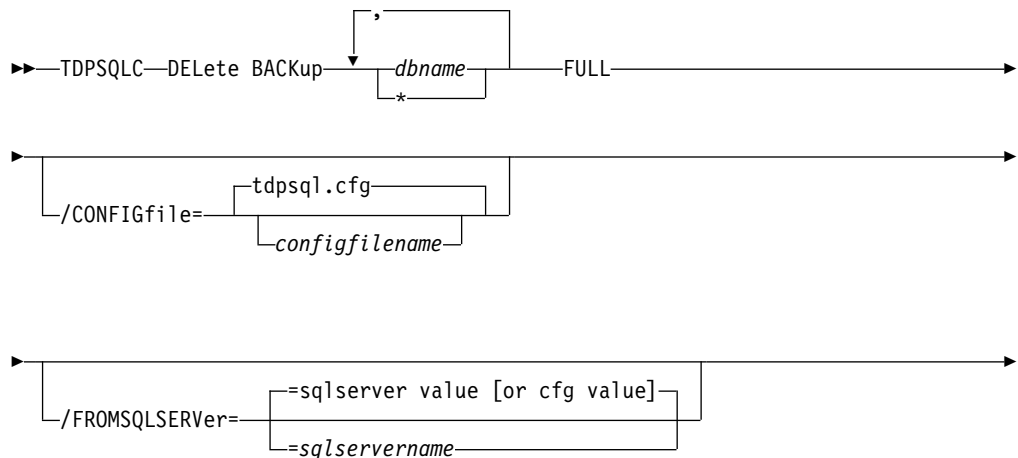
Use the **delete backup** command to delete a VSS backup of an SQL Server database.


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





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 blanks, enclose the database name in double quotation marks.

The following positional parameter specifies the type of delete backup to run:

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 option 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"
```

/FROMSQLSERVER=server-name

Use the **/fromsqlserver** parameter to specify the name of the SQL Server where the original backup is completed. 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.

- 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 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\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 you are using multiple simultaneous instances of Data Protection for SQL Server to complete operations, use the **/logfile** parameter to specify a different log file for each instance used. Logging for each instance is directed to a different log file, which prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPrune=numdays | No

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, a certain number of days of data are saved. By default, 60 days of log entries are saved. The option **No** can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.

- If you specify *no*, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is *60*.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, *60*, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

/OBJECT=objectname,...

Use the **/object** parameter to specify the names of backup objects that 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 Server databases and backup object type are deleted. The *objectname* variable specifies the names of the backup objects that 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 parameter does not affect the level of information that is 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 that 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 is not yet stored in the registry. In that case, the specified password is stored in the registry and used when you run this command.

If PASSWORDACCESS PROMPT is in effect, and you do not specify a password value on the command line, then you are prompted for a password.

The Tivoli Storage Manager password that Data Protection for SQL Server uses to log on to the Tivoli Storage Manager server can be up to 63 characters in length.

Related reference:

“Set positional parameters” on page 264

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 Server 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)

Tip: 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 Server 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 3.0
(C) Copyright IBM Corporation 1998, 2015.
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 3.0
 (C) Copyright IBM Corporation 1998, 2015.
 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=bufferizeinkb]             default: 1024 (or cfg value)
[/CONFIGFILE=configfilename]            default: tdpsql.cfg
[/DBONLY]
[/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]
[/RECOVER=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=numsqllibuffers]           default: 0 (or cfg value)
[/SQLBUFFERSIZE=sqlbufferizeinkb]        default: 1024 (or cfg value)
[/SQLPASSWORD=sqlpasswordname]          default: " "
[/SQLSERVER=[sqlprotocol:]sqlservername]
                                         default: local computer name (or cfg value)
                                         default sqlprotocol: "" (or cfg value)
                                         default: sa
[/SQLUSER=sqlusername]
[/STANDBY=undofile]
[/STRIPES=numstripes]                   default: 1 (or cfg value)
[/TSMNODE=tsmnodename]                   default: dsm.opt value
[/TSMOPTFILE=dsmoptfilename]             default: dsm.opt
[/TSMPPASSWORD=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 3.0
 (C) Copyright IBM Corporation 1998, 2015.
 All rights reserved.

```

TDPSQLC Restore *|dbname[,dbname,...] Log=*|logobjectname[,logobjectname,...]
[/BUFFers=numbuffers]           default: 3 (or cfg value)
[/BUFFERSize=bufferizeinkb]      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=sqlbufferizeinkb] 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=undofilefilename]
[/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 3.0
(C) Copyright IBM Corporation 1998, 2015.
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

REMOVEDSMAgentnode=nodename

SQLAUTHentication=[INTEgrated|SQLuserid]

SQLBUFFers=numsqbuffers (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 deactivate one or more active legacy backup objects on the Tivoli Storage Manager server.

Most backup objects are automatically deactivated as part of the normally scheduled backup processing. For those occasions when that processing is insufficient, you can issue 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

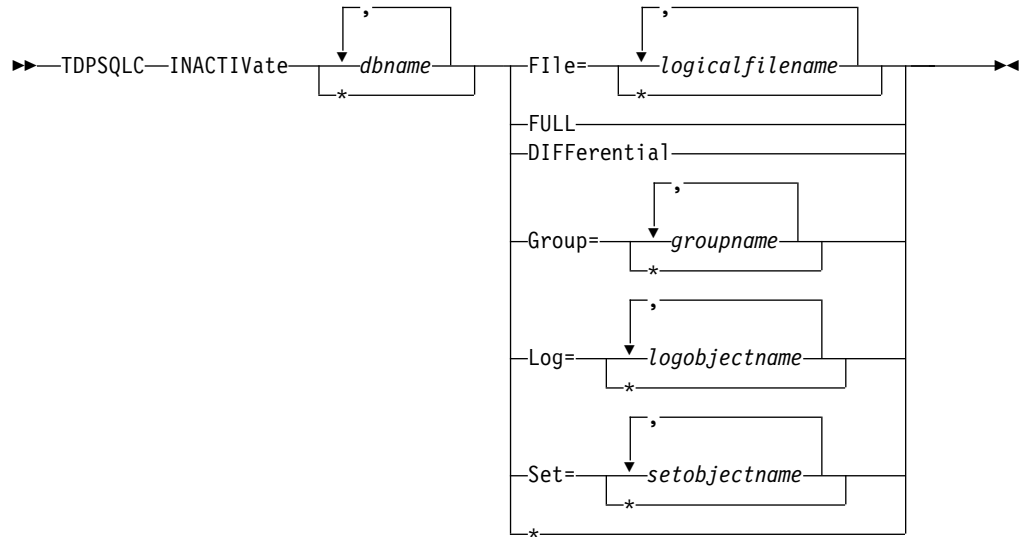
a backup object becomes inactive, the expiration processing defined in the object's management class determines exactly when the backup object is deleted.

Inactivate syntax


Use the **inactivate** command syntax diagrams as a reference to view available options and truncation requirements.

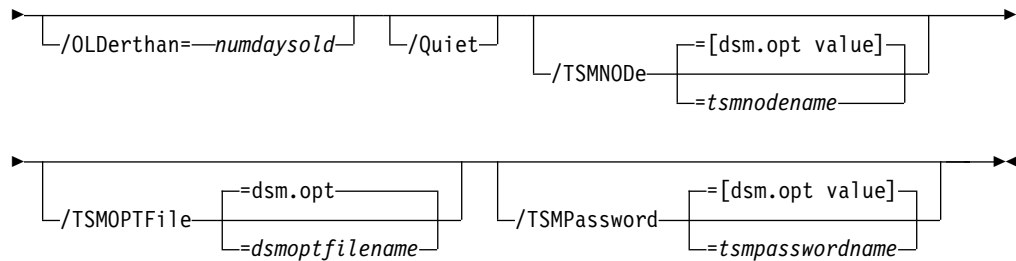
Syntax

TDPSQLC command



Inactivate Optional Parameters:





Inactivate positional parameters

Positional parameters immediately follow the **inactivate** command and precede the optional parameters.

File=* | *logicalfilename*,...

This option deactivates only the active file backup objects for the SQL Server databases you specify. The *logicalfilename* variable specifies the names of the SQL Server database logical files you want to deactivate.

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 deactivates only the active full database backup objects for the SQL Server databases you specify. Each SQL Server database backed up creates a separate backup object on the Tivoli Storage Manager server. A new full database backup object deactivates all prior active backup objects for the same SQL Server 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 deactivates only the active differential database backup object. Because each SQL Server database backup creates a separate backup object on the Tivoli Storage Manager server, a new differential database backup object deactivates any active differential backup object for the same SQL Server 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 deactivates only the active group database backup object for the SQL Server database you specify. The *groupname* variable specifies the names of the SQL Server database filegroups you want to deactivate.

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 deactivates only the active log database backup object for the SQL Server database you specify. This parameter takes the wildcard or *logobjectname* value. The *logobjectname* variable specifies the log backup objects to deactivate. 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 Server databases. You can specify this parameter more than once per command invocation.

Set or Set=* | *setobjectname*,...

This option deactivates only the active set database backup object for the SQL Server database you specify. This parameter takes the wildcard or *setobjectname* value. The *setobjectname* variable specifies the set backup objects to deactivate. 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 Server 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 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*.
- 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 deactivate 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. If the two SQL Server names are different, you must use this parameter even if **/fromsqlserver** was a non-clustered default instance.

/LOGFile=*logfile*

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 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

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, a certain number of days of data are saved. By default, 60 days of log entries are saved. The option **No** can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify **no**, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, 60, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

/OBJect=* | objectname,...

This parameter specifies that only particular backup objects for the specified SQL Server databases and backup object type (if specified) be deactivated. The *objectname* variable specifies the names of the backup objects you want to 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. 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 Server databases and backup object type.

/OLDERthan=numdaysold

This parameter specifies how old a backup object must be before the command can deactivate it.

Considerations:

- The *numdaysold* variable can range from 0 to 9999.
- If you specify 0, you deactivate all selected backup objects.
- If you specify 1, you deactivate 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 Server activity log.

/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 more information about the Tivoli Storage Manager options file, see Creating and modifying the client system-options file (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.4/client/t_cfg_crtmodoptunix.html).
- 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 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	<i>generate</i>	yes	<i>/tsmpassword</i> ignored
specified	<i>generate</i>	no	<i>/tsmpassword</i> used and stored
specified	<i>prompt</i>	—	<i>/tsmpassword</i> used
not specified	<i>prompt</i>	—	user is prompted

Related reference:

“Set command” on page 263

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 deactivates database backup objects. Once a backup object is deactivated, 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 7, Release 1, Level 3.0
(C) Copyright IBM Corporation 1998, 2015.
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 ..... 03/23/2014 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 ..... 20140323063104\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 ..... 03/20/2014 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 ..... 20140320053514\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 ..... 03/19/2014 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 ..... 20140319070139\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 ..... 03/23/2014 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 ..... 20140323063105\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 ..... 03/19/2014 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 ..... 20140319112615\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 ..... 03/17/2014 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 ..... 20140317011548\00001AC4  
Number of stripes in backup object ..... 1  
Assigned Management Class ..... DEFAULT
```

Completed

The user then decides to deactivate all *DB1_XIVmini_G_BAS* database objects older than two days (older than March 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 7, Release 1, Level 3.0
(C) Copyright IBM Corporation 1998, 2015.
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\20140320053514\00001AC4

Inactivating full backup DB1_XIVmini_G_BAS
Inactivating log backup DB1_XIVmini_G_BAS\20140319070139\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 7, Release 1, Level 3.0
(C) Copyright IBM Corporation 1998, 2015.
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 ..... 03/20/2014 05:35:14
Backup Size ..... 3,349
Database Object Name ..... 20140320053514\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 ..... 03/19/2014 07:01:39
Backup Size ..... 3,349
Database Object Name ..... 20140320053514\00001AC4
Number of stripes in backup object ..... 1

