

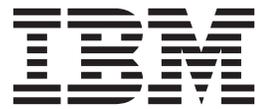
IBM Tivoli Storage Manager for Databases
Version 6.3

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



IBM Tivoli Storage Manager for Databases
Version 6.3

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



Note

Before using this information and the product it supports, read the information in "Notices" on page 249.

This edition applies to Version 6.3 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, 2011.**

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

Contents

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

Preface	ix
--------------------------	-----------

Who should read this guide	ix
Publications	x
Tivoli Storage Manager publications	x
Tivoli Storage FlashCopy Manager publications	xiii
Support information	xiii
Getting technical training	xiii
Searching knowledge bases	xiv
Contacting IBM Software Support	xvi
Conventions used in this book	xvii
Typeface conventions	xviii
Reading syntax diagrams	xviii

New in Version 6.3	xxi
-------------------------------------	------------

Additional Data Protection for SQL features	xxi
---	-----

Chapter 1. Data Protection for SQL

Overview	1
---------------------------	----------

Backup overview	1
Legacy Backup	1
VSS Backup	1
Restore methods	7
Legacy Restore	7
VSS Restore.	7
VSS Fast Restore	8
VSS Instant Restore	8
Restore types	9
SAN Volume Controller and Storwize V7000 FlashCopy support	10

Chapter 2. Planning for Data Protection for Microsoft SQL Server operations . . 19

How Tivoli Storage Manager server policy affects Data Protection for SQL	19
Security	21
Backup strategies	21
Using VSS and Legacy Backups together	24
Using VSS operations in a SQL Server Failover Cluster environment	25
Using Data Protection for SQL in a Microsoft Windows Failover Clustering environment	27
Using Data Protection for SQL in a Veritas Cluster Server environment	28
Back up to Tivoli Storage Manager storage versus back up to local shadow volumes	28
VSS operations in DS, SAN Volume Controller, and Storwize V7000 environments.	29
VSS operations in IBM N-series and NetApp environments.	30
Recommended Tivoli Storage Manager policy settings.	31
Domain.	31

Policy Set	31
Management Class	32
Copy Group	32
Storage Pool	34
Data Protection for SQL node name: Recommended settings.	34
Proxy node definitions (VSS Backups)	35
Required node names for basic VSS operations	35
Required node names for VSS off-loaded backups	36
Specifying Data Protection for SQL options.	37
Specifying Data Protection for SQL preferences	40
Setting automatic expiration (VSS and Legacy)	41

Chapter 3. Installing Data Protection for SQL Server 45

Quick installation and configuration	45
Installation prerequisites	47
Minimum hardware requirements	47
Minimum software and operating system requirements	48
Virtualization support	49
Install Data Protection for SQL on a local system	49
Installing Tivoli Storage FlashCopy Manager	50
Installing Data Protection for SQL in an MSCS or VCS environment	50
Install the language packs for Data Protection for SQL	50
Installing Data Protection for SQL silently	52
Silently installing Data Protection for SQL with the setup program	54
Silently installing Data Protection for SQL with MSI (msiexec.exe)	55
Installation problems: capturing a log of the installation	57
Creating the package on a DVD or a file server	58
Playing back the silent installation.	58
Setup error messages	59

Chapter 4. Configuring Data Protection for SQL 61

Manual configuration	61
1. Configure Data Protection for SQL manually	62
2. Configure the Tivoli Storage Manager server manually	63
3. Perform these tasks on the machine running the offloaded backups	63
4. Perform these tasks to verify your configuration	64

Chapter 5. Protecting Microsoft SQL Server data 67

Setting user preferences	67
Data Protection Properties	67
Managing policy.	75

Policy binding statements	76
How backups expire based on policy	76
Binding backups to a policy	77
Determining managed storage capacity	77
Using the Task Manager	78
Backing up SQL databases and files	78
Security	79
Backup types	79
Backing up SQL databases by using the VSS method	80
Backing up SQL databases by using the Legacy method	81
Backing up SQL groups or files by using the Legacy method	83
Restoring SQL databases and files	83
Restoring SQL server data	83
VSS Restore considerations	85
Restoring SQL file groups and files from Legacy backups	85
Restoring to an alternate machine	86
Restoring the master database	87
Restore options	88
Inactivating SQL databases (Legacy only)	90
Inactivate tree and list	91
Deleting SQL Server Backups	91
Using the Tivoli Storage Manager scheduler	91
Example scheduler procedure	92
Scheduler guidelines	97
Setting user mode	98
Automating tasks	99
Additional automation methods	99
Viewing, printing, and saving reports	100

Chapter 6. Troubleshooting Data Protection for SQL with VSS backup-restore support 101

Problem determination assistance	101
Resolving errors during Data Protection for SQL processing	102
Determining if the problem is a Tivoli Storage Manager or SQL issue	102
Determining if the problem is a Data Protection for SQL issue or a general VSS issue	103
Tracing the Data Protection client when using SQL VSS technology	106
Gathering SQL with VSS information before calling IBM	106
Gathering files from SQL with VSS before calling IBM	107
Troubleshooting Data Protection for SQL VSS and SAN Volume Controller, Storwize V7000, or DS6000/DS8000	108
Installation Problems: Creating an installation-log file	110
Emailing support files	110
Online IBM support	111
Viewing system information	111
Viewing trace and log files	112

Chapter 7. Performance 115

Chapter 8. Data Protection for Microsoft SQL Server reference 117

Command overview	117
Command line parameter characteristics	118
Backup command	121
Backup syntax	121
Backup positional parameters	124
Backup optional parameters	127
Legacy Backup output examples	134
VSS Backup output examples	143
Query command	146
Query syntax	147
Query positional parameters	149
Query optional parameters	151
Query output examples	156
Query Managedcapacity command	167
Restore command	168
VSS Restore command-line considerations	169
Restore syntax	170
Restore positional parameters	173
Restore optional parameters	175
Legacy Restore output examples	189
VSS Restore output examples	196
Restorefiles command	199
Restorefiles syntax	200
Restorefiles positional parameters	201
Restorefiles optional parameters	201
Restorefiles examples	204
Delete Backup command	204
Delete Backup syntax	204
Delete Backup positional parameters	205
Delete Backup optional parameters	205
Delete Backup example	208
Inactivate command (Legacy only)	208
Inactivate syntax	208
Inactivate positional parameters	210
Inactivate optional parameters	211
Inactivate output examples	214
Help command	218
Help syntax	218
Help positional parameters	219
Help output examples	219
Set command	222
Set syntax	223
Set positional parameters	223
Set optional parameters	229
Set output examples	229
Changetsmpassword command	230
Changetsmpassword	230
Changetsmpassword positional parameters	230
Changetsmpassword optional parameters	231
Changetsmpassword output examples	233
Query Policy command	233
Examples of Tivoli Storage Manager policy binding using include/exclude and VSSPOLICY statements	233
Transitioning SQL Server backups from Tivoli Storage FlashCopy Manager to Tivoli Storage Manager	237
Using the Tivoli Storage Manager server wizard	237
Implement these tasks on the Tivoli Storage Manager server	237

Implement these tasks on the workstation running the SQL Server	238	Notices	249
Appendix A. Frequently asked questions	241	Trademarks	250
Appendix B. Accessibility features for Data Protection for SQL	247	Glossary	251
		Index	273

Tables

1. Tivoli Storage Manager server publications	x	9. Backup strategy characteristics	25
2. Tivoli Storage Manager storage agent publications	xi	10. Required node names for basic VSS operations	35
3. Tivoli Storage Manager client publications	xi	11. Required node names for VSS off-loaded backups.	36
4. Tivoli Storage Manager data protection publications	xi	12. Silent installation options	52
5. IBM Tivoli Storage Manager troubleshooting and tuning publications	xii	13. Silent installation features (base client only)	53
6. Tivoli Storage FlashCopy Manager publications	xiii	14. Silent installation features (Language Packages only).	53
7. Configuration scenarios for Data Protection for SQL and SAN Volume Controller and Storwize V7000	13	15. Silent installation transforms.	53
8. Snapshot restore and delete behavior on SAN Volume Controller and Storwize V7000 space-efficient target volumes	15	16. Commands for creating an installation package	58
		17. Available property pages for SQL Server workloads	68
		18. Diagnostics modes and their usage.	70
		19. Data Protection for SQL commands	118
		20. Data Protection for SQL optional parameters	119

Preface

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

Data Protection for SQL is a comprehensive storage management software product. This book explains how to install, configure, and administer Data Protection for SQL.

Data Protection for SQL performs online backups of Microsoft SQL Server databases to Tivoli Storage Manager storage.

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

Throughout this document, the term *Windows* (unless otherwise specified) refers to the following operating systems:

- Windows Server 2003
- Windows Server 2003 R2
- Windows Server 2008
- Windows Server 2008 R2

Throughout this document, the term SQL Server (unless otherwise specified) refers to the following products:

- SQL Server 2005 (32-bit and x64)
- SQL Server 2008 (32-bit and x64)
- SQL Server 2008 R2 (32-bit and x64)

Throughout this document, the term *Windows VSS System Provider* (unless otherwise specified) refers to the standard Windows System provider.

Who should read this guide

This publication is intended for system installers, 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 2003
- Windows Server 2003 R2
- Windows Server 2008
- Windows Server 2008 R2

It is also assumed that you have an understanding of the following IBM storage system used for the database:

- IBM System Storage® Disk Storage Models DS3000, DS4000®, and DS5000
- IBM System Storage SAN Volume Controller (SVC)
- IBM Storwize® V7000 Disk System
- IBM XIV® Storage System
- IBM System Storage DS8000® (DS8100, DS8300, or DS8700)
- Any storage devices that implement the VSS provider interface. See “VSS Backup” on page 1 for more information.

Publications

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

To search all publications, go to the Tivoli Storage Manager information center at <http://publib.boulder.ibm.com/infocenter/tsminfo/v6r3>.

You can download PDF versions of publications from the Tivoli Storage Manager information center or from the IBM Publications Center at <http://www.ibm.com/shop/publications/order/>.

Go to Tivoli Documentation Central to find information centers that contain official product documentation for current and previous versions of Tivoli products, including the Tivoli Storage Manager product family. You can find Tivoli Documentation Central at <https://www.ibm.com/developerworks/wikis/display/tivolidoccentral/Home>.

You can also order some related publications from the IBM Publications Center website. The website provides information about ordering publications from countries other than the United States. In the United States, you can order publications by calling 1-800-879-2755.

Tivoli Storage Manager publications

The following tables list the publications that make up the Tivoli Storage Manager library.

Table 1. Tivoli Storage Manager server publications

Publication title	Order number
<i>IBM Tivoli Storage Manager for AIX Installation Guide</i>	GC23-9781
<i>IBM Tivoli Storage Manager for AIX Administrator's Guide</i>	SC23-9769
<i>IBM Tivoli Storage Manager for AIX Administrator's Reference</i>	SC23-9775
<i>IBM Tivoli Storage Manager for HP-UX Installation Guide</i>	GC23-9782
<i>IBM Tivoli Storage Manager for HP-UX Administrator's Guide</i>	SC23-9770
<i>IBM Tivoli Storage Manager for HP-UX Administrator's Reference</i>	SC23-9776
<i>IBM Tivoli Storage Manager for Linux Installation Guide</i>	GC23-9783
<i>IBM Tivoli Storage Manager for Linux Administrator's Guide</i>	SC23-9771
<i>IBM Tivoli Storage Manager for Linux Administrator's Reference</i>	SC23-9777

Table 1. Tivoli Storage Manager server publications (continued)

Publication title	Order number
<i>IBM Tivoli Storage Manager for Oracle Solaris Installation Guide</i>	GC23-9784
<i>IBM Tivoli Storage Manager for Oracle Solaris Administrator's Guide</i>	SC23-9772
<i>IBM Tivoli Storage Manager for Oracle Solaris Administrator's Reference</i>	SC23-9778
<i>IBM Tivoli Storage Manager for Windows Installation Guide</i>	GC23-9785
<i>IBM Tivoli Storage Manager for Windows Administrator's Guide</i>	SC23-9773
<i>IBM Tivoli Storage Manager for Windows Administrator's Reference</i>	SC23-9779
<i>IBM Tivoli Storage Manager for z/OS Media Installation and User's Guide</i>	SC27-4018
<i>IBM Tivoli Storage Manager Upgrade and Migration Guide for V5 Servers</i>	GC27-4017
<i>IBM Tivoli Storage Manager Integration Guide for Tivoli Storage Manager FastBack®</i>	SC27-2828

Table 2. Tivoli Storage Manager storage agent publications

Publication title	Order number
<i>IBM Tivoli Storage Manager for SAN for AIX Storage Agent User's Guide</i>	SC23-9797
<i>IBM Tivoli Storage Manager for SAN for HP-UX Storage Agent User's Guide</i>	SC23-9798
<i>IBM Tivoli Storage Manager for SAN for Linux Storage Agent User's Guide</i>	SC23-9799
<i>IBM Tivoli Storage Manager for SAN for Oracle Solaris Storage Agent User's Guide</i>	SC23-9800
<i>IBM Tivoli Storage Manager for SAN for Windows Storage Agent User's Guide</i>	SC23-9553

Table 3. Tivoli Storage Manager client publications

Publication title	Order number
<i>IBM Tivoli Storage Manager for UNIX and Linux: Backup-Archive Clients Installation and User's Guide</i>	SC23-9791
<i>IBM Tivoli Storage Manager for Windows: Backup-Archive Clients Installation and User's Guide</i>	SC23-9792
<i>IBM Tivoli Storage Manager Using the Application Programming Interface</i>	SC23-9793
<i>IBM Tivoli Storage Manager for Space Management for UNIX and Linux: User's Guide</i>	SC23-9794
<i>IBM Tivoli Storage Manager HSM for Windows Administration Guide</i>	SC23-9795

Table 4. Tivoli Storage Manager data protection publications

Publication title	Order number
<i>IBM Tivoli Storage Manager for Databases: Data Protection for Microsoft SQL Server Installation and User's Guide</i>	GC27-4010
<i>IBM Tivoli Storage Manager for Databases: Data Protection for Oracle for UNIX and Linux Installation and User's Guide</i>	SC27-4019

Table 4. Tivoli Storage Manager data protection publications (continued)

Publication title	Order number
<i>IBM Tivoli Storage Manager for Databases: Data Protection for Oracle for Windows Installation and User's Guide</i>	SC27-4020
<i>IBM Tivoli Storage Manager for Mail: Data Protection for Microsoft Exchange Server Installation and User's Guide</i>	GC27-4009
<i>IBM Tivoli Storage Manager for Mail: Data Protection for Lotus Domino® UNIX and Linux Installation and User's Guide</i>	SC27-4021
<i>IBM Tivoli Storage Manager for Mail: Data Protection for Lotus Domino for Windows Installation and User's Guide</i>	SC27-4022
<i>IBM Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP Installation and User's Guide for DB2</i>	SC33-6341
<i>IBM Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP Installation and User's Guide for Oracle</i>	SC33-6340
<i>IBM Tivoli Storage Manager for Virtual Environments Installation and User's Guide</i>	SC27-2898
<i>IBM Tivoli Storage Manager for Microsoft SharePoint Guide</i>	N/A

Table 5. IBM Tivoli Storage Manager troubleshooting and tuning publications

Publication title	Order number
<i>IBM Tivoli Storage Manager Problem Determination Guide</i>	GC23-9789
<i>IBM Tivoli Storage Manager Performance Tuning Guide</i>	GC23-9788
<i>IBM Tivoli Storage Manager Client Messages and Application Programming Interface Return Codes</i>	SC27-2878
<i>IBM Tivoli Storage Manager Server Messages and Error Codes</i>	SC27-2877
<i>IBM Tivoli Storage Manager for Mail: Data Protection for Microsoft Exchange Server Messages</i>	GC27-4011
<i>IBM Tivoli Storage Manager for Databases: Data Protection for Microsoft SQL Server Messages</i>	GC27-4012
<i>IBM Tivoli Storage Manager for Databases: Data Protection for Oracle Messages</i>	SC27-4014
<i>IBM Tivoli Storage Manager for Mail: Data Protection for Lotus Domino Messages</i>	SC27-4015
<i>IBM Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP Messages</i>	SC27-4016

Note: You can find information about IBM System Storage Archive Manager at http://publib.boulder.ibm.com/infocenter/tsminfo/v6r3/c_complydataretention_ovr.html.

Tivoli Storage FlashCopy Manager publications

The following table lists the publications that make up the Tivoli Storage FlashCopy Manager library.

Table 6. Tivoli Storage FlashCopy Manager publications

Publication title	Order number
<i>IBM Tivoli Storage FlashCopy Manager for UNIX and Linux Installation and User's Guide</i>	SC27-4005
<i>IBM Tivoli Storage FlashCopy Manager for Windows Installation and User's Guide</i>	SC27-4006
<i>IBM Tivoli Storage FlashCopy Manager for VMware Installation and User's Guide</i>	SC27-4007
<i>IBM Tivoli Storage FlashCopy Manager Messages</i>	GC27-4008

Support information

You can find support information for IBM products from various sources.

Start at the IBM Support Portal: <http://www.ibm.com/support/entry/portal/>. You can select the products that you are interested in and search for a wide variety of relevant information.

Getting technical training

Information about Tivoli technical training courses is available online.

Go to the following websites to sign up for training, ask questions, and interact with others who use IBM storage products.

Tivoli software training and certification

Choose from instructor led, online classroom training, self-paced Web classes, Tivoli certification preparation, and other training options at <http://www.ibm.com/software/tivoli/education/>

Tivoli Support Technical Exchange

Technical experts share their knowledge and answer your questions in webcasts at http://www.ibm.com/software/sysmgmt/products/support/supp_tech_exch.html.

Storage Management community

Interact with others who use IBM storage management products at <http://www.ibm.com/developerworks/servicemanagement/sm/index.html>

Global Tivoli User Community

Share information and learn from other Tivoli users throughout the world at <http://www.tivoli-ug.org/>.

IBM Education Assistant

View short "how to" recordings designed to help you use IBM software products more effectively at <http://publib.boulder.ibm.com/infocenter/ieduasst/tivv1r0/index.jsp>

Searching knowledge bases

If you have a problem with your Tivoli Storage Manager family product, there are several knowledge bases that you can search.

Begin by searching the Tivoli Storage Manager Information Center at <http://publib.boulder.ibm.com/infocenter/tsminfo/v6r3>. From this website, you can search the current Tivoli Storage Manager documentation.

Searching the Internet

If you cannot find an answer to your question in the IBM Tivoli Storage Manager information center, search the Internet for the information that might help you resolve your problem.

To search multiple Internet resources, go to the IBM support website at <http://www.ibm.com/support/entry/portal/>.

You can search for information without signing in. Sign in using your IBM ID and password if you want to customize the site based on your product usage and information needs. If you do not already have an IBM ID and password, click **Sign in** at the top of the page and follow the instructions to register.

From the support website, you can search various resources including:

- IBM technotes
- IBM downloads
- IBM Redbooks® publications
- IBM Authorized Program Analysis Reports (APARs)

Select the product and click **Downloads** to search the APAR list.

If you still cannot find a solution to the problem, you can search forums and newsgroups on the Internet for the latest information that might help you find problem resolution.

An independent user discussion list, ADSM-L, is hosted by Marist College. You can subscribe by sending an email to listserv@vm.marist.edu. The body of the message must contain the following text: `SUBSCRIBE ADSM-L your_first_name your_family_name`.

To share your experiences and learn from others in the Tivoli Storage Manager and Tivoli Storage FlashCopy Manager user communities, go to the following wikis:

Tivoli Storage Manager wiki

<http://www.ibm.com/developerworks/wikis/display/tivolistoragemanager>

Tivoli Storage FlashCopy Manager wiki

[https://www.ibm.com/developerworks/mydeveloperworks/wikis/home/wiki/Tivoli Storage FlashCopy Manager](https://www.ibm.com/developerworks/mydeveloperworks/wikis/home/wiki/Tivoli%20Storage%20FlashCopy%20Manager)

Using IBM Support Assistant

IBM Support Assistant is a complimentary software product that can help you with problem determination. It is available for some Tivoli Storage Manager and Tivoli Storage FlashCopy Manager products.

To learn about which products are supported, go to the IBM Support Assistant download web page at <http://www.ibm.com/software/support/isa/download.html>.

IBM Support Assistant helps you gather support information when you must open a problem management record (PMR), which you can then use to track the problem. The product-specific plug-in modules provide you with the following resources:

- Support links
- Education links
- Ability to submit problem management reports

You can find more information at the IBM Support Assistant website:

<http://www.ibm.com/software/support/isa/>

You can also install the stand-alone IBM Support Assistant application on any workstation. You can then enhance the application by installing product-specific plug-in modules for the IBM products that you use. Find add-ons for specific products at <http://www.ibm.com/support/docview.wss?uid=swg27012689>.

Finding product fixes

A product fix to resolve your problem might be available from the IBM software support website.

You can determine what fixes are available by checking the IBM software support website at <http://www.ibm.com/support/entry/portal/>.

- If you previously customized the site based on your product usage:
 1. Click the link for your product, or a component for which you want to find a fix.
 2. Click **Downloads**, and then click **Fixes by version**.
- If you have not customized the site based on your product usage, click **Downloads** and search for your product.

Receiving notification of product fixes

You can receive notifications about fixes, flashes, upgrades, and other news about IBM products.

To sign up to receive notifications about IBM products, follow these steps:

1. From the support page at <http://www.ibm.com/support/entry/portal/>, click **Sign in to create, manage, or view your subscriptions** in the **Notifications** pane.
2. Sign in using your IBM ID and password. If you do not have an ID and password, click **register now** and complete the registration process.
3. Click **Manage all my subscriptions** in the **Notifications** pane.
4. Click the **Subscribe** tab and then click **Tivoli**.
5. Select the products for which you want to receive notifications and click **Continue**.

6. Specify your notification preferences and click **Submit**.

Contacting IBM Software Support

You can contact IBM Software Support if you have an active IBM subscription and support contract and if you are authorized to submit problems to IBM.

To obtain help from IBM Software Support, complete the following steps:

1. Ensure that you have completed the following prerequisites:
 - a. Set up a subscription and support contract.
 - b. Determine the business impact of your problem.
 - c. Describe your problem and gather background information.
2. Follow the instructions in “Submitting the problem to IBM Software Support” on page xvii.

Setting up a subscription and support contract

Set up a subscription and support contract. The type of contract that you need depends on the type of product you have.

For IBM distributed software products (including, but not limited to, IBM Tivoli, Lotus®, and Rational® products, as well as IBM DB2® and IBM WebSphere® products that run on Microsoft Windows or on operating systems such as AIX or Linux), enroll in IBM Passport Advantage® in one of the following ways:

- **Online:** Go to the Passport Advantage website at <http://www.ibm.com/software/lotus/passportadvantage/>, click **How to enroll**, and follow the instructions.
- **By telephone:** You can call 1-800-IBMSERV (1-800-426-7378) in the United States. For the telephone number to call in your country, go to the IBM Software Support Handbook web page at <http://www14.software.ibm.com/webapp/set2/sas/f/handbook/home.html> and click **Contacts**.

Determining the business impact

When you report a problem to IBM, you are asked to supply a severity level. Therefore, you must understand and assess the business impact of the problem you are reporting.

Severity 1	Critical business impact: You are unable to use the program, resulting in a critical impact on operations. This condition requires an immediate solution.
Severity 2	Significant business impact: The program is usable but is severely limited.
Severity 3	Some business impact: The program is usable with less significant features (not critical to operations) unavailable.
Severity 4	Minimal business impact: The problem causes little impact on operations, or a reasonable circumvention to the problem has been implemented.

Describing the problem and gathering background information

When explaining a problem to IBM, it is helpful to be as specific as possible. Include all relevant background information so that IBM Software Support specialists can help you solve the problem efficiently.

To save time, know the answers to these questions:

- What software versions were you running when the problem occurred?
- Do you have logs, traces, and messages that are related to the problem symptoms? IBM Software Support is likely to ask for this information.
- Can the problem be re-created? If so, what steps led to the failure?
- Have any changes been made to the system? For example, hardware, operating system, networking software, and so on.
- Are you using a workaround for this problem? If so, be prepared to explain it when you report the problem.

Submitting the problem to IBM Software Support

You can submit the problem to IBM Software Support online or by telephone.

Online

Go to the IBM Software Support website at [http://www.ibm.com/support/entry/portal/Open_service_request/Software/Software_support_\(general\)](http://www.ibm.com/support/entry/portal/Open_service_request/Software/Software_support_(general)). Sign in to access IBM Service Requests and enter your information into the problem submission tool.

By telephone

For the telephone number to call in your country, go to the IBM Software Support Handbook at <http://www14.software.ibm.com/webapp/set2/sas/f/handbook/home.html> and click **Contacts**.

Conventions used in this book

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

This guide uses the following typeface conventions:

Bold

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

italics

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

bold italics

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

monospace

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

UPPER CASE

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

Typeface conventions

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.	▶—KEYWOrd—▶◀																		
Symbols: Enter these symbols exactly as they appear in the syntax diagram.	<table><tbody><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></tbody></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
<p>Variables:</p> <p>Italicized lowercase items (<i>var_name</i>) denote variables.</p> <p>In this example, you can specify a <i>var_name</i> when you enter the KEYWORD command.</p>	<p>The diagram shows a horizontal line with a double arrow on the left and a double arrow on the right. The text "KEYWORD" is written above the line, and "var_name" is written below the line.</p>
<p>Repetition:</p> <p>An arrow returning to the left means you can repeat the item.</p> <p>A character or space within the arrow means you must separate repeated items with that character or space.</p> <p>A footnote by the arrow references the number of times you can repeat the item.</p>	<p>The first diagram shows a horizontal line with a double arrow on the left and a double arrow on the right. The text "repeat" is written below the line. A curved arrow starts from the right side of the line and points back to the left side.</p> <p>The second diagram is identical to the first, but with a comma "," written inside the curved arrow.</p> <p>The third diagram is identical to the first, but with "(1)" written below the line.</p>
<p>Notes:</p> <p>1 Specify <i>repeat</i> as many as 5 times.</p>	
<p>Required Choices:</p> <p>When two or more items are in a stack and one of them is on the line, you <i>must</i> specify one item.</p> <p>In this example, you <i>must</i> choose A, B, or C.</p>	<p>The diagram shows a horizontal line with a double arrow on the left and a double arrow on the right. A vertical stack of three boxes labeled "A", "B", and "C" is positioned below the line. A bracket on the right side of the stack connects the three boxes to the line.</p>
<p>Optional Choice:</p> <p>When an item is below the line, that item is optional. In the first example, you can choose A or nothing at all.</p> <p>When two or more items are in a stack below the line, all of them are optional. In the second example, you can choose A, B, C, or nothing at all.</p>	<p>The first diagram shows a horizontal line with a double arrow on the left and a double arrow on the right. A single box labeled "A" is positioned below the line, with a bracket on its right side connecting it to the line.</p> <p>The second diagram is identical to the first, but with a vertical stack of three boxes labeled "A", "B", and "C" below the line, with a bracket on the right side connecting the stack to the line.</p>
<p>Defaults:</p> <p>Defaults are above the line. The default is selected unless you override it. You can override the default by including an option from the stack below the line.</p> <p>In this example, A is the default. You can override A by choosing B or C. You can also specify the default explicitly.</p>	<p>The diagram shows a horizontal line with a double arrow on the left and a double arrow on the right. A vertical stack of three boxes labeled "A", "B", and "C" is positioned below the line. A bracket on the right side of the stack connects the three boxes to the line. Additionally, a box labeled "A" is positioned above the line, with a bracket on its left side connecting it to the line.</p>
<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>The diagram shows a horizontal line with a double arrow on the left and a double arrow on the right. A vertical stack of three boxes labeled "A", "B", and "C" is positioned below the line. A bracket on the right side of the stack connects the three boxes to the line. A curved arrow starts from the right side of the line and points back to the left side.</p>

Syntax Diagram Description	Example
<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>►► The fragment name ◀◀</p> <p>The fragment name:</p> 

New in Version 6.3

The following features are new for Data Protection for SQL Version 6.3:

New MMC GUI

A new MMC GUI that enhances the user experience of Data Protection for SQL is available. The following features are available in the new MMC GUI:

- A new backup and restore interface with MMC GUI integration.
- Added filtering and refresh options.
- A new Task Manager to process query, backup, and restore operations as individual tasks. Multiple tasks can be submitted simultaneously.
- New configuration wizards to assist with setting up Data Protection for SQL.
- A new integrated interface for configuration properties.
- New dashboard views that provide graphical views of Data Protection for SQL activities.
- New local scheduling capabilities such as the schedule mode, in addition to default interactive mode. By using the schedule mode you can schedule operations as easily as interactively issuing the command in the MMC GUI.

Restore VSS backups to flat files

Use the **restorefiles** command to restore VSS backups to flat files without involving the SQL Server.

Data Protection for Microsoft SQL Server 2008 backup support

- Specify SQL native backup compression from the SQL Properties window.
- From the command line, use `SQLCOMPRESS` to set SQL native backup compression.
- On the MMC GUI, the SQL Workload column on the Recover tab shows the SQL compression status for legacy backups.

Enhanced Statistics

- The backup and restore statistics have been enhanced to include compression, LAN-free and client-side deduplication information
- The query backup commands have been enhanced to display information regarding backup encryption, compression and client-side deduplication status.

Additional Data Protection for SQL features

With Data Protection for SQL, you can back up and restore Microsoft SQL Server online databases to Tivoli Storage Manager server storage by using either the command-line or graphical user interface (GUI).

Data Protection for SQL helps you protect and manage SQL Server data with these tasks:

- Legacy full backups, differential backups, and transaction log backups of SQL Server databases.

- Legacy individual file group backups, individual file backups, and set backups (containing sets of file groups and files) of SQL Server data.
- Full VSS snapshot backups of SQL Server databases running on Windows Server.
- Full VSS offloaded snapshot backups of SQL Server databases running on Windows Server.
- Legacy restores of full, differential, and transaction log backups of SQL Server databases.
- Legacy restores of individual file groups, individual files, and set backups (containing sets of file groups and files) of SQL Server data.
- Partial Legacy restores of specific database file groups.
- VSS Restores of full snapshot backups of SQL Server databases (running on Windows Server) from Tivoli Storage Manager server storage or local VSS disks.
- Protect multiple SQL Server instances on the same system.

Note: You can access only one SQL Server per execution of Data Protection for SQL from the command line.

- Legacy backups that use data striping in parallel threads by using parallel sessions.
- Automate scheduled backups.
- Perform expanded restore operations on backup objects such as relocating, restoring to named marks, and partially restoring full backups.
- Restore Legacy backups for a point in time.
- Restore Legacy database backups to a different SQL Server.
- Retain, with a Legacy backup, the information required to recreate or move SQL databases or files, such as sort order, code page, and Unicode information, or file group and file logical and physical names. The meta object information is retained on the Tivoli Storage Manager server separately from the backup data objects.
- Inactivate all active Legacy backup objects, all objects of a particular Legacy backup type, or specific objects.
- Inactivate Legacy backup objects older than a specified number of days.
- Set automatic expiration of backup objects based on version limit and retention period. See “Recommended Tivoli Storage Manager policy settings” on page 31.
- Query any local SQL Server or any connected Tivoli Storage Manager server for database, status, and configuration information.
- Monitor results through the Data Protection for SQL activity log and automatically prune the activity log.
- Set Tivoli Storage Manager connection information options to Tivoli Storage Manager servers.
- Set Tivoli Storage Manager security and performance options.
- Participate in Microsoft Windows Failover Clustering (previously MSCS) and Veritas Cluster Server (VCS) failover clusters.
- Apply failover clustering without unclustering. Failover clustering is useful for both maintenance of and restoration of the master database.
- Obtain online context-sensitive task, and concept help.
- Globalization Support.
- Use the **restorefiles** command to restore VSS backups to flat files without involving the SQL Server.

Except where otherwise indicated, throughout this publication:

- SQL Server refers to SQL Server 2005, SQL Server 2008, or SQL Server 2008 R2.
- Windows Server refers to Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, or Windows Server 2008 R2.

Chapter 1. Data Protection for SQL Overview

Product overview information related to new features, backup methods, backup types, restore methods, restore types, performance, security, and available documentation is provided for Data Protection for SQL 6.3.

Backup overview

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

Legacy Backup

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

Data Protection for SQL provides selection mechanisms and the logic that are required to back up and restore SQL data. For example, when you initiate a backup operation, Data Protection for SQL:

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

Note:

1. Data Protection for SQL can compress SQL data before sending it to the Tivoli Storage Manager server.
2. **Metadata:**

When a backup is performed, Data Protection for SQL retains information about the SQL Server and database. This information is available for query and restore operations after the backup is completed. The information about the names and sizes of the database file groups and files is stored along with the database data, as a sub-object. This sub-object is referred to as metadata.

VSS Backup

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

A VSS Backup uses Microsoft Volume Shadow Copy Service technology to produce an online snapshot (point-in-time consistent copy) of SQL data that can be stored on local shadow volumes or on Tivoli Storage Manager server storage. Both of these storage destinations require that sufficient space be available for the snapshot. A VSS Backup means the SQL server is not in "backup mode" for an extended period of time because the length of time to perform the snapshot is usually measured in seconds and not hours. In addition, a VSS Backup allows a snapshot of large amounts of data at one time since the snapshot works at the volume level.

Optionally, VSS Backups can be stored locally on VSS shadow volumes that are directly accessible by the SQL system, as long as sufficient space is available for the snapshot. These types of backups are fast because data is not placed into Tivoli Storage Manager server storage. Restoring these backups is also fast because the SQL data is not transferred from Tivoli Storage Manager server storage over the network.

When performing VSS Backups and moving data to Tivoli Storage Manager server storage, sufficient space on local snapshot volumes is still required to hold the snapshot. For SQL data backed up to Tivoli Storage Manager server storage, the SQL data on the snapshot volume is sent to the Tivoli Storage Manager server. Once the data transfer to the server is complete, the snapshot volume is made available for reuse. If you are storing VSS Backups locally and the maximum number of local backup versions to be maintained (as specified by the Tivoli Storage Manager policy) is reached, the oldest backup version is expired in order to create the snapshot for the backup to Tivoli Storage Manager server storage. See “How Tivoli Storage Manager server policy affects Data Protection for SQL” on page 19 for details regarding how Tivoli Storage Manager proceeds in this situation.

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

For SQL data backed up to both destinations, a local snapshot backup is performed and the SQL data on the local snapshot volume is sent to the Tivoli Storage Manager server. The local snapshot volume is retained as a local backup.

For local VSS backups, you must have a licensed version of Tivoli Storage FlashCopy Manager installed on your system.

Characteristics of VSS Backups

The following characteristics are true of VSS Backup:

- Full backups only are supported. Log, differential, file, group, and set backup types are not supported. However, Legacy differential and Legacy log backups can be applied after a full VSS Backup has been restored.
- Backup granularity is at the database level only.
- Backups are managed through Tivoli Storage Manager server policy.
- Backups can be stored on local shadow volumes, Tivoli Storage Manager server storage, or both locations.
- Different policy settings can be defined for each storage location and backup method.
- Backups to Tivoli Storage Manager server storage can be off-loaded to an alternate machine as resource relief for production servers.
- Backups can be performed in a Microsoft Windows Failover Clustering or Veritas Cluster Server environment.

Planning requirements for VSS Backups

Consider the following requirements when planning for VSS Backups:

- When running VSS operations, make sure you have at least 200 megabytes of free disk space on your Windows System Drive. This space is used to hold the metadata files for Data Protection for SQL.
- Continue to schedule and perform Legacy backups in your strategy.

- Make sure you have a well-defined and tested recovery plan that meets your service level objectives.
- Use single hardware LUNs for each database or group of databases that will be backed up and restored together as a unit.
- Use basic disks.
- If you plan to keep some VSS snapshot backups on local shadow volumes only, make sure to consider the VSS provider-specific implementation and configuration options when setting up your strategy. For example, if your VSS hardware provider supports a full-copy snapshot versus a copy-on-write (COW) snapshot mechanism, be aware that full-copy type implementations have greater disk storage requirements but are less risky because they do not rely on the original volume to restore the data. COW implementations require much less disk storage but rely completely on the original volume to perform a restore. Since these implementations are entirely controlled by the VSS provider and not Data Protection for SQL, make sure to consult your VSS provider documentation for a complete understanding of your VSS implementation.
- If you must perform parallel VSS Backups, make sure to stagger the start of the backups for at least ten minutes. This action ensures the snapshot operations do not overlap. Failure to stagger the snapshots can result in errors.
- Do not place multiple volumes on the same LUN. Microsoft recommends that you configure a single volume/single partition/single LUN as 1 to 1 to 1.
- Do not set the ASNODENAME option in the dsm.opt file when using Data Protection for Microsoft SQL Server. Setting ASNODENAME can cause VSS backups and VSS restores to fail.

VSS Service overview

The VSS Service manages and directs three VSS software applications that are used during VSS operations.

VSS Writer

The VSS Writer is the Microsoft SQL Server. It is installed with the SQL Server software.

VSS Requestor

The VSS Requestor is Tivoli Storage Manager.

VSS Provider

The VSS provider manages the volumes where the SQL data resides. Configuration requirements are based upon the type of VSS provider used in your environment. For example:

- If you are using the standard Windows System provider (referred to as the Windows VSS System Provider throughout this book), no configuration is required.
- If you are using a VSS software provider, consult the documentation provided with your VSS software provider.
- If your SQL storage is configured on a disk storage subsystem and you want to leverage the hardware to manage the snapshots, you must install a VSS hardware provider.
- If you plan to perform VSS Instant Restores, be aware that IBM System Storage SAN Volume Controller, DS6000™, DS8000, Storwize V7000, and IBM XIV Storage System are the only storage subsystems that support VSS Instant Restores and require a VSS provider. Therefore, if you use IBM System Storage SAN Volume Controller, DS6000, DS8000, or Storwize V7000, you must install and configure IBM System Storage Support for Microsoft Virtual Disk and Volume Shadow Copy Services.

VSS provider support and documentation for IBM System Storage disk storage subsystems is available at the following Web site:
https://www.ibm.com/support/docview.wss?q1=ssg1*&dc=D400&rs=1114&uid=ssg1S4000372&context=HW2C2

If you use IBM XIV Storage System or later, you must install and configure IBM XIV Provider for Microsoft Windows Volume Shadow Copy Service. VSS provider support and documentation for IBM XIV Provider for Microsoft Windows Volume Shadow Copy Service is available by going to this website: <http://www.ibm.com/support/fixcentral/>. Then, search for an XIV VSS hardware provider that is suitable for the XIV storage system model that the Tivoli Storage Manager client uses. For example, specify the following values in these fields:

Product Group: Storage Systems
Product Family: Disk Systems
Product: XIV Storage System
Installed Version: 10.2.? or 11.0.0
Platform: Windows

For more information about VSS technology, see the Microsoft Technical Reference document "How Volume Shadow Copy Service Works" at the following website: [http://technet.microsoft.com/en-us/library/cc785914\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/cc785914(WS.10).aspx)

VSS System Provider:

A system provider assists with creating and maintaining copies on local shadow volumes. The VSS System Provider refers to the default VSS provider that is available with Windows Server.

If you are using the Windows VSS System Provider, no configuration tasks are required to perform VSS operations.

VSS Software or Hardware Provider:

A software or hardware provider (that is not the default Windows System provider) requires these configuration settings when planning for VSS Backups.

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

Backup types:

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

Data Protection for SQL provides six types of backup:

Full database backup (Legacy and VSS)

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

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

Differential backup (Legacy only)

Data Protection for SQL backs up only the data pages in a SQL Server database instance that have changed since the last full backup and a portion of the transaction log. This is equivalent to an incremental backup on the Tivoli Storage Manager backup-archive client.

Log backup (Legacy only)

Data Protection for SQL backs up only the contents of a SQL Server database transaction log since the last successful log backup. To do the first log backup, you need to have done a full backup or its equivalent first. Log backups normally follow full backups. The portion of the log included in full and differential backups is not equivalent to a log backup. Additionally, in full and differential backups, the log is not truncated as it is during a log backup. However, a log backup following a full or differential backup will include the same transactions as a full or differential. Log backups are not cumulative as are differential; they must be applied against a base backup and in the correct order.

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

File backup (Legacy only)

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

Group backup (Legacy only)

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

Set backup (Legacy only)

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

Examples

See “Legacy Backup output examples” on page 134 for samples of various backup types using the command line interface. Also, planning information regarding choosing the best strategy for your backup requirements is available in “Backup strategies” on page 21.

DS6000, DS8000, SAN Volume Controller, XIV, and Storwize V7000 requirements:

DS6000, DS8000, SAN Volume Controller, XIV, and Storwize V7000 require these configuration settings when planning for VSS Backups.

- Place databases files for each database or group of databases that will be backed up and restored together as a unit on their own dedicated logical volume.
- Place logs for each database or group of databases that will be backed up and restored together as a unit on their own logical volume.
- Do not place non-SQL data on storage volumes that are dedicated to SQL.
- When using hardware snapshot providers, do not share database LUNs with other databases or applications.
- If you delete a LOCAL snapshot that is stored on a SAN Volume Controller or Storwize V7000 Space Efficient volume (SEV) that has multiple dependent targets, you must delete them in the same order in which you created them. That is, you must delete the oldest one first, followed by the second oldest, and so on. Failure to delete them in this order can cause removal of other snapshots of the same source.
- (DS8000 only): In order to use incremental FlashCopy, the following two applications are also required for DS8000:
 - Tivoli Storage Manager backup-archive client Version 6.3 (or later)
 - IBM System Storage Support for Microsoft Virtual Disk and Volume Shadow Copy Services Version 4.1 (or later)

Note: You cannot restore from multiple Space-Efficient targets on DS8000.

XIV refers to XIV Gen 2 or later.

Offloaded VSS Backup

An offloaded backup uses another machine to move the data to the Tivoli Storage Manager server.

This type of backup shifts the backup load from the production machine to another machine. This frees the production system to serve the SQL server. This requires that a VSS hardware provider that supports transportable shadow copy volumes is installed on the production and secondary machines.

Offloaded VSS backups require a Tivoli Storage FlashCopy Manager license.

Restore methods

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

Legacy Restore

A Data Protection for SQL Legacy Restore obtains backup copies of all or part of one or more SQL databases and returns them to the SQL Server.

A complete restore of a database involves restoring a full backup or the equivalent thereof (from group, file, or set backups) and restoring *all* transaction logs since the last full backup.

For a Legacy Restore, Data Protection for SQL:

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

VSS Restore

A VSS Restore restores VSS Backups (SQL database files and log files) that reside on Tivoli Storage Manager server storage to their original location or to an alternate location.

The following characteristics are true of VSS Restores:

- You can only restore SQL Server VSS Backups to the same SQL Server instance.
- Only full backup types can be restored. Differential, individual file groups, individual files, and set backups are not supported by VSS and therefore, cannot be restored. However, Legacy differential and Legacy log backups can be applied after a full VSS Backup has been restored.
- Restore granularity is at the database level.
- Supports restoring one (or more) databases from a VSS snapshot backup located on Tivoli Storage Manager server storage.
- Supports restoring VSS Backups to an alternate location using the */relocatedir* option.
- Restores can be performed in a Microsoft Windows Failover Clustering or Veritas Cluster Server environment.

VSS Fast Restore

A VSS Fast Restore restores VSS Backups that reside on local shadow volumes.

In general, restore processing can conclude within minutes instead of hours in this situation. The following characteristics are true of VSS Fast Restore restores:

- Full backup types only can be restored. Differential, individual file groups, individual files, and set backups are not supported by VSS and therefore, cannot be restored. However, Legacy differential and Legacy log backups can be applied after a full VSS Backup has been restored.
- You can only restore SQL Server VSS Backups to the same SQL Server instance.
- Restore granularity is at the database level.
- VSS backups can be restored to an alternate location by using the `/relocatedir` option.
- Restores can be performed in a Microsoft Windows Failover Clustering or Veritas Cluster Server environment.

VSS Instant Restore

A VSS Instant Restore is when a set of target volumes (that contain a valid snapshot) are copied back to the original source volumes using hardware-assisted volume-level copy mechanisms. The application can return to normal operations as soon as the hardware-assisted volume-level copy has been started and the log replay is complete.

Be aware that a VSS Instant Restore is only possible when all of the data (from the database specified for restore) resides on storage subsystems supported by the VSS Instant Restore. If part of the data being restored (including the log files and full-text index files, if present) resides on a local disk, a VSS Instant Restore of this data is not supported. In this situation, a VSS Fast Restore is performed. In addition, you must make sure that any previous background copies (that involve the volumes being restored) are completed prior to initiating the VSS Instant Restore.

Although VSS Instant Restore is the default restore method when all SQL data specified for restore resides on storage subsystems supported by the VSS Instant Restore, a failover to VSS Fast Restore can occur when an error is detected early enough in the VSS Instant Restore process to trigger the failover. In this situation, an error is logged in the `dsmerror.log` file used by the `DSMAGENT`. However, a failover to VSS Fast Restore may not always be possible. For example, if an error occurs later in the restore process (such as a pending background copy on the storage subsystem, a failure to start the FlashCopy operation on the snapshot provider system, or other hardware error), VSS Instant Restore processing fails without a failover to VSS Fast Restore.

You can only restore SQL Server VSS Backups to the same SQL Server instance. Although VSS only supports full backups, Legacy differential and Legacy log backups can be applied after a full VSS Backup has been restored.

Be aware of the following considerations when planning for VSS Instant Restore:

- Requires IBM System Storage Support for Microsoft Volume Shadow Copy Service software.
- Backups must reside on the same DS6000, DS8000, SAN Volume Controller, XIV, or Storwize V7000 storage subsystem to which they are restored.

- Do not use version 2.2.2 of the IBM XIV Provider for Microsoft Windows Volume Shadow Copy Service with Tivoli Storage Manager. There are defects in that specific version that prevent VSS Instant Restore from working properly.

Restore types

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

Data Protection for SQL provides the same range of object types for restore as for backup. The following list contains descriptions of the available restore types:

Full database restore (Legacy and VSS)

The full database backup objects for the specified SQL databases are restored.

Differential restore (Legacy only)

Only the differential database backup objects for specified SQL databases are restored. Restore time is reduced as only the latest differential backup is restored (after its associated full backup is restored).

Log restore (Legacy only)

Log backup objects for the specified SQL databases are restored.

File restore (Legacy only)

The file backup objects needed from a full backup, file group backup, a file backup, or a set backup for the specified SQL databases are restored.

Group restore (Legacy only)

Only the group backup objects needed from a full backup, file group backup, a file backup, or a set backup for the specified SQL databases are restored.

Set restore (Legacy only)

Only set backup objects for the specified SQL databases are restored.

Depending on the backup strategy you choose, restoring a SQL database might involve restoring multiple backup objects from the Tivoli Storage Manager server. See “Backup strategies” on page 21.

Relocating files and other restores

In support of current SQL Server restore capabilities, Data Protection for SQL also provides the ability to relocate files during restore and to perform point-in-time restores, named-marks restores, or partial restores:

relocate

Allows you to move individual database files to a new location without having to first create the files.

point-in-time

Allows you to restore a transaction log backup to a specific SQL transaction date and time.

named-marks

Allows you to restore a transaction log backup to or before a named point, possibly after a specified point in time, and recover multiple related databases to the same named mark.

partial Allows you to restore just enough of a database into a temporary location to copy a specific table to the active database.

relocate dir

Allows you to move backed up SQL databases, logs, and SQL Server full-text index files to an alternate location.

Further Data Protection for SQL restore functions include the following:

- Restore a backup using the same number of data stripes used to create the backup, or fewer stripes.
- Restore with no recovery until the last restore with recovery.
- Restore from any available backup version created by Data Protection for SQL Version 5.5
- Replace an existing database with the restored database (or replace by relocating the restored database).
- Legacy restore to a different SQL Server or to a standby SQL Server.
- Automatically restore all backup objects needed to make a restore complete by using *smart selection* in the GUI.

Examples

See “Legacy Restore output examples” on page 189 for samples of various restore types using the command line interface.

SAN Volume Controller and Storwize V7000 FlashCopy support

Thin provisioning or the ability to allocate less physical storage than the declared size of a logical storage volume is available with SAN Volume Controller and Storwize V7000. A thinly provisioned volume is referred to as a space-efficient (SE) volume.

For more information about thinly provisioned volumes, see this website:
http://publib.boulder.ibm.com/infocenter/svc/ic/topic/com.ibm.storage.svc.console.doc/svc_spaceefficientvdisks_3r7ayd.html

Important: References to SAN Volume Controller in this section are referring to the following versions:

- 5.1.x (or later)
- 6.1.x (or later)
- 6.2.x (or later)

SAN Volume Controller and Storwize V7000 provide FlashCopy restore from SE target volumes as well as from fully allocated target volumes for which the background copy of the VSS backup has not yet completed. It is now possible to retain multiple FlashCopy images of a source volume as backup generations at a much reduced storage cost, since it is not necessary to allocate the full size of the source volume for each backup generation, while still maintaining the ability to do instant restore using FlashCopy from the SE target volumes.

SAN Volume Controller and Storwize V7000 minimize the demands required to maintain multiple snapshots of the same source volume by putting the target volumes into a cascade where each target is dependent on changes recorded in target volumes of subsequent snapshots. For example, if four VSS snapshots are created of a source volume, where S is the source and T1 through T4 are the targets, with T1 being the first chronologically and T4 the last, the following cascade occurs:

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

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

The following sections provide guidance and recommendations for effective use of Data Protection for SQL with SAN Volume Controller and Storwize V7000.

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

Data Protection for SQL exploitation of SAN Volume Controller and Storwize V7000 FlashCopy capabilities on Windows is dependent on the Volume Shadow Copy Service (VSS) hardware provider for SAN Volume Controller and Storwize V7000. Configuration of the VSS provider for SAN Volume Controller and Storwize V7000 controls what type of FlashCopy is performed when a VSS snapshot is requested, and the resultant behavior when you use VSS snapshots.

The VSS provider (4.1.x or later) that supports SAN Volume Controller and Storwize V7000 has the following characteristics:

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

In this case, deletion of the VSS snapshot removes the snapshot from the VSS inventory, but the FlashCopy relationship remains on SAN Volume Controller and Storwize V7000, so that a subsequent VSS snapshot of the same source volume will result in an incremental refresh of the target volume.

- When the VSS provider is configured to use SE target volumes - specifically, when the background copy rate is set to zero - the following is true:
 - Deletion of a VSS snapshot represented by a target volume in a cascade causes all target volumes dependent on the volume being deleted (in other words, the target volumes that were created earlier) also to be deleted. For example, deletion of a snapshot represented by target volume *T2* in the sample cascade *S -> T4 -> T3 -> T2 -> T1* causes *T2* and *T1* to be deleted, and the cascade *S -> T4 -> T3* to remain after the deletion.

Important: When you manually delete backups on SAN Volume Controller and Storwize V7000 space-efficient target volumes, and multiple backup versions exist, the backup being deleted as well as any older backups that contain the same volumes are deleted. Please take note that the deletion might not be performed until the next snapshot operation.

- A FlashCopy restore of the source volume from a target volume in a cascade of multiple target volumes is destructive to the target volume being restored, as well as to all newer targets in the cascade. For example, restore of a snapshot represented by target volume *T3* in the sample cascade *S -> T4 -> T3 -> T2 -> T1* causes *T4* and *T3* to be deleted, and the cascade *S -> T2 -> T1* to remain after the restore.

One exception to this pattern is that a FlashCopy restore from an SE target that is the only target in the cascade is not destructive.

- If an SE target volume runs out of space to hold the data from changed blocks on the source volume, that target volume and all target volumes dependent on that target volume go offline and render those backup versions unusable.

Note: An *SE backup version* is defined by a FlashCopy to an SE target volume that has a background copy rate of zero. Use of SE target volumes with "autoexpand" enabled and a background copy rate greater than zero does not create *SE backup versions*, because the target volumes grow to the allocated size of the source volumes when the background copy completes.

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

- Using a mix of SE and fully allocated target volumes is not supported. You must choose to use either SE or fully allocated volumes for FlashCopy targets, and set the VSS provider background copy rate parameter accordingly.
Making a transition from fully allocated targets to SE targets is accomplished by treating fully allocated targets as if they were SE when the background copy rate is set to 0.
- In order to determine how much storage space is required for each local backup, be aware that the backup LUNs require the same amount of storage space as the original LUNs. For example, if you have a 100GB database residing on a 200GB LUN, you will need a 200GB LUN for each backup version.
- When using SE backup versions:
 - Do not mix persistent and nonpersistent VSS snapshots. Use of a nonpersistent VSS snapshot following one or more persistent snapshots causes the older persistent snapshots to be deleted when the nonpersistent snapshot is deleted.
A VSS backup with *backupdestination* set to TSM creates a nonpersistent VSS snapshot. Therefore, do not follow a series of backups to local with *backupdestination* set to TSM. Instead, set *backupdestination* to both to send data to Tivoli Storage Manager while preserving local snapshot backup versions. Put another way, *backupdestination=LOCAL* and *backupdestination=TSM* are mutually exclusive settings. Do not use both in a backup strategy.
 - Enable *autoexpand* for the SE target volumes, to avoid out-of-space conditions.
 - Allocate enough space for SE target volumes to hold 120% of the data 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 are done every 6 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 space allocated to the SE target volumes should be $1.2 \times 5\% = 6\%$ 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.
 - Do not delete snapshots manually. Allow Data Protection for SQL to delete backup versions based on the defined policy, to ensure that deletion is done in the proper order. This avoids deletion of more backup versions than expected.

Data Protection for SQL and SAN Volume Controller and Storwize V7000 configurations

You can configure Data Protection for SQL and SAN Volume Controller and Storwize V7000 to meet your operational needs.

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

Table 7. Configuration scenarios for Data Protection for SQL and SAN Volume Controller and Storwize V7000

Use Cases or Objectives	SAN Volume Controller and Storwize V7000 Settings	VSS Provider Settings	Data Protection for SQL Settings	Additional information
<ul style="list-style-type: none"> Production application data resides on standard volumes Keep 14 snapshot backup versions Use minimum storage space for snapshot backup versions Full physical copy not required Perform 2 VSS backups per day 	<ul style="list-style-type: none"> Create 14 SE target volumes for each source volume to be protected Enable autoexpand for the SE target volumes Add the SE target volumes to the VSS free pool 	Set background copy rate = 0	<ul style="list-style-type: none"> Set policy to retain 14 local backup versions Schedule snapshot backups as preferred using backup destination=local 	When 14 VSS backups have been completed, the 15th VSS backup will cause the oldest backup to be deleted and will reuse that target set.
<ul style="list-style-type: none"> Production application data resides on standard volumes Keep 1 snapshot backup version Use minimum storage space for snapshot backup versions Full physical copy not required Perform one VSS backup per day and also send the backup to Tivoli Storage Manager 	<ul style="list-style-type: none"> Create 2 SE target volumes for each source volume to be protected Enable autoexpand for the SE target volumes Add the SE target volumes to the VSS free pool 	Set background copy rate = 0	<ul style="list-style-type: none"> Set policy to retain 2 local backup versions Schedule snapshot backups as preferred using backup destination=both 	Set 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 may not be available should a VSS backup fail after the prior backup was deleted
<ul style="list-style-type: none"> Production application data resides on standard volumes Keep 1 snapshot backup version Full physical copy is required Minimize overhead of background copies Perform 1 VSS backup per day and also send the backup to Tivoli Storage Manager 	<ul style="list-style-type: none"> Create 1 standard target vol for each source vol to be protected Add standard target volumes to the VSS free pool 	<ul style="list-style-type: none"> Use default background copy rate (50) Configure to use Incremental FC 	<ul style="list-style-type: none"> Set policy to retain 1 local backup version Schedule snapshot backups as preferred using backup destination=both 	When using INCR FC, the VSS provider will not delete the single snapshot target set even though Tivoli Storage FlashCopy Manager will delete the prior VSS snapshot before creating a new one

Table 7. Configuration scenarios for Data Protection for SQL and SAN Volume Controller and Storwize V7000 (continued)

Use Cases or Objectives	SAN Volume Controller and Storwize V7000 Settings	VSS Provider Settings	Data Protection for SQL Settings	Additional information
<ul style="list-style-type: none"> Production application data resides on standard volumes Keep 2 snapshot backup versions Full physical copies are required for local backup versions Perform VSS backups every 12 hours with one backup daily sent to Tivoli Storage Manager 	<ul style="list-style-type: none"> Create 3 standard target volumes for each source volume to be protected Add standard target volumes to the VSS free pool 	Use default background copy rate (50)	<ul style="list-style-type: none"> Set policy to retain 3 local backup versions Schedule VSS backups as follows: <ul style="list-style-type: none"> – backup destination=local at 11:00 – backup destination=both at 23:00 	Set policy for local snapshot backups to retain N+1 backup versions so that N snapshot backups are available for restore.
<ul style="list-style-type: none"> Production application data resides on standard volumes Keep 4 snapshot backup versions Use minimum storage space for snapshot backup versions Full physical copy not required Perform VSS backups every 6 hours with one backup daily sent to Tivoli Storage Manager 	<ul style="list-style-type: none"> Create 5 SE target volumes for each source volume to be protected Enable autoexpand for the SE target volumes Add SE target volumes to the VSS free pool 	Set background copy rate = 0	<ul style="list-style-type: none"> Set policy for local snapshot backups to retain 5 local backup versions Schedule VSS backups as follows: <ul style="list-style-type: none"> – backup destination=local at 06:00, 12:00 and 18:00 – backup destination=both at 00:00 	Set policy to retain N+1 backup versions so that N snapshot backups are available for restore
<ul style="list-style-type: none"> Production application data resides on SE volumes Keep 2 snapshot backup versions Full physical copies are required for local backup versions Perform VSS backups every 6 hours with one backup daily sent to Tivoli Storage Manager 	<ul style="list-style-type: none"> Create 3 SE target volumes for each source volume to be protected Allocate same percentage of real storage as for source volumes Add SE target volumes to the VSS free pool 	Use default background copy rate (50)	<ul style="list-style-type: none"> Set policy to retain 3 local backup versions Schedule VSS backups as follows: <ul style="list-style-type: none"> – backup destination=local at 06:00, 12:00 and 18:00 – backup destination=both at 00:00 	<ul style="list-style-type: none"> Set policy for local snapshot backups to retain N+1 backup versions so that N snapshot backups are available for restore Allows thin provisioning for both source and target volumes and lets them grow together

Using space-efficient target volumes with SAN Volume Controller and Storwize V7000

SAN Volume Controller and Storwize V7000 require special considerations when using space-efficient target volumes.

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

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

Important: When manually deleting backups on SAN Volume Controller and Storwize V7000 space-efficient target volumes and multiple backup versions exist, delete the backups in the same order that they were created. Otherwise, the FlashCopy mappings remain because of cascaded dependencies. This is a known limitation.

Table 8. Snapshot restore and delete behavior on SAN Volume Controller and Storwize V7000 space-efficient target volumes

Snapshots on space-efficient volumes	Snapshot to be restored	Snapshot deleted
s1, s2, s3, s4	s1	s1, s2, s3, s4
s1, s2, s3, s4	s4	s4
s1, s2, s3, s4	s2	s2, s3, s4
s1	s1	None

Note: The order of backup creation in the table is s1, s2, s3, s4.

Verifying snapshot creation

Use the IBM VSS provider `ibmvfcg` command to verify whether snapshots are created on SAN Volume Controller and Storwize V7000 space-efficient target volumes:

1. In a Windows command prompt, issue `ibmvfcg list infc -l` to display the FlashCopy mapping attributes. The `Tgt Type` column displays the FlashCopy mappings for the volumes attached to the current host. For example:

```
Tgt Type
```

```
Space Efficient
```

Space Efficient
Space Efficient
Space Efficient
Standard

2. Verify that the IBM VSS provider type states SVC Pegasus. For example:

```
Physical host.  
Provider Type is SVC Pegasus.
```

Migration considerations

Migration from Data Protection for SQL V5.5 is supported. After upgrading and configuring Data Protection for SQL V6.3, use VSS Restore for local VSS Backups that were originally created with Data Protection for SQL V5.5.

Guidelines for SAN Volume Controller and Storwize V7000 environments

Review the following guidelines before attempting backup operations. Further information is available at <http://publib.boulder.ibm.com/infocenter/svcic/v3r1m0/index.jsp>:

- Determine whether to use space-efficient or fully-allocated backup targets before issuing a backup operation. Provision enough target volumes in the SAN Volume Controller VSS_FREE volume group for as many of the backup versions you require. If using fully-allocated target volumes, their capacity size must match the size of the source volumes.
- If space-efficient virtual disks (VDisks) are used for backup targets, set the IBM VSS provider background copy value to zero by issuing the `ibmvfcg set backgroundCopy 0` command. To make the changes effective, restart the IBM VSS system service after issuing the command. For more details about configuring the IBM VSS Hardware Provider for space-efficient target volumes, make sure to read the appropriate VSS-related content in the SAN Volume Controller and Storwize V7000 documentation.
- Do not mix COPY and NOCOPY FlashCopy relationships from the same source volume or volumes.
- Do not mix fully-allocated and space-efficient VDisks (used for backup targets) in the VSS_FREE pool.
- If the protected data resides on SAN Volume Controller or Storwize V7000 volumes, and the VDisks in the VSS_FREE pool are space efficient, then VSS Instant Restore from multiple backups is possible. However, the VSS Instant Restore operation in this environment is destructive.
- The Windows host must be attached to an SAN Volume Controller or Storwize V7000 cluster. The volumes assigned to the Windows host must be participating in the SAN Volume Controller or Storwize V7000 cluster attached to a SAN Volume Controller.
- Make sure that IBM VSS hardware provider version 4.1 (or later) is installed. This provider must be configured to accommodate multiple backup versions on SAN Volume Controller or Storwize V7000 space-efficient target volumes.
- Multiple snapshots on SAN Volume Controller or Storwize V7000 space-efficient VDisks are not supported for single-copy cluster (SCC) environments. These SCC environments include Microsoft Windows Failover Clustering and Veritas Cluster Server environments. Multiple snapshots are not supported because VSS snapshots are not automatically identified as "cluster aware".

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

- You can remotely mount NOCOPY FlashCopy backups. However, SAN Volume Controller or Storwize V7000 is required to restore a NOCOPY FlashCopy backup.
- You can create a NOCOPY FlashCopy to a space-efficient target. However, protection from physical failures to the source volume is not provided.

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

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

Tip: Space-efficient target volumes go offline when their capacity limit is exceeded. As a result, the current backup and all older backups (which have not reached FULL_COPY status) are lost. To avoid this situation, use the AUTOEXPAND option when creating space-efficient targets. This option allocates additional physical storage in order to prevent space-efficient target volumes going offline.

Additional considerations when using SAN Volume Controller and Storwize V7000

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

Chapter 2. Planning for Data Protection for Microsoft SQL Server operations

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

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

Make sure to consider your production environment capabilities and backup objectives when planning for SQL Server data. Be sure to install and configure Data Protection for Microsoft SQL Server before attempting a backup operation.

How Tivoli Storage Manager server policy affects Data Protection for SQL

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

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

Tivoli Storage Manager policy can manage the VSS Backups that are placed on local shadow volumes as well as in Tivoli Storage Manager server storage pools. The Tivoli Storage Manager server is responsible for managing VSS Backups, whether the backup is stored on local shadow volumes or on the Tivoli Storage Manager server. Be aware that while a VSS snapshot (created for back up to Tivoli Storage Manager server storage) is deleted after the backup completes, a VSS snapshot (created for back up to local shadow volumes) remains active until the backup version is expired according to the policy settings for VSS Backups on local shadow volumes.

The number of local backup versions maintained by the Tivoli Storage Manager server is determined by the value specified by the Tivoli Storage Manager server *verexists* parameter (defined in the copy group of the management class to which the local backup belongs). The number of Target Volume sets allocated for local backups should be equal to the *verexists* parameter. 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 attempt will 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.

LUNs can also be reused for new backups (depending upon policy management settings) because a failed backup results in one less available backup version. For example, when *verexists*=3 and three backups have successfully completed, LUNs are reused from the oldest backup to accommodate a fourth backup operation. If

the backup fails, it is possible to have only two backup versions because the oldest version was deleted to make room for the new backup.

The policy management of local backups is responsible for reconciling the local backup repository with the information stored on the Tivoli Storage Manager server. For example, if Target Volume LUNs that were used for a local backup are removed from the storage subsystem, the information representing the backup on the Tivoli Storage Manager server must be reconciled. Likewise if the Tivoli Storage Manager server policy has determined that a local backup copy is no longer needed, the local backup manager must free the Target Volume LUNs to the storage subsystem so that these LUNs can be used for future backup operations. Tivoli Storage Manager automatically detects these situations and performs the reconciliation.

Storage space considerations for local shadow volumes

Tivoli Storage Manager requires that sufficient storage space be available to create shadow volumes required for VSS Backup processing. Even when the VSS Backup destination is the Tivoli Storage Manager server, storage space to create a shadow volume is still required (though on a temporary basis). Since the value of the *verexists* parameter (specified for your local backup policy) determines the number of backup versions to retain on local shadow volumes, a *verexists=1* setting will cause the deletion of an existing backup on local shadow volumes (during a VSS Backup to Tivoli Storage Manager server storage) in order to create enough temporary space for the new snapshot. Therefore, if you want to keep *N* backups on local shadow volumes and also perform VSS Backups to Tivoli Storage Manager server storage, make sure you provision enough storage space on local shadow volumes and specify *verexists=N+1*.

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

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

Policy considerations for VSS Backups

Be aware that the following issues impact your Tivoli Storage Manager policy for managing VSS Backups:

- Overall backup strategy.
- Length of time that VSS Backups will reside on Tivoli Storage Manager server storage.
- Number of VSS Backup versions to reside on Tivoli Storage Manager server storage.
- Types of VSS Backups to reside on Tivoli Storage Manager server storage.

- Number of VSS Backup versions to reside on local shadow volumes.
- Types of VSS Backups to reside on local shadow volumes.
- The amount of available target volume storage provisioned for VSS operations.

For information about Tivoli Storage Manager storage management policies, see this topic: http://publib.boulder.ibm.com/infocenter/tsminfo/v6r3/topic/com.ibm.itsm.client.doc/c_mgtc_understand.html.

Security

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

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

Three options are provided when specifying SQL Server logon information:

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

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

Backup strategies

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

Strategies defined by backup type

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

Full backup only (Legacy and VSS)

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

Full plus log backup (Legacy and VSS)

A full plus transaction log backup strategy is commonly used when the normal backup window or network capacity cannot support a full backup each time. In such cases, a periodic full backup followed by a series of log backups allows the backup window and network traffic to be minimized. For example, you can perform full backups on the weekend and log backups during the week. The full backups can be done during low usage times when a larger backup window and increased network traffic can be

tolerated. The restore process becomes more complex, however, because a full backup, as well as subsequent log backups, must be restored. It is also possible to do a point-in-time restore to restore a transaction log to a specified date and time.

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

Full plus differential backup (Legacy and VSS)

Perform this type of backup *between* full backups. A differential database backup can save both time and space — less space in that it consists of only the changed portions of a database since the last full backup (it is cumulative), and less time in that you can avoid applying all individual log backups within that time to the operation. This applies to restore operations as well; only the last differential backup (latest version) need be restored.

You can apply Legacy differential backups after a full VSS Backup has been restored. In order to do this, you must leave the database in a recovering state by specifying `/recovery=no` on the command-line interface or by making sure that the Recovery option is not selected when restoring the VSS Backup.

Full plus differential plus log backup (Legacy and VSS)

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

File or group backups (Legacy only)

Use a file backup strategy when it is impractical to backup an entire database due to its size and accompanying time and performance issues. Note that when performing restore operations for a file or file group, it is necessary to provide a separate backup of the transaction log.

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

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

Strategies defined by other considerations

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

Saving time:

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

Data striping (Legacy only):

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

Clustering:

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

Truncate log on checkpoint option:

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

Truncate log option:

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

Collocation:

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

Multiple SQL Servers:

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

Various Recommendations:

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

Using VSS and Legacy Backups together

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

Microsoft supports and recommends using both methods of backup in your complete backup strategy. Although VSS only supports full backups, Legacy differential and Legacy log backups can be applied after a full VSS Backup has been restored.

Also, be aware of the following best practice:

- Legacy and VSS Backups to Tivoli Storage Manager server storage are usually dictated by time, not versions.
- Backups to local shadow volumes are usually dictated by versions because of space limitations and provisioning of VSS storage.
- When running VSS operations, make sure you have at least 200 megabytes of free disk space on your Windows System Drive. This space is used to hold the meta data files for Data Protection for SQL.

Example strategy of using VSS and Legacy Backups together

Table 9. 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: <ul style="list-style-type: none"> • VSS Restore¹ • VSS Fast Restore² • VSS Instant Restore³ Legacy: <p>Restore to the production SQL Server or alternate SQL server</p>
Restore attributes	<ul style="list-style-type: none"> • FULL, LOG, DIFF, GROUP, FILE • Server, database, file group, and individual file-level restore granularity • Point-in-time recovery • Roll-forward recovery • Restore to alternate machine 	VSS: <ul style="list-style-type: none"> • FULL • Database level restore granularity • Point-in-time recovery⁴ • Roll-forward recovery⁴ Legacy: <ul style="list-style-type: none"> • FULL, LOG, DIFF, GROUP, FILE • Server, database, file group, and individual file-level restore granularity • Point-in-time recovery • Roll-forward recovery • Restore to alternate machine

Note:

1. Files are copied from the Tivoli Storage Manager server directly to the production source volumes.
2. Files are copied from local shadow volumes directly to the production source volumes.
3. Snapshot volumes are copied to the production source volumes. SAN Volume Controller, DS6000, DS8000, XIV, or Storwize V7000 is required to perform this type of restore.
4. In order to acquire these restore attributes, these backups must be Legacy log backups or Legacy differential backups that are applied to a full VSS Backup.