```

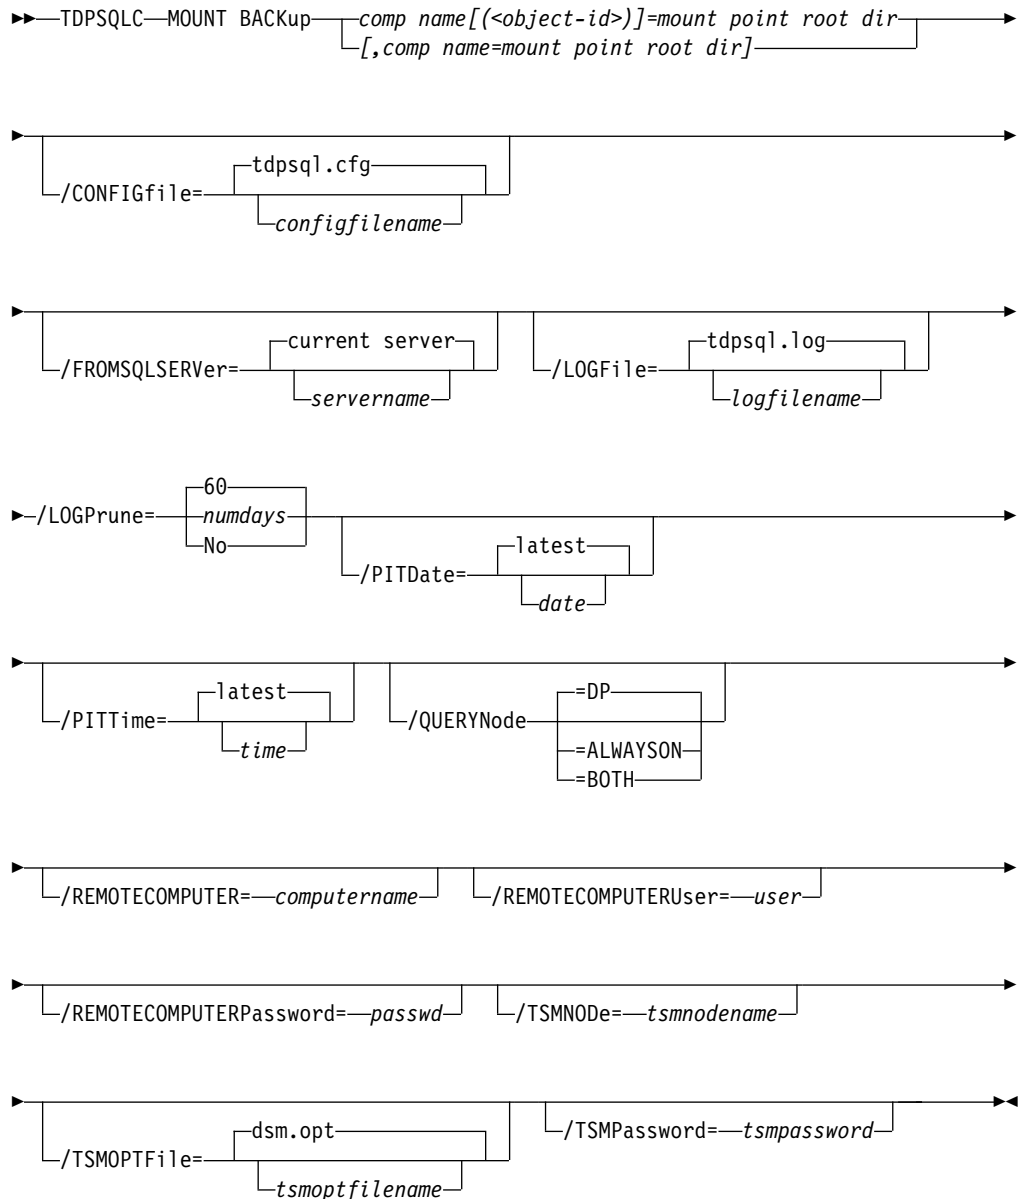
Mount Backup command

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 Server configuration file that contains the values to use for a **mount backup** operation.

The *configfilename* variable can include a fully qualified path. If the *configfilename* variable does not include a path, the Tivoli Storage FlashCopy Manager 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\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 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\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**

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, a certain number of days of data are saved. By default, 60 days of log entries are saved. The option **No** can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify *no*, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, 60, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

/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 Server 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 Server 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, and later versions, 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=password

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.

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 Server.

Considerations:

- Some of the information displays may have long text lines. You can redirect the informational output of the Data Protection for SQL Server 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

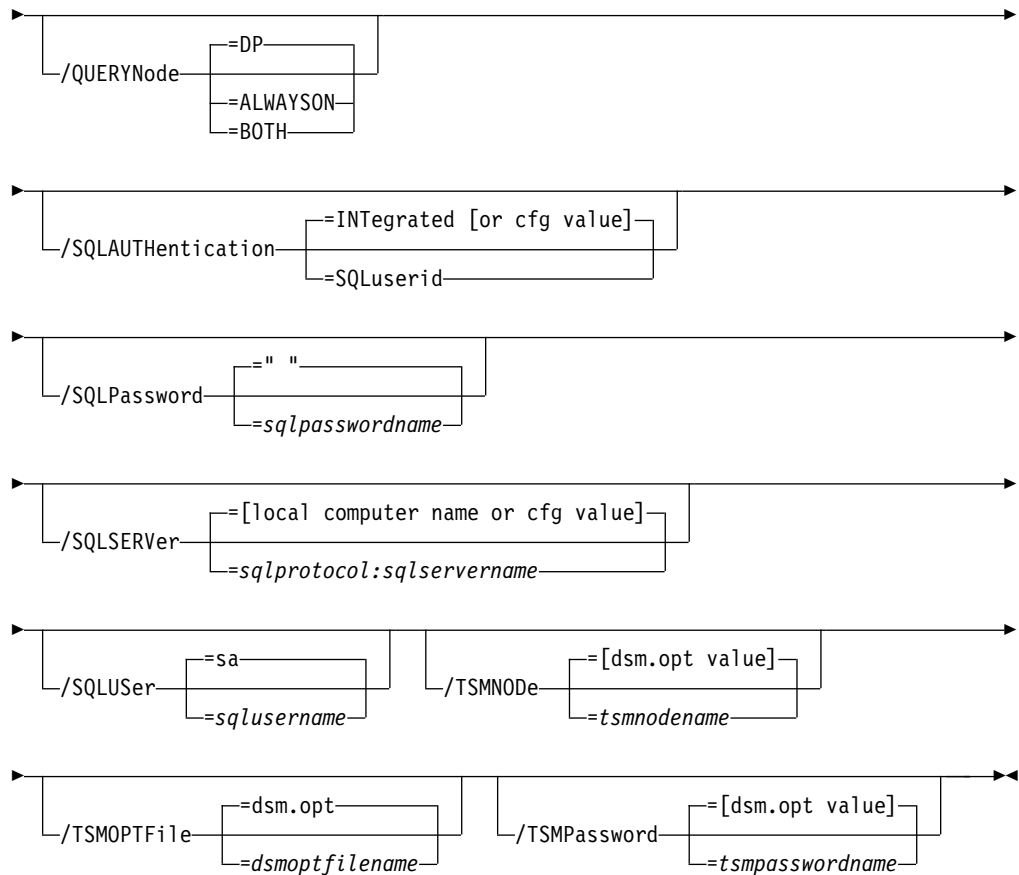
TDPSQLC SQL query



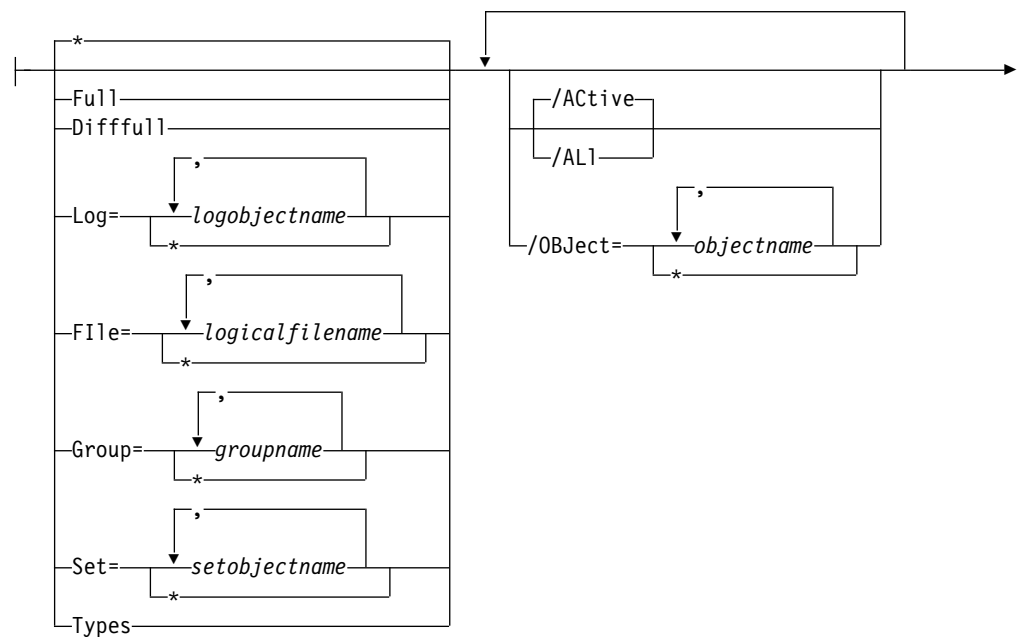
The syntax diagram of the Tivoli Storage Manager options corresponding to the letter A is shown following the Optional Parameters below.

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 you issue a Data Protection for SQL Server **query** command:

Query SQL *|*dbname*,...

This parameter displays information about the current SQL Server. The *dbname* variable specifies databases on the current SQL Server to display information about.

When you query an SQL Server, the following information is included:

- Server name
- Database name
- Database data space allocated
- Database space that is used
- Database log space allocated
- Database log space used
- Database options set (SELECT INTO / BULK COPY, TRUNCATE LOG ON CHECKPOINT)

If you specify the **/COMPATibilityinfo** parameter:

- Server version
- Server clustering state
- Database compatibility level

Query TDP

This parameter displays the Data Protection for SQL Server name and version information and the contents of the current Data Protection for SQL Server configuration file.

Query TSM *|*dbname*,...

This parameter 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 Server database. Use the *dbname* wildcard option to display information about all of the backup objects of one or more SQL Server databases.

When you query any backup object by using *dbname* parameter, the following information is included:

- SQL Server name
- SQL Server database name
- Backup object type
- Backup object active/inactive state
- Backup object Data Protection for SQL Server creation date and time
- Backup object Data Protection for SQL Server size

- Data Protection for SQL Server backup-object object name
- For legacy backups, whether the integrity of database and log files is verified by checksum processing
- 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 the **/COMPATibilityinfo** parameter:

- SQL Server version
- SQL Server clustering state
- Data Protection for SQL Server version that created the backup object
- SQL Server database compatibility level
- SQL Server database data space allocated
- SQL Server database data space used
- SQL Server database log space allocated
- SQL Server database log space used
- SQL Server database options

Note:

- You can also determine which backup objects to display through the **query TSM** optional parameters **/active** and **/all**.
- No information is displayed if no backup objects for a specified SQL Server database exist.

File=* | *logicalfilename,...*

This parameter displays information about file backup objects of one or more SQL Server databases from the current SQL Server that are on the current Tivoli Storage Manager server and node.

Full This parameter displays information about full backup objects of one or more SQL Server databases from the current SQL Server that are on the current Tivoli Storage Manager server and node.

Difffull

This parameter displays information about differential backup objects of one or more SQL Server databases from the current SQL Server that are on the current Tivoli Storage Manager server and node.

Group=* | *groupname,...*

This parameter displays information about one or more group backup objects of one or more SQL Server databases from the current SQL Server that are on the current Tivoli Storage Manager server and node.

Log=* | *logobjectname,...*

This parameter displays information about one or more log backup

objects of one or more SQL Server 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 Server databases.

Set=* | *setobjectname*,...

This parameter displays information about one or more set backup objects of one or more SQL Server 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 Server databases.

Types (Legacy backups only) This parameter displays a summary list of the backup objects by backup type, of one or more SQL Server 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 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.

/BUFFERSize=*buffer size in kb*

The **/buffersize** parameter specifies the size of each Data Protection for SQL Server buffer specified by the **/buffers** parameter. The *buffer size in kb* 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.

/COMPATibilityinfo

For **query** operations, this parameter displays information related to the compatibility of a backup object with an SQL Server. Certain SQL Server configuration options must be compatible before you can restore a backup object to an SQL Server. When you specify this parameter, Data Protection for SQL Server configuration information is listed to help determine if a backup object is correct for an 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 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*.
- 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 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.
- After you restore an SQL Server database to a different SQL Server, the logins of the SQL Server 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 Server login mismatches.

/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

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, a certain number of days of data are saved. By default, 60 days of log entries are saved. The option **No** can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify **no**, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, 60, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

/OBJect=*|*objectname*,...

For **restore** and **inactivate** operations, **/object** specifies that only particular backup objects for the specified SQL Server 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 Server. 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 Server databases and backup object type.

/QUERYNode=**DP** | **ALWAYSON** | **BOTH**

Specify whether you want to query standard databases from SQL Server 2012 and later versions that are 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 Server user id must have the SQL Server SYSADMIN fixed server role.