Using VSS operations in a SQL Server Failover Cluster environment

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

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

SQL Server Failover Cluster requirements for VSS

The following requirements must be met for VSS operations to perform successfully in a SQL Server Failover Cluster environment:

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

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

```
vssaltstagingdir d:\dir
```

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

```
vssaltstagingdir Q:\TSMVSS
```

- Make sure to specify the following options in each of the *dsm.opt* files that are used for the LOCALDSMAGENT and REMOTEDSMAGENT machines:

```
CLUSTERNODE NO  
CLUSTERDISKSONLY NO
```
- When running Tivoli Storage FlashCopy Manager in stand-alone mode (that is, not connected to the Tivoli Storage Manager server), if you are changing the *vssaltstagingdir* option, you must change its value to the same value in BOTH the *dsm.opt* file for the DSMAGENT and the *dsm.opt* file for Data Protection for SQL.

SQL Server Failover Cluster limitations for VSS

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

- When performing VSS Instant Restores, a best practice is to make sure that any previous background copies (that involve the volumes being restored) are completed prior to initiating the VSS Instant Restore. However, this check is not necessary for XIV or SAN Volume Controller or Storwize V7000 with space-efficient target volumes.
- All servers within the cluster must use the same levels of Tivoli Storage Manager, Windows, and other applicable software.
- Microsoft KB919117 is required to perform VSS Instant Restore in a cluster environment on Windows Server 2003 or later.
- Dynamic disks are not supported for VSS Instant Restore operations.
- For Windows Server 2003, VSS Backups that reside on local VSS shadow volumes can only be restored to the physical node that created the VSS Backup. For example, if NODE_A of a cluster created a VSS Backup and stored it on local shadow volumes, NODE_B of the cluster cannot restore that particular VSS Backup. NODE_B can only restore VSS Backups stored on the Tivoli Storage

Manager server or VSS Backups stored on local shadow volumes that were created by NODE_B. This is due to a limitation related to VSS cluster support and not to Tivoli Storage Manager.

- For Windows Server 2003, VSS Backups that reside on local VSS shadow volumes *can be deleted* in the event of a cluster failover. This means that if NODE_A of a cluster created a local VSS Backup and afterward, the SQL Server fails over to NODE_B, if a Data Protection for SQL VSS operation is performed on NODE_B of the cluster, the local VSS Backup created by NODE_A will be deleted. This is caused by the fact that the Microsoft VSS architecture is not cluster aware.
- The Tivoli Storage Manager Client Acceptor Daemon (CAD) must be installed on each cluster node so that it can continue operations in the event of a failover. Make sure the CAD service name is the same on all cluster nodes so that it can be started by a generic cluster service.
- It is recommended that the Local DSMAgent client node be a separate node from your normal backup-archive client, as this CAD service will need to be made a non-cluster option.
- The Remote DSMAgent client node does not require you to register a separate node for each server within the cluster as this server only acts as a secondary server.
- Use the Microsoft **vssadmin** and **vshadow** commands to verify the environment.
- A Data Protection for SQL configuration file should be configured for each node in the cluster. These files are almost identical, except that the *localdsmagentnode* parameter points to the corresponding local DSMAgent on each node.

SQL Server Failover Cluster VSS limitations for scheduled operations

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

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

Using Data Protection for SQL in a Microsoft Windows Failover Clustering environment

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

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

- References to the SQL Server made in this section pertain to the virtual SQL Server name in a Microsoft Windows Failover Clustering environment.
- You must install Data Protection for SQL on all nodes of the cluster. In addition, when installing Data Protection for SQL, you must install it on a disk local to each node (not on a shared cluster disk).
- You must specify *clusternode yes* in the Data Protection for SQL options file.
- Use identical configurations in the Data Protection for SQL options file when configuring Data Protection for SQL on each node of the cluster.

- If you are using the Tivoli Storage Manager scheduler for automating backups, you must install the scheduler service on each node of the cluster to enable fail-over support. See “Using the Tivoli Storage Manager scheduler” on page 91 for more information.
- The Tivoli Storage Manager treats backups as coming from a single server (the virtual server) regardless of which node of the cluster a backup was performed on.
- Multiple snapshots on SAN Volume Controller and Storwize V7000 space-efficient VDisks are not supported for single-copy cluster (SCC) environments. These SCC environments include Microsoft Windows Failover Clustering and Veritas Cluster Server environments. Multiple snapshots are not supported because VSS snapshots are not automatically identified as “cluster aware”.

Using Data Protection for SQL in a Veritas Cluster Server environment

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

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

- References to the SQL Server made in this section pertain to the virtual SQL Server name in an VCS environment.
- You must install Data Protection for SQL on all nodes of the cluster. In addition, when installing Data Protection for SQL, you must install it on a disk local to each node (not on a shared cluster disk).
- You must specify *clusternode yes* in the Data Protection for SQL options file.
- Use identical configurations in the Data Protection for SQL options file when configuring Data Protection for SQL on each node of the cluster.
- If you are using the Tivoli Storage Manager scheduler for automating backups, you must install the scheduler service on each node of the cluster to enable fail-over support. See “Using the Tivoli Storage Manager scheduler” on page 91 for more information.
- The Tivoli Storage Manager treats backups as coming from a single server (the virtual server) regardless of which node of the cluster a backup was performed on.

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

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

Tivoli Storage Manager storage

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

A Tivoli Storage Manager backup operation stores the backed up data on Tivoli Storage Manager server storage. Although this type of backup typically takes longer to process than a backup to local shadow volumes, a Tivoli Storage Manager backup is necessary when long term storage is needed such as saving SQL data on tape for archival purposes. 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 should backups on the local shadow volumes become corrupt or deleted.

Local shadow volumes

Backups to local shadow volumes are usually dictated by versions because of space limitations and provisioning of VSS storage.

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

VSS operations in DS, SAN Volume Controller, and Storwize V7000 environments

For information about storage space requirements for space-efficient volumes, see “Using Data Protection for SQL with SAN Volume Controller and Storwize V7000” on page 11.

VSS limitations for SAN Volume Controller and Storwize V7000

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

When performing two Data Protection for SQL VSS backups, back to back, and the SQL Server data resides on SAN Volume Controller or Storwize V7000 disks, if the volumes are large and/or the background copy rate is set a low number, it may appear that the second VSS backup is hanging. In fact, it is waiting for the background copy of the first backup to complete before proceeding with the second backup. SAN Volume Controller and Storwize V7000 will not allow two background copies of the same volume to occur at the same time. There is no indication that the second backup is waiting for the first background copy to complete. You may also see timeouts errors if the previous background copy takes too long. To work around this issue, schedule your VSS backups far enough apart to accommodate this situation. You can also try increasing the copy rate of the background copy.

VSS operations in IBM N-series and NetApp environments

Be aware that in environments that contain IBM N-series and NetApp systems, snapshots created using the IBM N-series and NetApp snapshot provider are stored on the same volume where the LUN resides. Disk space consumed by a local backup consists only of the blocks that have changed since the last local backup was created. The following formula can be used to help determine how much space is required for each local backup:

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

In addition, Write Anywhere File Layout (WAFL) reserves blocks equal to two times the specified size of the LUN to be used. This space reservation ensures writes for virtual disks. The following example demonstrates how to calculate the size of these volumes:

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

VSS limitations for NetApp FAS series or IBM N-series

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

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

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

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

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

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

- If the NetApp volume type is Traditional, VSS backups with backup destination set to Local must be bound to a management class that has verExists=1. This setting is not required if FlexVols are used.

- When performing VSS backups, you must ensure that previous backup has finished completely before starting a new backup. Any overlap of backups can result in undesirable side-effects on the Microsoft SQL Server, the VSS service, and, the NAS filer.

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

Taking above considerations into account, the following section describes a sample backup procedure that could be used to perform VSS backups utilizing both Tivoli Storage Manager and LOCAL backup destinations in an optimal manner. Note that the following assumptions apply to this sample backup procedure:

- The configuration requirements stated above are met.
- The VSS backup to Tivoli Storage Manager takes one hour to complete.
- The VSS backup to LOCAL takes five minutes to complete.

Your backup procedure could consist of the following backups:

- Daily VSS full backups to LOCAL every four hours - 12am, 4am, 8am, 12pm, 4pm, 8pm
- Daily VSS full backups to Tivoli Storage Manager storage by one of the following two methods:
 - Specify backupdestination set to BOTH at 12am. Note that this will create a 12am backup to local. Therefore, no separate 12am backup to local is required.
 - Full offloaded-backup at 1am. Note that no VSS local backup will be available to restore from between 1am and 4am when next VSS backup to local will take place.
- Perform weekly VSS-full backups to Tivoli Storage Manager (offloaded backup) 5am
- Perform weekly legacy full backups (or as needed).

Recommended Tivoli Storage Manager policy settings

Make sure the following policy items are defined with the recommended settings.

Consult your Tivoli Storage Manager administrator or see the *IBM Tivoli Storage Manager for Windows Administrator's Guide* and the *IBM Tivoli Storage Manager for Windows Administrator's Reference* for complete information on defining or updating these Tivoli Storage Manager policy items.

Domain

Create a policy domain on the Tivoli Storage Manager server to be used exclusively for Data Protection for SQL backups.

Policy Set

Policy sets contain management classes (which contain copy groups) that determine the rules by which Data Protection for SQL backups are performed and managed. Define the policy set to the policy domain to which Data Protection for SQL backups belong. Note that the policy set must be activated and only one policy set can be active in the policy domain.

Management Class

Define a management class for backups residing on local shadow volumes and a management class for backups residing 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 database (*VERExists=6*) while maintaining only two versions of the same database on Tivoli Storage Manager server storage (*VERExists=2*). In addition, you can create a separate management class for full backups for use in long term storage. Such policies can maximize storage resources and provide more control over your storage strategy.

Be aware that since Legacy backups on Tivoli Storage Manager server storage, VSS Backups on Tivoli Storage Manager server storage, and VSS Backups on local shadow volumes all have different Tivoli Storage Manager server naming and therefore, can each have their own management class, it is possible to have three active backups of the same database. Make sure your backup strategy is planned and well-defined before defining management classes.

Metadata considerations (Legacy only)

The management classes for Data Protection for SQL metadata should be identical to the corresponding management classes for database data except that the metadata management classes should not allow migration to removable media. If any Data Protection for SQL metadata is on removable media, queries may require media mounts, and backups or restores may require additional media mounts.

Data objects and their associated meta objects should have the same version limits and retention values. However, because meta objects may need to be restored as a result of a Data Protection for SQL query, you may want to consider storing meta objects in a disk-only storage pool so that a media mount is not necessary to resolve the query. To do this, you can:

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 using an *include* statement in the Data Protection for SQL 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 further details about binding backup objects to specific management classes, see “Setting automatic expiration (VSS and Legacy)” on page 41.

Copy Group

Define the copy group as a backup copy group and not an archive copy group. Since Data Protection for SQL 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 client system.

VERDeleted

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

RETEtra

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

RETOly

Determines the number of days to retain the last SQL Server database backup version of a database that has been deleted from the Data Protection for SQL client system. Be aware that log backups do not participate in expirations (due to version limit) because there is never more than one version of a log backup object. This is because log backups are always uniquely named. However, all Legacy backup objects for an SQL Server database are inactivated when a new full backup of that SQL Server database is performed (VSS backup objects remain active).

Because log and set objects are always uniquely named, they do not participate in expirations due to version limit. However, Data Protection for SQL inactivates all backup objects for a SQL database not otherwise inactivated whenever a new full database backup of that SQL database is performed. Therefore, the retention period defined through the ***RETOly*** parameter controls the expiration of log and set backups. There will never be more than one version of a log or set object. In addition, consider these guidelines:

- When selecting the ***RETOly*** value for log backups, ensure that it is at least as long as the value for the backup objects the logs are associated with. You can use the same management class for log backups and the backup objects that are retained the longest to be sure you use an adequate value. However, all Legacy backup objects for an SQL Server database are inactivated when a new Legacy full backup of that SQL Server database is performed (VSS backup objects remain active).
- Set backups are intended to be used in unusual one-of-a-kind situations.
- If you do not wish to wait for the next full database backup, you can explicitly inactivate any particular active object (if it is no longer needed) or any active objects older than a specified number of days using the “Inactivate command (Legacy only)” on page 208.

If you exploit data striping, each stripe of a backup must have the same version limits and retention values to ensure that some parts of a single logical backup object do not expire before others.

MODE, SERIALization, FREQuency

You can accept default values for these backup copy group parameters as they are not applicable to Data Protection for SQL.

It is recommended that you discuss these parameters with your Tivoli Storage Manager server administrator in order to accomplish your backup strategy.

Storage Pool

A single restore can require a full backup or a differential backup. It is recommended to use collocation if these backups are stored on removable media. Specify collocation by file space (**define stgpool COLlocate=FILEspace**) if you plan to restore multiple databases in parallel. This is recommended because individual data stripes stay on separate removable volumes. If you use data striping, use collocation by file space on sequential storage pools to maintain the stripes on separate storage volumes. This is necessary to allow concurrent parallel access to each of the stripes. If it happens that multiple stripes for the same object end up on the same sequential volume (because insufficient empty volumes are available), the Tivoli Storage Manager server **move data** command can be used to move the objects to a disk storage pool or to new sequential volumes added to the storage pool so that they can be accessed in parallel.

A single, complete restore may require a full database backup, a differential backup, and multiple log backups, or one or more group, file, or set backups and multiple log backups. It is recommended that you use collocation if these backups may be stored on removable media.

Data Protection for SQL node name: Recommended settings

The machine where Data Protection for SQL is installed must be registered to the Tivoli Storage Manager server with a node name. This node name owns and manages all Data Protection for SQL data that is backed up to the Tivoli Storage Manager server. Specify this node name with the *nodename* option in the *dsm.opt* options file located (by default) in the Data Protection for SQL installation directory. Note that in order to perform VSS operations, you may need to register node names for additional machines. See “Proxy node definitions (VSS Backups)” on page 35 for details about this task.

Be aware of the following Tivoli Storage Manager parameter conditions when registering your Data Protection for SQL node name (machine) to the Tivoli Storage Manager server:

BACKDElete

This parameter determines whether the Data Protection for SQL node can delete its own backup files from the Tivoli Storage Manager server. This parameter **MUST** have a value of *yes*.

MAXNUMMP

This parameter determines the maximum number of mount points a client node is allowed to use on the Tivoli Storage Manager server during a backup operation. This must be set to a number greater than the default value of 1 if you are to exploit SQL data striping with data going directly to a tape pool. For example, set this value to be at least the maximum number of stripes to be used for backup or restore when removable media such as tapes are used or if migration occurs during the backup or restore operation. If other backups or restores may occur at the same time, the value of this parameter must be large enough to allow for all of the needed mount points.

TXNGroupmax

This parameter determines the number of files transferred as a group between Data Protection for SQL and the Tivoli Storage Manager server between transaction commit points. This parameter **MUST** have a value of at least one more than the maximum number of stripes to be used for backup or restore operations regardless of media.

COMPression (Legacy only)

This parameter determines whether the Data Protection for SQL node compresses data before sending it to the Tivoli Storage Manager server during a backup operation. Specify *COMPression=Client* to allow the Data Protection for SQL node to make the decision whether to compress data via the value of the client *COMPRESSIon* option specified in the options file (dsm.opt) located in the Data Protection for SQL directory.

Note: If you are running Data Protection for SQL on a Microsoft Cluster Server, the node name cannot be the name of the local computer. Instead, the node name should match the SQL virtual server name.

See the *IBM Tivoli Storage Manager for Windows Administrator's Reference* for complete information regarding these parameters.

Proxy node definitions (VSS Backups)

Since Data Protection for SQL VSS Backup operations are implemented through the Tivoli Storage Manager backup-archive client, you must use node names specifically for VSS operations in addition to using a node name for where Data Protection for SQL is installed. As part of the configuration procedure, a proxy relationship is defined for these various node names. This proxy relationship allows node names to perform operations on behalf of another node name. When registering these nodes to the Tivoli Storage Manager server for VSS operations, do not specify the Tivoli Storage Manager *USerid=NONE* parameter. VSS operations will fail when this parameter is specified.

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

- **Target node:** A node name that controls backup and restore operations and that also owns the data on the Tivoli Storage Manager server. This is the node name specified in the Data Protection for SQL dsm.opt file.
- **Agent node:** A node name that performs operations *on behalf* of a target node. This is the node name specified in the Backup-Archive Client dsm.opt file.

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

```
GRANT PROXY TARGET=<dpsql node name> AGENT=<dsmagent node name>
```

Required node names for basic VSS operations

Data Protection for SQL node names are required for basic VSS operations.

The configuration wizards in the MMC GUI performs most of the tasks described here. Node name information is provided here for reference only so that you can become familiar with all of the options and their values in the respective environment.

To perform basic VSS operations, one target node and one agent node are required.

Table 10. Required node names for basic VSS operations

Proxy node type	Nodename	Where to specify
Target node	Data Protection for SQL node name	Use the <i>nodename</i> option in the Data Protection for SQL options file (dsm.opt)
Agent node	Local DSMAGENT Node	Use the <i>localdsmagentnode</i> parameter in the Data Protection for SQL configuration file (tdpsql.cfg)

Target node

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

Agent node

This is the node name where the backup-archive client is installed. This node is responsible for performing the VSS operations as Data Protection for SQL itself does not perform 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 configuration file (tdpsql.cfg by default). You can use the property pages of the MMC GUI or the **tdpsqlc set** command to specify this parameter.

Note: The agent node and target node will be on the same machine for basic VSS operations.

Required node names for VSS off-loaded backups

To perform VSS off-loaded backups, you must have one target node and two agent nodes:

Table 11. Required node names for VSS off-loaded backups

Proxy node type	Nodename	Where to specify
Target node	Data Protection for SQL node name	Use the <i>nodename</i> option in the Data Protection for SQL options file (dsm.opt)
Agent node	Local DSMAGENT Node	Use the <i>localdsmagentnode</i> parameter in the Data Protection for SQL configuration file (tdpsql.cfg)
Agent node	Remote DSMAGENT Node	Use the <i>remotedsmagentnode</i> parameter in the Data Protection for SQL configuration file (tdpsql.cfg)

Target node

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

Agent node

This is the node name where the backup-archive client and VSS provider are installed. This node is responsible for performing the VSS operations as Data Protection for SQL itself does not perform 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 configuration file (tdpsql.cfg by default). You can use the property pages of the MMC GUI or the **tdpsqlc set** command to specify this parameter.

Agent node

The node name of a separate machine that must also have the backup-archive client and VSS provider installed. This machine is responsible for performing the movement of VSS snapshot data from local shadow volumes to the Tivoli Storage Manager server. This node name is referred to as the Remote DSMAGENT Node and is specified with the *remotedsmagentnode* parameter in the Data Protection for SQL

configuration file (`tdpsql.cfg` by default). You can use the property pages of the MMC GUI or the **tdpsqlc set** command to specify this parameter. The choice of available machines depends on whether the machines have access to the local shadow volumes that contain the VSS snapshot backups. This node name is only valid for VSS environments that support transportable shadow copies. It is not supported if you are using the default VSS system provider. Refer to your VSS provider documentation for details.

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

Specifying Data Protection for SQL options

Once Data Protection for SQL is registered to a Tivoli Storage Manager server, several Data Protection for SQL parameters need to be configured. The Tivoli Storage Manager administrator should have provided you with the node name, password, and the communications method with the appropriate parameters to connect to the Tivoli Storage Manager server.

These values, together with other parameters, are stored in an options file in the Data Protection for SQL directory. The default options file name is `dsm.opt`. To modify the initial `dsm.opt` file, display `dsm.opt` using a text editor.

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

nodename

The Tivoli Storage Manager *nodename* is the unique name by which Tivoli Storage Manager knows the machine running Data Protection for SQL.

commmethod

The communication protocols link the Data Protection for SQL node with the Tivoli Storage Manager server. Data Protection for SQL supports the same set of communication protocols supported by other Tivoli Storage Manager clients.

- For Legacy backups, specify the *commmethod* option in the Data Protection for SQL 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.

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

passwordaccess

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

A utility program named `dsmcutil.exe` allows you to manage (update or display) the password as stored in the registry. This utility program is distributed with the Tivoli Storage Manager Backup-Archive Client package. For more information on using the `dsmcutil` program, see the `dsmcutil.hlp` file or the `dsmcutil.txt` file which are distributed with the Tivoli Storage Manager Backup-Archive Client package.

compression

A setting of *compression on* instructs the Tivoli Storage Manager API to compress data before sending it to the Tivoli Storage Manager server; this reduces traffic and storage requirements. If you enable compression, it affects performance in two ways:

- CPU utilization is higher on the machine on which Data Protection for SQL is running.
- Network bandwidth utilization is lower because fewer bytes are sent.
- Storage usage on the Tivoli Storage Manager server is reduced.

You may want to turn compression on if any of the following conditions exist:

- The network adapter has a data overload.
- Communications between the Data Protection for SQL and Tivoli Storage Manager server are over a low bandwidth connection.
- There is heavy network traffic.

Considerations:

- For Legacy backups, specify the *compression* option in the Data Protection for SQL 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.

It may be better to set *compression off* in the following cases:

- The computer running Data Protection for SQL has a CPU overload; the added CPU usage can impact other applications including SQL Server. You can monitor CPU and network resource utilization using the Performance Monitor program shipped with Windows.
- You are not constrained by network bandwidth; in this case, you can achieve the best performance by setting *compression off* and enabling hardware compaction on the tape drive, which also reduces storage requirements.

Note: The Tivoli Storage Manager administrator can override the compression option setting for the Data Protection for SQL node when registering or updating the node by specifying, on the Tivoli Storage Manager server side, that a particular node:

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

See “How to enable SQL Server 2008 backup compression” on page 44

clusternode

A setting of *clusternode yes* directs the Tivoli Storage Manager API and Data Protection for SQL to be cluster-aware when running in a Microsoft Windows Failover Clustering or Veritas Cluster Server environment. This option *must* be specified for Data Protection for SQL to function properly in either of these cluster environments.

deduplication

Client-side data deduplication is used by the Tivoli Storage Manager API, to remove redundant data during backup and archive processing before the data is transferred to the Tivoli Storage Manager server. Specify whether the Tivoli Storage Manager API deduplicates data before sending it to the Tivoli Storage Manager server. You can specify *Yes* or *No*. The default value is *No*. The value of the deduplication option for Data Protection for SQL applies only if the Tivoli Storage Manager administrator allows client-side data deduplication.

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

enablelanfree

To perform a LAN-free Legacy backup with Data Protection for SQL, a Tivoli Storage Manager Storage Agent must be installed on the same machine and *enablelanfree yes* must be specified in the Data Protection for SQL options file. To perform a LAN-free VSS Backup with Data Protection for SQL, specify *enablelanfree yes* in the DSMAGENT (VSS Requestor) options file. See *Managed System for SAN Storage Agent User's Guide* for detailed information about LAN-free environments.

enableclientencryptkey

This option encrypts SQL databases during backup and restore processing. One random encryption key is generated per session and is stored on the Tivoli Storage Manager server with the object in the server database. Although Tivoli Storage Manager manages the key, a valid database must be available in order to restore an encrypted object. Specify *enableclientencryptkey yes* in the Data Protection for SQL options file. In addition, assign the type of encryption to use by specifying the *encryptiontype* option in this same options file. You can specify *DES56* (56-bit) or *AES128* (128bit). The default is *AES128*. In this same file, you must also specify the databases you want encrypted by adding an include statement with the *include.encrypt* option.

- For Legacy backups, specify these encryption options in the Data Protection for SQL 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 example, edit the appropriate options file by adding the following three options:

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

3. (Legacy backups only) Add your include statements with the *include.encrypt* option. For example, to encrypt all SQL data, specify the following:

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

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

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

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

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

Note:

1. If you are running Data Protection for SQL on a Microsoft Windows Failover Clustering or Veritas Cluster Server, the options files on each node of the cluster must be identical.
2. You can create additional Data Protection for SQL options files to point to other Tivoli Storage Manager servers. You might also want to create more than one options file, each with different parameters, to use with a single Tivoli Storage Manager server.

Specifying Data Protection for SQL preferences

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

You can set the values of the Data Protection for SQL configuration parameters in two ways:

- In the MMC GUI, set the value in Properties.
- The **tdpsqlc set** command in the Data Protection for SQL Command Line Interface. See “Set command” on page 222.

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

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

See “Set command” on page 222 for descriptions of available configuration parameters.

Set policy for VSS backups by specifying the VSSPOLICY statement in your Data Protection for SQL configuration file, or by selecting **Properties > VSS Policy Binding** in the MMC GUI, and then configuring the policy.

Note that you must specify this statement manually. You cannot specify it using the **tdpsqlc set** command. See “Setting automatic expiration (VSS and Legacy)” for detailed information and examples.

Setting automatic expiration (VSS and Legacy)

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

Setting automatic expiration for VSS

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

VSSPOLICY statements are processed from the bottom up and processing stops at the first match. To ensure that more specific specifications are processed at all, the more general specification should be listed before the more specific ones, so as to be processed after the more specific specifications. Otherwise, the more general specification will match the target before the more specific specifications are seen.

Specify the following information in the VSSPOLICY statement:

```
VSSPOLICY <srvname> <dbname> <backuptype> <backupdest> <mgmtcls>
```

The statement contains the following values:

srvname

Name of the SQL Server or wildcard character (*)

dbname

Name of database or wildcard character (*)

backuptype

Backup type: FULL or wildcard character (*)

Note: Since VSS supports full backup types only, specifying a wildcard character for <backuptype> performs a full backup only.

backupdest

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

mgmtcls

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

See “Examples of Tivoli Storage Manager policy binding using include/exclude and VSSPOLICY statements” on page 233 for examples of the VSSPOLICY statement.

Setting automatic expiration for Legacy

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

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

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

The general include/exclude syntax for object naming is:

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

where:

objectNameSpecification is:

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

dataType is:

meta | data

typeInfo is:

LogicalFileName (for **file** backup type)

GroupName (for **group** backup type)

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

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

backupType is:

full | diff | log | group | file | set

Considerations:

- The wildcard character * matches zero or more characters. The wildcard character ? matches any one character.
- The wildcard character * within a qualifier replaces zero or more characters only within that qualifier. The qualifier itself must exist in the matching object name.
- To match zero or more qualifiers, use ellipses: \...\
- All specifications must end with the wildcard character *. This is required because the specification must match both object names and temporary names. Temporary names are used to enable rolling back a backup transaction if an error occurs. Temporary names are object names with a unique string appended to the *backupType* qualifier.
- An *objectNameSpecification* should be placed within double quotation marks. If the specification includes spaces or special characters, the double quotes are required.
- For exclude statements, **meta** should be a match in the specification, either explicitly, or by wildcard or ellipses. Excluding **meta** excludes the entire object.

- Include statements can specify either **meta** or **data** separately and explicitly, or both by wildcard or ellipses.
- You may specify both data and meta objects in options file include lists in order to assign management classes. However, when you use exclude statements, you should specify only the meta objects. If a data object is not backed up, its meta object will not be created.
- **Log** and **set** object names are always unique. The *typeInfo* contains the qualifiers whose values make them unique. Because they are generated at the time of the backup, they are not predictable and you cannot specify them.
- Include/exclude lists are processed from the bottom up, and processing stops at the first match. To ensure that more specific specifications are processed at all, you should list the more general specifications before the more specific ones so that they will be processed after the specific. Otherwise, the more general will match the target before the more specific are seen.
 - When a match is found, processing of the list stops and the statement that matches is examined. If it is an exclude statement, the matching object name is not backed up. If it is an include statement, the matching object name is backed up. If the include statement contains a ManagementClassName, that management class is associated with the object name for this backup and for all backups of the same name on the current node.
 - If a match is not found, the object is backed up using the default management class for the current node.
 - If a match is found for an include that specifies a management class, but the specified management class is not valid for the current node, the default management class for the current node is used.
- Include/exclude processing does not produce error messages for invalid specifications. Therefore, you should thoroughly test all include/exclude lists. Specifying an invalid management class name will generate an error message in the dserror.log.
- In regard to case-sensitivity, the Windows Tivoli Storage Manager API currently assumes the specifications are for a Windows file system and ignores case. However, because case may be honored in the future, you should always use the correct case. Specifically,
 - Use correct case for SQL names (server, database, group, or file names) as displayed by the **query sql** or **query tsm** commands.
 - Use lower case for the Data Protection for SQL constants: **meta**, **data**, **full**, **diff**, **log**, **group**, **file**, and **set**.

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

SqlServerNames:

SQL2000, SQL2005

InstanceNames:

INST1, INST2

DatabaseNames:

Db1, Db2, Db3

GroupNames:

g1, g2, g3

LogicalFileNames:

f1, f2, f3

For complete sample illustrations of *include* and *exclude* statements using the syntax outlined above, refer to “Examples of Tivoli Storage Manager policy binding using include/exclude and VSSPOLICY statements” on page 233.

How to enable SQL Server 2008 backup compression

Support for SQL Server 2008 backup compression is available on Data Protection for SQL. You can use either the MMC GUI or the command line to enable this feature.

- From the MMC GUI, specify SQL native backup compression from the SQL Properties window. After you have set this option, the **SQL Workload** column on the Recover tab shows the SQL compression status for legacy backups.
- From the command line, add this statement to the SQL configuration file (tdpsql.cfg). Edit the file and enter the command as follows:
SQLCOMPRESSION Yes | No

The default value is No.

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

SQL Server 2008 backup compression is generally faster and more effective than using it together with Tivoli Storage Manager compression. Tivoli Storage Manager compression is also effective at reducing the backup window needed to protect SQL Server and is a good choice for SQL Server 2005 (or earlier).

Chapter 3. Installing Data Protection for SQL Server

Learn about prerequisite information before installing Data Protection for SQL and review the installation instructions.

Quick installation and configuration

Install and set up Data Protection for SQL to help you start protecting your SQL server data as quickly as possible.

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

1. Install Data Protection for SQL.
 - a. Log on as an administrator.
 - a. Insert the IBM Tivoli Storage Manager for Databases: Data Protection for Microsoft SQL Server product DVD into your DVD drive. If autorun is enabled, the setup wizard starts automatically when the DVD loads. Otherwise, click **Start > Run**, and at the prompt, specify: `x:\setupfcm.exe`, where `x:` is your DVD drive. Click **OK**.
 - b. Follow the installation instructions that are displayed on the screen.
 - c. If prompted, restart your system before the installation completes.
 - d. Click **Finish** to complete the installation of Data Protection for SQL.
 - e. If you plan to use VSS operations, you must install the most recent version of the Tivoli Storage Manager backup-archive client. The backup-archive client is also the VSS Requestor and is available separately.
2. Configure Data Protection for SQL.
 - a. Start the Management Console (MMC GUI) by clicking **Start > All Programs > Tivoli Storage Manager > Data Protection for Microsoft SQL Server > DP for SQL Management Console**. If you have not previously configured Data Protection for SQL, the Tivoli Storage Manager configuration wizard starts automatically.
 - b. If the Tivoli Storage Manager configuration wizard does not start automatically, click **Manage > Configuration > Wizards** in the tree view, select the wizard, and click **Start** in the Actions pane.
 - c. Complete the following pages of the wizard:

Data Protection Selection

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

Requirements Check

Click any **Failed** or **Warnings** links for help on resolving potential issues.

TSM Node Names

Enter the Data Protection for SQL Tivoli Storage Manager node name. If you are using VSS, enter the VSS Requestor node name.

TSM Server Settings

Specify the Tivoli Storage Manager server address, and choose whether to have the wizard configure the Tivoli Storage Manager server.

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

This page shows the status of the configuration. If you are using Windows Server 2008 or later, and you plan to use VSS operations, check the **VSS Diagnostics** check box to begin VSS verification.

Important: If you did not choose to have the wizard configure the Tivoli Storage Manager server, the Tivoli Storage Manager administrator must configure the Tivoli Storage Manager server before verification can be done. If the wizard does not configure the Tivoli Storage Manager server, it provides a link to a macro that can be provided to the Tivoli Storage Manager administrator as an example of one way to configure the Tivoli Storage Manager server.

3. Verify the configuration.
 - a. If you are using VSS operations on Windows Server 2008 or later, verify that VSS is working correctly.

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

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

Complete the following pages in the VSS Diagnostics wizard:

Snapshot Volume Selection

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

VSS Snapshot Tests

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

Completion

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

- b. Verify that Data Protection for SQL is configured properly:
 - 1) Click the **Automate** tab to open the integrated command-line interface.
 - 2) At the bottom part of the window, click the Open folder icon, and select the `verify_sql.txt` file. Then click **Open**.
 - 3) These commands are displayed in the command-line panel:

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

With the cursor in the command-line panel, press **Enter** to run the commands to verify your configuration. The configuration is verified when these commands run without warnings or errors.
 - 4) When verification is complete, you can use Data Protection for SQL to back up and restore SQL server data.

- 5) Back up and restore a set of test data.
4. Customize Data Protection for SQL.
After Data Protection for SQL is configured and verified successfully, customize your settings by defining your policy settings and scheduled operations to ensure that your business requirements are satisfied.

For detail informations about the installation and configuration procedures, or if you want to perform these tasks manually, see Chapter 3, “Installing Data Protection for SQL Server,” on page 45 and Chapter 4, “Configuring Data Protection for SQL,” on page 61.

Related concepts

“Managing policy” on page 75

Installation prerequisites

Before you install IBM Tivoli Storage Manager for Databases: Data Protection for SQL, ensure that your system meets the minimum hardware, software, and operating system requirements.

The following sections provide an overview of the minimum hardware and software requirements for the V6.3 release of Data Protection for SQL.

Details of the hardware and software requirements for Data Protection for SQL can evolve over time due to maintenance updates and the addition of operating system, application, and other software currency support.

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

Minimum hardware requirements

Before installing Data Protection for SQL, ensure that your system meets the minimum hardware requirements.

The hardware required to install Data Protection for SQL is listed in this section:

Hardware for an x86 platform

Compatible hardware supported by the Windows operating system and SQL Server in use.

Hardware for an x64 platform

Compatible hardware supported by the Windows operating system and SQL Server in use.

Minimum software and operating system requirements

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

The operating systems supported for the x86 platform are listed in this section:

- 32-bit Windows Server 2003 SP2, or later Service Pack levels: Standard, Enterprise, or Data Center editions
- 32-bit Windows Server 2003 R2 SP2, or later Service Pack levels: Standard, Enterprise, or Data Center editions
- 32-bit Windows Server 2008 SP2, or later Service Pack levels: Standard, Enterprise, or Data Center editions.

The operating systems supported for the x64 platform are listed:

- 64-bit Windows Server 2003 SP2, or later Service Pack levels: Standard x64, Enterprise x64, or Data Center x64 editions
- 64-bit Windows Server 2003 R2 SP2, or later Service Pack levels: Standard x64, Enterprise x64, or Data Center x64 editions
- 64-bit Windows Server 2008 SP2, or later Service Pack levels: Standard x64, Enterprise x64, or Data Center x64, editions
- 64-bit Windows Server 2008 R2, or later Service Pack levels: Standard x64, Enterprise x64, or Data Center x64 editions

Additional operating environments that are supported are listed in this section:

- Microsoft Cluster Server (MSCS) and Windows Failover Clustering
- Veritas Cluster Server (VCS) environments.

Microsoft SQL levels supported that are supported for x86 platforms are listed in this section:

- Microsoft SQL Server 2005 SP3, or later Service Pack levels: Standard or Enterprise editions
- Microsoft SQL Server 2008 SP1, or later Service Pack levels: Standard or Enterprise editions
- Microsoft SQL Server 2008 R2, or later Service Pack levels: Standard, Enterprise, or Data Center editions.

Microsoft SQL levels supported that are supported for x64 platforms are listed in this section:

- Microsoft SQL Server 2005 SP3, or later Service Pack levels: Standard x64 or Enterprise x64 editions
- Microsoft SQL Server 2008 SP1, or later Service Pack levels: Standard x64 or Enterprise x64 editions
- Microsoft SQL Server 2008 R2, or later Service Pack levels: Standard, Enterprise, or Data Center editions.

Virtualization support

Information for virtualization environments supported by Data Protection for SQL is available at the *IBM Tivoli Storage Manager (TSM) guest support for Virtual Machines and Virtualization* web page at: <http://www.ibm.com/support/docview.wss?uid=swg21239546>

Install Data Protection for SQL on a local system

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

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

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

Licensed package

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

Maintenance update (fix pack or interim fix package)

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

See the README.FTP file for instructions regarding 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.

Follow these instructions to install Data Protection for SQL from a DVD:

1. Install Data Protection for SQL using the setup wizard. This installs the product and any prerequisites such as the .NET Framework and Report Viewer.
 - a. Log on as an administrator.
 - b. Insert the IBM Tivoli Storage Manager for Databases: Data Protection for Microsoft SQL Server product DVD into your DVD drive. If autorun is enabled, the installation dialog starts automatically when the DVD loads. Otherwise, select **Start > Run**, and at the prompt, specify: `x:\setupfcm.exe`, where `x:` is your DVD drive, and click **OK**.
 - c. Follow the installation instructions displayed on the screen.
 - d. If prompted, restart your system before the installation completes.
 - e. Click **Finish** to complete the installation of Data Protection for SQL.

Note: The MMC GUI is shared among, Data Protection for Microsoft SQL Server, Data Protection for Microsoft Exchange Server, and Tivoli Storage FlashCopy Manager. If one of these products is installed in a non-default location, then the setup wizard defaults to the existing installation directory. Use the same directory when installing any of these products on the same computer. The default base directory is `c:\program files\tivoli`.

2. If you are installing Data Protection for SQL in a Microsoft Windows Failover Clustering environment or Veritas Cluster server environment, repeat the installation procedure on all nodes of your cluster.

3. To install additional language packs, see “Install the language packs for Data Protection for SQL.”
4. If you plan to back up and restore local snapshots or perform VSS offloaded backup operations, perform the tasks described in “Installing Tivoli Storage FlashCopy Manager.” If not, see Chapter 4, “Configuring Data Protection for SQL,” on page 61 for important configuration information.

Installing Tivoli Storage FlashCopy Manager

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

Install IBM Tivoli Storage FlashCopy Manager if you want to use advanced functionality with Tivoli Storage Manager such as:

- Offloaded backups
- Create and manage local (persistent) VSS snapshots
- VSS Fast Restore
- VSS Instant restore
- Query managed capacity

For information about how to install Tivoli Storage FlashCopy Manager, see http://publib.boulder.ibm.com/infocenter/tsminfo/v6r3/com.ibm.itsm.fcm.win.doc/c_fcm_installing.html.

After installing Data Protection for SQL and Tivoli Storage FlashCopy Manager, see Chapter 4, “Configuring Data Protection for SQL,” on page 61 for important configuration information.

Installing Data Protection for SQL in an MSCS or VCS environment

Instructions are provided for installing the Data Protection for SQL in a cluster environment.

Installing Data Protection for SQL in a MSCS or VCS environment requires the following:

1. Install Data Protection for SQL on both nodes of your MSCS or VCS if you are installing Data Protection for SQL for a clustered SQL Server.
2. Follow the instructions in “Install Data Protection for SQL on a local system” on page 49 for all nodes of your cluster.

Install the language packs for Data Protection for SQL

Use these instructions for installing Data Protection for SQL language packs.

Make sure that Data Protection for SQL has been successfully installed before attempting to install the Language Packs.

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

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

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

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

Data Protection for SQL Management Console language packs

```
x:\fcm\aaa\mmc\3100\bbb\setup.exe
```

DP for SQL language packs

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

Where *x:* is your DVD drive, *aaa* is either *x86* or *x64*, and *bbb* is the three-letter country code that is associated with that language.

3. Select setup.exe and click **OK**. Follow the installation instructions contained in the prompt windows.
4. Click **Finish** to complete the installation.

After installing the language pack, activate the language by updating the Data Protection for SQL configuration file (tdpsql.cfg by default) by using either of these methods:

- Use the **set** command with the **language** parameter to specify the language that you want. For example:
tdpsqlc set lang=fra

See the description of the **language** parameter in “Set positional parameters” on page 223 for a list of available languages and their three-letter country codes.

- Use the property pages to set the language by doing the following steps:
 1. Select the SQL server instance in the tree view.
 2. Click **Properties** in the Actions pane.
 3. Select the Regional property page.
 4. Click **Regional and Language Options** to ensure that system settings match the language that you want to use. The MMC GUI uses system language settings.
 5. Select the language from the list of installed language packs. The Data Protection components use language settings from a configuration file.
 6. For the best results and correct operation, select the language that matches the system settings. Click **Match MMC language** to automatically update the language to match the system.

Installing Data Protection for SQL silently

Administrators can install Data Protection for SQL by using silent installation. A silent installation runs on its own without any intervention so that administrators are freed from the task of monitoring the installation and providing input to dialogs.

This method is especially useful when Data Protection for SQL must be installed on a number of different computers with identical hardware. For example, a company might have 25 SQL servers installed across 25 different sites. To ensure a consistent configuration and to avoid having 25 different people enter Data Protection for SQL parameters, an administrator can choose to produce an unattended installation package and make it available to the 25 sites. The installation package can be placed on a DVD and sent to each of the remote sites, or the package can be placed on a file server for distribution to the sites.

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

Setup Program

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

Microsoft Installer (MSI)

Use `msiexec.exe` to install the MSI package.

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

Table 12. Silent installation options

Option	Description
<code>/i</code>	Specifies the program is to install the product.
<code>/l*v</code>	Specifies verbose logging.
<code>/qn</code>	Runs the installation without running the external user interface sequence.
<code>/s</code>	Specifies silent mode.
<code>/v</code>	Specifies the Setup Program to pass the parameter string to the call it makes to the MSI executable program (<code>msiexec.exe</code>). Note the following syntax requirements when invoking the <code>/v</code> option: <ul style="list-style-type: none">• A backslash (<code>\</code>) must be placed in front of any quotation marks (<code>" "</code>) that reside within existing quotation marks.• Do not include a space between the <code>/v</code> command line option and its arguments.• Multiple parameters entered with the <code>/v</code> command line option must be separated with a space.• You can create a log file by specifying the directory and filename at the end of the command. The directory must exist at the time a silent installation is performed.
<code>/x</code>	Specifies the program is to uninstall the product.
<code>addlocal</code>	Specifies features to install.
<code>allusers</code>	Specifies which users can use the installation package.
<code>installdir</code>	Specifies the directory where Data Protection for SQL is to be installed.

Table 12. Silent installation options (continued)

Option	Description
reboot	Specifies whether or not to prompt the user to reboot the system after silent installation. Force Always prompts user to reboot after silent installation. Suppress Suppresses prompt to reboot after silent installation. ReallySuppress Suppresses all reboots and prompts to reboot after silent installation.
rebootyesno	Specifies whether or not to reboot the system after silent installation. Specify <i>Yes</i> to reboot the system after silent installation. Specify <i>No</i> not to reboot the system after silent installation.
transform	Specifies language to install.

Note: Setting the **rebootyesno** option to *No* applies only to the installation of the Data Protection for SQL software. The installation package includes a number of prerequisites that will be installed by Data Protection for SQL if they have not been installed as prerequisites onto the system. 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.

Tip: Details of prerequisites are included in the Hardware and Software Requirements document for the release level. These requirements are linked to the All Requirements Document, see <http://www.ibm.com/support/docview.wss?uid=swg21218747>.

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

Table 13. Silent installation features (base client only)

Feature	Description
Client	Data Protection for SQL code

Table 14. Silent installation features (Language Packages only)

Feature	Description
LanguageFiles	Language specific files

The following transforms are used in this procedure.

Table 15. Silent installation transforms

Transform	Language
1028.mst	CHT Chinese (Traditional)
1031.mst	DEU German
1033.mst	ENG English
1034.mst	ESP Spanish
1036.mst	FRA French

Table 15. Silent installation transforms (continued)

Transform	Language
1040.mst	ITA Italian
1041.mst	JPN Japanese
1042.mst	KOR Korean
1046.mst	PTB Portuguese
2052.mst	CHS Chinese (Simplified)

Silently installing Data Protection for SQL with the setup program

Use the setup program (setup.exe) to silently install Data Protection for SQL.

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

Data Protection for SQL Management Console setup program

- (32-bit) x:\fcm\x86\mmc\3100\enu\setup.exe
- (64-bit) x:\fcm\x64\mmc\3100\enu\setup.exe

Data Protection for SQL setup program

- (32-bit) x:\fcm\x86\sql\6300\enu\setup.exe
- (64-bit) x:\fcm\x64\sql\6300\enu\setup.exe

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

The following commands are examples of the silent installation. You must substitute the appropriate feature when installing a language other than English. See Table 14 on page 53 for more information.

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

```
x:\fcm\aaa\mmc\3100\enu\setup.exe /s /v/qn
x:\fcm\aaa\sql\6300\enu\setup.exe /s /v/qn
```

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

- Here are examples of commands that specify the target directory, the features, language transform, boot suppression, and logging. Specify each command on a single line.

```
x:\fcm\x64\mmc\3100\enu\setup.exe /s
/v"INSTALLDIR="C:\Program Files\Tivoli\"
ADDLOCAL="Client" TRANSFORM=1033.mst
REBOOT=ReallySuppress /qn /!v "C:\Temp\DpSqlMmcSetupLog.txt"
```

```
x:\fcm\x64\sql\6300\enu\setup.exe /s
/v"INSTALLDIR="C:\Program Files\Tivoli\tsm\"
ADDLOCAL="Client" TRANSFORM=1033.mst
REBOOT=ReallySuppress /qn /!v "C:\Temp\DpSqlSetupLog.txt"
```

Note:

- You must place a backslash (\) before each quotation mark that is within an outer set of quotation marks (").
- 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 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 software. The installation package includes a number of prerequisites that will be installed by Data Protection for SQL if they have not been installed as prerequisites onto the system. Ensure that all the prerequisites are installed before starting the silent installation, then set the *rebootyesno* option to *No* so that no system restart is required after the silent installation process finishes.

Creating batch files

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

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

```
@echo off
rem =====
rem sample silent install script
rem
call x:\fcm\x64\mmc\3100\enu\setup.exe /s
/v"INSTALLDIR="C:\Program Files\Tivoli\" ADDLOCAL="Client" TRANSFORM=1033.mst
REBOOT=ReallySuppress /qn /! *v "C:\Temp\DpSqlMmcSetupLog.txt\"
rem
call x:\fcm\x64\sql\6300\enu\setup.exe /s
/v"INSTALLDIR="C:\Program Files\Tivoli\tsm\" ADDLOCAL="Client"
TRANSFORM=1033.mst REBOOT=ReallySuppress /qn /! *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 Data Protection for SQL with MSI (msiexec.exe)

You can silently install Data Protection for SQL by using the Microsoft Installer program, also known as MSI (msiexec.exe).

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

Important: Unlike setup.exe, using msiexec.exe does not install any prerequisites. When using msiexec.exe, you must install all prerequisites manually:

- 64-bit Data Protection for SQL Management Console requires the following prerequisite:

Microsoft .NET Framework 4 Client Profile

```
x:\fcm\x64\mmc\3100\enu\ISSetupPrerequisites\{29FC0BB0-95A7-4420-
B46C-68F628FB5FC4}\dotNetFx40_Client_x86_x64.exe
```

32-bit Microsoft Visual C++ 2010 Redistributable Package

```
x:\fcm\x64\mmc\3100\enu\ISSetupPrerequisites\{270b0954-35ca-4324-
bbc6-ba5db9072dad}\vc redistrib_x86.exe
```

64-bit Microsoft Visual C++ 2010 Redistributable Package

x:\fcm\x64\mmc\3100\enu\ISSetupPrerequisites\{7f66a156-bc3b-479d-9703-65db354235cc}\vcredist_x64.exe

Microsoft ReportViewer 2010 SP1 Redistributable

x:\fcm\x64\mmc\3100\enu\ISSetupPrerequisites\{E79DF561-C1FA-48ED-9BB0-6C9DC84ABAAE}\ReportViewer.exe

- 32-bit Data Protection for SQL Management Console requires the following:

Microsoft .NET Framework 4 Client Profile

x:\fcm\x86\mmc\3100\enu\ISSetupPrerequisites\{29FC0BB0-95A7-4420-B46C-68F628FB5FC4}\dotNetFx40_Client_x86_x64.exe

32-bit Microsoft Visual C++ 2010 Redistributable Package

x:\fcm\x86\mmc\3100\enu\ISSetupPrerequisites\{270b0954-35ca-4324-bbc6-ba5db9072dad}\vcredist_x86.exe

Microsoft ReportViewer 2010 SP1 Redistributable

x:\fcm\x86\mmc\3100\enu\ISSetupPrerequisites\{E79DF561-C1FA-48ED-9BB0-6C9DC84ABAAE}\ReportViewer.exe

- 64-bit Data Protection for SQL Server requires:

Microsoft .NET Framework 2.0 (not required if .NET 4 is installed)

x:\fcm\x64\sql\6300\enu\ISSetupPrerequisites\{552A6A80-5255-42AD-AF28-9660532ED1F6}\NetFx20SP2_x64.exe

Microsoft Core XML Services (MSXML) 6.0

x:\fcm\x64\sql\6300\enu\ISSetupPrerequisites\{726F97A8-63B9-4A58-ACFB-B8A56B383740}\msxml6_x64.msi

Microsoft SQL Server 2008 Native Client

x:\fcm\x64\sql\6300\enu\ISSetupPrerequisites\{82189DE7-71FD-4A69-AE68-E343859BA5E5}\sqlncli_x64.msi

Microsoft SQL Server System CLR Types

x:\fcm\x64\sql\6300\enu\ISSetupPrerequisites\{B1404676-6041-45EB-8BA7-65EFBA193C88}\SQLSysClrTypes_x64.msi

Microsoft SQL Server 2008 Management Objects

x:\fcm\x64\sql\6300\enu\ISSetupPrerequisites\{FEDA1805-E765-4496-A06D-26CFDAE590ED}\SharedManagementObjects_x64.msi

- 32-bit Data Protection for SQL Server requires:

Microsoft .NET Framework 2.0 (not required if .NET 4 is installed)

x:\fcm\x86\sql\6300\enu\ISSetupPrerequisites\{9277CA1D-4C87-4A9F-928F-0E95A021FA49}\NetFx20SP2_x86.exe

Microsoft Core XML Services (MSXML) 6.0

x:\fcm\x86\sql\6300\enu\ISSetupPrerequisites\{726F97A8-63B9-4A58-ACFB-B8A56B383740}\msxml6_x86.msi

Microsoft SQL Server 2008 Native Client

x:\fcm\x86\sql\6300\enu\ISSetupPrerequisites\{82189DE7-71FD-4A69-AE68-E343859BA5E5}\sqlncli_x86.msi

Microsoft SQL Server System CLR Types

x:\fcm\x86\sql\6300\enu\ISSetupPrerequisites\{B1404676-6041-45EB-8BA7-65EFBA193C88}\SQLSysClrTypes_x86.msi

Microsoft SQL Server 2008 Management Objects

x:\fcm\x86\sql\6300\enu\ISSetupPrerequisites\{FEDA1805-E765-4496-A06D-26CFDAE590ED}\SharedManagementObjects_x86.msi

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

Language packs can be installed in a similar way. MSI files for the language packs are located in the language folders that are associated with each component. For language packs use `ADDLOCAL="LanguageFiles"` instead of `ADDLOCAL="Client"`. See Table 14 on page 53 for more information.

- Install the Data Protection for SQL Management Console:

```
msiexec /i"x:\fcm\aaa\mmc\3100\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.

- Install Data Protection for SQL:

```
msiexec /i"x:\fcm\aaa\sql\6300\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.

Tip: Details of prerequisites are included in the Hardware and Software Requirements document for the release level. These requirements are linked to the All Requirements Document, see <http://www.ibm.com/support/docview.wss?uid=swg21218747>.

Note:

- You must place quotation marks (") around the following:
 - 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 listed in a custom installation must be specified after the *addlocal* option.

Installation problems: capturing a log of the installation

If a silent installation fails, you must record the symptoms and environment information for the failing installation and contact customer support with that information.

The following environmental information can be helpful:

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

Before contacting support, check for the following items:

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

Assuming that all looks correct, gather a detailed log of the failing installation into a file called `setup.log`. To generate a log file ensure that `/! *v \ "filename\"` is used on the command-line interface. For example, issue the following command (on a single line) to generate a log file named `C:\Temp\DpSqlSetupLog.txt`:

```
x:\fcm\x64\sql\6300\enu\setup.exe /s /v"INSTALLDIR=\ "C:\Program
Files\Tivoli\tsm\ " ADDLOCAL=\ "Client\ " TRANSFORM=1033.mst
REBOOT=ReallySuppress /qn /! *v \ "C:\Temp\DpSqlSetupLog.txt\ "
```

Creating the package on a DVD or a file server

The administrator can make the Data Protection for SQL installation package available in different ways.

You can burn a DVD or place the package in a shared directory on a file server. Typically, the package contains the Data Protection for SQL code distribution files and a batch file for silent install.

Creating a silent installation package

Follow these instructions to create a silent installation package.

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

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

Table 16. Commands for creating an 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 Data Protection for SQL DVD distribution files to the staging directory
<code>copy c:\setup.bat</code>	Replace the existing <code>setup.bat</code> with the one created in the previous step

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

Playing back the silent installation

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

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

From a silent installation package on DVD:

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

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

From a distribution directory:

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

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

In either case, the silent installation begins.

Setup error messages

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

In most cases, you, the administrator, will encounter these messages when a severe error occurs. Rarely will your users see these messages. When you get an error message, it is displayed in a message box. Every error message has a number. These are system error messages and there is no way to suppress them in your script.

If you encounter an error, you can go to this website: <http://support.installshield.com/default.asp>. Then, use the search facility to obtain information about the error.

Chapter 4. Configuring Data Protection for SQL

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

When you have completed the installation, you can begin the configuration.

1. Start the Management Console by clicking **Start > All Programs > Tivoli Storage Manager > Data Protection for Microsoft SQL Server > DP for SQL Management Console**.
2. If the Tivoli Storage Manager configuration wizard does not start automatically, start it by selecting the **Manage** node in the tree view and then selecting **Configuration > Wizards**.
3. Select the Configuration wizard and click **Start** in the Actions pane.
4. Enter the requested information on each wizard page and complete the wizard.

Important: If any warnings or error links appear, click them to get information about how to resolve the issues.

If you plan to automate your backup operations by using the Tivoli Storage Manager central scheduler, the Tivoli Storage Manager backup-archive client scheduler must also be installed and configured.

You can verify your configuration by doing the following steps:

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

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

Related tasks

“Quick installation and configuration” on page 45

Manual configuration

If you ever need to manually configure Data Protection for SQL, follow these instructions.

1. Configure Data Protection for SQL manually

Follow these instructions to manually configure Data Protection for SQL.

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

1. (Legacy only) Specify your Data Protection for SQL node name and communication method in the `dsm.opt` file located (by default) in the Data Protection for SQL installation directory. Additional options are also available. See “Specifying Data Protection for SQL options” on page 37 for additional information.
2. (Legacy only) Using the **set** command, specify your Data Protection for SQL preferences (language, date format, log file) in the `tdpsql.cfg` file located in the Data Protection for SQL installation directory. See “Specifying Data Protection for SQL preferences” on page 40 and “Set positional parameters” on page 223 for additional information.
3. (VSS only) Specify your **VSSPOLICY** statement in your Data Protection for SQL configuration file. See “Specifying Data Protection for SQL preferences” on page 40 and “Set positional parameters” on page 223 for additional information.
4. (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**. Note that the node name for this machine is referred to as the **Local DSMAGENT Node** and is specified with the *localdsmagentnode* parameter in the Data Protection for SQL configuration file (`tdpsql.cfg`). See *IBM Tivoli Storage Manager for Windows Backup-Archive Client Installation and User's Guide* and “Proxy node definitions (VSS Backups)” on page 35 for additional information.
5. (VSS only) Install and configure the Tivoli Storage Manager Client Acceptor Daemon (CAD) Service 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 proceeding to Step 6.
6. (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**. You can use the existing DSMAGENT if one is already installed and configured.
7. (VSS only) If you want to manage local persistent VSS snapshots, including VSS backups to LOCAL, VSS Instant Restores, and offloaded backups, you must install IBM Tivoli Storage FlashCopy Manager.
8. (VSS only) Install and configure a VSS provider. Consult the VSS provider documentation for information regarding configuration of that software. There is no installation or configuration required if you are using the default Windows VSS System Provider.
9. (VSS only) Change the SQL Server VSS Writer from Manual to Automatic and start the service.
10. (VSS only) Define storage space to hold VSS Backups that will reside on local shadow volumes. Make sure you define enough space to hold all copies of the VSS Backups as designated by your policies. Provisioning storage space to

manage VSS snapshots is dependent on the VSS provider that you use. Consult the VSS Provider documentation for more details.

See “Back up to Tivoli Storage Manager storage versus back up to local shadow volumes” on page 28 for recommendations regarding sufficient disk storage space.

2. Configure the Tivoli Storage Manager server manually

Make sure that Tivoli Storage Manager server is available before performing this task.

Perform these steps on the Tivoli Storage Manager server.

1. Define the policy domains, policy sets, management classes, copy groups, and storage pools needed to meet your Data Protection for SQL backup and restore requirements. For VSS operations, Tivoli Storage Manager server authentication must be on.
2. Register your Data Protection for SQL node name and password with the Tivoli Storage Manager **register node** command. For example, for VSS operations, this node is the Target Node. When registering nodes to the Tivoli Storage Manager server specifically for VSS operations, do not specify the Tivoli Storage Manager **USerid=NONE** parameter. VSS operations fail when this parameter is specified.
3. (VSS only) If not already defined, register your Tivoli Storage Manager backup-archive client node name and password for the machine where the SQL Server is installed. For example: For VSS operations, this agent node is the Local DSMAGENT Node.
4. (VSS only) If you plan to run offloaded backups from a particular machine, first register the Tivoli Storage Manager backup-archive client node name and password for the machine. For example, the agent node is the Remote DSMAGENT Node. *BAOFF* is used here (and in Step 5) to differentiate between this Remote DSMAGENT Node and the Local DSMAGENT Node (Step 3). You can replace *BAOFF* with the node name of your backup-archive client, and remove the *BAOFF* from the **grant proxynode** command.
5. (VSS only) Define the proxy node relationship (for the Target Node and agent nodes) with the Tivoli Storage Manager **grant proxynode** command. For example:

```
grant proxynode target=DP agent=BAnodename,BAOFF
```

Note: Always complete and fix any warnings that are flagged during the configuration process. Some warnings have a link to a macro that you can use to configure the Tivoli Storage Manager, and others have links to websites where you can download the packages you need to complete the configuration successfully.

3. Perform these tasks on the machine running the offloaded backups

This task is for VSS operations only.

Perform these steps on the machine running the offloaded backups:

1. Configure the Tivoli Storage Manager backup-archive client (if it is not already configured). If the backup-archive client is already configured, you can use existing client services. In the backup-archive client GUI menu, select **Utilities->Setup Wizard->Help me configure the TSM Backup Archive Client**.

The node name for this machine is called the Remote DSMAGENT Node and is specified with the *remotedsmagentnode* parameter in the Data Protection for SQL configuration file (tdpsql.cfg).

2. Install and configure the Tivoli Storage Manager Client Acceptor Daemon (CAD) Service and the Remote Client Agent Service (DSMAGENT) if they are not already installed. You can use an existing client CAD Service if one is already installed and configured. Use the backup-archive client Setup Wizard to guide you through the CAD installation process by selecting **Utilities->Setup Wizard->Help me configure the TSM Web Client**.
3. Install and configure a VSS provider (if you are not using the default system VSS provider). Consult the VSS provider documentation for information regarding configuration of that software.

4. Perform these tasks to verify your configuration

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

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

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

1. Click **Start > All Programs > Tivoli Storage Manager > Data Protection for Microsoft SQL Server > SQL Client - Command Line**.

2. Enter the following commands:

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

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

Verify that a SQL Server is ready to perform VSS operations

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

After all of these tests complete without errors, you are ready for the Tivoli Storage Manager component installations. For Windows 2003, before issuing the **VSHADOW** commands you must download the VShadow tool. For Windows 2008 and later, the DiskShadow tool is preloaded. You can run the **VSHADOW** commands, but you must download the VShadow tool in order to run them.

Using the VSHADOW command (Windows 2003)

1. Test the creation and the deletion of a non-persistent shadow copy:
 - a. From the command line, issue the Microsoft Windows VSHADOW command:

```
VSHADOW k: l:
```

where the variable "k:" is the SQL Server database volume and the variable "l:" is the SQL Server log volume.

- b. Repeat the preceding step four times.
- c. Verify that the Windows Event Log contains no errors.

2. Test the creation and the deletion of a persistent shadow copy:
 - a. From the command line, issue the Microsoft Windows VSHADOW command:

```
VSHADOW -p k: l:
```

where the variable "k:" is the SQL Server database volume and the variable "l:" is the SQL Server log volume.

Important: If you do not have enough space to run the command, you might want to delete shadow copies on the server. To delete *all* shadow copies on a server, issue the Microsoft Windows VSHADOW command: VSHADOW -da

- b. Repeat the preceding step four times.
 - c. Verify that the Windows Event Log contains no errors.
3. If you want to perform OFFLOAD VSS Backup operations, test the creation and the deletion of a non-persistent transportable shadow copy (VSS Hardware Provider environments only):

- a. From the command line, issue the Microsoft Windows VSHADOW command:

```
VSHADOW -t=export.xml k: l:
```

where the variable "k:" is the SQL Server database volume and the variable "l:" is the SQL Server log volume.

- b. On the server that you have designated for offloaded backup, from a command-line prompt, issue the Microsoft Windows VSHADOW command as follows:

```
VSHADOW -i=export.xml
```

- c. Verify that the Windows Event Log contains no errors.

Using the DISKSHADOW command (Windows 2008 and later)

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

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

```
DISKSHADOW>begin backup
DISKSHADOW>add volume f: (Database volume)
DISKSHADOW>add volume g: (Log volume)
DISKSHADOW>create
DISKSHADOW>end backup
DISKSHADOW>list shadows all (this may take a few minutes)
DISKSHADOW>delete shadows all
```

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

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

```

DISKSHADOW>set context persistent
DISKSHADOW>begin backup
DISKSHADOW>add volume f: (Database volume)
DISKSHADOW>add volume g: (Log volume)
DISKSHADOW>create
DISKSHADOW>end backup
DISKSHADOW>list shadows all (This may take a few minutes)
DISKSHADOW>delete shadows all

```

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

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

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

```

DISKSHADOW>set context persistent
DISKSHADOW>set option transportable
DISKSHADOW>begin backup
DISKSHADOW> add volume f: (Database volume)
DISKSHADOW> add volume g: (Log volume)
DISKSHADOW>set metadata c:\metadata\sqlmeta.cab (specify the path
where you want the metadata stored)
DISKSHADOW> create
DISKSHADOW>end backup

```

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