/SQLPassword=*sqlpasswordname*

This parameter specifies the SQL Server 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 Server 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, port*

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. With this variable, you can specify an *sqlservername* and, in the case of the *tcp* protocol, you can also define a port. You can choose from the following protocols:

Table 24.

Protocol Name	Description	Example Usage
lpc	Shared Memory	/sqlserver=lpc: <servername>\<instancename>

Table 24. (continued)

Protocol Name	Description	Example Usage
np	Named Pipes	<p>/sqlserver=np:<servername>\pipe</p> <p>Note: You can optionally specify a specific named pipe. For example, /sqlserver=np:\hostname\pipe\pipe name</p> <p>By default, the pipe name is <i>sql\query</i>. If you connect to a named instance, the pipe name is typically in the following format: \\<servername>\pipe\MSSQL\$<instancename>\SQL\query</p>
tcp	Transmission Control	<p>/sqlserver=tcp:<servername>[\<instancename>],<port></p> <p>Note: For tcp protocols only, you have the option of defining a <i>port</i>. If you do not define a port, the default port value is 1443.</p>
via Note: SQL Server only supports this protocol through SQL Server 2008 R2.	Virtual Interface Architecture	/sqlserver=via:<virtualservername>[\instancename]

If you do not specify a protocol, 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) That is, 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?
<i>local-computername</i>	default	no	no
<i>local-computername\instancename</i>	named	no	yes
<i>virtualservername</i>	default	yes	yes
<i>virtualservername\instancename</i>	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 Server user id for this password must both be configured for SQL Server authentication.
- The SQL Server 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 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 more information about the Tivoli Storage Manager options file, see Creating and modifying the client system-options file (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.4/client/t_cfg_crtmodoptunix.html).
- 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 more information about the Tivoli Storage Manager options file, see Creating and modifying the client system-options file (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.4/client/t_cfg_crtmodoptunix.html).

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	<i>generate</i>	yes	<i>/tsmpassword</i> ignored
specified	<i>generate</i>	no	<i>/tsmpassword</i> used and stored
specified	<i>prompt</i>	—	<i>/tsmpassword</i> used
not specified	<i>prompt</i>	—	user is prompted

Related reference:

“Set command” on page 263

“Set command” on page 263

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 3.0
(C) Copyright IBM Corporation 1997, 2015.
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 3.0
(C) Copyright IBM Corporation 1997, 2015.
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 Server 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 3.0
(C) Copyright IBM Corporation 1997, 2015.
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 3.0
(C) Copyright IBM Corporation 1997, 2015.
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 3.0
 (C) Copyright IBM Corporation 1997, 2015.
 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 3.0
(C) Copyright IBM Corporation 1997, 2015.
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 3.0
(C) Copyright IBM Corporation 1997, 2015.
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 3.0
 (C) Copyright IBM Corporation 1997, 2015.
 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.1df
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 3.0
(C) Copyright IBM Corporation 1997, 2015. 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 Server databases.

Use this command to restore all or part of one or more SQL Server databases from Tivoli Storage Manager storage to an SQL Server.

- You cannot restore SQL Server databases that are in use. By placing SQL Server 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 Server databases or server must be the same user that Data Protection for SQL Server 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 Server 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 Server databases.

During SQL Server database restore processing, the SQL Server prepares the database files after first restoring a minimal amount of metadata. For large SQL Server 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):

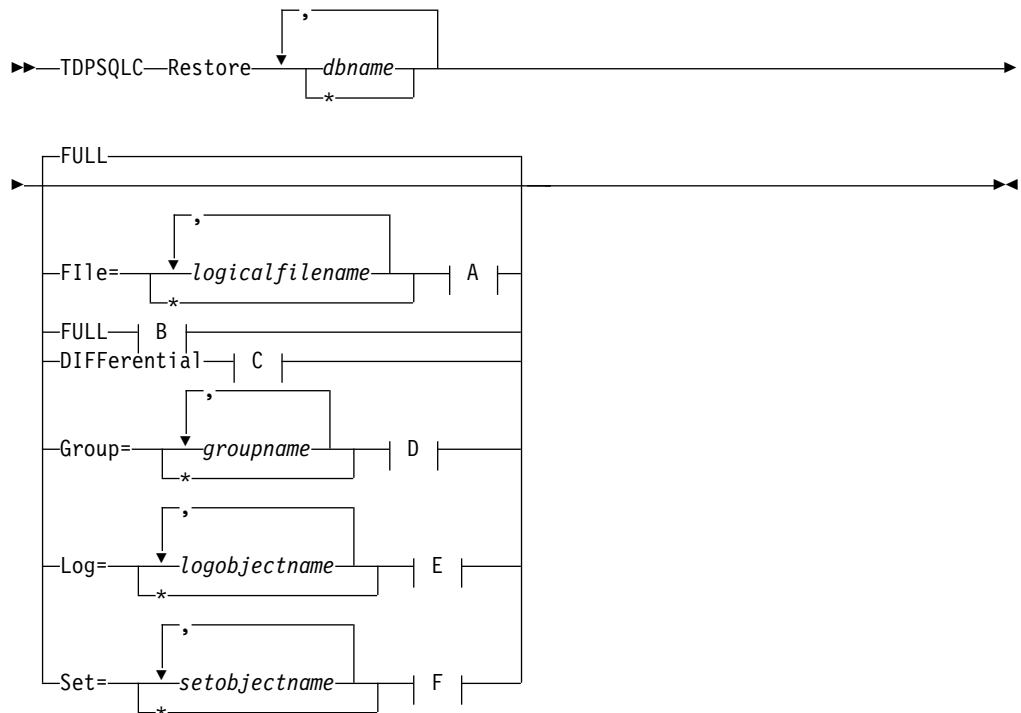
- A VSS instant restore overwrites the entire contents of the source volumes. However, you can avoid overwriting the source volumes by specifying **/instantrestore=no**. This parameter setting bypasses volume-level copy and uses file-level copy instead to restore the files from a VSS backup that resides on local shadow volumes. The source volume contains only the SQL Server 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.

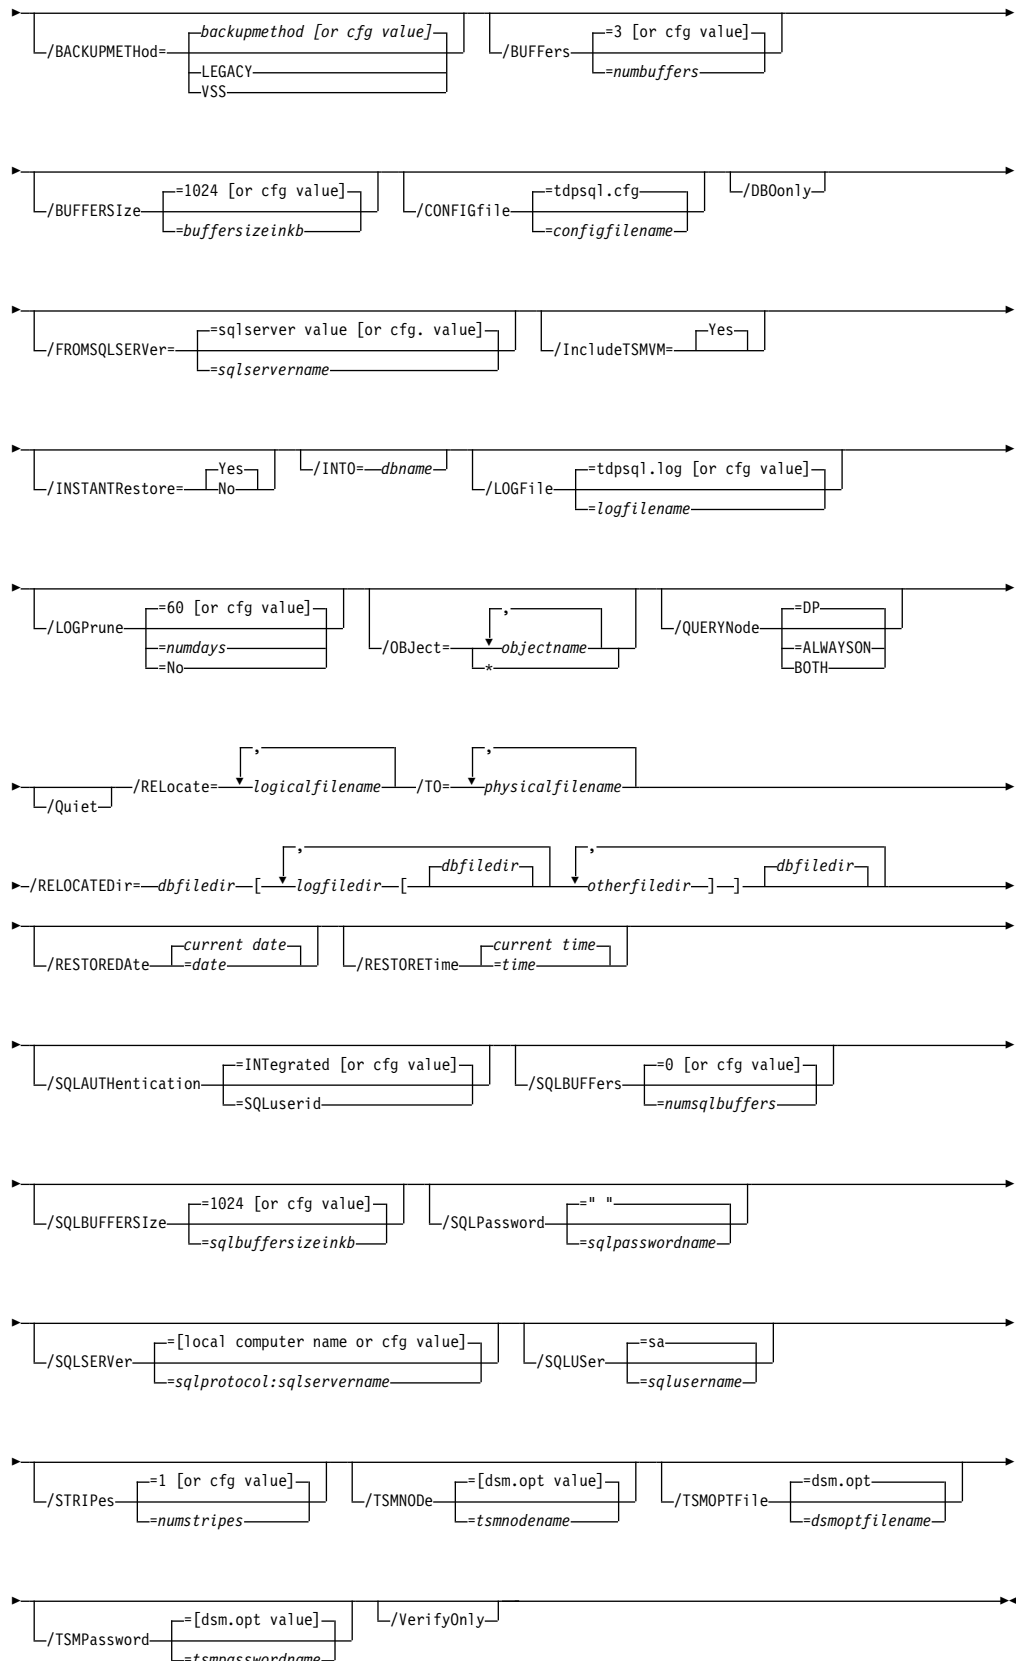
Syntax



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:






A Restore File Options:



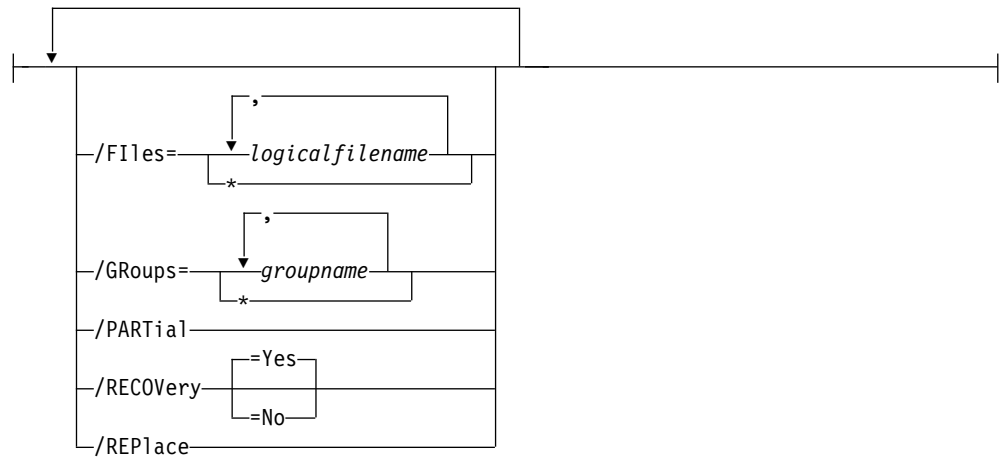
B Restore Full Options:



B1 Restore Full Options 1:



B2 Restore Full Options 2:



C Restore Diff Options:



The diagram illustrates the structure of the `/REplace` command. It shows a sequence of arguments: `/Files=`, `logicalfilename`, and `*`. A comma (,) is placed between `logicalfilename` and `*`. The entire sequence is enclosed in a box, with a vertical line on the left and a horizontal line on the right. A vertical line also separates the `/Files=` argument from the rest of the sequence.

```

graph TD
    Start(( )) --> Recover[RECOVER]
    Recover --> RecoverYes{=Yes}
    RecoverYes --> Standby[STANDBY=undofilename]
    RecoverYes --> Stopat[STOPAT=datetime]
    RecoverYes --> Stopatmark[STOPATMark=markname]
    RecoverYes --> Stopbeforemark[STOPBEFOREMark=markname]
    Standby --> After1[After=datetime]
    Stopat --> After1
    Stopatmark --> After2[After=datetime]
    Stopbeforemark --> After2
    After1 --> End(( ))
    After2 --> End
  
```

The flowchart illustrates the logic for the RECOVER command. It starts with a decision point: if the condition is "Yes", it proceeds to execute STANDBY=undofilename, STOPAT=datetime, STOPATMark=markname, and STOPBEFOREMark=markname. If the condition is "No", it proceeds to execute After=datetime. Both paths eventually lead to the end of the process.

```

graph TD
    Start(( )) --> Loop
    subgraph Loop [ ]
        direction TB
        FileNames["/Files="]
        Groups["/Groups="]
        Replace["/REPlace"]
        FileNames --> SetFile["setFile logicalfilename"]
        Groups --> SetGroup["setGroup groupname"]
        SetFile --> SetGroup
        SetGroup --> Replace
    end
    Replace --> End(( ))

```

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 Server database due to available backup time and space or due to performance requirements. This option restores file backup objects for the SQL Server databases you specify. The *logicalfilename* variable specifies the names of the SQL Server database logical files you want to restore to.

- 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 Server 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 Server 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 Server 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 Server 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 Server 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 Server

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 Server 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 option 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 completed. 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 operation is completed with the legacy API. This option is the default if no value is specified in the Data Protection for SQL Server preferences file (tdpsql.cfg).

VSS The restore operation is completed 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 also increase storage use. Each buffer is the size that is specified in the **/BUFFERSize** parameter.
- The default value is the value that is specified by the buffers configurable option in the Data Protection for SQL Server configuration file. This value is initially 3.

- If you specify **/BUFFers**, its value is used instead of the value that is 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=bufferizeinkb

The **/BUFFERSize** parameter specifies the size of each Data Protection for SQL Server buffer that is specified by the **/BUFFers** parameter. The *bufferizeinkb* variable refers to the size of data buffers in KB. The number can be in the range 64 - 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 that is specified by the buffers configurable option in the Data Protection for SQL Server configuration file. This value is initially 1024.
- If you specify **/BUFFERSize**, its value is used instead of the value that is 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 *bufferizeinkb*, 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 includes 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, the XIV system, or IBM Storwize V7000 storage system 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 Server database files and log files are the only data overwritten on the source volumes.

When you are running VSS instant restore on IBM System Storage DS8000 series and Storwize family, ensure that any previous background copies that involve the volumes you are restoring, complete before you initiate the VSS instant restore.

/IncludeTSMVM=Yes

Set the **/IncludeTSMVM** to Yes to view all database backups that includes Virtual Environment backup SQL Server 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, the XIV system, or IBM Storwize V7000 storage system 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 Server database files and log files are the only data overwritten on the source volumes.

When you are running VSS instant restore on IBM System Storage DS8000 series and Storwize family, 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

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, a certain number of days of data are saved. By default, 60 days of log entries are saved. The option **No** can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify **no**, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, 60, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat**

- or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
- Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

/MOUNTWait=Yes|No

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 might eventually be either waiting for its own or another stripe's media to become available. In this case, it might become necessary to end the Data Protection for SQL Server command from a prolonged wait. This action can be completed by closing 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 ends the command with an error message. The same outcome happens 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 value is initially yes. Specifying this parameter does not change the value in the configuration file.
- If you specify **/MOUNTWait**, but do not specify either yes, or 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 Server 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 Server databases and backup object type.

/PARTial

The **/PARTial** parameter restores only part of an SQL Server database. You can complete 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 Server database. After the partial restore, differential database restores and transaction log restores can return the subset SQL Server database to a point where the required data exists or is undamaged. You can then copy the required data from the subset SQL Server database to the original SQL Server database. You can also use partial restores whenever you need a subset of an SQL Server 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 might be used to complete a recovery. The statistics that are displayed reflect the amount of data that is restored from the Tivoli Storage Manager server only, not the amount of data that is 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 Server groups, even if you did not specify all SQL Server files in an SQL Server group with the files option.
 - The primary group is always included.
 - SQL Server groups that are not restored are marked offline and are not accessible.
- If you are restoring the subset SQL Server database to a location where it was backed up, you must use the **/RELocate** and **/to** parameters.
- Microsoft Management Console (MMC) does not support the **/RELocate** and **/to** parameters. You must use the command-line interface when you are completing 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 and later versions that are 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 you want to restore more SQL Server database that are 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 an SQL Server database and the current restore is the final restore in the sequence, or is the only restore to an SQL Server database. The SQL Server detects that the restore operation is complete and ready for incomplete transactions to be rolled back.

No Whenever you make a sequence of restores to an SQL Server 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**. You can specify a list of logs without having to specify the final log in a separate command.

The following is a sample scenario:

Sequence of restore operations	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, all subsequent restore objects for that database in a single restore command are skipped regardless of the **/RECOVery** or **/STANDBy** settings.

/RELocate=logicalfilename,... /T0=physicalfilename,...

For **restore** operations, the **/RELocate** and **/T0** parameters as a pair specify the new location of an SQL Server database file. You must use this parameter for every SQL Server 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 Server database file you want to relocate. The *physicalfilename* variable specifies the new physical Windows

file name where you want to relocate the SQL Server database file. This parameter is available when you are restoring legacy backups only.

Considerations:

- You cannot specify more than one database name as the value for the restore command when you are specifying **/RELocate**.
- **/RELocate** and **/TO** can each take a list of values and can be specified more many times. 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`
- MMC does not support the **/RELocate** and **/TO** parameters. You must use the command-line interface when you are completing 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 quotation marks.
- 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 might 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 Server databases, logs, and SQL Server full-text index files. FILESTREAM files are included for SQL Server 2008, SQL Server 2008 R2, or later versions. This parameter is available when you are restoring VSS backups or legacy backups.

The *dbfiledir* variable specifies the directory location of the SQL Server database you want to relocate. 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 Server log files you want to relocate. If the *logfiledir* variable is not specified, the SQL Server 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, SQL Server 2008 R2, or later versions). If the *otherfiledir* variable is not specified, the SQL Server full-text index files and FILESTREAM files (SQL Server 2008, SQL Server 2008 R2, or later versions) are restored to the directory specified by *dbfiledir*.

/REPlace

For **restore** operations, the **/REPlace** parameter specifies that you want existing SQL Server files to be overwritten when they otherwise would not be. You might have to use this parameter in the following instances:

- You are completing a full database restore, and one of the following is true:
 - You are using the **/into** parameter, and the **/into** database exists on the SQL Server.
 - The database exists on the SQL Server, and one of the following is also true:
 - The number of SQL Server files in the existing database differs from the number of SQL Server files in the full database backup object.
 - The names of one or more SQL Server files in the existing database are not the names of any of the SQL Server files in the full database backup object.
- You are completing a file, group, or set restore, and one or more of the SQL Server files 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 that is 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 be specified only when you are 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 that is 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 be specified only when you are 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 that is used when logging on to the SQL Server. The **INTEgrated** value specifies Windows authentication. The user ID that 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 that is specified by the *sqluserid* parameter is the ID that you use to log on to the SQL Server. Any SQL Server 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 Server determines how many buffers must be used.

Considerations:

- The default value is the value that is specified by the SQL Server buffers configurable option in the Data Protection for SQL Server configuration file. This value is initially 0.
- If you specify **/SQLBUFFers**, its value is used instead of the value that is 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 KB. The number be in the range 64 - 4096. The default value is 1024.

Considerations:

- The default value is the value that is specified by the SQL Server buffers configurable option in the Data Protection for SQL Server configuration file. This value is initially 1024.
- If you specify **/SQLBUFFERSize**, its value is used instead of the value that is 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 Server 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 Server 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, port

The **/SQLSERVER** parameter specifies the SQL Server that Data Protection for SQL Server logs on to. For **restore** operations, this server 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. With this variable, you can specify an *sqlservername* and, in the case of the *tcp* protocol, you can also define a port. You can choose from the following protocols:

Table 25.

Protocol Name	Description	Example Usage
lpc	Shared Memory	/sqlserver=lpc: <servername>\<instancename>
np	Named Pipes	/sqlserver=np:<servername>\pipe Note: You can optionally specify a specific named pipe. For example, /sqlserver=np: \\hostname\pipe\pipe name By default, the pipe name is <i>sql\query</i> . If you connect to a named instance, the pipe name is typically in the following format: \\<servername>\pipe\ MSSQL\$<instancename>\SQL\ query
tcp	Transmission Control	/sqlserver=tcp:<servername>[\<instancename>],<port> Note: For tcp protocols only, you have the option of defining a <i>port</i> . If you do not define a port, the default port value is 1443.
via Note: SQL Server only supports this protocol through SQL Server 2008 R2.	Virtual Interface Architecture	/sqlserver=via:<virtualservername>[\<instancename>]

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 that is specified by the SQL Server configurable option in the Data Protection for SQL Server configuration file. This value 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) That is, 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?
<i>local-computername</i>	default	no	no
<i>local-computername\instancename</i>	named	no	yes
<i>virtualservername</i>	default	yes	yes
<i>virtualservername\instancename</i>	named	yes	yes

local-computername

The network computer name of the computer where the SQL Server and Data Protection for SQL Server is located. The TCP/IP host name might not always be the same.

instancename

The name that is given to the named instance of SQL Server that is specified during installation of the instance.

virtualservername

The name that is given to the clustered SQL Server specified during clustering service setup. This name is not the cluster or node name.

/SQLSer=sqlusername

The **/SQLSer** 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 Server user ID must have the SQL Server SYSADMIN fixed server role.
- If you do not specify **/SQLSer**, the default is *sa*.
- If you specify **/SQLSer**, 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 more restore operations 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 that is specified by the %TEMP% environment variable.
- If *undofilename* includes spaces, you must enclose it in double quotation marks.
- If the specified undo file does not exist, SQL Server creates it. If the file exists but was not used for the same SQL Server database, SQL Server overwrites it.
- If you do not specify either **/recovery** or **/STANDby**, the default is **/recovery=yes**.

/STOPAT=datetime

For **restore** operations, **/STOPAT** specifies the point in time that you restore an SQL Server database to. Only transaction logs that are written before the point in time is applied to the SQL Server database. The *datetime*

variable specifies both the date and time, separated by a space. Use any valid date and time format that is accepted by SQL Server.

Considerations:

- This parameter applies only to transaction log restores, but the base restore that the transaction logs apply to must be 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 quotation marks.
- 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 Server database remains 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. You can specify a time 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 can be found at or after the specified point in time) are applied to the SQL Server database. The *markname* variable specifies the name of an SQL Server transaction. The SQL Server transaction might 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 Server database in the distributed transaction.

markname is the transaction name, not the description that follows the **MARK** keyword in an **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 that is accepted by SQL Server.

Considerations:

- This parameter applies only to transaction log restores. The base restore that the transaction logs apply to must be 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 Server 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 quotation marks.
- You cannot 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 Server database remains 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. You can specify a later point in time if you specify the **/AFTER** option. Only

transaction log records written before and not including the named transaction (which can be found at or after the specified point in time) are applied to the SQL Server database. The *markname* variable specifies the name of an SQL Server transaction. The SQL Server transaction might 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 Server database in the distributed transaction.

markname is the transaction name, not the description that follows the MARK keyword in an 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 that is accepted by SQL Server.

Considerations:

- This parameter applies only to transaction log restores. The base restore that the transaction logs apply to must be 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 Server 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 quotation marks.
- You cannot 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 Server database remains 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 value that is used in the backup operation.
- If you specify **/STRIPes** but not *numstripes*, the stored value is used.
- You can use up to the number used to create the backup. You can determine the number of data stripes that are 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 that you require.