```

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

```

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

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

Diagnose the cause of common errors returned from VSS operations

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

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

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

ANS1532E (RC5722) Proxy Rejected: Proxy authority has not been granted to this node. This is displayed when the Tivoli Storage Manager server has not been configured for the proxy nodes correctly.

Chapter 5. Protecting Microsoft SQL Server data

Tasks required for backing up and restoring SQL Server data are provided.

Before you begin, Data Protection for Microsoft SQL Server must be installed and configured on your system and a SQL Server must be available.

Setting user preferences

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

The property pages described in this section customize preferences such as activity logging or how languages and information are displayed.

Be aware of the backup strategy, resource needs, policy settings, and hardware environment capabilities so that you configure these preferences to values that enhance Data Protection for SQL features.

1. In the tree view of the Management Console, select the SQL instance for which you would like to edit preferences.
2. Click **Properties** in the Action pane. A properties window appears, displaying the name of the selected SQL instance in the window title bar.
3. Select the property page that you would like to edit. Available property pages for a workload vary depending on whether it is configured for stand-alone snapshot support or Tivoli Storage Manager support.
For information about the available property pages, see “Data Protection Properties.”
4. Edit the property page and click **OK** (or **Apply**) to save your changes and close the window.

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

Data Protection Properties

Property pages are available for you to customize your configuration preferences.

The available property pages vary depending on whether it is configured for stand-alone snapshot support or Tivoli Storage Manager support.

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

The following table shows which property pages are available for which snapshot configuration. A stand-alone environment is only supported when Tivoli Storage FlashCopy Manager is installed.

Table 17. Available property pages for SQL Server workloads

Properties for the stand-alone environment	Properties for the Tivoli Storage Manager environment
Server Information	Server Information
Server Password	Server Password
Policy Management	Policy Management
VSS Policy Binding	VSS Policy Binding
Managed Capacity	Managed Capacity
Diagnostics	Diagnostics
SQL Login	SQL Login
Logging	General
Regional	Logging
VSS Backup	Regional
Custom Settings	VSS Backup
	Custom Settings
	Performance

Server Information

This property page displays information about the server that you contact for backup services.

Different information is displayed depending on whether the product is configured for stand-alone snapshot support or for Tivoli Storage Manager support.

Node name

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

TSM API version

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

Server name

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

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

Server Network Host name

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

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

Server type

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

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

Server version

The version of the Tivoli Storage Manager server.

Compression mode

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

Domain name

The policy domain that your node belongs to. A policy domain contains one or more policy sets.

Active Policy Set

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

Default Management Class

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

Server Password

Use this property page to change the password for accessing the Tivoli Storage Manager server. This property sheet applies only to Tivoli Storage Manager configurations.

Old password

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

New password

Type a new password. The password must be 1 - 63 characters in length, and can include any alphanumeric character, underscore (_), period (.), hyphen (-), plus (+), or ampersand (&).

Confirm new password

Type the new password again. Ensure that you click **OK** (or **Apply**) to save your changes.

VSS Policy Binding

Use this 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 bottom to the top and processing stops at the first match. To ensure that more specific statements are processed at all, the more general specification should be listed before the more specific ones, so as to be processed after the more specific specifications. Otherwise, the more general specification will match the target before the more specific specifications are processed.

Click a field to edit its content. Click **Add** to add a policy binding statement. Click **Delete** to delete a statement. Click **Up** and **Down** to modify the processing order. Click **OK** (or **Apply**) to save or apply your changes immediately.

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

Managed Capacity

Use this property page to track the capacity of currently managed storage. This information is helpful when you renew your product license.

The total managed capacity is displayed. Typically there is a difference between the capacity used by SQL Server data and the capacity of the volume that contains that data. For example, a set of SQL Server databases might require a capacity of 1 GB and occupy a 10 GB volume. When a snapshot of the volume is performed, the managed capacity measurement is 10 GB.

Click **Show Details** to view a list of the volumes that contain backups and their respective managed capacity.

Diagnostics

Use this 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 you want to use by clicking **Normal**, **Complete**, or **Custom**. Then click **Begin** to start the trace. Close the property page. Recreate 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 perform tracing only for MMC GUI.

Diagnostics 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 the MMC GUI only. Only MMC tracing can be performed here.

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, along with the specific trace flags, and guidance on when to use each mode.

Table 18. Diagnostics modes and their usage

Mode	Components traced along with trace flags used	When to use
Normal	MMC, DP (service), API (service,api_detail)	Use for Legacy operations, 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 this button to collect trace and log files for Legacy operations.

Complete

Click this button to collect trace and log files for VSS operations.

Custom

Click this button, then click the checkmark icon next to the button 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

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

Set Default Trace Flags

Click this button to set the most commonly requested trace flags.

Enable **Data Protection** tracing

Check this box to enable tracing of Data Protection for Microsoft SQL Server support. Click **Review** to view the trace file. Add or update trace flags in the field. Your service representative can tell you which trace flags to use.

Enable **DSM Agent** tracing

Check this box to enable tracing of the Tivoli Storage Manager client node. You must restart the TSM Client Acceptor service before starting the trace. Click **Review** to view the trace file. Add or update trace flags in the field. Your service representative can tell you which trace flags to use.

Enable **API** tracing

Check this box to enable tracing of the Tivoli Storage Manager API. Click **Review** to view the trace file. Add or update trace flags in the field. Your service representative can tell you which trace flags to use.

Event log entries

Click this button 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.

E-mail Select diagnostic files and click this button to send a diagnostic email to an IBM service representative with the selected files attached.

Screenshot

This button 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**.

When the tool opens, click **Add New Screenshot** to add a screen capture to the FlashCopyManager\ProblemDetermination folder. The screen capture can be selected with other diagnostic data.

Select All

Click **Select All** to select all files available in the diagnostic results window.

Copy Select the diagnostic files, click **Copy** to open a Browse For Folder dialog, and select a location to copy the selected diagnostic files.

View Click **View** to open the selected diagnostic file.

Delete Click **Delete** to delete the selected diagnostic files.

Tracing details for each component

All trace files are stored in the flashcopymanager folder, which is C:\Program Files\Tivoli\flashcopymanager by default. When the **End** diagnostics button is clicked 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 the MMC user settings file:

TraceFm.trc

Data Protection

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

TraceFileSql.trc

TraceFileExc.trc

TraceFileFs.trc

Agent Tracing options are stored in the VSS requestor dsm.opt file:

TraceFileAgent.trc

API Tracing options are stored in the respective Data Protection dsm.opt file:

TraceFileSqlAPI.trc

TraceFileExcAPI.trc

TraceFileFsAPI.trc

General

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

SQL Server

Specify the unique name that identifies the SQL Server instance.

From Server

Use this field when you want to restore a SQL database or database component from another SQL Server. Enter the name of the other SQL Server from which you want to restore backups. By default, this field displays the same name for the **SQL Server** and the **From SQL Server**.

Wait for tape mounts for backup or restore

Check this box when you want Data Protection for SQL 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, it is likely that 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

Check 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 using SQL Server compression

Check 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.

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 if sufficient storage space is available for the backup. The value that is specified here becomes the default value for all differential backups.

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

Logging

Use this property page to specify activity log preferences.

Log File Name

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

Enable pruning

Select this option to automatically delete older entries from the log. By default, log pruning is activated and performed daily.

Number of days to keep old entries

Specify 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 button to prune the activity log for one command run.

Regional

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

Regional and Language options

Click this button to set preferences for the Management Console. The Management Console 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**), and year (**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**).

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.

Match MMC Language

Click this button to change the MMC regional settings to match the system's regional and language options. Clicking this button also matches the number, date, and time formats to the default formats of the selected language.

VSS Backup

Use this property page to configure preferences used during VSS backup operations.

Default Backup Destination

Select the default storage location for your backups. You can select from these storage locations:

TSM The backup is stored on Tivoli Storage Manager server storage only. This selection is the default.

Local The backup is stored on local disk only.

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

Local DSMAGENT Node name

Specify the node name (the DSM Agent node) of the local client system that creates the VSS backups. This parameter must be specified for VSS operations to succeed.

Remote DSMAGENT Node name

Specify the node name of the machine 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.

Custom Settings

Check the box to display **Show Refresh Options** in the toolbar in the Recover view.

This box is not selected by default. It is useful when used in environments with many thousands of objects stored on a Tivoli Storage Manager server. The administrator can use the **Refresh Options** button and toolbar to switch between manual and automatic refresh mode.

Automatic and manual refresh modes differ in the following manner:

- In automatic refresh mode, the first time a view is selected, it is automatically refreshed. If there are tens of thousands to millions of objects on the server, the refresh can take a long time to complete.
- In manual refresh mode, no automatic refresh takes place. A name filter is available on the **Refresh Options** toolbar, which can be used to narrow down the selection of objects. After you enter a name pattern, you can click **Refresh**. Using manual refresh mode can greatly reduce the amount of information that is returned from the server, and in turn can greatly speed up the completion time. You can also specify a wildcard character (*) in the name pattern to assist your filtering effort.

Performance

Use this property page to set preferences that affect performance.

DP Buffers

Specify a number from 2 to 8 that specifies the number of communication data buffers that Data Protection for SQL uses when transferring data to the Tivoli Storage Manager server. Each buffer is the size that is specified by the **DP Buffer Size** option. This option applies to Legacy backups only.

DP Buffer Size

Specify a number from 64 to 8192 that specifies the size of the buffers that

are used by Data Protection for SQL to transfer data to the Tivoli Storage Manager server. This option applies to Legacy backups only.

SQL Buffers

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

SQL Buffer Size

Specify a number from 64 to 4096 that specifies the size of the buffers that are used by Data Protection for SQL to transfer data from the SQL Server to Data Protection for SQL. This option applies to Legacy backups only.

Stripes

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

SQL Login

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

Use Windows Authentication

Click this radio button to use a trusted connection and allow Microsoft Windows to authenticate the logon.

Use SQL Server Authentication

Click this radio button to use SQL User ID security. With this type of security, the administrator provides the logon ID and the password to logon to the Microsoft SQL Server.

User name

Type the SQL User ID.

Password

Type the password to logon to the Microsoft SQL Server.

Managing policy

Data Protection for SQL uses policy to determine how backups are retained.

Although Tivoli Storage Manager policy determines how Data Protection for SQL backups are managed on Tivoli Storage Manager storage, backup retention on local shadow volumes is dictated by version and time-based policies. Sufficient local storage space must be available on local shadow volumes for a VSS backup strategy to be successful. Ensure that there is enough available storage space assigned to the volumes to accommodate your backup operations. The shadow copy volume that is the storage destination of a snapshot must have sufficient space for the snapshot. Environment and storage resources also affect how many backup versions are maintained on local shadow volumes. The amount of space required is dependent on the VSS provider that is used.

For information about Tivoli Storage Manager policy, see “How Tivoli Storage Manager server policy affects Data Protection for SQL” on page 19.

Policy binding statements

Policy binding statements associate SQL backups to a management policy.

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

To avoid making errors, use the GUI to specify policy binding statements.

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

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

How backups expire based on policy

Backups are expired based on Data Protection for SQL policy.

Expiration is the process by which SQL Server backup objects are identified for deletion because their expiration date has passed or the maximum number of backup versions to be retained is reached. The value of this data is dependent on the business needs as identified by the recovery point objective (RPO) and the recovery time objective (RTO). For example, legal, operational, and application requirements impact how data must be protected to meet these RPO and RTO demands. To support such requirements, Data Protection for SQL allows you specify the number of backups to retain and the length of time to retain them. Expiration is how Data Protection for SQL implements this function.

Expiration of backups occurs during the first query, backup, or restore operation of a Data Protection for SQL session. Expiration of backups might also occur during any backup operation.

If an operation occurs when the maximum number of backup versions to be retained (as specified by the Data Protection for SQL policy) is reached, the oldest backup version is expired and deleted before creating, restoring, or displaying information about a backup.

If an operation occurs when the maximum number of days to retain a backup (as specified by the Data Protection for SQL policy) is reached, the *inactive* backup versions older than the number of days specified are expired before creating, restoring, or displaying information about a backup.

Binding backups to a policy

Add, update, delete, or change the processing order of existing binding statements. Policy determines how backups are managed and retained.

1. Start the Management Console.
2. Select the **SQL Server** instance from the tree view.
3. In the **Protect** tab, click **Properties** in the **Action** pane. A dialog appears that displays properties for the selected instance.
4. Select **VSS Policy Binding** from the list of available property pages. The existing bindings are displayed.
5. Add, update, delete, or change the processing order of existing binding statements.

Tip: Click a field to edit it.

- Any field can take a wildcard character (*) to mean "all". For example, specify a wildcard character (*) in the **Server** field to bind the policy to all SQL servers.
 - All fields, other than the **Server** field have drop-down menus of available options.
6. Use **Move Up** and **Move Down** to modify the processing order. Policies are processed from the bottom up and processing stops at the first match. To ensure that more specific statements are processed at all, the more general specification should be listed before the more specific ones, so as to be processed after the more specific specifications. Otherwise, the more general specification will match the target before the more specific specifications are seen.
 7. Save any new or changed binding statement by clicking **Save statements**.
 8. Verify new or updated policies and bindings.
 - a. Run one or more test backups.
 - b. In the **Recover** tab, verify the management classes that are bound to your test backups.

Determining managed storage capacity

Tracking the capacity of currently managed storage assists during license renewal.

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

To determine managed storage capacity:

1. Select an SQL instance from the GUI.
2. In the **Protect** tab, Click **Properties** in the **Action** pane. A dialog appears that displays properties for the selected instance.
3. Select **Managed Capacity** from the list of available property pages. The managed capacity is calculated and displayed.
4. Click **Show Details** to view a list of the volumes (that contain backups) and their respective managed capacity.
5. Click **OK** to close this dialog.

Using the Task Manager

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

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

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

Use the **Task List** view to complete these tasks:

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

Backing up SQL databases and files

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

Security

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

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

Three options are provided when specifying SQL Server logon information:

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

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

Backup types

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

Data Protection for SQL provides six types of backup:

Full database backup (Legacy and VSS)

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

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

Differential backup (Legacy only)

Data Protection for SQL backs up only the data pages in a SQL Server database instance that have changed since the last full backup and a portion of the transaction log. This is equivalent to an incremental backup on the Tivoli Storage Manager backup-archive client.

Log backup (Legacy only)

Data Protection for SQL backs up only the contents of a SQL Server database transaction log since the last successful log backup. To do the first log backup, you

need to have done a full backup or its equivalent first. Log backups normally follow full backups. The portion of the log included in full and differential backups is not equivalent to a log backup. Additionally, in full and differential backups, the log is not truncated as it is during a log backup. However, a log backup following a full or differential backup will include the same transactions as a full or differential. Log backups are not cumulative as are differential; they must be applied against a base backup and in the correct order.

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

File backup (Legacy only)

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

Group backup (Legacy only)

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

Set backup (Legacy only)

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

Examples

See “Legacy Backup output examples” on page 134 for samples of various backup types using the command line interface. Also, planning information regarding choosing the best strategy for your backup requirements is available in “Backup strategies” on page 21.

Backing up SQL databases by using the VSS method

Perform these tasks to back up SQL server data by using Microsoft Volume Shadow Copy Service (VSS) technology.

Before you begin, see “Security” on page 21 for the settings to use in a secure environment.

If you want to manage local VSS backups or perform offloaded backups to Tivoli Storage Manager server storage, you must have Tivoli Storage FlashCopy Manager configured in your environment.

1. Start the MMC GUI.
2. If you are going to utilize offloaded backups, make sure a **Remote DSMAGENT Node name** is specified. An offloaded backup uses another machine (specified with the **Remote DSMAGENT Node name** parameter) to move SQL data to Tivoli Storage Manager server storage. This might reduce the load on network, I/O, and CPU resources during backup processing.

To verify or specify a remote **Remote DSMAGENT Node name**, select the **SQL Server** instance in the tree view, and click **Properties** in the Actions pane. Then select the VSS Backup property page. If the **Remote DSMAGENT Node name** is blank, enter a node name.

3. In the **Protect** tab of the SQL instance, specify the type of SQL data to back up:
 - Select **View: Databases** for a list of discovered SQL databases that are available for backup.
 - Select **View: Files** for a list of discovered SQL database files that are available for backup.

Use the results pane to browse and select the databases or files to back up.

Tip: Fine tune the list of available databases in the results pane by entering a keyword in the **Search** field.

4. Verify backup options. If the backup options are not currently displayed, click **Show Backup Options**.
 - If you want to use offloaded backups, select **True** in the **Offload** field. This field applies only to VSS backups.
5. In the Actions pane, click **Backup Method** and select **VSS**.
6. In the Actions pane, click the **Backup Destination** and select where to store the backup:
 - **Local**
 - **TSM**
 - **Both**
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 will start, complete with the appropriate command that is required to complete the task.
8. Create the backup by clicking **Full Backup** in the Actions pane. You can also right-click a database, and select the backup action you want from the context-menu.

Backing up SQL databases by using the Legacy method

This task describes how to complete a Legacy back up of SQL databases using the MMC GUI.

Three types of backup are supported:

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

Differential

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

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

To perform a Legacy Backup, the IBM Tivoli Storage Manager for Databases Data Protection for Microsoft SQL Server license file must be installed.

Before you begin, see “Security” on page 21 for the settings to use in a secure environment.

Follow these steps to perform a Legacy backup of your data:

1. Start the Management Console.
2. Select the SQL Server instance in the tree.
3. In the **Protect** tab for the SQL instance, make sure that the **Databases** view is selected.
4. Select one or more databases to back up.
5. Verify backup options. If the backup options are not currently displayed, click **Show Backup Options**.
 - Use the **Stripes** option to specify the number of data stripes to use in a backup or restore operation. The *numstripes* variable can range from 1 to 64. The default value is 1.
 - Use the **DiffEstimate** option to specify the estimated fraction of the database that has changed since its last full database backup. This estimate is needed because SQL Server does not provide a way to determine the size of a differential backup, and because the Tivoli Storage Manager server requires an accurate size estimate to efficiently allocate space and place objects. The Tivoli Storage Manager server uses this value to determine whether there is enough space in the primary storage pool to contain the backup. The default value is 20.
 - Use the **LogEstimate** option to specify the estimated the fraction of a SQL database that has changed due to non-logged operations since the last log backup. The default value is 0.
 - Use the **Truncate** option to specify whether to dispose of entries you no longer need in the SQL database transaction log after you back up the log. In general, you do not want to truncate the log when rebuilding a corrupted database. This option enables the server to back up the transaction log but does not try to touch the data in any way. It writes all transaction log entries from the time of the last log backup to the point of database corruption. If you do not truncate the transaction log, you might be able to back up the transaction log of a damaged, suspect, or unrecoverable SQL Server database. The default value is **False**.
6. In the Actions pane, click **Backup Method** and select **Legacy**.
7. In the Actions pane, click the **Backup Destination** and select where to store the backup:
 - **Local**
 - **TSM**
 - **Both**
8. Create the backup by clicking one of these actions in the Actions pane:
 - If your backup destination is **TSM**, select one of these actions:
 - **Full Backup**
 - **Differential Backup to TSM**
 - **Log Backup to TSM**
 - If your backup destination is **Local** or **Both**, click **Full Backup**.

Backing up SQL groups or files by using the Legacy method

This task describes how to complete a Legacy back up of SQL groups or files using the MMC GUI.

The following three types of backup are supported.

Group Backs up the contents of the specified file group.

File Backs up the contents of the specified logical file.

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

Attention: You must back up the transaction logs after completing a Group, File, or Set backup operation. See “Backup types” on page 5 for a description of a log backup.

Before you begin, see “Security” on page 21 for the settings to use in a secure environment.

Follow these steps to perform a Legacy backup of SQL groups or files.

1. Start the Management Console.
2. Select the SQL Server instance in the tree.
3. In the **Protect** tab for the SQL instance, make sure that the **Files** view is selected.
4. Select one or more groups or files to back up.
5. Verify backup options. If the backup options are not currently displayed, click **Show Backup Options**.
 - Use the **Stripes** option to specify the number of data stripes to use in a backup or restore operation. The *numstripes* variable can range from 1 to 64. The default value is 1.
6. Create the backup by clicking one of the following **Actions**.
 - **Group Backup to TSM**
 - **File Backup to TSM**
 - **Set Backup to TSM**

Restoring SQL databases and files

Restore SQL databases and files with Data Protection for SQL.

Restoring SQL server data

Follow these steps to restore SQL server data.

Before you begin, see “Security” on page 21 for the settings to use in a secure environment.

Restore Databases allows you to restore databases or parts of databases only from **full**, **differential**, and **log** backups. VSS only supports full backups. Legacy differential and Legacy log backups can be applied after a full VSS Backup has been restored.

- Make sure to review the “VSS Restore considerations” on page 85 before attempting any type of VSS Restore.

- A Legacy restore or VSS Restore of the master database requires special attention. See “Restoring the master database” on page 87 for step-by-step instructions.

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

1. Start the Management Console.
2. Select the SQL Server instance in the tree.
3. In the **Recover** tab for the SQL instance, specify the type of SQL data to restore:
 - Select **View: Databases** for a list of discovered SQL databases that are available for restore.
 - Select **View: Files** for a list of discovered SQL database files that are available for restore.
4. Use the results pane to browse and select the databases or files to restore. The following features are available:
 - Search: Fine tune the list of available databases or files in the results pane by entering a keyword in the **Search** field.
 - Filter: Use the filter options to narrow the list of items in the result pane.
 - a. Click **Show Filter Options** and **Add Row**.
 - b. Click the down arrow in the **Column Name** field and select an item to filter.
When you click **Select All**, all rows that reflect the filter specifications are selected.
 - c. Select an operator in the **Operator** field.
 - d. Specify a value to filter on in the **Value** field.
 - e. In you want to filter on additional items, click **Add Row**.
 - f. Click **Apply Filter** to filter the items on the list.
 - Backups: You can click **Active Backups** to show only active backups, or click **All Backups** to show both active and inactive backups.
 - Refresh: Click **Refresh** to update the view with your changes.

If you applied a filter, the objects on the server that match the filter or search criteria are listed in the **Recover** tab. The status area indicates the number of items that match the criteria n of x displayed, where n equals the number of objects that match the filter criteria, and x is the number of objects that are retrieved from the server. For example, "5 of 20 displayed." If you specify refresh options to further narrow your results, and click **Refresh** again, the objects on the server that match the filtered and refresh options are displayed. Each time you click **Refresh**, another query is run against the Tivoli Storage Manager server.

5. Verify restore options. If the restore options are not currently displayed, click **Show Restore Options**.
 - Use the **Restore Into** option to specify a target location for the restore operation.
Attention: Any type of **Restore Into** function automatically disables VSS Instant Restore.
 - Use the **Instant Restore** option to turn Instant Restore on or off. Disable Instant Restore if you want to use Fast Restore.
Attention: Instant Restore will overwrite all files on the destination file system.

- See the following sections for additional information:
 - “Restore options” on page 88
 - “VSS Instant Restore” on page 8
 - “VSS Fast Restore” on page 8
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 will start, complete with the appropriate command that is required to complete the task.
 7. Click **Restore** in the Action pane to begin the restore operation.

VSS Restore considerations

Be aware of 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):

- If you plan to perform a VSS Restore of the master database, see “Restoring the master database” on page 87 for step-by-step instructions.
- A VSS Instant Restore overwrites the entire contents of the source volumes. However, you can avoid overwriting the source volumes by specifying **InstantRestore False** in the MMC GUI. This option bypasses volume-level copy and uses file-level copy instead to restore the files from a VSS Backup that resides on local shadow volumes. It is recommended that the source volume contain only the SQL database.
- Be aware that when a VSS restore from local shadow volumes is performed, the bytes transferred will display "0". That is because no data ("0") is restored from the Tivoli Storage Manager server.
- When performing VSS Instant Restores, a best practice is to make sure that any previous background copies (that involve the volumes being restored) are completed prior to initiating the VSS Instant Restore. However, this check is not necessary for XIV, SAN Volume Controller, or Storwize V7000 with space-efficient target volumes.

Restoring SQL file groups and files from Legacy backups

This task describes how to restore a Legacy backup of SQL file groups and files using the MMC GUI.

Before you begin, see “Security” on page 21 for the settings to use in a secure environment.

With the Restore Groups/Files option, you can restore databases or parts of databases from **group**, **file**, **set**, **log** and **full** Legacy backups.

Restoring parts of a database from a **full** Legacy backup is referred to 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:

1. Use the **Files** view on the Recover tab to select and restore the full backup object. Make sure the **Recovery** option is set to **False**.
2. If you plan to apply a log restore with point-in-time, use the Restore Databases tab and the **Point in Time...** option to restore the log. Make sure the **Recovery** option is set to **True**.

3. If you plan to apply a differential restore, use the Restore Databases tab to perform a differential restore. Make sure the **Recovery** option is set to **True**.

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

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

1. Start the MMC GUI.
2. Select the SQL Server instance in the tree.
3. In the **Recover** tab for the SQL instance, make sure that the **Files** view is selected.
4. Select one or more groups, files, or sets to restore.
5. Verify restore options. If the restore options are not currently displayed, click **Show Restore Options**.
 - Use the **Restore Into** option to specify a target location for the restore operation.
Attention:
Any type of **Restore Into** function automatically disables VSS Instant Restore.
 - See “Restore options” on page 88 for descriptions of additional restore options.
6. Click **Restore** in the Action pane to begin the restore operation.

Restoring to an alternate machine

This procedure demonstrates how to restore a SQL database backup to an alternate SQL Server machine (or database) using the MMC GUI. Note that Data Protection for SQL must be installed on both machines. Note that unlike Legacy backups, VSS Backups cannot be restored into a SQL Server that has a different name.

This procedure uses the following terms:

Source machine

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

Target machine

The alternate machine to which the backup will be restored.

1. Copy the Data Protection for SQL options file (dsm.opt) from the source machine to the target machine. By default, the dsm.opt file is located in the C:\Program Files\Tivoli\TSM\TDPSql directory. If *passwordaccess generate* is specified in this dsm.opt file, it may be necessary to reset the password for this node on the Tivoli Storage Manager server.
2. Start the Management Console.
3. In the **Recover** tab for the SQL instance, select one or more databases to restore.
4. Verify restore options. If the restore options are not currently displayed, click **Show Restore Options**.
 - a. Click **Wait for Tape Mounts for Restore**.
 - b. Click **Wait for Tape Mounts for File Information**.
 - c. You might have to click **Show Active and Inactive** in order to view the database to be restored.
 - d. If the database to be restored is going to replace an existing database on the target machine, click **Replace**.

- e. Use the **Restore Into** option to specify a target location for the restore operation.
Attention: Any type of **Restore Into** function automatically disables VSS Instant Restore.
- f. Use the **Instant Restore** option to turn Instant Restore on or off. Disable Instant Restore if you want to use Fast Restore.
Attention: Instant Restore will overwrite all files on the destination file system.

See “Restore options” on page 88 for descriptions of additional restore options.

5. Click **Restore** in the **Action** pane to begin the restore operation.

Restoring the master database

A damaged master database can result in the SQL Server failing to start, as well as a number of other error conditions. A special procedure is required to restore the master database. Note that during the process of rebuilding the master database, the SQL Server setup program drops and then recreates the **msdb** database so it must be restored along with the master database.

In general, the following steps are required:

1. Run the SQL Server setup program to rebuild the master database. You must rebuild using the same character set and sort order as the master database backup that will be restored.
2. Start the SQL Server in single-user mode. This can be done at a command prompt. See also Note 1 under “Setting user mode” on page 98.
3. Use Data Protection for SQL to restore the master database.

Note: When the master database has finished restoring, the SQL Server shuts itself down. As a result, an error message is generated stating that the connection was lost to the SQL Server. This is expected.

4. Restart the SQL Server normally (in multi-user mode).
5. Manually reapply any changes that were made to the master database *after* the date of the database backup used to do the restore operation.
6. Use Data Protection for SQL to restore the **msdb** database.

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

- ALTER DATABASE
- CREATE DATABASE
- DISK INIT
- DISK RESIZE
- DISK MIRROR
- DISK UNMIRROR
- DISK REMIRROR
- Various DBCC options such as SHRINKDB

- System stored procedure such as: sp_dropremotelogin, sp_addumpdevice, sp_dropdevice, sp_addlogin, sp_droplogin, sp_addserver, sp_dropserver, sp_addremotelogin

Restore options

Descriptions of the options available in the Data Protection for SQL GUI restore windows are provided.

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

AutoSelect

Set this option to **True** to enable auto-selection. With auto-selection, when you select the most recent backup to restore, all other necessary backups are automatically selected for you, up to the previous full backup.

Stripes

Under Performance, the number of **Stripes** is listed. You can specify the number of data stripes to use in a restore operation. A maximum of 64 data stripes is allowed. The default value is 1. To change the Stripes value on the default 1 and type in the new number. The value you enter should correspond to the value set for SQL buffers. Note that this option is always enabled for Legacy backups. Stripes are not available for VSS backups.

Database Owner Only

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

Instant Restore

Under Restore Behavior, **InstantRestore** is listed with a default value of **True**. You can disable Instant Restores by setting the value to **False**, which bypasses volume-level copy and uses file-level copy to restore the files from a local VSS Backup. If this option is set to **True**, the volume level snapshot restore is used for local VSS Backups if the backup exists on SAN-attached volumes. The default value is to use **volume level snapshot restore** if it is supported. This option is available for VSS operations only. When performing VSS Instant Restores, a best practice is to make sure that any previous background copies (that involve the volumes being restored) are completed prior to initiating the VSS Instant Restore.

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

Recovery

Under Restore Behavior, **Recovery** is listed with a default value of False. Change this option to True to allow the Data Protection for SQL to run a recovery after the restore operation has succeeded.

Replace

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

Smart Select

Under Performance, **SmartSelect** is listed with a default value of False. Change this value to True if you want to enable the smart select capability.

Stand By Undo File Name

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

Wait for Tape Mounts for Restore

Under Tape, you can specify whether or not the Data Protection for SQL restore operation waits for the Tivoli Storage Manager server to mount removable media such as tapes or DVDs. This information is retrieved from Tivoli Storage Manager when you click the Recover tab, or select the Refresh button. The default value is True.

Wait for Tape Mounts for File Information

Under Tape, **WaitForTapeMountsForFileInformation** is listed with a default value of True. When querying Tivoli Storage Manager for file information, you can specify whether or not Data Protection for SQL waits for the Tivoli Storage Manager server to mount removable media. This option applies to Legacy restores only.

From the **Restore Databases** view, the following additional options are available:

Point in Time

You can specify a point in time to which to restore a database if desired by clicking the **Point in Time** button. This button is enabled only when you select for restore a **full** backup object and at least one **log** backup.

Point in Time Dialog

Clicking on the Point in Time button displays a dialog box with the following options:

- No point in time
- Stop at
- Stop at mark
- Stop before mark

The **stop** radio buttons allow you to specify a date and time. With **Stop at mark** and **Stop before mark**, you can name a mark for the restore and include the date and time to help locate the mark.

To clear a point in time that is set, select the **No point in time** radio button.

When **point in time** is in use, a static field is enabled to display the results of the action.

Shortcut Menu: You can display additional restore options by right-clicking a selected item in the list control. From the **Restore Groups/Files** tab, this menu is available only when you highlight a database in the tree. All of its backup objects will be displayed in the list control, and the menu will be available for any selected objects. The right-click popup menu contains the following items:

Restore Into

Use this option to specify the database to restore a backup object to. Click **Restore Into** to display an edit box. If you have selected several databases to be restored, the **Restore into** name you specify applies only to the selected backup object that you right-clicked. If other selected backups require the *Restore into* parameter, you will have to specify them one at a time, but you can do this in one restore operation.

Relocate

Use the Relocate dialogs to specify new destination locations in which to restore backed up SQL databases, logs, and SQL Server full-text index files (SQL Server 2005) or FILESTREAM files (SQL Server 2008 and SQL Server 2008 R2):

- *Relocate All Files Into a Directory*: Select this option to restore the SQL data files, logs, and other related files into a location different from where the data was originally backed up.
 - *Relocate Log Files Into*: Check this box to restore the log files into a location different from where the SQL database and other related files are being restored.
 - *Relocate Other Files Into*: Check this box to restore SQL Server full-text index files (SQL Server 2005) or FILESTREAM files (SQL Server 2008 and SQL Server 2008 R2) into a location different from where the SQL database and logs are being restored.
- *Relocate Files Individually*: Select this option to restore each SQL database, log, and SQL Server full-text index file (SQL Server 2005) or FILESTREAM files (SQL Server 2008 and SQL Server 2008 R2) individually. This is available for Legacy backups only.

Standby Server Undo File

Use this option to specify the undo file for a Legacy restore to a standby SQL database. If the target SQL database is not already in standby mode, it will be placed in standby mode. This menu item appears only in the **Restore Databases** window and is available for full, differential, and log backup types, but only for one database at a time. Click this option to display an edit box for the undo file name. Once you specify this for a database, it applies to all backup objects for that database. Likewise, once you remove this option for a backup object, it is removed for all.

Note that the MMC GUI does *not* support the */relocate.../to* or */relocatedir* parameters for partial restores. You must use the command line interface when performing a partial restore that requires these parameters.

Inactivating SQL databases (Legacy only)

Use these tasks to inactivate an existing legacy backup of SQL databases.

Use the **Inactivate** action to make an SQL database backup **Inactive** on the Tivoli Storage Manager server. Typical backups do not require this command as Tivoli Storage Manager inactivates an SQL database backup as a part of Tivoli Storage Manager policy management. As a result, backup objects are typically inactivated as part of the scheduled backup processing.

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

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

Note: The SQL database that you want to inactivate must be a Legacy backup. VSS backups cannot be inactivated by using this method. The **Inactivate** action in the Actions pane is not available for VSS backups.

To inactivate backup objects:

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

Inactivate tree and list

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

Deleting SQL Server Backups

Perform these steps to delete an SQL Server backup that was created with the VSS backup method.

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

To delete an SQL Server backup:

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

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

Using the Tivoli Storage Manager scheduler

This section provides an example of how to use the Tivoli Storage Manager scheduler with Data Protection for SQL to automate online backups of SQL Server databases. This example illustrates the use of the Tivoli Storage Manager scheduler and creating a full backup of your SQL databases.

Note: To ensure that this example works, obtain and install the latest Tivoli Storage Manager backup-archive client.

When Data Protection for SQL has been registered to a Tivoli Storage Manager server and installed on the SQL Server, the procedure consists of these steps:

1. **On the Tivoli Storage Manager server:**

- a. Define a schedule to run a Windows command file. This schedule must be defined in the policy domain to which Data Protection for SQL is registered.
 - b. Associate the Data Protection for SQL node to the defined schedule.
2. **On the machine where the SQL Server and Data Protection for SQL are installed:**
- a. Install the Tivoli Storage Manager server scheduler client as a Windows service for Data Protection for SQL. If a scheduler exists for the regular Tivoli Storage Manager backup client, install another one for Data Protection for SQL.
 - b. Define a command file that contains the Data Protection for SQL commands to do the desired backup.

Note: See the Windows documentation for details on creating a Windows command file.
 - c. If you are running in a SQL Server Failover Cluster environment:
 - 1) Install the Tivoli Storage Manager server scheduler client as a Windows service on both cluster nodes.
 - 2) Create a cluster resource that represents the Tivoli Storage Manager scheduler. Verify that the cluster resource is started. This service is dependent on the SQL Server cluster resource.
 - d. Start the scheduler installed in step 2a.

Example scheduler procedure

This example assumes the following environment:

- Data Protection for SQL is registered to a Tivoli Storage Manager server:
 - The node name is *mynode*.
 - The password is *mypassword*.
 - The policy domain is *mydomain*.
- The event to be scheduled:
 - A daily full backup of all of the SQL databases.
 - The backups begin between 9:00 and 9:15 pm.

This method is flexible because you can define a command file with any set of commands you choose. This allows you to use the same Tivoli Storage Manager schedule for other Data Protection for Microsoft SQL Server applications on Windows as well.

Complete these tasks on the Tivoli Storage Manager server

You must set up a scheduler service on the machine where the Backup-Archive Client is installed before performing this procedure.

1. Create a command file called `c:\sqlfull.cmd`. A sample command file (`sqlfull.smp`) is provided in the directory where Data Protection for SQL is installed. This sample file contains commands necessary to perform a scheduled full backup of all SQL Server databases to Tivoli Storage Manager storage. You must specify COMPLETE PATHNAMES in the command file for all file names and non-system commands.

Note: Perform the following if you are setting up the scheduler for a SQL Server running in a Single Copy Cluster (SCC) serverenvironment:

- Your command file must reside on the SQL Server File Share. The schedule you define on the Tivoli Storage Manager server needs to match this command file. In Step 1 and Step 2, the command file `c:\sqlfull.cmd` could be `x:\sqlfull.cmd` where `x` is the SQL Server File Share.
 - The *tsmoptfile* and *logfile* options specified in your command file must reflect the location of the options file and log file on the SQL Server File Share.
2. Enter the following command to define the schedule. You can enter this command on the server console or from an administrative client. The administrative client does not have to be running on the same system as the Tivoli Storage Manager server.

```
def sched mydomain my_schedule desc="SQL Daily Full Backup" action=command objects="c:\sqlfull.cmd" priority=2 starttime= 21:00 duration=15 duru=minutes period=1 perunits=day dayofweek=any
```

Tivoli Storage Manager displays this message:

```
ANR2500I Schedule MY_SCHEDULE defined in policy domain MYDOMAIN.
```

3. Issue the following command to associate Data Protection for SQL to this schedule:

```
define association mydomain my_schedule mynode
```

Tivoli Storage Manager displays this message:

```
ANR2510I Node MYNODE associated with schedule MY_SCHEDULE in policy domain MYDOMAIN.
```

A schedule is now defined on the Tivoli Storage Manager server with the following attributes:

- It runs a command file called `c:\sqlfull.cmd`.
- It begins at 9:00 pm.
- It is performed daily and can start on any day of the week.

You can use the Tivoli Storage Manager administrative commands **query schedule** and **query association** to confirm that the schedule and association are set correctly.

Complete these tasks on the SQL Server

These tasks are performed on the SQL Server.

This section of the procedure assumes the following environment:

- The Tivoli Storage Manager Backup-Archive client is installed on the SQL Server in the `d:\Program Files\Tivoli\TSM\baclient` directory.
- Data Protection for SQL is installed on the SQL Server in the `d:\Program Files\Tivoli\TSM\TDPSql` directory.
- The communication options in the `dsm.opt` option files located in these directories point to the Tivoli Storage Manager server to which the SQL databases are to be backed up.

The options file that is defined for Data Protection for SQL is used by the scheduler when validating the node and password. The options file is also used when contacting the Tivoli Storage Manager server for schedule information.

If this message is displayed:

```
A communications error occurred connecting to the Tivoli Storage Manager Server
```

- make sure the communication options in the `dsm.opt` file points to the correct Tivoli Storage Manager server.

- make sure the Tivoli Storage Manager server is running.

All cluster references in this procedure refer to a SQL Server Failover Cluster setup.

Perform the following steps on the SQL Server:

1. Login by using a Windows account that has administrative privileges.
2. Open a Windows command prompt window.
3. In the window, issue the following command:

```
cd d:\Program Files\Tivoli\TSM\baclient
```

If the path you want to use has a space in it, you can place quotation marks around the section of the path name that contains a space (d:\Program Files\Tivoli\TSM\baclient). You can also use the short form of a path name by placing a tilde (~) and unique identifier after the first six characters in the path.

An example of the short form of the path is as follows:

```
d:\Progra~1\Tivoli\TSM\baclient
```

Note: If a Tivoli Storage Manager scheduler is already installed on your machine (for the regular backups of the Windows system), you need to install another scheduler (with a unique name) to run the schedules defined for Data Protection for SQL. The Tivoli Storage Manager scheduler must have a different node name from the regular Tivoli Storage Manager Backup-Archive client.

4. In the window, issue the following command: The Windows ID that is used to start the scheduler service may require that you specify the client **dsmcutil** options *ntdomain*, *ntaccount*, and *ntpassword* in this command.

```
dsmcutil INSTAll scheduler /name:"Data Protection for SQL Scheduler"
/node:mynode /password:mypassword /autostart:yes
/clientdir:"d:\Program Files\Tivoli\TSM\baclient"
/optfile:"d:\Program Files\Tivoli\TSM\TDPSql\dsm.opt"
/startnow:no
```

Note: If you are setting up the scheduler for an SQL Server running in a cluster environment:

- a. Change the */autostart* option to *no*. For example:


```
/autostart:no
```
- b. Move the SQL virtual server to the secondary node of the cluster to create the scheduler service. Make sure that the secondary node of the cluster has ownership of the SQL virtual server.
- c. The primary node of the cluster must contain the command file on the file share used to create the scheduler service.
- d. Copy the options file (dsm.opt in the Step 4 example) to a shared drive associated with the virtual server. For example:

```
/optfile:"x:\dsm.opt"
```

Tivoli Storage Manager displays this output:

```
TSM Windows Client Service Configuration Utility
Command Line Interface - Version 6, Release 3, Level 0.0
(C) Copyright IBM Corporation, 1990, 2011, All Rights Reserved.
TSM Api Version 6.3.0
```

```
Command: Install TSM Client Service
Machine: TDPSQ1(Local Machine)
```

Installing TSM Client Service:

Machine : MYNODE

Service Name : Data Protection for SQL Scheduler

Client Directory : D:\Program Files\Tivoli\TSM\baclient

Automatic Start : yes

Logon Account : LocalSystem

The service was successfully installed.

Creating Registry Keys ...