- When you use data striping, you must 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. SQL Server allows a maximum of 64 data stripes.

/TSMNODE=tsmnode

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 parameter 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 node name in the options file. Otherwise, you can change PASSWORDACCESS to PROMPT to use the **/TSMNODE** parameter. For more information about the Tivoli Storage Manager options file, see Creating and modifying the client system-options file (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.4/client/t_cfg_crtmodoptunix.html).
- If you do not specify **/TSMNODE**, the default value is that specified by the node name 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 operation 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 more information about the Tivoli Storage Manager options file, see Creating and modifying the client system-options file (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.4/client/t_cfg_crtmodoptunix.html).

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 quotation marks.
- 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	<i>generate</i>	yes	/TSMPassword ignored
specified	<i>generate</i>	no	/TSMPassword used and stored
specified	<i>prompt</i>	—	/TSMPassword used
not specified	<i>prompt</i>	—	user is prompted

/VerifyOnly

The **/VerifyOnly** parameter specifies whether the integrity of a legacy database backup is verified. Before you restore a backup, you can use this parameter to evaluate whether the backup is complete and can be read. By default, the integrity of a legacy database backup is not verified.

Restriction: The **/VerifyOnly** parameter is available only for legacy database backups. This parameter is only a command optional parameter, and it cannot be set as a configuration option.

Related reference:

“Set command” on page 263

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 3.0
(C) Copyright IBM Corporation 1998, 2015.
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 3.0
(C) Copyright IBM Corporation 1998, 2015.
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 3.0
(C) Copyright IBM Corporation 1998, 2015.
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 3.0
(C) Copyright IBM Corporation 1998, 2015.
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 3.0
(C) Copyright IBM Corporation 1998, 2015.
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 3.0
(C) Copyright IBM Corporation 1998, 2015.
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 backups 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 3.0
(C) Copyright IBM Corporation 1998, 2015.
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 3.0
(C) Copyright IBM Corporation 1998, 2015.
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 3.0
(C) Copyright IBM Corporation 1998, 2015.
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 3.0
(C) Copyright IBM Corporation 1998, 2015.
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 3.0
(C) Copyright IBM Corporation 1998, 2015.
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 3.0  
(C) Copyright IBM Corporation 1998, 2015.  
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 3.0
(C) Copyright IBM Corporation 1998, 2015.
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 *svtodb* from Tivoli Storage Manager server storage to directory *m:\svtodb* using the new optional parameter, **/relocatedir**. All SQL logs and full-text index files associated with database *svtodb* are also restored and relocated.

Command:

```
tdpsqlc restore svtodb full /relocatedir=m:\svtodb /backupdestination=tsm
/backupmethod=vss
```

Output:


```

IBM Tivoli Storage Manager for Databases:
Data Protection for Microsoft SQL Server
Version 7, Release 1, Level 3.0
(C) Copyright IBM Corporation 1998, 2015.
All rights reserved.

Connecting to SQL Server, please wait...

Querying TSM Server for Backups ....

Starting Sql database restore...

Beginning VSS restore of 'svtldb'...

Preparing for restore of 'svtldb' 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 *svtldb*, its logs, and its full-text index files into their own respective locations, the following command is issued:

```
tdpsqlc restore svtldb full /relocatedir=m:\svtldb,e:\svtldb,f:\svtldb
/backupdestination=tsm /backupmethod=vss
```

The **/relocatedir** values in this command are as follows:

- *m:\svtldb*: The directory where only the *svtldb* database is relocated.
- *e:\svtldb*: The directory where only the *svtldb* logs are relocated.
- *f:\svtldb*: The directory where only the *svtldb* 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). It is not possible to use a mapped network drive as a destination directory.
- 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:\tempstore`
- Use /RELOCATEDIR to restore a database that 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





The syntax diagrams of the backup object type contain options corresponding to the letters *A*, *B*, *C*, *D*, *E*, *F*. Follow the Optional Parameters for the **restorefiles** command.

Restorefiles positional parameters

Positional parameters immediately follow the **restorefiles** command and precede the optional parameters.

The following positional parameters specify the object to restore:

tdpsqlc restorefiles

*** | componentname1, ..., componentnameNFULL**

***** 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 Server **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 Microsoft Management Console (MMC) 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. It is not possible to use a mapped network drive as a destination directory.

The *dbfiledir* variable specifies the directory location of the SQL Server 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 Server log files you want to relocate. Note that if the *logfiledir* variable is not specified, the SQL Server 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.

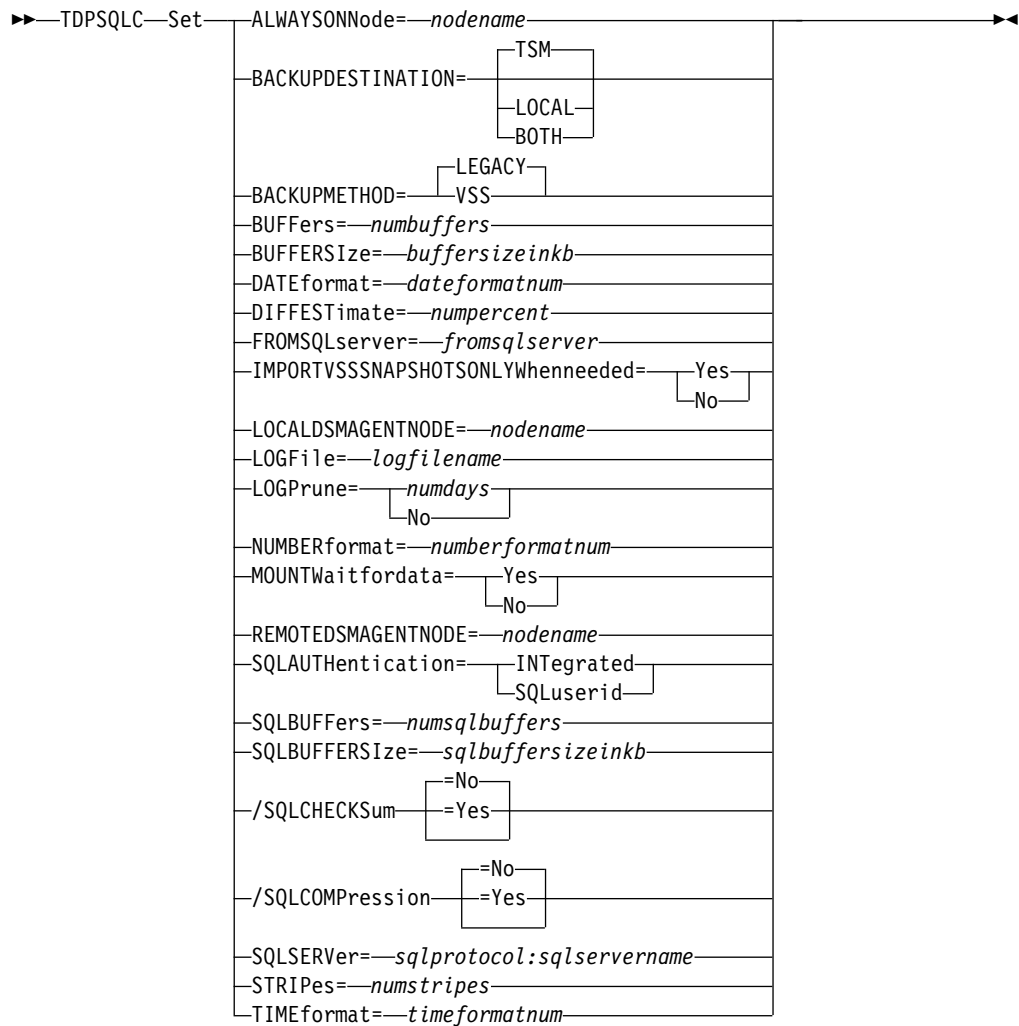
The values are saved in a configuration file. The default file is **tdpsql.cfg**. Configuration values can also be set in the Data Protection Properties window in the GUI.

If a configuration file is not specified, the values in the **tdpsql.cfg** 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



Set Optional Parameters



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 you issue a **set** command.

/ALWAYSOnNode=*nodename*

Specify the Tivoli Storage Manager node name that is used to back up AlwaysOn availability databases with SQL Server 2012 and later versions. This parameter is required when you are configuring Data Protection for SQL Server with SQL Server 2012 and later versions. 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 By default, the backup is stored on Tivoli Storage Manager server storage only.

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 By default, Data Protection for SQL Server uses the legacy API to complete the backup.

VSS Data Protection for SQL Server uses VSS to complete the backup.

BUFFERS=numbuffers

The **BUFFERS** parameter specifies the number of data buffers that are 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 also increase storage use. Each buffer is the size that is specified by the **BUFFERSize** parameter. The *numbuffers* variable refers to the number of data buffers to use. The number can be in the range 2 - 8. The initial value is 3.

BUFFERSize=bufferizeinkb

The **BUFFERSize** parameter specifies the size of each Data Protection for SQL Server buffer that is specified by the **BUFFERS** parameter. The *bufferizeinkb* variable refers to the size of data buffers in KB. The number can be in the range 64 - 8192. The default is initially 1024.

DATEformat=dateformatnum

The **DATEformat** parameter selects the format that you want to use to display dates.

The *dateformatnum* variable can be in the range 1 - 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 completing one of the following steps:

- After you change the value of the **DATEformat** parameter, make a copy of the existing log file before you run Data Protection for SQL Server.
- Specify a new log file with the **LOGFile** parameter.

DIFFESTimate=numpercent

For differential database backups by using the Data Protection for SQL Server **backup** command, **DIFFESTimate** specifies the estimated fraction of an entire SQL Server 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 enough space exists in the primary storage pool to contain the SQL Server database backup. Because a separate backup object is created for each specified SQL Server database, this estimate applies to each specified SQL Server database individually. The *numpercent* variable can be in the range 1 - 99. Because a differential backup backs up database pages, this number is the percent of database pages that changed since the last full database backup. The initial value is 20.

Considerations:

- If the estimate is smaller than the actual quantity of changes, the Tivoli Storage Manager server might 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 larger than the actual quantity of changes, the server might 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

Use the **/IMPORTVSSSNAPSHOTSONLYWhenneeded** parameter to specify whether Data Protection for SQL Server automatically imports VSS snapshots to the Windows system where the snapshots are created.

Specify one of the following values:

- | | |
|------------|---|
| Yes | Import VSS snapshots to the Windows system where the snapshots are created. The option is the default. During backup processing, transportable snapshots are automatically created and imported to storage systems when the snapshots are required. This option is the default. |
| No | Do not create transportable VSS snapshots during backup processing, and do not automatically import the snapshot to storage systems after the backup is completed. |

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 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 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.

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 7, Release 1, Level 3.0
(C) Copyright IBM Corporation 1997, 2015. All rights reserved.

AC05054I 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 7, Release 1, Level 3.0
(C) Copyright IBM Corporation 1997, 2015. All rights reserved.

Connecting to SQL Server, please wait...

AC05054I The configuration option was set successfully.
```

The following statement is added to the `tdpsql.cfg` configuration file:

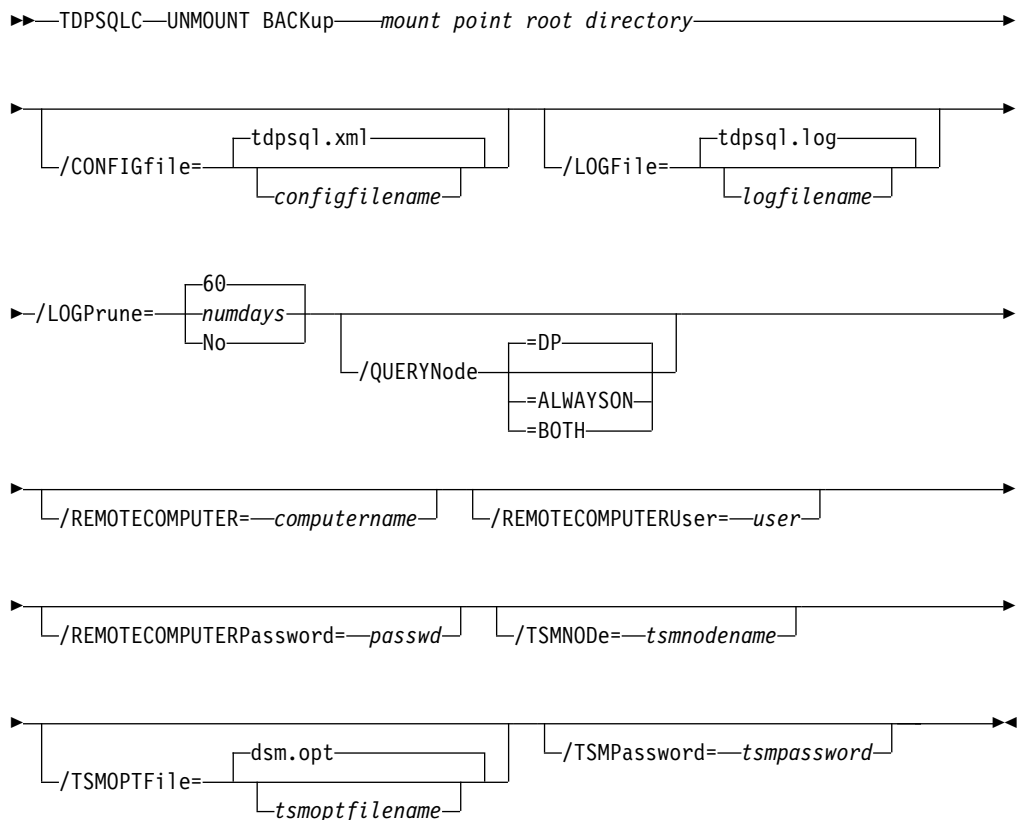
Unmount Backup command

Use the **unmount backup** command to unmount backups that have been previously mounted, and are managed by Tivoli Storage FlashCopy Manager for SQL Server.

Unmount Backup syntax

Use the **unmount 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

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**

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, a certain number of days of data are saved. By default, 60 days of log entries are saved. The option **No** can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify **no**, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or **no**; in this case, the default value, 60, is used.

- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

/QUERYNode=DP|ALWAYSON|BOTH

Specify whether you want to query standard databases from SQL Server 2012 and later versions 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=password

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.

Appendix A. Frequently asked questions

Review the answers to the following frequently asked questions about Data Protection for SQL Server.

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.

For more information about the **compression** option, see “Specifying configuration parameters for Tivoli Storage Manager” on page 47.

- You can specify SQL Server backup compression from the SQL Properties windows in Microsoft Management Console (MMC), or you can use the **sqlcompression** option from the command line to set SQL Server native backup compression for Data Protection for SQL Server backups. For more information, see “Enabling SQL Server backup compression” on page 99.

Backup compression is only available with legacy backups. You can run backup compression only on Enterprise Edition. You can run SQL Server 2008 R2 backup compression 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.

For more information about the **enableclientencryptkey** and **encryptiontype** options, see “Specifying configuration parameters for Tivoli Storage Manager” on page 47.

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 configuration parameters for Tivoli Storage Manager” on page 47.

Can I restore an individual table from an 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.

Can I restore an SQL Server database backup to an alternate SQL Server machine or instance?

Yes, but only for legacy backups. You cannot restore VSS backups directly to an alternate SQL Server machine or instance.

Can I restore VSS backups to alternate locations within the same SQL Server instance?

You can restore VSS backups to alternate locations within the same SQL Server instance as follows:

- Use the **/relocatedir** parameter from the command line.
- Use the Relocate option in the Restore Databases window from the GUI.

Can I restore VSS backups to alternate database names within the same SQL Server instance?

You can restore VSS backups to alternate database names as follows:

- Use the **/into** parameter from the command line.
- Use the Restore Into option in the Restore Databases window from the GUI.

Can I use Data Protection for SQL Server to back up SQL Server 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 an SQL Server database?

First, use the same node name as the primary SQL Server 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:

- “Specifying Data Protection for SQL Server node name parameters” on page 50
- “Specifying configuration parameters for Tivoli Storage Manager” on page 47

Can I perform VSS operations in a clustered SQL Server environment?

Yes, you can run Data Protection for SQL Server VSS operations in a clustered SQL Server environment. For more information, see “AlwaysOn Availability Groups with Data Protection for SQL Server in a Windows Failover Cluster environment” on page 21.

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 “Prerequisites” on page 29.

How can I use VSS and legacy backups together in a common backup strategy?

For more information, see “VSS and legacy backup methods” on page 89.

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 **/recover=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. With VSS, you can run only full backups. You cannot run log, differential, individual filegroups, individual files, and set backups with VSS. For more information, see “VSS and legacy backup methods” on page 89.

When restoring very large SQL Server 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 operations work?

In a VSS instant restore operation, a volume-level hardware-assisted copy that includes target volumes (that contain the snapshot) are copied back to the original source volumes. A SAN Volume Controller, Storwize V7000, XIV, or IBM System Storage DS8000 series storage system is required to complete VSS instant restore operations. For more information, see “VSS instant restore processing” on page 19.

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 policy affects backup management on Data Protection for SQL Server” on page 8 and “VSS and legacy backup methods” on page 89.

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 7, “Performance tuning,” on page 151
- “Specifying configuration parameters for Tivoli Storage Manager” on page 47
- “**/buffers** and **/buffersize** parameters” (with the **backup** command) on “Backup optional parameters” on page 162.
- “**/buffers** and **/buffersize** parameters” (with the **restore** command) on “Restore optional parameters” on page 232.
- “**/buffers** and **/buffersize** parameter” (with the **set** command) on “Set optional parameters” on page 267.

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 MMC 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:

- “AlwaysOn Availability Groups with Data Protection for SQL Server in a Windows Failover Cluster environment” on page 21

How do I know if my backup ran successfully?

A message displays that states the backup completed successfully. In

addition, the Task Manager in MMC 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 node name parameters” on page 50.

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 7, “Performance tuning,” on page 151.

For more information, see <http://www.redbooks.ibm.com/abstracts/sg246148.html>.

Appendix B. 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 Operations Center help system is enabled for accessibility. For more information, click the question mark icon on the help system menu bar.

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, MD-NC119
Armonk, NY 10504-1785
United States of America*

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
US*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

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 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:

Portions of this code are derived from IBM® Corp. Sample Programs.

© Copyright IBM® Corp. _enter the year or years_. All rights reserved.

Trademarks

IBM, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Adobe is a registered trademark of Adobe Systems Incorporated in the United States, and/or other countries.

Linear Tape-Open, LTO, and Ultrium are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

Intel and Itanium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java™ and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights Except as expressly granted in this permission, no other permissions,

licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Privacy policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at <http://www.ibm.com/privacy> and IBM's Online Privacy Statement at <http://www.ibm.com/privacy/details> in the section entitled "Cookies, Web Beacons and Other Technologies," and the "IBM Software Products and Software-as-a-Service Privacy Statement" at <http://www.ibm.com/software/info/product-privacy>.

Glossary

A glossary is available with terms and definitions for the IBM Tivoli Storage Manager family of products.

See Tivoli Storage Manager glossary.

To view glossaries for other IBM products, see IBM Terminology.

Index

Special characters

- /alwaysonpriority parameter
 - and backup command 163
- /querynode parameter
 - and query command 207
 - and restore command 232
- /RELocate and /TO parameters 238
- /usealwaysonnode parameter
 - and backup command 172

A

- access to databases, restricting 233
- accessibility features 279
- active/inactive state
 - affected by full backup 160, 190
 - in restore operations 237
- adjustkbtstimate parameter 162
- adjustpercentestimate parameter 163
- AlwaysOn node
 - transitioning standard databases to 85
- alwaysonnode parameter
 - and set command 264
- APAR 149
- API, Tivoli Storage Manager 205
- authorization mode, setting
 - using the CLI 167, 168, 170, 210, 212, 240, 241, 243
- auto select option, GUI 110
- automated failover
 - overview 26
- automatic expiration policy, setting 63
- availability database restores
 - overview 23, 24, 92

B

- backing up SQL availability databases
 - by using the legacy method 102
- backing up SQL databases
 - by using the legacy method 102
 - on Windows Server Core 106
- backing up SQL databases with Data Protection for SQL Server
 - on Windows Server Core 124
- backing up SQL Server availability databases
 - by using the VSS method 100
- backing up SQL Server groups or files
 - by using the legacy method 105
- backup
 - copy-only full 89
 - file backup 89
 - full 89
 - full plus differential plug log 89
 - full plus differentials 89
 - full plus log 89
 - Legacy 7
- backup command
 - and /alwaysonpriority parameter 163
 - and /backupdestination parameter 163
 - and /backupmethod parameter 163

- backup command (*continued*)
 - and /logfile parameter 181
 - and /offload parameter 167
 - and /quiet parameter 182
 - and /usealwaysonnode parameter 172
 - optional parameters 162
 - positional parameters 159
- backup methods 6
- backup object types 4
 - copyfull 231
 - COPYFull 160
 - differential 160, 190, 231
 - file 190, 230
 - File 159
 - for query Data Protection for SQL Server 206
 - full 190, 231
 - FULL 160
 - group 190, 231
 - Group 160
 - log 191, 231
 - Log 161
 - set 191, 231
 - Set 162
- backup objects
 - compatibility with server 208
 - query of 205
- backup operations
 - time-saving strategy 89
 - using the GUI
 - backup databases tab 100
 - backup groups/files tab 105
- backup strategy
 - copy-only full 89
 - file backup 89
 - full backup 89
 - full plus differential plug log 89
 - full plus differentials 89
 - full plus log 89
 - group backup 89
 - Tivoli Storage Manager versus local shadow volumes 89
 - VSS cluster 62, 141
- backupdestination parameter
 - and backup command 163
 - and delete backup command 180
 - and restore command 232
- BACKUPDESTINATION parameter
 - and set command 265
- backupmethod parameter
 - and backup command 163
 - and restore command 232
 - and restorefiles command 259
- BACKUPMETHOD parameter
 - and set command 265
- backups of availability databases
 - overview 22
- binary sort order 208
- buffering data 233
 - Data Protection for SQL Server performance 151
 - for Data Protection for SQL Server 164, 207, 232, 265
 - for SQL Server 167, 168, 240, 241
 - for TDP 164, 207

- buffering data (*continued*)
 - used with stripes 167, 240
- buffers parameter 164, 207, 232
- BUFFers parameter 265
- buffer size parameter 164, 207, 233

C

- capacity
 - determining managed storage 97
- changetsm password command
 - optional parameters 176
 - positional parameters 176
- changing configuration values
 - on Windows Server Core 79, 125
- checksum
 - SQL Server 99
- CHECKSum
 - SQL Server 99
- cluster
 - VSS 62, 141
- clustering
 - strategy 22
- clustering state
 - querying 205
- code page ID
 - querying 206
- command
 - policy 225
- command line parameters
 - /alwaysonnode
 - and set 264
 - /alwaysonpriority
 - and backup 163, 172
 - /backupdestination
 - and backup 163
 - and restore 180, 232
 - /backupmethod
 - and backup 163
 - and restore 232
 - /instantrestore
 - and restore 234
 - /INSTANTRestore
 - and restore 234
 - /offload
 - and backup 167
 - /pitdate
 - and mount backup 200
 - /pittime
 - and mount backup 200
 - /querynode
 - and query 207
 - and restore 232
 - /quiet
 - and backup 182
 - BACKUPDESTination
 - and set 265
 - BACKUPMETHod
 - and set 265
- command-line considerations
 - VSS restore 226
- command-line interface
 - overview 153
- command-line parameters
 - /backupdestination
 - and restorefiles 259
 - command-line parameters (*continued*)
 - /configfile
 - and delete backup 180
 - and mount backup 199
 - and restorefiles 259
 - and unmount backup 269
 - /fromsqlserver
 - and delete backup 181
 - and mount backup 199
 - and restorefiles 260
 - /IMPORTVSSSNAPSHOTSONLYWhenneeded
 - and set 266
 - /into
 - and restorefiles 261
 - /logfile
 - and backup 181
 - and mount backup 199
 - and restorefiles 260
 - and unmount backup 269
 - /logprune
 - and restorefiles 260
 - /mountwait
 - and restorefiles 261
 - /object
 - and delete backup 182
 - and restorefiles 261
 - /quiet
 - and restorefiles 261
 - /tsmnode
 - and mount backup 201
 - and restore 182
 - and restorefiles 262
 - and unmount backup 270
 - /tsmoptfile
 - and mount backup 201
 - and restore 182
 - and restorefiles 262
 - and unmount backup 270
 - /tsmpassword
 - and mount backup 202
 - and restore 183
 - and restorefiles 262
 - and unmount backup 270
 - and local 259
 - and tsm 259
 - and vss 259
- commands
 - query managedcapacity 224
- commands, Data Protection for SQL
 - tdpsqlc changetsm password 176
- commands, Data Protection for SQL Server
 - tdpsqlc help 184
 - tdpsqlc set 264
- communication protocol option 47
- compatibility level
 - querying 205
- compatibilityinfo
 - query of Data Protection for SQL Server 206
- compatibilityinfo parameter 208
- compression
 - SQL Server 99
- compression option 48
- configfile parameter 164, 176, 191, 208, 233, 267
 - and delete backup command 180
 - and mount backup command 199
 - and restorefiles command 259
 - and unmount backup command 269