```
Updated registry value 'ImagePath' .
Updated registry value 'EventMessageFile' .
Updated registry value 'TypesSupported' .
Updated registry value 'Data Protection for SQL Scheduler' .
Updated registry value 'ADSMClientKey' .
Updated registry value 'OptionsFile' .
Updated registry value 'EventLogging' .
Updated registry value 'ClientNodeName'.
```

Generating registry password ...

Authenticating password with Tivoli Storage Manager for node MYNODE

Connecting to Tivoli Storage Manager via client options file
'd:\Program Files\Tivoli\TSM\TDPSql\dsm.opt' ...

Password authentication successful.

The Registry password for node MYNODE has been updated .

Note: To make corrections after installing a service:

a. Issue the following command to remove the service:

```
dsmcutil remove /name:"Data Protection for SQL Scheduler"
```

b. Issue the command in Step 4 again to install a new service.

Note: If you are setting up the scheduler service for a SQL Server running in a cluster environment, repeat Step 1 through Step 4 on the secondary node of the cluster. **IMPORTANT!** The name of the scheduler service created by the **dsmcutil** command in Step 4 and the cluster service must have the same logon authority.

5. The Tivoli Storage Manager scheduler is now installed but has not started. To start the scheduler IN A NON-CLUSTER ENVIRONMENT, issue the following command in the Windows command prompt window:

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

This output is displayed:

```
The Data Protection for SQL Scheduler service is starting.
The Data Protection for SQL Scheduler service was started
successfully.
```

Note that because */autostart:yes* is specified, the Tivoli Storage Manager scheduler automatically starts each time the Windows system is rebooted.

Important:

- If you are creating the scheduler service in a *non-cluster* environment, proceed directly to Step 13.

- If you are creating the scheduler service in a *cluster* environment, perform Step 6 through Step 13.

The newly created scheduler service is tied to a cluster group. The Tivoli Storage Manager scheduler can correctly fail over between the nodes and also manage automatic password changes.

6. Start the Cluster Administrator.
7. Select the SQL Server Cluster Group and create a Resource to represent the Data Protection for SQL scheduler (**File→New→Resource**). Make sure that the following parameters are specified:

Resource Type

Specify *Generic Service*.

Group Specify the SQL Virtual Server.

Possible Owners

Make sure that both node machines are listed. Add them if they are not listed.

Resource Dependencies

Make sure the SQL Virtual server and the shared drives where the options file is located are listed.

Generic Service→Service Name

Specify the exact name of the scheduler service.

8. Select the new resource and modify its properties. Under the **Registry Replication→Root Registry Key** parameter, add the exact key where the Data Protection for SQL node name is listed. For example:

```
SOFTWARE\IBM\ADSM\CurrentVersion\BackupClient\NODES\\ServerName
```

Replace <nodename> with your Data Protection for SQL node name. For example:

```
SOFTWARE\IBM\ADSM\CurrentVersion\BackupClient\NODES\MYNODE\ServerName
```

9. In the Windows command prompt window, change to the Data Protection for SQL installation directory.

```
cd /d d:\Program Files\Tivoli\TSM\TDPSql
```

10. Verify that the *clusternode* option in the dsm.opt file is set to *yes*. After this setting is verified, enter a command that connects with the Tivoli Storage Manager server. Specify the */tsmpassword=yourpassword* parameter to ensure that the correct password is stored in the registry. The scheduler can properly connect automatically to the Tivoli Storage Manager server.

```
tdpsqlc query tsm /tsmpassword=myspassword
```

11. From the Cluster Administrator, select the new resource and bring it online (**File→ Bring Online**).

12. After the new resource is online, move the Group from the secondary node to the primary node of the cluster. Verify that the Data Protection for SQL scheduler service on the primary node has started.

13. Enter the following command to start the client scheduler on the machine where the Backup-Archive Client is installed:

```
dsmc sched
```

An example of the output is displayed below:

```
Tivoli Storage Manager
Command Line Backup/Archive Client Interface - Version 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1990, 2011 All Rights Reserved.
```

```

Querying server for next scheduled event.
Node Name: MYNODE
Session established with server TYPP_CLINT: Win 2003

Server Version 6, Release 3, Level 0.0

Server date/time: 09/12/2011 17:55:37 Last access: 09/12/2011 17:43:51

Next operation scheduled:
-----
Schedule Name: DATA PROTECTION FOR SQL SCHEDULER
Action: Command
Objects: C:\sqlfull.cmd
Options:
Server Window Start: 13:40:30 on 09/13/2011
-----
Schedule will be refreshed in 1 hour.

Time remaining until execution: 00:59:58

```

Your system is now ready to run automatic daily full backups of the SQL databases.

Scheduler guidelines

The guidelines assist when planning scheduled operations.

Be aware of the following guidelines when defining a Tivoli Storage Manager schedule:

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

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

- If any changes that affect the scheduler are made to the Data Protection for SQL options file, the scheduler has to be restarted in order to pick up the changes. An example of this is the Tivoli Storage Manager server address, the schedule mode, or the client TCP address or port. This can be done by issuing the following commands:

```

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

```

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

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

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

You can override this file name by specifying the *shedlogname* option in the Data Protection for SQL options file.

- Data Protection for SQL creates its own log file with statistics about the backed up database objects when the *logfile* parameter is specified during the **tdpsqlc** command. In the sample file (sqlfull.smp), the log file is sqlsch.log. This file is different from the Tivoli Storage Manager scheduler log file and must also be different from the file to which the **tdpsqlc** command output is redirected. In the example above, this file is sqlfull.log.

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

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

```
Scheduled event eventname completed successfully
```

This is merely an indication that Tivoli Storage Manager successfully issued the scheduled command associated with the *eventname*. No attempt is made to determine the success or failure of the command. You should assess the success or failure of the command by evaluating the return code from the scheduled command in the scheduler log. The scheduler log entry for the command's return code is prefaced with the following text:

```
Finished command. Return code is:
```

If any scheduled backups fail, the scheduler script will exit with the same error code as the failed backup command. A non-zero error code means that backup failed.

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

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

Setting user mode

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

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

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

ALTER DATABASE DBNAME SET SINGLE_USER WITH ROLLBACK AFTER N SECONDS

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

ALTER DATABASE DBNAME SET MULTI_USER

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

Note:

1. You can set a SQL server to single-user mode by using the `-m SQL SERVER` startup option when restarting the SQL server.
2. You can use the SQL stored procedure `SP_WHO` to determine which users are using the databases.

Automating tasks

This section explains how to use the Automate view to work with commands. It shows how to save commands and schedule the running of commands.

You can use the Automate view to create, save, store, and schedule commands. Open the Automate view by selecting a workload that you want to work with and clicking the **Automate** tab. When you enter commands in the bottom details pane, the output of the command is displayed in the results pane at the top. You can click the **Tips** button to display tips on how to use the Automate view.

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

The commands can be entered without `tdpsql c`. For example, for each selected workload instance, you can enter a single command or multiple commands, such as:

```
q tsm
q sql
```

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

Additional automation methods

This section explains how to automate your commands from other views in the Management Console.

You can automate your commands from the Protect, Recover, Schedule, and Task List views.

1. Start the Management Console and select the SQL Server in the tree view.
2. Click the appropriate tab for the task (**Protect** or **Recover**).
3. Automate the command by using one of the following methods:
 - Result Pane
Select the databases for your task in the result pane, then select **Run Scheduled** in the toolbar drop-down menu. Click the appropriate task in the **Action** pane. When the schedule wizard starts, enter the information for each prompt in order to create a scheduled task.
 - Task List Pane

When a task has been 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 in order to create a scheduled task.

Viewing, printing, and saving reports

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

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

1. Select **Reporting** in the tree view, under **Manage**. A list of available reports appears. Each report has a description of what data the report contains.
2. Select a report from the list. The selected report appears.
3. To print or save the current report, click the appropriate icon at the top of the report.

Chapter 6. Troubleshooting Data Protection for SQL with VSS backup-restore support

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

Problem determination assistance

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

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

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

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

```
Scheduled event eventname completed successfully
```

This is merely an indication that Tivoli Storage Manager successfully issued the scheduled command associated with the *eventname*. No attempt is made to determine the success or failure of the command. You can assess the success or failure of the command by evaluating the return code from the scheduled command in the scheduler log. The scheduler log entry for the command's return code is prefaced with the following text:

```
Finished command. Return code is:
```

- Windows Event Log.
- For VSS operations, view the `dsmerror.log` file in the backup-archive client installation directory.

Resolving errors during Data Protection for SQL processing

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

Perform the following steps as a first attempt to resolve the problem:

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

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

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

For Legacy operations:

- Use the Backup or Restore utility provided in the SQL Server administrator program to see if the problem can be recreated.
- If the error message "AC05350E An unknown SQL API error has occurred" is displayed, the SQL server encountered an unexpected situation. Microsoft assistance may be needed if the problem continues.
- Data Protection for SQL error messages occasionally contain an HRESULT code. Use this code to search Microsoft documentation and the Microsoft Knowledge Base for resolution information.

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

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

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

You can also search for helpful diagnostic information in the Tivoli Storage Manager information center: <http://publib.boulder.ibm.com/infocenter/tsminfo/v6r3>

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

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

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

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

- SQL Server status
- Databases
- VSS components

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

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

IBM Tivoli Storage Manager for Databases:
Data Protection for Microsoft SQL Server
Version 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1997, 2011. All rights reserved.

AC05057I The C:\Program Files\Tivoli\tsm\TDPSql\tdpsql.log log file
was pruned successfully.

Connecting to SQL Server, please wait...

SQL Server Information
-----

SQL Server Name ..... JAZZVM2EXCH2007\JAZZVM2SQL2K8R2
SQL Server Version ..... 10.50.1617 (SQL Server 2008 R2)

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

Writer Name           : SqlServerWriter
Local DSMAGENT Node   : jazzvm2
Remote DSMAGENT Node  :
Writer Status         : Online
Selectable Components : 7
```

If the **TDPSQLC QUERY SQL** command does not return all of this information, you might have a proxy configuration problem. Contact the Tivoli Storage Manager server administrator to have the correct server **GRANT PROXY** commands issued to enable proxy authority for nodes. If all of the information returned to you seems correct, proceed to the next step.

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

vssadmin

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

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

Restriction: The **VSSADMIN LIST SHADOWS** command does not list shadows of SAN-attached volumes for Windows 2008 and later. The **vssadmin** utility uses Microsoft Software Shadow Copy provider to list the shadows that are created. For more information on the **vssadmin** utility, see <http://technet.microsoft.com/en-us/library/cc754968%28WS.10%29.aspx>.

diskshadow

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

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

```
diskshadow>begin backup
diskshadow>add volume f: (database volume)
diskshadow>add volume g: (log volume)
diskshadow>create
diskshadow>end backup
diskshadow>list shadows all
diskshadow>delete shadows all
```

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

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

```
diskshadow>set context persistent
diskshadow>begin backup
diskshadow>add volume f: (database volume)
diskshadow>add volume g: (log volume)
diskshadow>create
diskshadow>end backup
diskshadow>list shadows all (This may take a few minutes)
diskshadow>delete shadows all
```

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

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

vshadow

A utility included with the Microsoft Volume Shadow Copy Services software developers kit (SDK) that can be used to exercise most of the VSS infrastructure, such as creating/querying/deleting shadow copies. You can also use vshadow to create both persistent and nonpersistent shadow copies, transportable snapshots, as well as assign a drive letter or mount point to a shadow copy. See the <http://msdn.microsoft.com/en-us/library/> site.

Note: You can use the **vshadow** utility on Windows Server 2003 (and 2003 R2). If you are using Windows 2008, you can use the **diskshadow** tool.

- The following items can be determined by using the vssadmin or vshadow utility:
 - Verify VSS Provider configurations and setup
 - Rule out any possible VSS problems before running the Tivoli Storage Manager VSS functions
 - That you might have a VSS configuration problem or a real hardware problem if an operation does not work with vshadow/vssadmin
 - That you might have a Tivoli Storage Manager problem if an operation works with vshadow/vssadmin but not with the Tivoli Storage Manager
- Perform the following tests to ensure that VSS is working correctly:

Test non-persistent shadow copy creation and deletion

- a. Run “VSHADOW *k*: *l*:” where *k*: and *l*: are the SQL Server database and log volumes.
- b. Repeat the step above four times.
- c. Inspect the Windows Event Log to ensure that things look appropriate.

Test persistent shadow copy creation and deletion

- a. Run “VSHADOW -p *k*: *l*:” (where *k*: and *l*: are the SQL Server database and log volumes. You might have to run “VSHADOW -da” to remove this if you do not have enough space.
- b. Repeat the previous step four times.
- c. Inspect the Windows Event Log to ensure that things look appropriate.

Test non-persistent transportable shadow copy creation and deletion (VSS Hardware Provider environments only)

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

If any of these tests fail repeatedly, there is hardware configuration problem or a real VSS Problem. Consult your hardware documentation for known problems or search Microsoft Knowledge Database for any information.

If all tests pass, continue to Step 3.

3. Recreate your specific problem by using vshadow. If you can only recreate your problem through a series of steps (for example: a backup fails only when you perform two consecutive local backups), try to perform those same tests by using vshadow.
 - SQL VSS backups to Local are simulated by running a vshadow persistent snapshot.
 - SQL VSS backups to the Tivoli Storage Manager are simulated by running a vshadow nonpersistent snapshot.
 - SQL VSS backups to Local and to the Tivoli Storage Manager are simulated by running a vshadow persistent snapshot.
 - Offloaded SQL VSS backups to the Tivoli Storage Manager are simulated by running a vshadow nonpersistent, transportable snapshot.

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

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

Tracing the Data Protection client when using SQL VSS technology

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

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

Data Protection for SQL trace

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

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

Tivoli Storage Manager API trace

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

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

DSMAGENT trace

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

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

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

Gathering SQL with VSS information before calling IBM

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

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

See "Emailing support files" on page 110 for more information on collecting diagnostic information to send to IBM.

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

- The exact level of the Windows operating system, including all service packs and hotfixes that were applied.
- The exact level of the SQL Server, including all service packs and hotfixes that were applied.
- The exact level of Data Protection for SQL with Virtual Shadow Copy Service (VSS) backup-restore support.
- The exact level of the Tivoli Storage Manager API.
- The exact level of the Tivoli Storage Manager server.
- The exact level of the Tivoli Storage Manager backup-archive client.
- The exact level of the Tivoli Storage Manager storage agent (if LAN-free environment).
- The Tivoli Storage Manager server platform and operating system level.
- The output from the Tivoli Storage Manager server **QUERY SYSTEM** command.
- The output from the Data Protection for SQL **TDPSQLC QUERY SQL** command.
- The device type (and connectivity path) of the SQL databases and logs.
- (SAN only) The specific hardware that is being used. For example: HBA, driver levels, microcode levels, SAN Volume Controller or Storwize V7000 levels, DS8000 hardware details.
- Permissions and the name of the user ID being used to run backup and restore operations.
- The name and version of antivirus software.
- (SAN only) The VSS hardware provider level.
- The VSS hardware provider log files. See the documentation of the specific VSS hardware provider on how to enable tracing and collect the trace log files.
- (SAN only) The IBM CIM agent level for DS8000, SAN Volume Controller, or Storwize V7000.
- A list of other applications running on the system.
- A list of the steps needed to recreate the problem (if the problem can be recreated).
- If the problem can not be recreated, list the steps that caused the problem.
- Is Data Protection for SQL running in a Microsoft Cluster Server (MSCS) environment?
- Is it an Exchange Server 2010 DAG or Exchange Server 2007 CCR, LCR, or SCR environment?
- Is the problem occurring on other SQL servers?

Gathering files from SQL with VSS before calling IBM

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

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

See “Emailing support files” on page 110 for more information on collecting diagnostic information to send to IBM.

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

- The contents of the C:\adsm.sys\vss_staging directory and subdirectories. Or gather the appropriate directories if you are using the VSSALTSTAGINGDIR option.
- The Data Protection for SQL configuration file. The default configuration file is tdpsql.cfg.
- The Data Protection for SQL Tivoli Storage Manager application programming interface (API) options file. The default options file is dsm.opt.
- The Tivoli Storage Manager registry hive export.
- The SQL Server registry hive export.
- The Tivoli Storage Manager Server activity log. The Data Protection client logs information to the server activity log. A Tivoli Storage Manager administrator can view this log for you if you do not have a Tivoli Storage Manager administrator user ID and password.
- If the Data Protection client is configured for LAN-free data movement, also collect the options file for the Tivoli Storage Manager storage agent. The default name for this file is dsmsta.opt.
- Any screen capture or command-line output of failures or problems.

Log files can indicate the date and time of a backup, the data that is backed up, and any error messages or completion codes that could help to determine your problem. The following files are the Tivoli Storage Manager log files to gather:

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

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

- System Provider (Windows Event Log)
- IBM System Storage SAN Volume Controller, Storwize V7000, and DS8000 - Program Files\IBM\Hardware Provider for VSS\IBMVss.log.
- NetApp - Program Files\SnapDrive*.log
- XIV - zip up all of the files in the C:\Windows\Temp\xProvDotNet directory

You can use the Data Protection for SQL console to list the events originated by Data Protection for SQL. Select **Dashboard > ServerName > Diagnostics > System Information** and double-click the dpevents.ps1 script in the PowerShell section of the System Information page.

Troubleshooting Data Protection for SQL VSS and SAN Volume Controller, Storwize V7000, or DS6000/DS8000

The troubleshooting tips included here are designed to help you accelerate your problem determination task.

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

- **CIMOM (Common Information Model Object Manager) Connectivity issues**
To verify connectivity to the CIMOM, perform the following steps:

1. Refer to your SAN Volume Controller, Storwize V7000, or DS8000 documentation.
2. Run the **IBMVCFG LIST** command. The default location is D:\Program Files\IBM\Hardware Provider for VSS-VDS.
3. Issue the **IBMVCFG SHOWCFG** command to view the provider configuration information.
4. Check that the CIMOM is properly configured. Run `verifyconfig.bat -u username -p password` on the Master Console.
5. Check the username and password. If there is a problem with the truststore, follow the procedure in the documentation to generate a new truststore.

- **CIMOM operational issues**

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

```
Wed Jan 11 17:34:34.793 - Calling AttachReplicas
Wed Jan 11 17:34:35.702 - AttachReplicas: 909ms
Wed Jan 11 17:34:35.702 - returnValue: 34561
Wed Jan 11 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 11 17:34:35.779 - IBMVSS: AbortSnapshots
```

A return value of 0 means that it was successful. To determine why it failed, look at the log files generated by the command line interface (CLI) or graphical user interface (GUI), depending on how you run your operation. These might provide more information on the failure.

- **Host configuration issues**

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

- **Collecting logs in this environment**

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

- Information listed in the Tivoli Storage Manager diagnostic information section
- HBA type, firmware and driver levels
- SDD version
- SAN Volume Controller microcode version (if applicable)
- DS8000 microcode version (if applicable)
- Storwize V7000 microcode version (if applicable)
- SAN Volume Controller Master Console version (if applicable)
- For DS8000, the CIM Agent version (if applicable)
- `IBMVSS.log`
- `IBMVDS.log`
- Application Event Log
- System Event Log

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

Installation Problems: Creating an installation-log file

In the event a silent installation fails, gather the following information to assist Customer Support when evaluating your situation:

- Operating system level
- Service pack
- Hardware description
- Installation package (DVD or electronic download) and level
- Any Windows event log relevant to the failed installation
- Windows services active during the failed installation (for example, anti-virus software)
- Whether you are logged on to the local console (not through a terminal server)
- Whether you are logged on as a local administrator, not a domain administrator (Tivoli does not support cross-domain installations)

You can create a detailed log file (setup.log) of the failed installation. Run the setup program (setup.exe) in the following manner:

```
setup /v"1*v setup.log"
```

Emailing support files

Send diagnostic information to IBM support personnel.

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.

Follow these steps to send diagnostic information to IBM support personnel:

1. Click **Start > Tivoli Storage FlashCopy Manager > Management Console**.
2. Click **Diagnostics** in the results pane of the welcome page. Click the **E-Mail Support files** icon in the action pane.
3. Enter the required information in the various fields and click **Done**. The information is sent to the designated support personnel and the dialog closes.

Return to the Tivoli Storage FlashCopy Manager Management Console and begin backup operations.

Online IBM support

Integrated web content is provided.

Search for the most current information regarding Tivoli Storage FlashCopy Manager product support at this website: <http://www.ibm.com/software/tivoli/products/storage-flashcopy-mgr/>

Enter the search term, such as an authorized program analysis report (APAR) number, release level, or operating system to narrow the search criteria for your support need.

Viewing system information

View scripts that provide information on system components such as Data Protection for SQL-related Windows Services, Windows Event Log entries, and Volume Shadow Copy Service (VSS) information.

The System Information view is extensible. You can take advantage of this flexibility to add and share customize scripts.

To work with scripts, follow these steps:

1. Open the System Information view by doing the following steps:
 - a. Click **Diagnostics** in the start page.
 - b. Double-click **System Information** in the results pane. A list of scripts is displayed in the results pane of the System Information view. The types of scripts that are displayed are PowerShell scripts, Windows Management Instrumentation scripts, and Tivoli Storage Manager scripts.
2. Add, update, or delete your scripts.
 - To add your own scripts, click **New** in the Actions pane. You can also copy your scripts directly to the scripts folder in the installation directory. The file type extension is used to determine how to run the script. As a result, make sure that your scripts follow these extension requirements:
 - PowerShell scripts: *filename.ps1*
 - Windows Management Instrumentation (WMI) scripts: *filename.wmi*
 - Tivoli Storage Manager scripts: *filename.tsm*
 - To view or edit an existing script:
 - a. From the list of script files in the results pane, select the name of a script that you want to view or edit.

Tip: The name of the script is displayed in the Actions pane. Click the name of the script in the Actions pane to reveal or hide a list of actions to perform.
 - b. Click **Command Editor** in the Actions pane to open the script file for viewing or editing.
 - c. View or edit the script. Click **OK** to save your changes, or click **Cancel** to exit the System Information Command Editor without saving any changes.
 - To delete a script:
 - a. From the list of script files in the results pane, select the name of a script that you want to delete.

Tip: The name of the script is displayed in the Actions pane. Click the name of the script in the Actions pane to reveal or hide a list of actions to perform.

- b. Click **Delete** in the Actions pane.

Viewing trace and log files

View files used during troubleshooting tasks.

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

- Tivoli Storage FlashCopy Manager
 - Installation directory: `c:\program files\tivoli\flashcopymanager`
 - `dsierror.log`
 - `fcm.log`
 - `TraceFm.trc`
 - `TraceUx.trc`
 - `TraceManagedCapacityHistory.trc`
 - `TraceSchedLaunch.trc`
 - `VssProvisioning.log`

Note: If the `fcm.log` is defined in a path other than the default `c:\program files\tivoli\flashcopymanager\fcm.log`, the reports will 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 `fcm.log`.

- Data Protection for SQL
 - Installation directory: `C:\Program Files\Tivoli\TSM\TDPSql`
 - `dsierror.log`
 - `tdpsql.log`
 - `TraceFileSql.trc`

Note: If the `tdpsql.log` is defined in a path other than the default `c:\program files\tivoli\TSM\TDPSql\tdpsql.log` the reports will not include task completion, type of data protection activity, and amount of data protection activity for scheduled backup and restore operations. The charts and reports display only information that is present in the default log file `tdpsql.log`.

- VSS Requestor
 - Installation directory: `C:\Program Files\Tivoli\TSM\baclient`
 - `dsmerror.log`
- IBM VSS
 - `IBMVDS.log`
 - `IBMVss.log`

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

You can collect trace and log files in the Diagnostics property page for a workload. For more information about collecting trace and log files, see [Diagnostics](#).

Chapter 7. Performance

Data Protection for SQL provides certain parameters that can be tuned for optimum performance.

Many factors can affect the backup and restore performance of Data Protection for SQL, such as hardware configuration, network type, and capacity. These factors are not within the scope of this document. However, some parameters that are related to Data Protection for SQL can be tuned for optimum performance.

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

Buffering (Legacy only)

Data Protection for SQL is a multi-threaded application that uses asynchronous execution threads to transfer data between the SQL and Tivoli Storage Manager servers. To accomplish this, multiple data buffers are used to allow one thread to receive data from one side, while another thread sends data to the other side. For example, one thread can be reading data from a SQL Server while another is sending data to the Tivoli Storage Manager server. The number of buffers that Data Protection for SQL allocates to these threads can be specified in the */buffers* and */sqlbuffers* parameters of the command line interface. The size of these buffers can be specified in the */buffersize* and */sqlbuffersize* parameters.

Data Striping (Legacy only)

In addition to multi-threading to maximize throughput on a single session, Data Protection for SQL uses separate threads to support SQL data striping, which allows use of multiple parallel sessions to backup and restore a single database. This is another method to maximize data throughput. If a single session cannot fully exploit available bandwidth, multiple parallel sessions can yield improved data throughput, especially if the database is spread across multiple physical volumes.

If you use one data stripe per physical volume for both the SQL Server and the Tivoli Storage Manager server, the performance (measured as the amount of time necessary to backup or restore a particular SQL database) should show an improvement over the unstriped case (approximately proportional to the number of data stripes used, given the constraints of the devices and the network used, and striping independent overhead in SQL Server, Tivoli Storage Manager server, and Data Protection for SQL).

You can specify the number of stripes to use with the */STRIPes* parameter on the command-line interface. You can also specify the number of stripes to use from the MMC GUI, by changing the number in the **Stripes** field in the Backup options or Restore options panel.

Note:

- Additional striping does not necessarily improve performance and may even decrease performance if system constraints involving real and paged memory, processors, network interface cards, networks, device reads and writes, and RAID become saturated or exceed capacity.
- If you use striping in conjunction with SQL buffers, be certain that the number of SQL buffers specified is equal to or greater than the number of stripes.
- The default values that Data Protection for SQL assigns to buffers, buffer size, and stripes can be changed in the Data Protection for SQL configuration file. Use the **set** command or the Performance property page in the MMC GUI to modify the configuration file.

LAN-free environment (Legacy and VSS)

Running Data Protection for SQL in a LAN-free environment if you are equipped to do so avoids network constraints.

- For Legacy backups, specify *enablelanfree yes* in the Data Protection for SQL options file.
- For VSS Backups, specify *enablelanfree yes* in the DSMAGENT (VSS Requestor) dsm.opt file only.

For information on setting up a LAN-free environment, refer to the Tivoli publication: *IBM Tivoli Storage Manager for SAN for Windows Storage Agent User's Guide*.

Chapter 8. Data Protection for Microsoft SQL Server reference

Reference information for Data Protection for Microsoft SQL Server is provided.

Command overview

The name of the Data Protection for SQL command line interface is **tdpsqlc.exe**. This executable is located in the directory where Data Protection for SQL is installed.

Using the Data Protection for SQL command line interface from the GUI

Follow these steps to launch the Data Protection for SQL command-line interface:

1. Start the MMC GUI.
2. Expand the **Protect and Recover Data** node.
3. In the tree view, select an SQL Server node.
4. Select the **Automate** tab on the center display, an integrated command line is available in the bottom of the task window for inputting commands. The top section of the window displays the command output.

Command-line parameter characteristics

The command-line parameters have the following characteristics:

- positional parameters do not include a leading slash (/) or dash (-)
- optional parameters can appear in any order after the required parameters
- optional parameters begin with a forward slash (/) or a dash (-)
- minimum abbreviations for keywords are indicated in upper case text
- some keyword parameters require a value
- for those keyword parameters that require a value, the value is separated from the keyword with an equal sign (=)
- if a parameter requires more than one value after the equal sign, the values are separated with commas
- each parameter is separated from the others by using spaces
- if a parameter's value includes spaces, the value must be enclosed in double quotation marks
- a positional parameter can appear only once per command invocation

Data Protection for SQL uses the following command line syntax:

```
tdpsqlc <command> <positional parameter> <0 or more optional parameters>
```

The **tdpsqlc** executable is followed by high level operations called *commands*. Each command accepts various command line parameters. These parameters consist of *positional parameters* and *optional parameters*. Positional parameters must precede other options in the command line. In the following case, the backup command with its database name *xyz*, the object to back up, is followed by the type of backup, **full**, a positional parameter, and finally by an optional parameter, **/sqlbuffers**.

```
tdpsqlc backup xyz full /sqlbuffers=2
```

For output samples of the Data Protection for SQL commands illustrating various positional and optional parameters, see the "Output Examples" section following each command's syntax.

Note: You can display a complete list of Data Protection for SQL commands and all their parameters by simply entering `tdpsqlc` or the `tdpsqlc help |?` command. See "Help command" on page 218.

Command-line interface help

Issue the `tdpsqlc ?`, `tdpsqlc help`, or `tdpsqlc` command to display help for the command-line interface.

Available commands

Table 19. Data Protection for SQL commands

Command	Description	Page
Backup	Backs up all or part of one or more SQL databases to Tivoli Storage Manager server.	"Backup command" on page 121
Query	Displays information about servers, databases, backup objects, and Data Protection for SQL configuration.	"Query command" on page 146
Restore	Restores all or part of one or more SQL databases to SQL server.	"Restore command" on page 168
RESTOREFiles	Restores flat files (.mdf, .ldf) from a VSS backup to a specified directory on a fixed or network mapped drive	" Restorefiles command" on page 199
INACTIVate	Inactivates one or more active backup objects on the Tivoli Storage Manager server.	"Inactivate command (Legacy only)" on page 208
Help	Displays the syntax of Data Protection for SQL commands.	"Help command" on page 218
Set	Changes the values of configuration parameters.	"Set command" on page 222
CHANGETSMPassword	Changes the Tivoli Storage Manager password used by Data Protection for SQL.	"Changetsmpassword command" on page 230

Command line parameter characteristics

As seen in the preceding command line syntax, there are several characteristics to take note of in the Data Protection for SQL command line:

- Do not include a slash or dash before positional parameters.
- Begin optional parameters with a forward slash (/) or a dash (-).
- You may place multiple optional parameters per command invocation in any order *after* positional parameters.
- You may abbreviate keywords. Minimum abbreviations are indicated in upper case in the syntax diagrams.
- All SQL names of databases or parts of databases are case-sensitive.
- Separate parameters with at least one space.

- Some keyword parameters may require a value; separate values from their keywords with an equal sign. (=).
- If a parameter's value includes spaces or special characters, enclose the value in double quotes.
- You can use most positional and optional parameters only once per command invocation. The following exceptions allow lists of values or repetition of the parameter:
 - File=
 - Group=
 - Log=
 - Set=
 - /Files=
 - /GRoups=
 - /RELocate=
 - /RELOCATEDir=
 - /TO=

For example: /files=a,b or /files=a /files=b

Note: Multiple instances of optional parameters do not have to be contiguous.
For example: /files=a /groups=y /files=b /groups=z

Where repeatable syntax appears, separate multiple values with commas as indicated in the following:



Use the wildcard asterisk (*) following the command to select all instances on the server of database names or file names.

For help in reading syntax diagrams, refer to “Reading syntax diagrams” on page xviii.

Data Protection for SQL parameters available by backup method

Table 20. Data Protection for SQL optional parameters

Optional Parameters	Legacy	VSS
/ACTIVE	Yes	Yes
/ADJUSTKBtsmestimate	Yes	No
/ADJUSTPERcentsmestimate	Yes	No
/ALI	Yes	Yes
/BACKUPDESTination	Yes	Yes
/BACKUPMETHod	Yes	Yes
/BUFFers	Yes	No
/BUFFERSIze	Yes	No
/COMPATibilityinfo	Yes	Yes
/CONFIGfile	Yes	Yes

Table 20. Data Protection for SQL optional parameters (continued)

Optional Parameters	Legacy	VSS
/DBOonly	Yes	No
/DIFFESTimate	Yes	No
/EXCLUDEDB	Yes	Yes
/FILEInfo	Yes	No
/Files	Yes	No
/GRoups	Yes	No
/INSTANTRestore	No	Yes
/INTO	Yes	Yes
/LOGESTimate	Yes	No
/LOGFile	Yes	Yes
/LOGPrune	Yes	Yes
/MOUNTWait	Yes	No
/OBJECT	Yes	Yes
/OFFLOAD	No	Yes
/OLDERthan	Yes	No
/PARTial	Yes	No
/Quiet	Yes	Yes
/RECOVery	Yes	Yes
/RELOCATEDir	Yes	Yes
/RELocate /TO	Yes	No
/REPlace	Yes	No
/SQLAUTHentication	Yes	Yes
/SQLBUFFers	Yes	No
/SQLBUFFERSIze	Yes	No
/SQLPassword	Yes	Yes
/SQLSERVer	Yes	Yes
/SQLUSer	Yes	Yes
/STANDBy	Yes	No
/STOPAT	Yes	No
/STOPATMark /AFTER	Yes	No
/STOPBEFOREMark /AFTER	Yes	No
/STRIPes	Yes	No
/TRUNCate	Yes	No
/TSMNODE	Yes	Yes
/TSMOPTFile	Yes	Yes
/TSMPassword	Yes	Yes

Data Protection for SQL allows you to perform online backups and restores of Microsoft SQL Server databases to Tivoli Storage Manager server storage using either command-line or graphical user interfaces (GUI).

Backup command

Use the **backup** command to back up all or part of one or more SQL databases from the SQL Server to Tivoli Storage Manager storage on the Tivoli Storage Manager server.

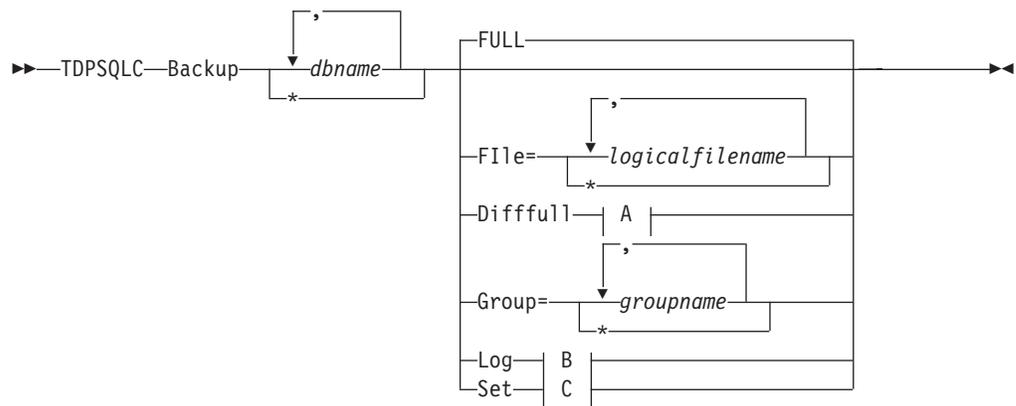
You can enter the * character to backup all databases. You can specify more than one database at once for multiple database and transaction log backups.