- configuration
 - IBM SAN Volume Controller and IBM Storwize V7000 86
 - manual procedure
 - offloaded backups 70
 - SQL Server 68
 - Tivoli Storage Manager server 70
 - options 47
- configuration file, Data Protection for SQL Server
 - query TDP 205
 - setting values, CLI 164, 176, 191, 208, 233, 264, 267
- configuration files
 - non-default locations 51
- configuration information, listing 208
- configuration tasks
 - Data Protection for SQL Server on Windows Server Core 73
- configuring
 - backup priority of replicas 81
 - binding
 - policy 13
 - Data Protection for SQL Server 45, 65
 - policy 13
 - quick instructions 30
 - where scheduled backups are run 81
- configuring Data Protection for SQL Server
 - for Windows Server configuration 77
- configuring the backup-archive client
 - for Windows Server configuration 76
- considerations
 - considerations 3
 - VSS restore 3
- copyfull parameter
 - described 231
- COPYFull parameter
 - described 160
- create index operation 160, 161
- creating a client node
 - for Windows Server Core configuration 73
- custom settings 59

D

- data objects
 - in object naming 15
- data protection
 - Exchange with VSS backup/restore support
 - gathering files before calling IBM 146
 - tracing when using VSS 145
 - general help 138
 - SQL Server with VSS backup/restore support
 - general help 133
 - SQL with VSS backup-restore support
 - gathering information before calling IBM 145
 - SQL with VSS backup/restore support
 - determining the issue 134
 - troubleshooting 142
- Data Protection for Microsoft Exchange Server
 - silent installation with batch file 38
- Data Protection for SQL
 - backup types 4
 - options file
 - automatic expiration policy 63
- Data Protection for SQL Server
 - backups 6
 - buffer size 164, 207, 233
 - BUFFERSize 265
 - commands 153

- Data Protection for SQL Server *(continued)*
 - configuration file query 205
 - configuration file, setting
 - using the CLI 164, 176, 191, 208, 233, 267
 - configuration parameters 63
 - configuring 45, 65
 - configuring options 47
 - creating an installation package on a DVD 41
 - installing on a local system 34
 - installing on Windows Server Core 35
 - number of buffers 164, 207, 232, 265
 - operating environment 1
 - overview 1, 89
 - policy settings 8
 - quick configuration 30
 - quick installation 30
 - reference 153
 - registering 50
 - restore types 19
 - silent installation 36
 - silent installation on Windows Server Core (setup.exe) 37
 - version query 206
 - VSS planning 2
- Data Protection for SQL Server GUI
 - starting 89, 94
- Data Protection for SQL Server on Windows Server Core
 - overview 123
- Data Protection for SQL Server on Windows Server Core
 - configuration tasks 73
- Data Protection for SQL Server scripts
 - adding 147
 - editing 147
 - viewing 147
- Data Protection for SQL Server silent installation
 - playing back the installation 41
 - setup error messages 42
- Data Protection for SQL Server tasks
 - automating 130
 - Data Protection for SQL Server tasks 130
- Data Protection for SQL Server trace and log files
 - viewing 144
- Data Protection for SQL Server VSS backup
 - policy binding 13
- Data Protection for VMware Recovery Agent 118
- data stripes
 - buffering 164, 207, 232, 265
 - Data Protection for SQL Server performance 151
 - querying 205
 - used with /sqlbuffers 167, 170, 240, 245
 - using the CLI 170, 245
- data striping
 - buffering 164, 207, 233
- database
 - options set 205
 - querying 205
 - restoring master 117
 - restoring to alternate 115
 - restoring using the GUI 122
 - restoring with full-text catalogs and indexes 118
- database backups, verifying 247
- database integrity checking 247
- database name
 - restorefiles
 - command line 259
- database owner option, GUI 110
- DATEformat parameter 265
- dboonly parameter 233

- deactivate operations
 - using the GUI 107
- delete backup
 - storage group
 - command line 180
- delete backup command
 - and /backupdestination parameter 180
 - and /configfile parameter 180
 - and /fromsqlserver parameter 181
 - and /object parameter 182
 - optional parameters 180
 - overview 179
 - syntax diagram 179
- deleting SQL Server VSS backups 107
- described
 - set parameter 207
- developerWorks wiki 149
- diagnostics properties 54
- differential backup
 - overview 4
- differential parameter
 - described 160, 190, 206, 231
- differential restore
 - using the GUI 112, 122
- differential versus log backup strategy 89
- DIFFESTimate parameter 266
- disability 279
- dsm.opt file 47
 - communication protocol 47
 - compression 48
 - enableclientencryptkey 49
 - enablelanfree 49
 - encryptiontype 49
 - include.encrypt 49
 - nodename 48

E

- enableclientencryptkey option 49
- enablelanfree option 49
- ENABLEREplacementchars parameter 165
- encryption 49
- encryptiontype option 49
- error log files 133
- example
 - restorefiles command 263
- excfull.log 131
- excludedb parameter 165
- excsched.log 131
- expiration policy, setting 63
- expiring VSS backups
 - policy 8

F

- failover
 - overview 26
- FAQ 273
- file backup
 - overview 4
 - strategy 89
- file parameter
 - described 190, 206, 230
- File parameter
 - described 159

- file restore
 - using the GUI 112, 122
- files
 - Data Protection for SQL Server options 183, 262
 - dsm.opt 47
 - excfull.log 131
 - excsched.log 131
 - tdpexc.cfg
 - and delete backup command 180
 - tdpexc.log
 - and delete backup command 181
 - tdpsql.cfg
 - and mount backup command 199
 - and restorefiles command 259
 - and unmount backup command 269
 - vsspolicy option 63
 - tdpsql.log 277
 - and mount backup command 199
 - and restorefiles command 260
 - and unmount backup command 269
 - tdpsqlc.exe 153
 - Tivoli Storage FlashCopy Manager options 202, 270
- frequently asked questions 273
- from SQL Server option, GUI 111
- fromsqlserver parameter 191, 208, 233, 266
 - and delete backup command 181
 - and mount backup command 199
 - and restorefiles command 260
- full backup
 - overview 4
 - strategy 89
- full parameter
 - described 190, 206, 231
- FULL parameter
 - described 160
- full plus differential backup
 - strategy 89
- full plus differential plus log backup
 - strategy 89
- full plus log backup
 - strategy 89
- full restore
 - using the GUI 112, 122

G

- general properties for SQL Server 57
- graphical user interface
 - backup groups/files tab 105
- graphical user interface (GUI)
 - backup databases tab 100
 - inactivating SQL Server databases 107
 - restore groups/files tab 122
 - restore options 110
 - restoring SQL Server databases 122
- group backup
 - overview 4
 - strategy 89
- group parameter
 - described 190, 206, 231
- Group parameter
 - described 160
- group restore
 - using the GUI 112, 122
- GUI
 - SQL Server VSS backup 98
 - SQL VSS restore 109

H

help command
described 184

I

IBM Knowledge Center x
IBM SAN Volume Controller and IBM Storwize V7000
configurations 86
using Data Protection for SQL Server 20
IBM System Storage DS8000 series
requirements 21
inactivate command
optional parameters 191
positional parameters 190
include/exclude
syntax 63
INCLUDE/EXCLUDE
sample statements 15
include.encrypt option 49
IncludeTsmVm 111
indexes and tables
backing up 89
installation
configuring options 47
registering Data Protection for SQL Server 50
installing
creating an installation package on a DVD 41
on a local system 34
quick instructions 30
silently with batch file 38
silently with msixec.exe 39
installing Data Protection for SQL Server
on multiple servers (silent) 36
on Windows Server core 35
unattended (silent) 36
installing Tivoli Storage Manager client on Windows Server
Core
on multiple servers (silent) 37
unattended (silent) 37
instantrestore parameter
and restore command 234
integrated user id mode 167, 210, 240
integrity checking, database backups 247
into parameter 234
and restorefiles command 261

K

keyboard 279
Knowledge Center x

L

LAN-free
Data Protection for SQL Server performance 152
Legacy backup
overview 7
local backup policy
setting 12
log backup
overview 4
log files
using for problem determination 133

log parameter
described 191, 206, 231
Log parameter
described 161
log restore
using the GUI 112, 122
logfile parameter 165, 191, 209, 235
and delete backup command 181
and mount backup command 199
and restorefiles command 260
and unmount backup command 269
logging properties 57
login settings
using the CLI 167, 210, 240
logprune parameter 166, 177, 181, 192, 199, 209, 235, 269
and restorefiles command 260

M

managed storage
determining capacity 97
management class
INCLUDE statements 15
meta and data objects 15
object naming 63
master database, restoring 117
media migration 63
meta objects
in object naming 15
metadata policy, setting 15
migration 43
MMC GUI
starting 89, 94
mount backup command
and /configfile parameter 199
and /fromsqlserver parameter 199
and /logfile parameter 199
and /pitdate parameter 200
and /pittime parameter 200
and /tsmnode parameter 201
and /tsmoptfile parameter 201
and /tspassword parameter 202
overview 198
syntax diagram 198
mountwait parameter 166, 236
and restorefiles command 261
msixec.exe
used for silent installation 39
multiple-user mode 109

N

New in Data Protection for SQL Server Version 7.1.6 xiii
node name
Data Protection for SQL Server 50
offloaded backup 46
proxy nodes 45
VSS 45
nodename option 48
number of buffers
for Data Protection for SQL Server 164, 207, 232, 265
for SQL Server 167, 240
used with stripes 167, 240

O

- object parameter 192, 210, 236
 - and delete backup command 182
 - and restorefiles command 261
- offload parameter
 - and backup command 167
- offloaded backup
 - manual configuration procedure 70
 - node names 46
- offloaded VSS backup
 - overview 7
- olderthan parameter 193
- operating environment 1
- optional parameters
 - backup command 162
 - changetmpspassword command 176
 - delete backup command 180
 - inactivate command 191
 - query command 207
 - restore command 232
 - set command 267
- options file, Data Protection for SQL
 - INCLUDE/EXCLUDE statements 15
- options files
 - non-default locations 51
- overview 1, 89
 - availability database restores 23, 24, 92
 - backups of availability databases 22
 - Data Protection for SQL Server on Windows Server
 - Core 123
 - Legacy backup 7
 - offloaded VSS backup 7
 - SQL Server AlwaysOn Availability Groups 21
 - VSS backup 2, 6
 - VSS restore 3
 - VSS service 1

P

- parameters
 - /alwaysonnode
 - and set command 264
 - /alwaysonpriority
 - and backup command 163
 - /backupdestination
 - and backup command 163
 - and delete backup command 180
 - and restore command 232
 - and restorefiles command 259
 - /backupmethod
 - and backup command 163
 - and restore command 232
 - /configfile
 - and delete backup command 180
 - and mount backup command 199
 - and restorefiles command 259
 - and unmount backup command 269
 - /fromsqlserver
 - and delete backup command 181
 - and mount backup command 199
 - and restorefiles command 260
 - /IMPORTVSSSNAPSHOTSONLYWhenneeded
 - and set command 266
 - /instantrestore
 - and restore command 234

parameters (*continued*)

- /INSTANTRestore
 - and restore command 234
- /into
 - and restorefiles command 261
- /logfile
 - and delete backup command 181
 - and mount backup command 199
 - and restorefiles command 260
 - and unmount backup command 269
- /logprune
 - and restorefiles command 260
- /mountwait
 - and restorefiles command 261
- /object
 - and delete backup command 182
 - and restorefiles command 261
- /offload
 - and backup command 167
- /pitdate
 - and mount backup command 200
- /pittime
 - and mount backup command 200
- /querynode
 - and query command 207
 - and restore command 232
- /quiet
 - and delete backup command 182
 - and restorefiles command 261
- /tsmnode
 - and mount backup command 201
 - and restore command 182
 - and restorefiles command 262
 - and unmount backup command 270
- /tsmoptfile
 - and mount backup command 201
 - and restore command 182
 - and restorefiles command 262
 - and unmount backup command 270
- /tsmpsassword
 - and mount backup command 202
 - and restore command 183
 - and restorefiles command 262
 - and unmount backup command 270
- /usealwaysonnode
 - and backup command 172
- BACKUPDESTination
 - and set command 265
- BACKUPMETHod
 - and set command 265

parameters, described

- /to 238
- optional
 - /adjustkbtstestimate 162, 163
 - /buffers 164, 207, 232
 - /buffersize 164, 207, 233
 - /compatibilityinfo 208
 - /configfile 164, 176, 191, 208, 233, 267
 - /dboonly 233
 - /excludedb 165
 - /fromsqlserver 191, 208, 233
 - /into 234
 - /logfile 165, 191, 209, 235
 - /logprune 166, 177, 181, 192, 199, 209, 235, 269
 - /mountwait 166, 236
 - /object 192, 210, 236
 - /olderthan 193

parameters, described (*continued*)

optional (*continued*)

- /partial 237
- /quiet 193, 237
- /recovery 238
- /relocate 238
- /relocatedir 239
- /replace 239
- /restoredate 240
- /restorettime 240
- /sqlauthentication 167, 210, 240
- /sqlbuffers 167, 240
- /sqlbuffersize 168, 241
- /sqlpassword 168, 210, 241
- /sqlserver 169, 210, 241
- /sqluser 170, 212, 243
- /standby 243
- /stopat 243
- /stopatmark 244
- /stopbeforemark 244
- /stripes 170, 245
- /tsmnode 171, 178, 193, 212, 246
- /tsmoptfile 171, 178, 193, 212, 246
- /tsmpassword 171, 194, 213, 246
- /VerifyOnly 247
- ENABLEREPlacementchars 165
- SQLCHECKSum 168
- SQLCOMPression 168

positional

- backup object types 206
- copyfull 231
- COPYFull 160
- differential 160, 190, 231
- file 190, 230
- File 159
- for changetsmpassword command 176
- for set command 264
- full 190, 231
- FULL 160
- group 190, 231
- Group 160
- log 191, 231
- Log 161
- set 191, 231
- Set 162

partial parameter 237

password, Tivoli Storage Manager

- changing
 - using the CLI 176

performance issues 159, 190, 230

performance properties 61

pitdate parameter

- and mount backup command 200

pittime parameter

- and mount backup command 200

planning requirementsVSS backup 2

point in time named marks restore

- using the CLI 244

policy 14

- binding 13
- binding Data Protection for SQL Server VSS backups 13
- configuring 13
- expiring VSS backups 8
- setting local policy 12

policy command

- overview 225

policy management properties 53

policy settings

- Data Protection for SQL Server and Tivoli Storage Manager 8

positional parameters

- backup command 159
- inactivate command 190
- query command 205
- restore command 230

printing reports 126

problem determination 133, 208

product support 149

properties

- custom settings 59
- diagnostics 54
- general SQL Server 57
- logging 57
- performance 61
- policy management 53
- regional settings 58
- SQL login 56
- VSS 58

protecting SQL Server data with Data Protection for SQL Server

- on Windows Server Core 123

- Windows Sever Core

- protecting SQL Server data with Data Protection for SQL Server 123

proxy nodes 45

publications x

Q

query command

- and /querynode parameter 207
- optional parameters 207
- positional parameters 205

query managedcapacity command

- overview 224

query operations

- query SQL 205
- query TDP 205
- query TSM 205
- sample output 213

query TSM options 204, 206

querying backup objects 205

quiet parameter 193, 237

- and delete backup command 182

- and restorefiles command 261

R

recovery model

- querying 205

recovery option, GUI 111

recovery parameter 238

reference

- Data Protection for SQL Server 153

regional properties 58

registration 50

relocatedir parameter 239

removable media 63

replace option, GUI 111

replace parameter 239

reports

- viewing, printing, and saving 126

requirements 21

- requirements (*continued*)
 - IBM System Storage DS8000 series 21
 - SAN Volume Controller 21
 - Storwize V7000 21
 - XIV 21
- restore 19
 - database 19
 - databases 118
 - restorefiles command 19
 - transaction log 19
 - types 19
 - VSS 3
- restore command
 - and /backupdestination parameter 232
 - and /backupmethod parameter 232
 - and /instantrestore parameter 234
 - and /INSTANTRestore parameter 234
 - and /querynode parameter 232
 - and /tsmnode parameter 182
 - and /tsmoptfile parameter 182
 - and /tsmpassword parameter 183
 - optional parameters 232
 - positional parameters 230
- restore operations
 - master database 117
 - named marks 244
 - of inactive objects 237
 - point in time 243, 244
 - SQL Server databases with full-text catalogs and indexes 118
 - to alternate instance 115
 - using the GUI 122
 - auto select option 110
 - database owner option 110
 - from SQL Server option 111
 - instant restore 112
 - recovery option 111
 - replace option 111
 - restore options 110
 - standby server undo file option 111
 - stripes option 110
 - verify only option 111
 - Wait for Tape Mounts for File Information 112
 - wait for tape mounts options 112
- restoredate parameter 240
- restorefiles command
 - and /backupmethod parameter 259
 - and /configfile parameter 259
 - and /fromsqlserver parameter 260
 - and /into parameter 261
 - and /logfile parameter 260
 - and /logprune parameter 260
 - and /mountwait parameter 261
 - and /object parameter 261
 - and /quiet parameter 261
 - and /tsmnode parameter 262
 - and /tsmoptfile parameter 262
 - and /tsmpassword parameter 262
 - example 263
- restoretme parameter 240
- restoring SQL databases with Data Protection for SQL Server
 - on Windows Server Core 124
- restoring SQL Server availability databases 112
- restoring SQL Server databases
 - on Windows Server Core 121
- restoring SQL Server file groups and files
 - from availability databases 122
- restoring VE databases 112
- restricting database access 233

S

- sample output
 - query command 213
 - set command 267
- saving reports 126
- scheduler
 - guidelines 131
- server, SQL
 - querying 205
- server, Tivoli Storage Manager
 - querying 205
- Service Management Console 149
- set backup
 - overview 4
- set command
 - and /alwaysonnode parameter 264
 - and /IMPORTVSSSNAPSHOTSONLYWhenneeded parameter 266
 - and BACKUPDESTination parameter 265
 - and BACKUPMETHOD parameter 265
 - optional parameters 267
 - positional parameters 264
 - sample output 267
- set parameter
 - described 191, 231
- Set parameter
 - described 162
- set restore
 - using the GUI 112, 122
- setting up a proxy node for offloaded VSS backups
 - for Windows Server configuration 75
- setup.exe
 - used for silent installation 37
- silent installation
 - with setup.exe 37
- silent installation (setup.exe)
 - Data Protection for SQL Server on Windows Server Core 37
 - silent installation of Data Protection for SQL Server 36
 - silent installation on Windows Server Core of Tivoli Storage Manager client 37
- single-user mode 109
- size of buffers 233, 265
 - for SQL Server 168, 241
 - for TDP 164, 207
- sort order ID
 - querying 206
- space allocated and used
 - querying 205, 206
- space-saving considerations 159, 160, 190, 230, 231
- space, saving
 - strategy 89
- SQL availability databases
 - backing up with the legacy method 102
- SQL login properties 56
- SQL restore
 - DAG environment 109
 - VSS
 - GUI 109
- SQL Server
 - number of buffers 167, 240
 - size of buffers 168, 241
 - user id 167, 210, 240

- SQL Server 2000
 - differential strategy 89
 - query of 205
- SQL Server AlwaysOn Availability Groups
 - overview 21
- SQL Server availability databases
 - backing up with VSS 100
 - restoring 112
 - restoring SQL Server file groups and files 122
- SQL Server backup
 - VSS
 - GUI 98
- SQL Server compression 99
- SQL Server databases with full-text catalogs and indexes, restoring 118
- SQL Server groups or files
 - backing up with the legacy method 105
- SQL Server integrity checking
 - checksum 99
- SQL Server VSS backup
 - deleting 107
- sqlauthentication parameter 167, 210, 240
- sqlbuffers parameter 167, 240
- sqlbuffersize parameter 168, 241
- SQLCHECKSum parameter 168
- SQLCOMPression parameter 168
- sqlpassword parameter 168, 210, 241
- sqlserver parameter 169, 210, 241
- sqluser parameter 170, 212, 243
- standby parameter 243
- standby server undo file option, GUI 111
- starting
 - Data Protection for SQL Server GUI 89, 94
 - MMC GUI 89, 94
- stopat parameter 243
- stopatmark parameter 244
- stopbeforemark parameter 244
- storage
 - determining managed capacity 97
- storage group
 - delete backup
 - command line 180
 - VSS backup
 - GUI 98
 - VSS restore
 - GUI 109
- storage management, policy 7
- Storwize V7000
 - requirements 21
- stripes option, GUI 110
- stripes parameter 170, 245
- stripes, data
 - Data Protection for SQL Server performance 151
 - used with /sqlbuffers 167, 170, 240, 245
 - using the CLI 170, 245
- support files
 - sending to IBM with email 148
- syntax diagrams
 - delete backup command 179
 - mount backup command 198
 - reading x
 - unmount backup command 268
- sysadmin fixed server role 225

T

- tables and indexes
 - backing up 89
- task manager 97
- tdpexc.cfg file
 - and delete backup command 180
- tdpexc.log file
 - and delete backup command 181
- tdpsql.cfg file
 - and mount backup command 199
 - and restorefiles command 259
 - and unmount backup command 269
- tdpsql.cfg, setting values
 - using the CLI 164, 176, 191, 208, 233, 267
- tdpsql.log file
 - and mount backup command 199
 - and restorefiles command 260
 - and unmount backup command 269
- tdsqlc.exe
 - overview 153
- throughput, improving 164, 207, 208, 232, 233
- time-saving considerations 159, 160, 190, 230, 231
- Tivoli Storage FlashCopy Manager
 - transitioning backups 82
- Tivoli Storage Manager
 - policy settings 8
- Tivoli Storage Manager API 205
- Tivoli Storage Manager client
 - silent installation on Windows Server Core 37
- Tivoli Storage Manager server
 - querying 205
- transact-SQL command 109
- transaction log
 - querying 205
 - restore 19
- transitioning standard SQL Server databases to the AlwaysOn node 85
- tsmnode parameter 171, 178, 193, 212, 246
 - and mount backup command 201
 - and restore command 182
 - and restorefiles command 262
 - and unmount backup command 270
- tsmoptfile parameter 171, 178, 193, 212, 246
 - and mount backup command 201
 - and restore command 182
 - and restorefiles command 262
 - and unmount backup command 270
- tsmpassword parameter 171, 194, 213, 246
 - and mount backup command 202
 - and restore command 183
 - and restorefiles command 262
 - and unmount backup command 270
- TSMVM 111
- types parameter
 - described 207
 - syntax 204

U

- Unicode information
 - querying 206
- unmount backup command
 - and /configfile parameter 269
 - and /logfile parameter 269
 - and /tsmnode parameter 270
 - and /tsmoptfile parameter 270

- unmount backup command (*continued*)
 - and /tsmpassword parameter 270
 - overview 268
 - syntax diagram 268
- user mode, setting 109
- using
 - syntax diagrams x
- using command-line help
 - on Windows Server Core 95
- using Data Protection for SQL Server
 - with IBM SAN Volume Controller and IBM Storwize V7000 20
- using the backup command
 - on Windows Server Core 106, 124
- using the help command
 - on Windows Server Core 95
- using the restore command
 - on Windows Server Core 121, 124
- using the set command
 - on Windows Server Core 79, 125

V

- Verify Only option, GUI 111
- verifyonly parameter 247
- viewing reports 126
- viewing system information for Data Protection for SQL Server 147
- volume mount 63
- VSS
 - cluster 62, 141
 - N-series and NetApp
 - storage 4
 - node names 45
 - overview 1
 - proxy nodes 45
- VSS backup 2
 - overview 6
 - planning requirements 2
- VSS fast restore
 - method 19
- VSS instant restore
 - method 19
- VSS planning 2
- VSS properties 58
- VSS provider 1
- VSS Requestor 1
- VSS restore
 - command-line considerations 226
- VSS writer 1
- vsspolicy option 63
- VSSPOLICY, statements 14

W

- wait for tape mounts for file information, GUI 112
- wait for tape mounts options, GUI 112
- Windows authentication mode, setting
 - using the CLI 167, 210, 240
- Windows Server Core
 - backing up SQL databases 106
 - backing up SQL databases with Data Protection for SQL Server 124
 - changing configuration values 79, 125
 - restoring SQL databases with Data Protection for SQL Server 124

- Windows Server Core (*continued*)
 - restoring SQL Server databases 121
 - using command-line help 95
 - using the backup command 106, 124
 - using the help command 95
 - using the restore command 121, 124
 - using the set command 79, 125
- Windows Server Core configuration
 - configuring Data Protection for SQL Server 77
 - configuring the backup-archive client 76
 - creating a client node 73
 - setting up a proxy node for offloaded backups 75

X

- XIV
 - requirements 21



Product Number: 5608-E04

Printed in USA