Considerations:

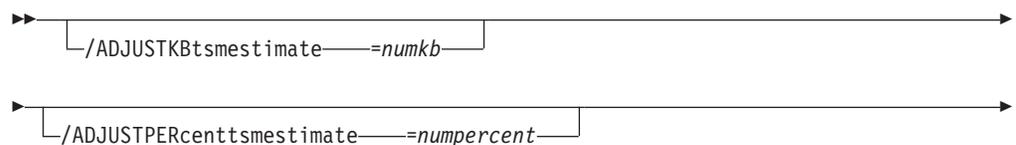
- Simple recovery model databases are automatically excluded from log backups.
- The master database is automatically excluded from log and differential backups.
- You cannot back up or restore the **tempdb** database because it is created by SQL server each time the server is started.
- Although full and differential backups include a part of the transaction log, that part is only what is required to make a restore consistent. It is not a log backup and does not truncate the log.
- The user id used by Data Protection for SQL to log on to the SQL server must have the SQL Server SYSADMIN fixed server role.
- You can use the TRANSACT-SQL database consistency checker statement DBCC CHECKDB ('DBNAME') to verify the integrity of the SQL databases before you back them up.

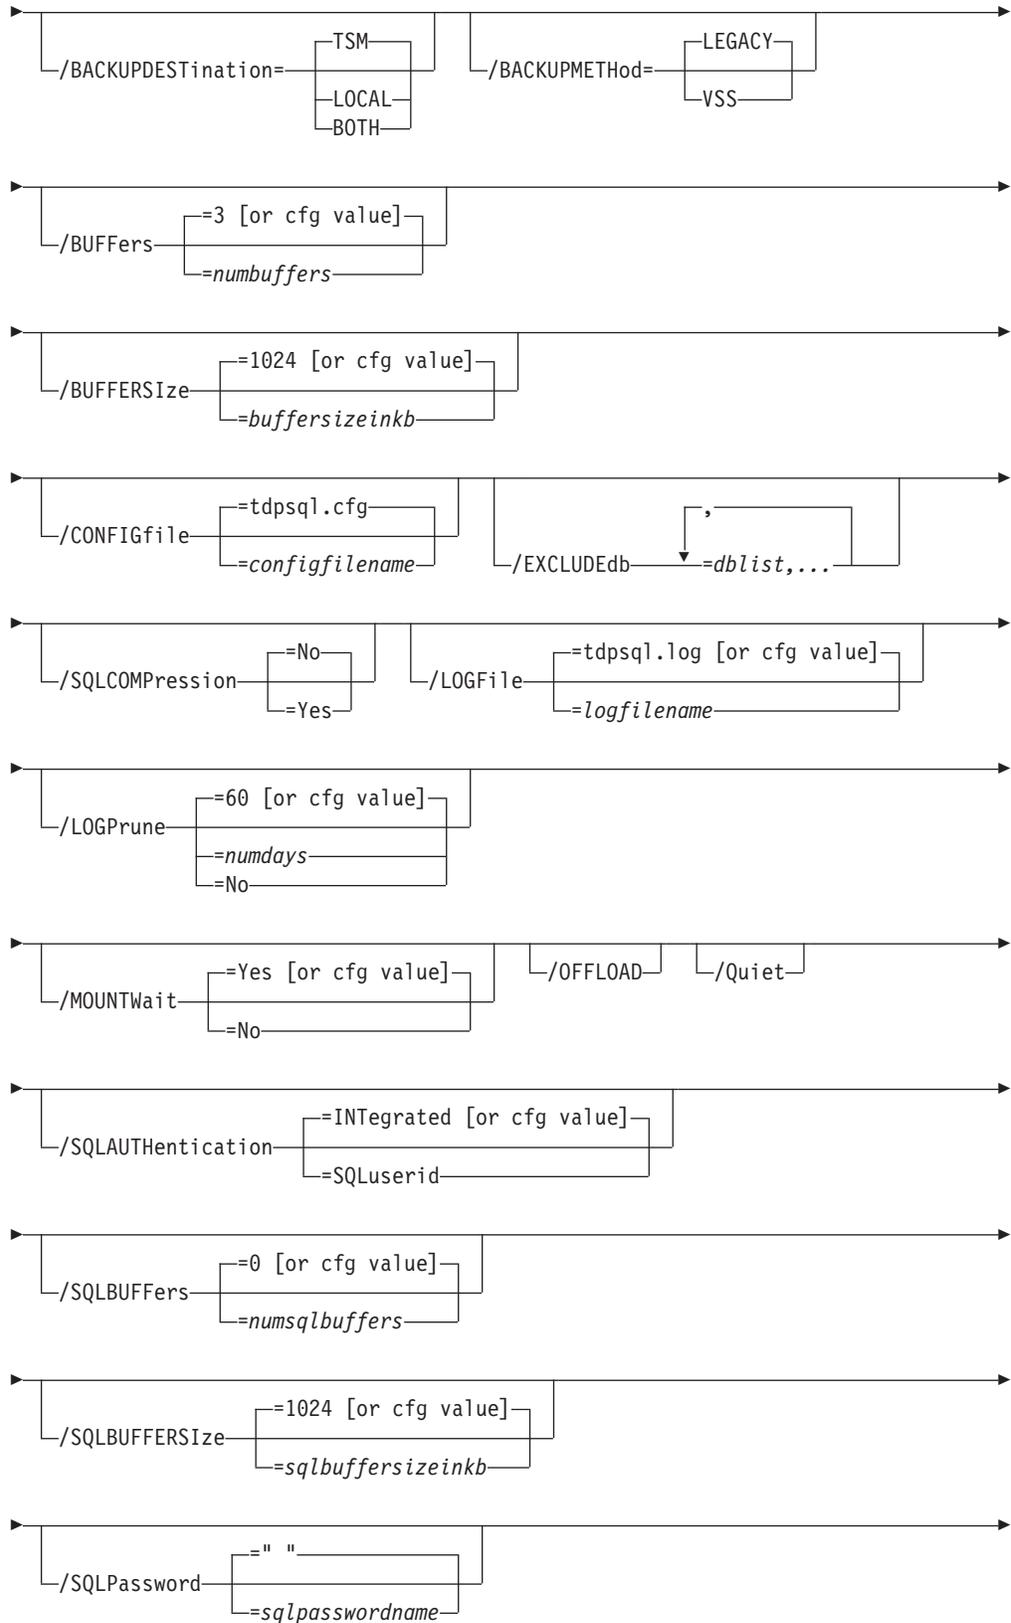
Backup syntax

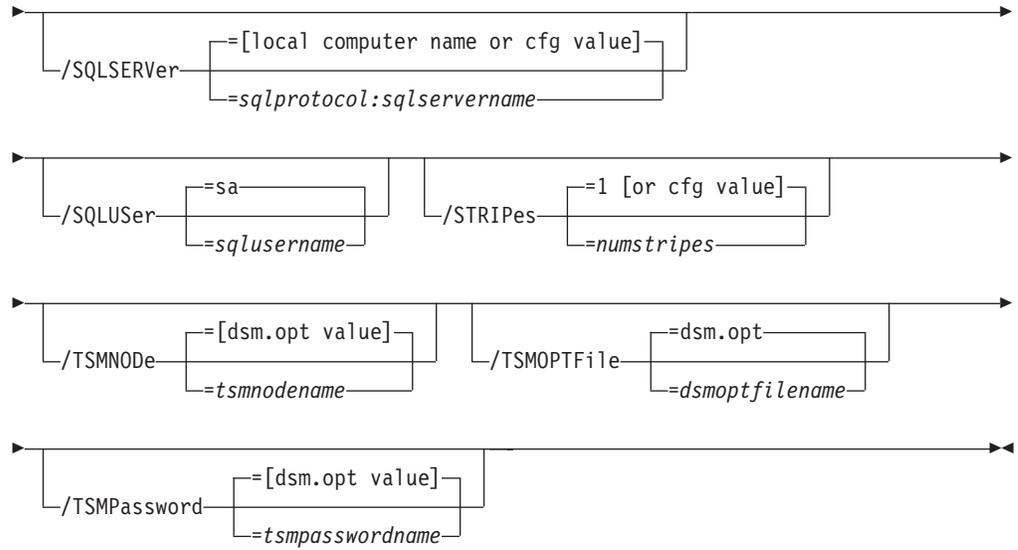
Use the **backup** command syntax diagrams as a reference to view available options and truncation requirements.



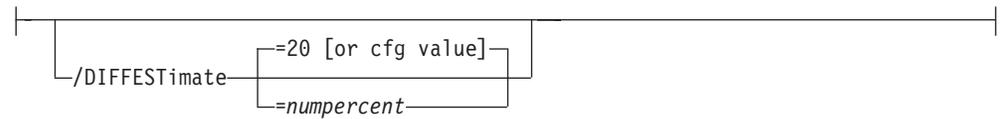
Backup Optional Parameters:



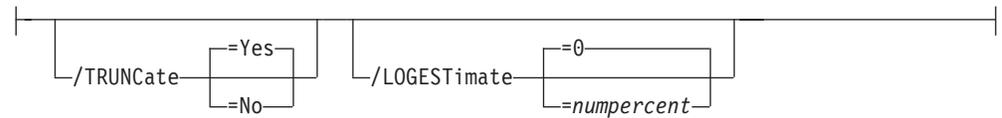




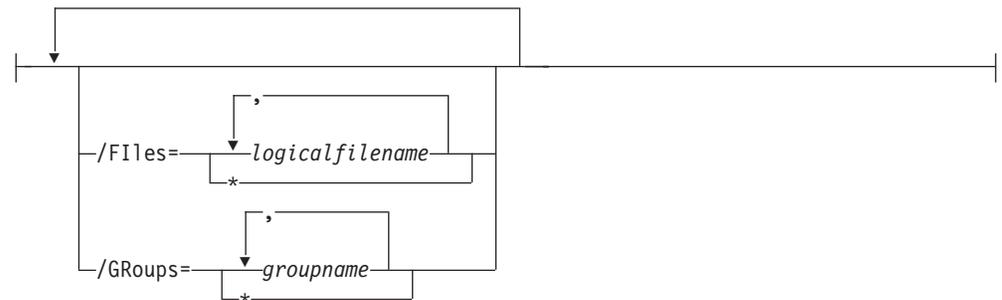
A Diffull Options:



B Log Options:



C Set Options:



Backup positional parameters

Positional parameters immediately follow the **backup** command and precede the optional parameters.

File=*|*logicalfilename*,...

A **file** backup contains only the contents of the SQL server logical file you specify. You can use this option when it is not practical to back up an entire SQL database due to available backup time and space or due to performance requirements. The *logicalfilename* variable specifies the names of the SQL server database logical files you want to back up or restore to.

Considerations:

- You should follow file backups with transaction log backups for all SQL databases you back up.
- You can specify this parameter more than once per command invocation.
- A new backup object inactivates any active backup object of the same name in the same SQL database.
- Use * as a wildcard character in *logicalfilename* to replace zero or more characters for each occurrence. Specifying only the wildcard character indicates all logical files in the SQL server database.

Note: 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.
- A SQL server Create Index operation requires that you back up all affected file groups as a unit. You cannot back up a file in the affected file groups until you perform the unit backup. A SQL server error message will indicate which file groups are affected. You must perform a full database backup or a set backup of at least the affected file groups before the file backup succeeds.

FULL A **full** Legacy database backup contains all of the contents of a SQL server database plus enough of the database's transaction log to make a restore consistent. A **full** VSS database backup contains all of the contents of a SQL server database (database files, log files, and full-text index files).

Each SQL database backed up using the Legacy backup method creates a separate backup object on the Tivoli Storage Manager server. A new full database backup object inactivates all prior Legacy active backup objects for the same SQL database. This inactivation includes any active full backup object as well as any active file, group, set, differential, and log backup objects. For additional policy information, including VSS aspects, see "How Tivoli Storage Manager server policy affects Data Protection for SQL" on page 19, "Management Class" on page 32, and "Setting automatic expiration (VSS and Legacy)" on page 41.

Difffull

A **difffull** (differential) database backup contains only the parts of a SQL server database changed since the latest full backup plus enough of the SQL database's transaction log to make a restore consistent. As such, a

differential backup usually takes up less space than a full backup. Use this option so that all individual log backups since the last full database backup do not need to be applied.

Group=*|*groupname*,...

A **group** backup contains only the contents of the SQL server file group you specify. A group backup is useful when selected SQL database table or indexes have been assigned to a file group and only those tables or indexes need backing up. Specifically:

- You can save backup time by not backing up other tables or indexes in the SQL database that do not change as often.
- You can save restore time if, for example, the file group is on a different volume from the rest of the SQL database's file groups and that volume needs to be restored. You need restore only that file group for that SQL database.

The *groupname* variable specifies the names of the SQL server database file groups you want to back up.

Considerations:

- You can specify this parameter more than once per command invocation.
- A new group backup object inactivates any active group backup object of the same name in the same SQL database.
- Use * as a wildcard character in the *groupname* variable to replace zero or more characters for each occurrence.
- Specifying only the wildcard character indicates all file groups in the SQL server database.

Note: Since 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 file group in the SQL server database.

- If the *groupname* variable includes spaces or special characters, enclose it in double quotes.
- The *groupname* variable is case-sensitive.
- You should follow group backups with transaction log backups for all SQL databases you back up.
- You cannot perform group backups for the following SQL databases:
 - Those with the SQL Server attribute TRUNCATE LOG ON CHECKPOINT.
 - Those using the SIMPLE recovery model.
- You cannot specify the */recovery* parameter with **restore group** operations.
- A SQL Server Create Index operation requires that you back up all affected file groups as a unit. You cannot back up a single file group of the affected file groups until you perform the unit backup. A SQL Server error message will indicate which file groups are affected. You must perform a full database backup or a set backup of at least the affected file groups before the group backup succeeds.

Log or Log=*|*logobjectname*,...

A log backup contains the contents of the transaction log for an active SQL server database since the latest successful log backup. This option can save backup time by requiring fewer SQL database backups. For **backup**

operations, **log** takes no values. Use * as a wildcard character in *logobjectname* to replace zero or more characters for each occurrence. Specifying only the wildcard character indicates all log backup objects for the SQL databases.

Considerations:

- You can control the size of a transaction log by allowing a log backup to truncate the inactive part of the transaction log. This is the default.
- By using the */truncate=no* parameter, you may be able to backup the transaction log of a damaged, suspect, or unrecovered SQL Server database.
- Each log backed up creates a separate backup object with a unique name on the Tivoli Storage Manager server. A new log backup object does *not* inactivate any active backup objects (unlike the other backup types except **set** backups). Log backup objects do not participate in Tivoli Storage Manager server automatic expiration processing except when full database backup objects inactivate all active backup objects for a SQL database. Therefore, you can inactivate log backup objects using the **inactivate** command if full database backups are not performed frequently or at all.
- You cannot perform log backups for the following SQL databases:
 - Those with the SQL Server attribute TRUNCATE LOG ON CHECKPOINT.
 - Those using the SIMPLE recovery model.

Set or Set=* | setobjectname,...

A **set** backup contains the contents of the SQL server file groups and files you specify with the */files* and */groups* parameters. For **backup** operations, **set** takes no values. Use * as a wildcard character in *setobjectname* to replace zero or more characters for each occurrence. Specifying only the wildcard character indicates all set backup objects for the SQL databases.

Considerations:

- Use this option for unusual circumstances or special, one-time backups. One such case is when SQL Server requires that certain file groups be backed up as a unit and a full database backup is not practical. See the description of the **file**, and **group** parameters in this section, specifically in reference to the Create Index operation.
- Each SQL database backed up creates a separate backup object on the Tivoli Storage Manager server. All of the files and file groups backed up as part of a set backup for the same SQL server database are contained in a *single* backup object. Note that this is different from group and file backups, which create a separate backup object of *each* file and file group even if they are part of the same SQL server database.
- A new set backup object does *not* inactivate any active backup objects (unlike the other backup types except **log** backups). Set backup objects do not participate in Tivoli Storage Manager server automatic expiration processing except when full database backup objects inactivate all active backup objects for a SQL database. Therefore, you can inactivate set backup objects using the **inactivate** command if full database backups are not performed or not performed frequently.
- You should follow set backups with transaction log backups for all SQL databases you back up.

- The **file**, **group**, **log**, and **set** parameters can take a list of values (repeatable syntax) and may be specified more than once. 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.

/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 backup, the size of transaction logs can increase beyond the original backup estimate and cause the backup to fail. Use this parameter to customize the number of kilobytes in the backup estimate and avoid possible backup failures.

/ADJUSTPERcenttsmestimate=*numpercent*

The **/adjustpercenttsmestimate** parameter specifies the percentage number to add to the size of the backup estimate. The *numpercent* variable refers to the percentage number to add. The number can range from 0 to 99. The default is 0. Increasing the percentage estimate may be necessary when the backup estimate (generated by the SQL Server) may be too low as the disk storage pool has cache enabled. For example, if maintenance is performed on the production server during a Data Protection for SQL 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.

/BACKUPDESTination=TSM|LOCAL|BOTH

Use the **/backupdestination** parameter to specify the location where the backup is stored.

You can specify:

TSM The backup is stored on Tivoli Storage Manager server storage only. This is the default.

LOCAL

The backup is stored on local shadow volumes only. This is only valid when the **/backupmethod** parameter specifies **VSS**.

BOTH The backup is stored on Tivoli Storage Manager server storage and local shadow volumes. This is only valid when the **/backupmethod** parameter specifies **VSS**.

/BACKUPMETHod=LEGACY|VSS

Use the **/backupmethod** parameter to specify the manner in which the backup is performed.

You can specify:

LEGACY

The backup is performed with the legacy API. This is the SQL

streaming backup and restore API as used in previous versions of Data Protection for SQL. This is the default.

VSS The backup is performed with VSS.

/BUFFERS=*numbuffers*

The **/buffers** parameter specifies the number of data buffers used for each data stripe to transfer data between Data Protection for SQL 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 configuration file. This is initially 3.
- If you specify **/buffers**, its value is used instead of the value stored in the Data Protection for SQL configuration file. Specifying this parameter does not change the value in the configuration file.
- If you specify **/buffers** but not *numbuffers*, the default value 3 is used.

/BUFFERSIZE=*buffersizeinkb*

The **/buffersize** parameter specifies the size of each Data Protection for SQL buffer specified by the **/buffers** parameter. The *buffersizeinkb* variable refers to the size of data buffers in kilobytes. The number can range from 64 to 8192. The default is 1024.

Considerations:

- Though increasing the number of buffers can improve throughput, it also increases storage use as determined by this parameter.
- The default value is the value specified by the buffers configurable option in the Data Protection for SQL configuration file. This is initially 1024.
- If you specify **/buffersize**, its value is used instead of the value stored in the Data Protection for SQL configuration file. Specifying this parameter does not change the value in the configuration file.
- If you specify **/buffersize** but not *buffersizeinkb*, the default value 1024 is used.

/CONFIGFILE=*configfilename*

The **/configfile** parameter specifies the name of the Data Protection for SQL configuration file, which contains the values for the Data Protection for SQL configurable options. See “Set command” on page 222 for details on the file's contents.

Considerations:

- *configfilename* can include a fully qualified path. If *configfilename* does not include a path, it uses the directory where Data Protection for SQL is installed.
- If *configfilename* includes spaces, place it in double quotes.
- If you do not specify **/configfile**, the default value is *tdpsql.cfg*.

/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.

/SQLCOMPression=No | Yes

The **SQLCOMPression** parameter specifies if SQL compression is applied. If you do not specify **/sqlcompression**, the default value No is used. This parameter is only applicable on systems running SQL Server 2008 or later. For SQL Server 2008, backup compression is only supported on the Enterprise Edition. For SQL Server 2008 R2, backup compression is supported on Standard, Enterprise, and Datacenter editions.

/LOGFile=logfilename

The **/logfile** parameter specifies the name of the activity log that is generated by Data Protection for SQL. This activity log records significant events such as completed commands and error messages. The Data Protection for SQL activity log is distinct from the SQL Server error log. The *logfilename* variable identifies the name to be used for the activity log generated by Data Protection for SQL.

Considerations:

- If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file.
- The file name can include a fully-qualified path; however, if you specify no path, the file is written to the directory where Data Protection for SQL is installed.
- You cannot turn Data Protection for SQL activity logging off. If you do not specify **/logfile**, log records are written to the default log file. The default log file is *tdpsql.log*.
- When using multiple simultaneous instances of Data Protection for SQL to perform operations, use the **/logfile** parameter to specify a different log file for each instance used. This directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPrune=numdays | No

The **/logprune** parameter prunes the Data Protection for SQL activity log and specifies how many days of entries are saved. By default, log pruning is enabled and performed once each day Data Protection for SQL is executed; however, this option allows you to disable log pruning or explicitly request a prune of the log for one command run even if the log file has already been pruned for the day. The *numdays* variable represents the number of days to save log entries. By default, 60 days of log entries are saved in the prune process.

Considerations:

- If you specify *numdays*, it can range from 0 to 9999. A value of 0 deletes all entries in the Data Protection for SQL activity log file except for the current command entries.
- If you specify **/logprune**, its value is used instead of the value stored in the Data Protection for SQL configuration file. Specifying this parameter does not change the value in the configuration file.
- Changes to the value of the **timeformat** or **dateformat** parameter can result in an undesired pruning of the Data Protection for SQL log file. If you are running a command that may prune the log file and the value

of the **timeformat** or **dateformat** parameter has changed, perform one of the following to prevent undesired pruning of the log file:

- Make a copy of the existing log file.
- Specify a new log file with the **/logfile** parameter or **logfile** setting.

/MOUNTWait=Yes | No

If the Tivoli Storage Manager server is configured to store backup data on removable media such as tapes or optical media, it is possible that the Tivoli Storage Manager server may indicate to Data Protection for SQL 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 cannot complete waiting until the initial media for all stripes are available, although Data Protection for SQL 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 command from a prolonged wait. This can be done *only* by terminating the Data Protection for SQL 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 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 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 all of 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 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 ONLY valid when **/backupmethod=VSS** and **/backupdestination=TSM**. Note that this parameter requires a VSS provider that supports transportable shadow copies. It is not supported with the default Windows VSS System Provider.

/SQLAUTHentication=INTEgrated | SQLuserid

This parameter specifies the authorization mode used when logging on to the SQL server. The **integrated** value specifies Windows authentication. The user id you use to log on to Windows is the same id you will use to log on to the SQL server. This is the default value.

Use the **sqluserid** value to specify SQL Server user id authorization. The user id specified by the **/sqluserid** parameter is the id you will use to log on to the SQL server. Any SQL user id must have the SQL Server SYSADMIN fixed server role.

/SQLBUFFers=numsqbuffers

The **/sqlbuffers** parameter specifies the total number of data buffers SQL Server uses to transfer data between SQL Server and Data Protection for SQL. The *numsqbuffers* variable refers to the number of data buffers to use. The number can range from 0 to 999. The initial value is 0. When **/sqlbuffers** is set to 0, SQL determines how many buffers should be used.

Considerations:

- The default value is the value specified by the SQL buffers configurable option in the Data Protection for SQL 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 configuration file. Specifying this parameter does not change the value in the configuration file.
- If you specify **/sqlbuffers** but not *numsqbuffers*, 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. The *sqlbuffersizeinkb* variable refers to the size of data buffers in kilobytes. The number can range from 64 to 4096. The default is 1024.

Considerations:

- The default value is the value specified by the SQL buffers configurable option in the Data Protection for SQL 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 configuration file. Specifying this parameter does not change the value in the configuration file.
- If you specify **/sqlbuffersize** but not *sqlbuffersizeinkb*, the default value 1024 is used.

/SQLPassword=sqlpasswordname

This parameter specifies the SQL password that Data Protection for SQL uses to log on to the SQL server that objects are backed up from or restored to.

Considerations:

- Using this parameter means that you are using SQL Server authentication. The SQL Server and the SQL user id for this password must both be configured for SQL Server authentication.
- If you do not specify **/sqlpassword**, the default value is blank ("").
- If you specify **/sqlpassword** but not *sqlpasswordname*, the default is also blank ("").
- This parameter is ignored if you use the **/sqlauth=integrated** parameter with it.

/SQLSERVer=sqlprotocol:sqlservername

The **/sqlserver** parameter specifies the SQL server that Data Protection for SQL logs on to. The *sqlprotocol* variable specifies the communication protocol to use. You can specify one of the following protocols:

- *lpc*: Use Shared Memory protocol.
- *np*: Use Named Pipes protocol.
- *tcp*: Use Transmission Control protocol.
- *via*: Use Virtual Interface Architecture protocol.

If no protocol is specified, Data Protection for SQL logs on to the SQL server according to the first protocol that becomes available.

Considerations:

- The default value is the value specified by the SQL server configurable option in the Data Protection for SQL configuration file. This is initially the local computer name.
- If you specify */sqlserver* but not *sqlservername*, the local computer name is used.
- The following two shortcuts are accepted as the local computer name: . (local) These are a period or the word *local* within parentheses.
- 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 reside on. The TCP/IP host name may not always be the same.

instancename

The name given to the named instance of SQL Server specified during installation of the instance.

virtualservername

The name given to the clustered SQL Server specified during clustering service setup. This is not the cluster or node name.

/SQLUSER=sqlusername

The */sqluser* parameter specifies the name that Data Protection for SQL uses to log on to the SQL server.

Considerations:

- Using this parameter means that you are using SQL Server authentication. The SQL Server and the SQL user id for this password must both be configured for SQL Server authentication.
- The SQL user id must have the SQL server SYSADMIN fixed server role.
- If you do not specify */sqluser*, the default is **sa**.

- If you specify */sqluser* but not *sqlusername*, the default is also **sa**.
- This parameter is ignored if you use the */sqlauth=integrated* parameter with it.

/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 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 command: `query tsm dbname backup_object`
- You must use the MAXNUMMP parameter on a Tivoli Storage Manager REGISTER NODE or UPDATE NODE command to allow a node to use multiple sessions to store data on removable media (which requires you to allocate multiple mount points to that node). The MAXNUMMP value must be equal to or less than the maximum number of stripes you desire.
- When you use data striping, you should use Tivoli Storage Manager server file space collocation to try to keep each stripe on a different storage volume.
- The maximum number of data stripes you can use is one less than the value of the Tivoli Storage Manager server TXNGROUPMAX option in the *dsmserv.opt* file.

/TSMNODE=*tsmnodename*

The **/tsmnode** parameter specifies the Tivoli Storage Manager node name that Data Protection for SQL uses to log on to the Tivoli Storage Manager server. This identifies which Tivoli Storage Manager client is requesting services. You can also store the node name in the options file. The command line parameter overrides the value in the options file.

Considerations:

- You cannot use the **/tsmnode** parameter if PASSWORDACCESS GENERATE is specified in the Tivoli Storage Manager options file. You must specify the nodename in the options file. Otherwise, you can change PASSWORDACCESS to PROMPT to utilize the **/tsmnode** parameter. For details about the Tivoli Storage Manager options file, see the reference manual *IBM Tivoli Storage Manager for Windows Backup-Archive Client Installation and User's Guide*.
- If you do not specify **/tsmnode**, the default value is that specified by the nodename option in the Tivoli Storage Manager options file. Specifying this parameter does not change the value in the options file.

/TSMOPTFile=*dsmoptfilename*

The **/tsmoptfile** parameter specifies the Tivoli Storage Manager options file to use. This is similar to selecting a Tivoli Storage Manager server from the server list in the GUI. The Tivoli Storage Manager options file contains the configuration values for the Tivoli Storage Manager API. For details about

the Tivoli Storage Manager options file, see the reference manual *IBM Tivoli Storage Manager for Windows Backup-Archive Client Installation and User's Guide*.

Considerations:

- The *tsmoptfilename* variable can include a fully qualified path. If you do not include a path, the directory where Data Protection for SQL is installed is used.
- If *tsmoptfilename* includes spaces, you must enclose it in double quotes.
- If you do not specify **/tsmoptfile**, the default value is *dsm.opt*.
- If you specify **/tsmoptfile** but not *tsmoptfilename*, the default is also *dsm.opt*.

/TSMPassword=tsmpasswordname

The **/tsmpassword** parameter specifies the Tivoli Storage Manager password that Data Protection for SQL uses to log on to the Tivoli Storage Manager server. This parameter and the option **PASSWORDACCESS** in the Tivoli Storage Manager options file interact in the following ways:

/tsmpassword	PASSWORDACCESS in Tivoli Storage Manager options file	Password already stored in registry?	Result
specified	<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

Legacy Backup output examples

These output examples provide a sample of the text, messages, and process status that displays when using the **backup** command.

Legacy Backup 1: full

Legacy Backup 1 performs a Legacy full backup of two databases, *model* and *msdb*, to Tivoli Storage Manager server storage. Two optional parameters, */sqlbuffers* and */stripes*, are included.

Command:

```
tdpsqlc backup model,msdb full /sqlbuffers=2 /stripes=2
```

Output:

```

IBM Tivoli Storage Manager for Databases:
Data Protection for Microsoft SQL Server
Version 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011.
All rights reserved.

Starting SQL database backup...

Beginning full backup for database model, 1 of 2.
Full: 0   Read: 1218832   Written: 1218832   Rate: 322.48 Kb/Sec

Backup of model completed successfully.

Beginning full backup for database msdb, 2 of 2.
Full: 0   Read: 8250640   Written: 8250640   Rate: 2,463.24 Kb/Sec

Backup of msdb completed successfully.

Total SQL backups selected:           1
Total SQL backups attempted:          1
Total SQL backups completed:          1
Total SQL backups excluded:           0
Total SQL backups inactivated:        0
Total SQL backups deduplicated:       0

Throughput rate:                       907.63 Kb/Sec
Total bytes inspected:                  2,185,984
Total bytes transferred:                2,185,984
LanFree bytes transferred:              0
Total bytes before deduplication:       0
Total bytes after deduplication:        0
Data compressed by:                     0%
Deduplication reduction:                0.00%
Total data reduction ratio:             0.00%
Elapsed processing time:                 2.35 Secs

```

Legacy Backup 2: full

Legacy Backup 2 performs a Legacy full backup of the *test2* database with no output due to the */quiet* parameter. In addition, the default Windows authentication mode has been overridden by */sqlauthentication*.

Command:

```
tdpsqlc backup test2 full /quiet /sqlauth=sql
```

Output:

```

IBM Tivoli Storage Manager for Databases:
Data Protection for Microsoft SQL Server
Version 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011.
All rights reserved.

```

Legacy Backup 3: differential

Legacy Backup 3 performs a Legacy differential backup of the previous full backup of the *test2* database. An estimate of the changed portion of *test2* is included.

Command:

```
tdpsqlc backup test2 difffull /diffest=10
```

Output:

```

IBM Tivoli Storage Manager for Databases:
Data Protection for Microsoft SQL Server
Version 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011.
All rights reserved.

Starting SQL database backup...

Beginning diffull backup for database test2, 1 of 1.
Full: 0   Read: 5341016   Written: 5341016   Rate: 1,120.96 Kb/Sec

Backup of test2 completed successfully.

Total SQL backups selected:           1
Total SQL backups attempted:          1
Total SQL backups completed:          1
Total SQL backups excluded:           0
Total SQL backups inactivated:        0
Total SQL backups deduplicated:       0

Throughput rate:                       907.63 Kb/Sec
Total bytes inspected:                 2,185,984
Total bytes transferred:               2,185,984
LanFree bytes transferred:             0
Total bytes before deduplication:      0
Total bytes after deduplication:       0
Data compressed by:                   0%
Deduplication reduction:              0.00%
Total data reduction ratio:            0.00%
Elapsed processing time:                2.35 Secs

```

Note: Differential backup objects are denoted as "**diffull**" in CLI output and in the GUI tree and list.

Legacy Backup 4: log

Legacy Backup 4 performs a Legacy log backup of the previous full backups of *test2* and *model*. The default to truncate the logs is overridden.

Command:

```
tdpsqlc backup test2,model log /truncate=no
```

Output:

```

IBM Tivoli Storage Manager for Databases:
Data Protection for Microsoft SQL Server
Version 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011.
All rights reserved.

Starting SQL database backup...

Beginning log backup for database model, 1 of 2.
Full: 0   Read: 80656   Written: 80656   Rate: 67.55 Kb/Sec

Backup of model completed successfully.

Beginning log backup for database test2, 2 of 2.
Full: 0   Read: 88920   Written: 88920   Rate: 73.28 Kb/Sec

Backup of test2 completed successfully.

Total SQL backups selected:           1
Total SQL backups attempted:          1
Total SQL backups completed:          1
Total SQL backups excluded:           0
Total SQL backups inactivated:         0
Total SQL backups deduplicated:        0

Throughput rate:                       907.63 Kb/Sec
Total bytes inspected:                  2,185,984
Total bytes transferred:                 2,185,984
LanFree bytes transferred:               0
Total bytes before deduplication:        0
Total bytes after deduplication:         0
Data compressed by:                     0%
Deduplication reduction:                 0.00%
Total data reduction ratio:              0.00%
Elapsed processing time:                 2.35 Secs

```

Legacy Backup 5: group

Legacy Backup 5 performs a Legacy backup of all file groups belonging to the *netapp_db2* database.

Command:

```
tdpsqlc backup netapp_db2 Group=*
```

Output:

```

IBM Tivoli Storage Manager for Databases:
Data Protection for Microsoft SQL Server
Version 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011.
All rights reserved.

Connecting to SQL Server, please wait...

Starting SQL database backup...

Connecting to TSM Server as node 'TREELO_SQL'...

Beginning group backup for database netapp_db2, 1 of 1.
Full: 0   Read: 2182784   Written: 2182784   Rate: 2,135.90 Kb/Sec
Backup of netapp_db2 completed successfully.

Total SQL backups selected:           1
Total SQL backups attempted:          1
Total SQL backups completed:          1
Total SQL backups excluded:           0
Total SQL backups inactivated:         0
Total SQL backups deduplicated:        0

Throughput rate:                       907.63 Kb/Sec
Total bytes inspected:                  2,185,984
Total bytes transferred:                2,185,984
LanFree bytes transferred:              0
Total bytes before deduplication:        0
Total bytes after deduplication:         0
Data compressed by:                     0%
Deduplication reduction:                0.00%
Total data reduction ratio:             0.00%
Elapsed processing time:                 2.35 Secs

```

Legacy Backup 6: file

Legacy Backup 6 performs a Legacy file backup of all files belonging to *test2* using the wildcard character (*). This consists of three files within two groups within one database.

Command:

```
tdpsqlc backup test2 file=*
```

Output:

```

IBM Tivoli Storage Manager for Databases:
Data Protection for Microsoft SQL Server
Version 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011.
All rights reserved.

Starting SQL database backup...

Beginning file backup for database test2, 1 of 3.
Full: 0   Read: 669832   Written: 669832   Rate: 569.31 Kb/Sec

Backup of test2 completed successfully.

Beginning file backup for database test2, 2 of 3.
Full: 0   Read: 669832   Written: 669832   Rate: 388.90 Kb/Sec

Backup of test2 completed successfully.

Beginning file backup for database test2, 3 of 3.
Full: 0   Read: 13958280   Written: 13958280   Rate: 2,911.39 Kb/Sec

Backup of test2 completed successfully.

Total SQL backups selected:           1
Total SQL backups attempted:         1
Total SQL backups completed:         1
Total SQL backups excluded:          0
Total SQL backups inactivated:       0
Total SQL backups deduplicated:      0

Throughput rate:                      907.63 Kb/Sec
Total bytes inspected:                 2,185,984
Total bytes transferred:               2,185,984
LanFree bytes transferred:             0
Total bytes before deduplication:      0
Total bytes after deduplication:       0
Data compressed by:                   0%
Deduplication reduction:              0.00%
Total data reduction ratio:           0.00%
Elapsed processing time:                2.35 Secs

```

Legacy Backup 7: set

Legacy Backup 7 performs a Legacy set backup of one file group and two separate files (jointly as a single backup object) from the *test2* database. The */groups* and */files* parameters specify which items constitute this set backup.

Command:

```

tdpsqlc backup test2 set /groups=primary /files=test2_2data,
test2_3data

```

Output:

```

IBM Tivoli Storage Manager for Databases:
Data Protection for Microsoft SQL Server
Version 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011.
All rights reserved.

Starting SQL database backup...

Beginning set backup for database test2, 1 of 1.
Full: 0   Read: 14706896   Written: 14706896
Rate: 3,404.98 Kb/Sec

Backup of test2 completed successfully.

Total SQL backups selected:           1
Total SQL backups attempted:          1
Total SQL backups completed:          1
Total SQL backups excluded:           0
Total SQL backups inactivated:        0
Total SQL backups deduplicated:       0

Throughput rate:                       907.63 Kb/Sec
Total bytes inspected:                 2,185,984
Total bytes transferred:              2,185,984
LanFree bytes transferred:             0
Total bytes before deduplication:      0
Total bytes after deduplication:       0
Data compressed by:                   0%
Deduplication reduction:              0.00%
Total data reduction ratio:           0.00%
Elapsed processing time:               2.35 Secs

```

Legacy Backup 8: differential

Legacy Backup 8 performs a Legacy differential backup of all available databases using the wildcard character (*). However, the new */excludedb* parameter is used to exclude the *master* and *msdb* databases from being backed up.

Command:

```
tdpsqlc backup * difffull /excludedb=master,msdb
```

Output:

```

IBM Tivoli Storage Manager for Databases:
Data Protection for Microsoft SQL Server
Version 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011.
All rights reserved.

Connecting to SQL Server, please wait...

Starting SQL database backup...

Connecting to TSM Server as node 'TREELO_SQL'...

Beginning diff full backup for database AdventureWorks, 1 of 5.
Full: 0 Read: 1140480 Written: 1140480 Rate: 904.02 Kb/Sec
Backup of AdventureWorks completed successfully.

Beginning diff full backup for database AdventureWorksDW, 2 of 5.
Full: 0 Read: 1135360 Written: 1135360 Rate: 1,447.45 Kb/Sec
Backup of AdventureWorksDW completed successfully.

Beginning diff full backup for database model, 3 of 5.
Full: 0 Read: 1050485 Written: 1050485 Rate: 801.04 Kb/Sec
Backup of model completed successfully.

Beginning diff full backup for database netapp_db1, 4 of 5.
Full: 0 Read: 1227360 Written: 1227360 Rate: 1,109.82 Kb/Sec
Backup of netapp_db1 completed successfully.

Beginning diff full backup for database test1, 5 of 5.
Full: 0 Read: 1135360 Written: 1135360 Rate: 1,447.45 Kb/Sec
Backup of test1 completed successfully.

Total SQL backups selected:          1
Total SQL backups attempted:         1
Total SQL backups completed:         1
Total SQL backups excluded:          0
Total SQL backups inactivated:       0
Total SQL backups deduplicated:      0

Throughput rate:                     907.63 Kb/Sec
Total bytes inspected:                2,185,984
Total bytes transferred:              2,185,984
LanFree bytes transferred:            0
Total bytes before deduplication:     0
Total bytes after deduplication:      0
Data compressed by:                   0%
Deduplication reduction:              0.00%
Total data reduction ratio:           0.00%
Elapsed processing time:               2.35 Secs

```

Legacy Backup 9: full

Legacy Backup 9 performs a Legacy full backup of all available databases using the wildcard character (*). However, the new */excludedb* parameter is used to exclude the *master* and *msdb* databases from being backed up.

Command:

```
tdpsqlc backup * full /excludedb=master,msdb
```

Output:

```

IBM Tivoli Storage Manager for Databases:
Data Protection for Microsoft SQL Server
Version 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011.
All rights reserved.

Connecting to SQL Server, please wait...

Starting SQL database backup...

Connecting to TSM Server as node 'STRINGVM1_SQL'...

Beginning full backup for database DB1_XIVmini_G_BAS, 1 of 2.
Full: 0   Read: 0   Written: 0   Rate: 0.00 Kb/Sec
Full: 0   Read: 2185984   Written: 2185984   Rate: 987.85 Kb/Sec

Backup of DB1_XIVmini_G_BAS completed successfully.

Beginning full backup for database model, 2 of 2.
Full: 0   Read: 2188032   Written: 2188032   Rate: 2,189.29 Kb/Sec

Backup of model completed successfully.

Total SQL backups selected:           4
Total SQL backups attempted:         2
Total SQL backups completed:         2
Total SQL backups excluded:          2
Total SQL backups inactivated:       0
Total SQL backups deduplicated:      0

Throughput rate:                      1,353.88 Kb/Sec
Total bytes inspected:                4,374,016
Total bytes transferred:              4,374,016
Total LanFree bytes transferred:      0
Total bytes before deduplication:     0
Total bytes after deduplication:      0
Data compressed by:                   0%
Deduplication reduction:              0.00%
Total data reduction ratio:           0.00%
Elapsed processing time:               3.15 Secs

Completed

```

Legacy Backup 10: full

Legacy Backup 10 performs a Legacy full backup of the *test1* database. The new *ladjustkbtsmestimate* parameter is used to customize the number of kilobytes.

Command:

```
tdpsqlc backup test1 full /adjustkbtsmestimate=25
```

Output:

```

IBM Tivoli Storage Manager for Databases:
Data Protection for Microsoft SQL Server
Version 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011.
All rights reserved.

Connecting to SQL Server, please wait...

Starting SQL database backup...

Connecting to TSM Server as node 'TREELO_SQL'...

Beginning full backup for database test1, 1 of 1.
Full: 0   Read: 2183936   Written: 2183936   Rate: 1,949.50 Kb/Sec
Backup of test1 completed successfully.

Total SQL backups selected:           1
Total SQL backups attempted:         1
Total SQL backups completed:         1
Total SQL backups excluded:          0
Total SQL backups inactivated:        0
Total SQL backups deduplicated:       0

Throughput rate:                      907.63 Kb/Sec
Total bytes inspected:                 2,185,984
Total bytes transferred:               2,185,984
LanFree bytes transferred:             0
Total bytes before deduplication:      0
Total bytes after deduplication:       0
Data compressed by:                   0%
Deduplication reduction:               0.00%
Total data reduction ratio:            0.00%
Elapsed processing time:                2.35 Secs

```

VSS Backup output examples

These output examples provide a sample of the text, messages, and process status that displays when using the **backup** command.

VSS Backup 1: full local

VSS Backup 1 performs a VSS full backup of database *test1* to local shadow volumes using the new optional parameters, */backupdestination* and */backupmethod*.

Command:

```
tdpsqlc backup DB1_XIVmini_G_BAS full /backupdestination=local
/backupmethod=vss
```

Output:

```

IBM Tivoli Storage Manager for Databases:
Data Protection for Microsoft SQL Server
Version 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011.
All rights reserved.

Connecting to SQL Server, please wait...

Connecting to TSM Server as node 'STRINGVM1_SQL'...
Connecting to Local DSM Agent 'STRINGVM1'...

Starting SQL database backup...

Beginning VSS backup of 'DB1_XIVmini_G_BAS'...

Performing Snapshot Operation...

VSS Backup operation completed with rc = 0.

    Elapsed Processing Time: 29.09 seconds

Completed

```

VSS Backup 2: full local and Tivoli Storage Manager server

VSS Backup 2 performs a VSS full backup of database *model* to local shadow volumes and Tivoli Storage Manager server storage using *lbackupmethod*.

Command:

```
tdpsqlc backup model full /backupmethod=vss
```

Output:

```

IBM Tivoli Storage Manager for Databases:
Data Protection for Microsoft SQL Server
Version 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011.
All rights reserved.

Connecting to SQL Server, please wait...

Connecting to TSM Server as node 'STRINGVM1_SQL'...
Connecting to Local DSM Agent 'STRINGVM1'...

Starting SQL database backup...

Beginning VSS backup of 'model'...

Performing Snapshot Operation...

    Files Examined/Completed/Failed: [ 5 / 0 / 0 ]   Total Bytes: 2262039
    Files Examined/Completed/Failed: [ 5 / 5 / 0 ]   Total Bytes: 3959984

VSS Backup operation completed with rc = 0
Files Examined           : 5
Files Completed          : 5
Files Failed             : 0
Files Deduplicated       : 0
Total Bytes Inspected    : 3957509
Total Bytes              : 3959984
Total LanFree Bytes      : 0
Total Bytes Before Deduplication : 0
Total Bytes After Deduplication : 0
Files Compressed By      : 0%
Deduplication Reduction : 0.00%
Total Data Reduction Ratio : 0.00%

Completed

```

VSS Backup 3: full local

VSS Backup 3 performs a VSS full backup of all available databases to local shadow volumes using the wildcard character (*). Using the parameter */excludedb* enables you to exclude the *master* and *msdb* databases from being backed up in this example.

Command:

```
tdpsqlc backup * full /backupdestination=local /backupmethod=vss  
/exclude=master,msdb
```

Output:

```
IBM Tivoli Storage Manager for Databases:  
Data Protection for Microsoft SQL Server  
Version 6, Release 3, Level 0.0  
(C) Copyright IBM Corporation 1998, 2011.  
All rights reserved.  
  
Connecting to SQL Server, please wait...  
  
AC05459W Database master excluded from backup.  
AC05459W Database msdb excluded from backup.  
Connecting to TSM Server as node 'STRINGVM1_SQL'...  
Connecting to Local DSM Agent 'STRINGVM1'...  
  
Starting SQL database backup...  
  
Beginning VSS backup of 'DB1_XIVmini_G_BAS', 'model'...  
  
Performing Snapshot Operation...  
  
VSS Backup operation completed with rc = 0.  
  
Elapsed Processing Time: 44.31 seconds  
  
Completed
```

VSS Backup 4: full local

VSS Backup 4 performs a VSS full backup of the *test1* database.

Command:

```
tdpsqlc backup test1 full /backupdestination=local /backupmethod=vss
```

Output:

```

IBM Tivoli Storage Manager for Databases:
Data Protection for Microsoft SQL Server
Version 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011.
All rights reserved.

Connecting to SQL Server, please wait...

Connecting to TSM Server as node 'TREELO_SQL'...
Connecting to Local DSM Agent 'tree_lo_agent'...

Starting SQL database backup...

Beginning VSS backup of 'test1'...

Preparing to backup using snapshot.
Files Examined/Completed/Failed: [ 5 / 5 / 0 ] Total Bytes: 3175517

VSS Backup operation completed with rc = 0
Files Examined           : 5
Files Completed          : 5
Files Failed             : 0
Files Deduplicated       : 0
Total Bytes Inspected    : 3175517
Total Bytes              : 3175517
Total LanFree Bytes     : 0
Total Bytes Before Deduplication : 0
Total Bytes After Deduplication : 0
Files Compressed By     : 0%
Deduplication Reduction : 0.00
Total Data Reduction Ratio : 0.00%

```

Query command

Use the **query** command to display information about the SQL server and its databases, about the Tivoli Storage Manager server and its backup objects, and about Data Protection for SQL.

Considerations:

- Some of the information displays may have long text lines. You can redirect the informational output of the Data Protection for SQL query command to a text file using the Windows command output redirection syntax (command prompt):

```
TDPcommandstatement > [[drive:]path\]filename.ext
```

This creates or replaces the file.

```
TDPcommandstatement >> [[drive:]path\]filename.ext
```

This appends to the file.

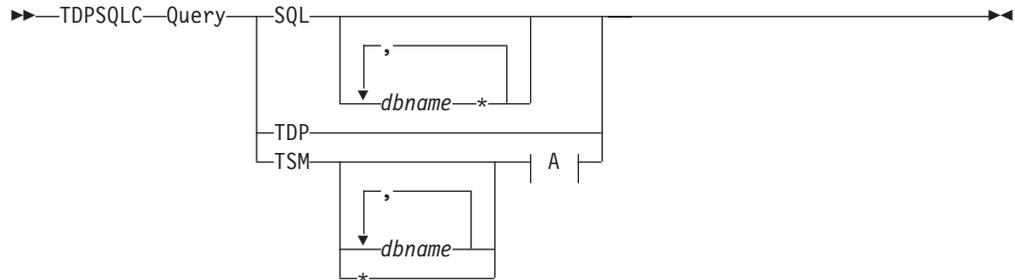
You can then browse or edit the file.

- You can use the Windows **more** filter command (command prompt) to display the informational output one screen at a time, in conjunction with the Windows command pipe character: `TDPcommandstatement | more`

Query syntax

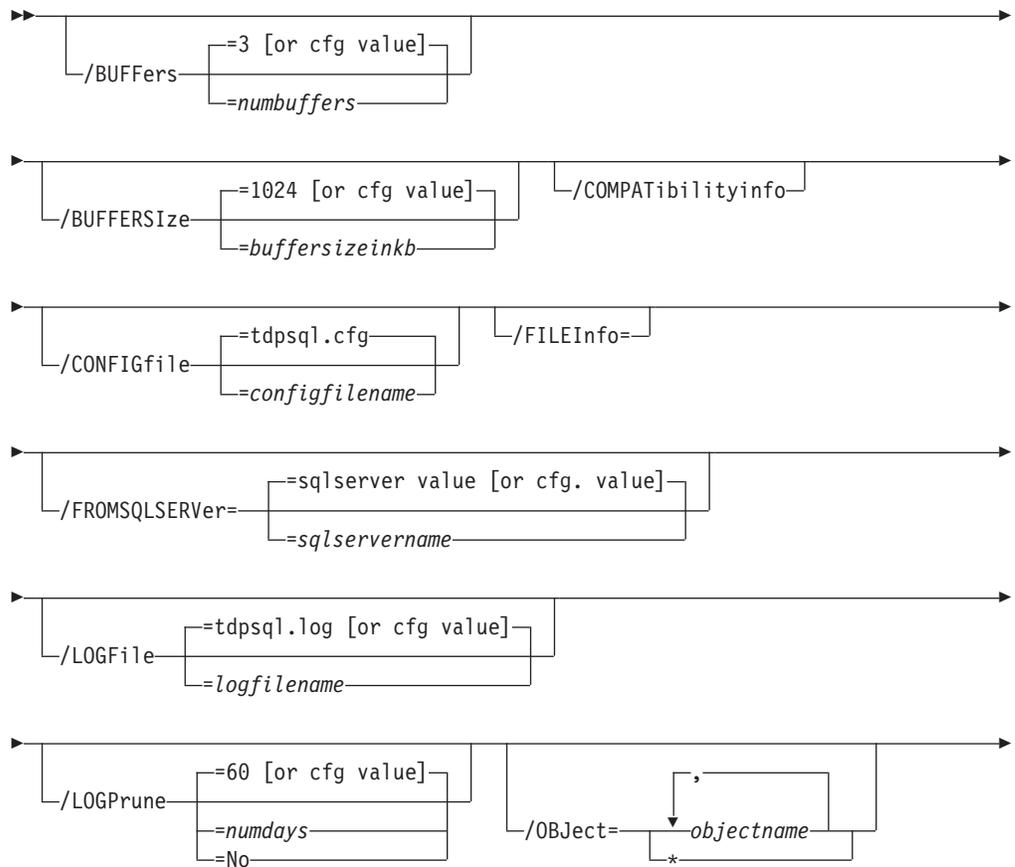
Use the **query** command syntax diagrams as a reference to view available options and truncation requirements.

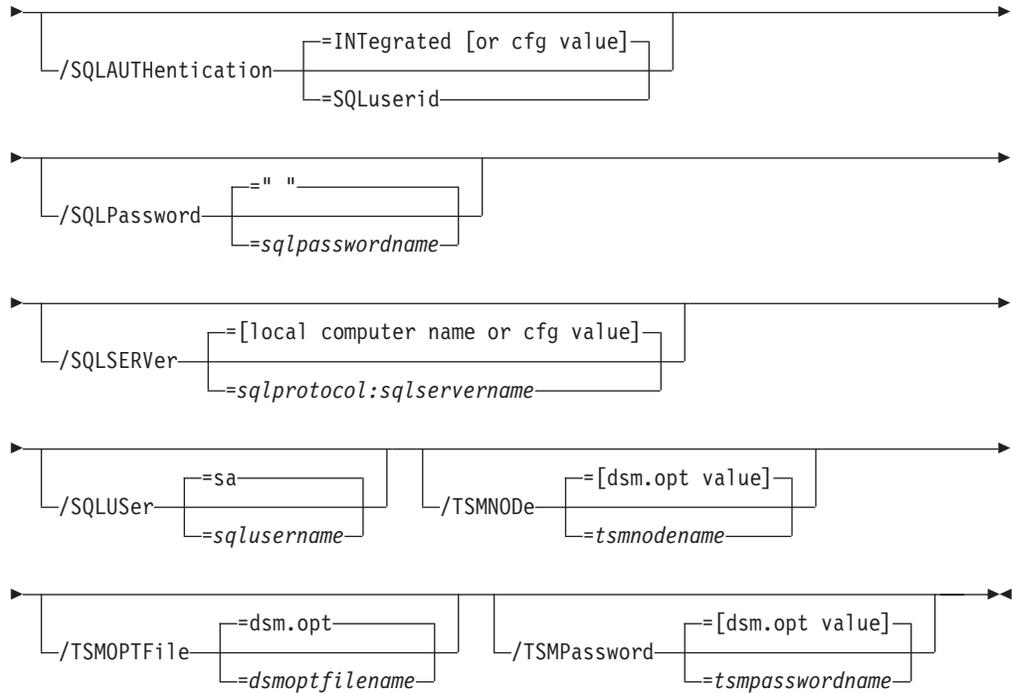
Syntax



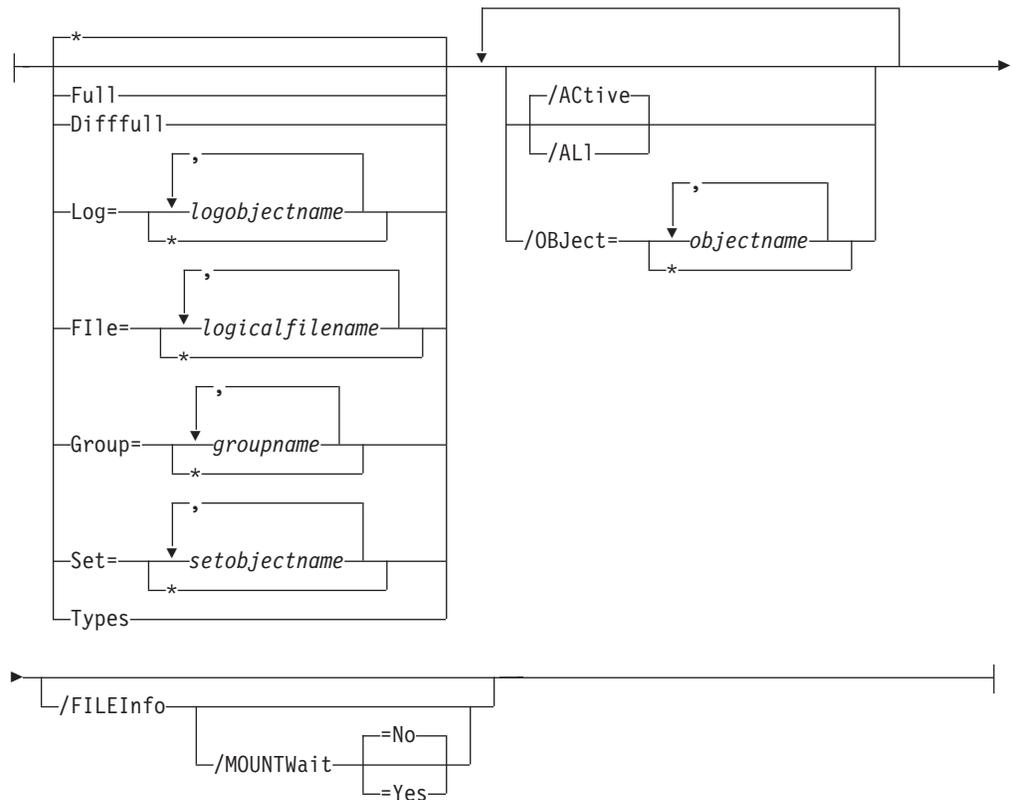
The syntax diagram of the Tivoli Storage Manager options corresponding to the letter above (A) follows the Optional Parameters below.

Query Optional Parameters:





A Query TSM Options:



Query positional parameters

Positional parameters immediately follow the **query** command and precede the optional parameters.

Specify one of the following when issuing a Data Protection for SQL **query** command:

Query SQL *|*dbname*,...

This displays information about the current SQL server. The *dbname* variable specifies databases on the current SQL server to display information about.

When querying a SQL Server, the following information is included:

- Server name
- Database name
- Database data space allocated
- Database space used
- Database log space allocated
- Database log space used
- Database options set (SELECT INTO / BULK COPY, TRUNCATE LOG ON CHECKPOINT, and so on)

If you specify */COMPATibilityinfo*:

- Server version
- Server clustering state
- Database compatibility level

Query TDP

This displays the Data Protection for SQL name and version information and the contents of the current Data Protection for SQL configuration file.

Query TSM *|*dbname*,...

This displays the Tivoli Storage Manager API and Tivoli Storage Manager server version information. The *dbname* variable names the specified databases from the current SQL server that have backup objects on the current Tivoli Storage Manager server and node. No name is displayed if specified objects do not exist as backup objects in the SQL database. Use the *dbname*,...* wildcard option to display information about *all* of the backup objects of one or more SQL databases.

When querying any backup object using **TSM** *dbname*, the following information is included:

- SQL server name
- SQL database name
- Backup object type
- Backup object active/inactive state
- Backup object Data Protection for SQL creation date and time
- Backup object Data Protection for SQL size
- Data Protection for SQL backup-object object name
- SQL compressed
- Backup compressed
- Backup encryption type
- Backup deduplicated

- Number of data stripes in backup object
- For VSS only, whether the backup supports Instant Restore

The following is included if you specify */compatibilityinfo*:

- SQL server version
- SQL Server clustering state
- Data Protection for SQL version that created the backup object
- SQL database compatibility level
- SQL database data space allocated
- SQL database data space used
- SQL database log space allocated
- SQL database log space used
- SQL database options

Note:

- You can also determine which backup objects to display through the **query TSM** optional parameters */active* and */all*.
- No information will be displayed if there are no backup objects for a specified SQL database.

File=* | logicalfilename,...

This displays information about file backup objects of one or more SQL databases from the current SQL server that are on the current Tivoli Storage Manager server and node.

Full This displays information about full backup objects of one or more SQL databases from the current SQL server that are on the current Tivoli Storage Manager server and node.

Difffull

This displays information about differential backup objects of one or more SQL databases from the current SQL server that are on the current Tivoli Storage Manager server and node.

Group=* | groupname,...

This displays information about one or more group backup objects of one or more SQL databases from the current SQL server that are on the current Tivoli Storage Manager server and node.

Log=* | logobjectname,...

This displays information about one or more log backup objects of one or more SQL databases from the current SQL server that are on the current Tivoli Storage Manager server and node. The *logobjectname* variable specifies which log backup objects to display information about. Use * as a wildcard character in *logobjectname* to replace zero or more characters for each occurrence. Specifying only the wildcard character indicates all log backup objects for the SQL databases.

Set=* | setobjectname,...

set parameter This displays information about one or more set backup objects of one or more SQL databases from the current SQL server that are on the current Tivoli Storage Manager server and node. The *setobjectname* variable specifies which set backup objects to display information about. Use * as a wildcard character in

setobjectname to replace zero or more characters for each occurrence. Specifying only the wildcard character indicates all set backup objects for the SQL databases.

Types This displays a summary by backup type of the backup objects of one or more SQL databases from the current SQL server that are on the current Tivoli Storage Manager server and node. Only backup types with one or more backup objects are displayed. If the */all* optional parameter is specified, the number of inactive backup objects is included. You cannot specify either the */compatibility* or the */fileinfo* optional parameter with the *types* parameter.

Query optional parameters

Optional parameters follow the **query** command and positional parameters.

The following are detailed descriptions of each of the optional parameters:

/BUFFers=*numbuffers*

The **/buffers** parameter specifies the number of data buffers used for each data stripe to transfer data between Data Protection for SQL 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 configuration file. This is initially 3.
- If you specify **/buffers**, its value is used instead of the value stored in the Data Protection for SQL configuration file. Specifying this parameter does not change the value in the configuration file.

/BUFFERSize=*bufferizeinkb*

The **/buffersize** parameter specifies the size of each Data Protection for SQL 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 configuration file. This is initially 1024.
- If you specify **/buffersize**, its value is used instead of the value stored in the Data Protection for SQL configuration file. Specifying this parameter does not change the value in the configuration file.

/COMPATibilityinfo

For **query** operations, this parameter displays information related to the compatibility of a backup object with a SQL server. Certain SQL Server configuration options must be compatible before you can restore a backup object to a SQL server. When you specify this parameter, SQL and Data

Protection for SQL configuration information is listed to help determine if a backup object is correct for a SQL server, or to help in problem determination.

Considerations:

- You cannot specify this parameter with the **types** parameter on a **query TSM** command.
- Compatible generally means identical. However, if you use a binary sort order for both the SQL server and the backup object, the code pages may be different, although the interpretation of individual character values may result in different characters being displayed or printed.

/CONFIGfile=*configfilename*

The **/configfile** parameter specifies the name of the Data Protection for SQL configuration file, which contains the values for the Data Protection for SQL configurable options. See “Set command” on page 222 for details on the file's contents.

Considerations:

- *configfilename* can include a fully qualified path. If *configfilename* does not include a path, it uses the directory where Data Protection for SQL is installed.
- If *configfilename* includes spaces, place it in double quotes.
- If you do not specify **/configfile**, the default value is *tdpsql.cfg*.
- If you specify **/configfile** but not *configfilename*, the default value *tdpsql.cfg* is used.

/FROMSQLSERVER=*sqlservername*

For **restore**, the **/fromsqlserver** parameter specifies the SQL server that backup objects were backed up from. This parameter is necessary only when the name of the SQL server to restore to, as determined by the **/sqlserver** parameter, is different from the name of the SQL server that the backup objects were created from. The default value is the **/sqlserver** value or the value set in the Data Protection for SQL configuration file.

Considerations:

- If the two SQL server names are different, you must use this parameter even if **/fromsqlserver** was a non-clustered default instance.
- After you restore a SQL database to a different SQL server, the logins of the SQL database may not match the logins for the different SQL server. If appropriate, you can use the SQL stored procedure **SP_CHANGE_USERS_LOGIN** to find and correct such SQL login mismatches.

/LOGFile=*logfile*

The **/logfile** parameter specifies the name of the activity log that is generated by Data Protection for SQL. This activity log records significant events such as completed commands and error messages. The Data Protection for SQL activity log is distinct from the SQL Server error log. The *logfile* variable identifies the name to be used for the activity log generated by Data Protection for SQL.

Considerations:

- If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file.

- The file name can include a fully-qualified path; however, if you specify no path, the file is written to the directory where Data Protection for SQL is installed.
- You cannot turn Data Protection for SQL activity logging off. If you do not specify **/logfile**, log records are written to the default log file. The default log file is *tdpsql.log*.
- When using multiple simultaneous instances of Data Protection for SQL to perform operations, use the **/logfile** parameter to specify a different log file for each instance used. This directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPrune=numdays | No

The **/logprune** parameter prunes the Data Protection for SQL activity log and specifies how many days of entries are saved. By default, log pruning is enabled and performed once each day Data Protection for SQL is executed; however, this option allows you to disable log pruning or explicitly request a prune of the log for one command run even if the log file has already been pruned for the day. The *numdays* variable represents the number of days to save log entries. By default, 60 days of log entries are saved in the prune process.

Considerations:

- If you specify *numdays*, it can range from 0 to 9999. A value of 0 deletes all entries in the Data Protection for SQL activity log file except for the current command entries.
- If you specify **no**, the log file is not pruned during this command.
- If you do not specify **/logprune**, the default value is that specified by the **logprune** configurable option in the Data Protection for SQL configuration file. This is initially 60.
- If you specify **/logprune**, its value is used instead of the value stored in the Data Protection for SQL configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/logprune** without specifying *numdays* or **no**; in this case, the default 60 is used.
- Changes to the value of the **timeformat** or **dateformat** parameter can result in an undesired pruning of the Data Protection for SQL log file. If you are running a command that may prune the log file and the value of the **timeformat** or **dateformat** parameter has changed, perform one of the following to prevent undesired pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/logfile** parameter or **logfile** setting.

/OBJect=* | objectname,...

For **restore** and **inactivate** operations, **/object** specifies that only particular backup objects for the specified SQL databases and backup object type (if specified) be restored or inactivated. For **query** operations, **/object** includes particular objects and object types in the display. The *objectname* variable specifies the names of the backup objects you want to restore or inactivate. The object name uniquely identifies each backup object and is created by Data Protection for SQL. Use **query** to view the names of backup objects. You can use * as a wildcard character in *objectname* to replace zero or more

characters for each occurrence. Specifying only the wildcard character indicates all backup objects of the specified SQL databases and backup object type.

/SQLAUTHentication=INTEgrated | SQLuserid

This parameter specifies the authorization mode used when logging on to the SQL server. The **integrated** value specifies Windows authentication. The user id you use to log on to Windows is the same id you will use to log on to the SQL server. This is the default value. Use the **sqluserid** value to specify SQL Server user id authorization. The user id specified by the **/sqluserid** parameter is the id you will use to log on to the SQL server. Any SQL user id must have the SQL Server SYSADMIN fixed server role.

/SQLPassword=sqlpasswordname

This parameter specifies the SQL password that Data Protection for SQL uses to log on to the SQL server that objects are backed up from or restored to.

Considerations:

- Using this parameter means that you are using SQL Server authentication. The SQL Server and the SQL user id for this password must both be configured for SQL Server authentication.
- If you do not specify **/sqlpassword**, the default value is blank (" ").
- If you specify **/sqlpassword** but not *sqlpasswordname*, the default is also blank (" ").
- This parameter is ignored if you use the **/sqlauth=integrated** parameter with it.

/SQLSERVer=sqlprotocol:sqlservername

The **/sqlserver** parameter specifies the SQL server that Data Protection for SQL logs on to. Use **/sqlserver** for the **query SQL** command, but use **/fromsqlserver** for the **query TSM** command. The *sqlprotocol* variable specifies the communication protocol to use. You can specify one of the following protocols:

- *lpc*: Use Shared Memory protocol.
- *np*: Use Named Pipes protocol.
- *tcp*: Use Transmission Control protocol.
- *via*: Use Virtual Interface Architecture protocol.

If no protocol is specified, Data Protection for SQL logs on to the SQL server according to the first protocol that becomes available.

Considerations:

- The default value is the value specified by the SQL server configurable option in the Data Protection for SQL configuration file. This is initially the local computer name.
- If you specify **/sqlserver** but not *sqlservername*, the local computer name is used.
- The following two shortcuts are accepted as the local computer name: . (local) These are a period or the word *local* within parentheses.
- 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 reside on. The TCP/IP host name may not always be the same.

instancename

The name given to the named instance of SQL Server specified during installation of the instance.

virtualservername

The name given to the clustered SQL Server specified during clustering service setup. This is not the cluster or node name.

/SQLUSER=sqlusername

The **/sqluser** parameter specifies the name that Data Protection for SQL uses to log on to the SQL server.

Considerations:

- Using this parameter means that you are using SQL Server authentication. The SQL Server and the SQL user id for this password must both be configured for SQL Server authentication.
- The SQL user id must have the SQL server SYSADMIN fixed server role.
- If you do not specify **/sqluser**, the default is **sa**.
- If you specify **/sqluser** but not *sqlusername*, the default is also **sa**.
- This parameter is ignored if you use the **/sqlauth=integrated** parameter with it.

/TSMNODE=tsmnodename

The **/tsmnode** parameter specifies the Tivoli Storage Manager node name that Data Protection for SQL uses to log on to the Tivoli Storage Manager server. This identifies which Tivoli Storage Manager client is requesting services. You can also store the node name in the options file. The command line parameter overrides the value in the options file.

Considerations:

- You cannot use the **/tsmnode** parameter if PASSWORDACCESS GENERATE is specified in the Tivoli Storage Manager options file. You must specify the nodename in the options file. Otherwise, you can change PASSWORDACCESS to PROMPT to utilize the **/tsmnode** parameter. For details about the Tivoli Storage Manager options file, see the reference manual *IBM Tivoli Storage Manager for Windows Backup-Archive Client Installation and User's Guide*.
- If you do not specify **/tsmnode**, the default value is that specified by the nodename option in the Tivoli Storage Manager options file. Specifying this parameter does not change the value in the options file.

/TSMOPTFile=dsmoptfilename

The **/tsmoptfile** parameter specifies the Tivoli Storage Manager options file to use. This is similar to selecting a Tivoli Storage Manager server from the server list in the GUI. The Tivoli Storage Manager options file contains the configuration values for the Tivoli Storage Manager API. For details about the Tivoli Storage Manager options file, see the reference manual *IBM Tivoli Storage Manager for Windows Backup-Archive Client Installation and User's Guide*.

Considerations:

- The *tsmoptfilename* variable can include a fully qualified path. If you do not include a path, the directory where Data Protection for SQL is installed is used.
- If *tsmoptfilename* includes spaces, you must enclose it in double quotes.
- If you do not specify **/tsmoptfile**, the default value is *dsm.opt*.
- If you specify **/tsmoptfile** but not *tsmoptfilename*, the default is also *dsm.opt*.

/TSMPassword=tsmpasswordname

The **/tsmpassword** parameter specifies the Tivoli Storage Manager password that Data Protection for SQL uses to log on to the Tivoli Storage Manager server. This parameter and the option **PASSWORDACCESS** in the Tivoli Storage Manager options file interact in the following ways:

/tsmpassword	PASSWORDACCESS in Tivoli Storage Manager options file	Password already stored in registry?	Result
specified	<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

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 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011.
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 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011.
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)

MSCS Cluster ..... No

SQL Database Information
-----

SQL Database Name ..... DB1_XIVmini_G_BAS
SQL Database Data Space Allocated ..... 3,145,728
SQL Database Data Space Used ..... 1,376,256
SQL Database Log Space Allocated ..... 2,097,152
SQL Database Log Space Used ..... 393,216
SQL Database Compatibility level..... 100
SQL Database Options .....

Completed

```

Query 3–TDP (Legacy)

Query 3 queries Data Protection for SQL for configuration file information. Note that this configuration is for Legacy operations only as *BACKUPDESTINATION TSM*, *BACKUPMETHOD LEGACY*, and the *LOCALDSMAgentnode* and *REMOTEDSMAgentnode* are not set.

Command:

```
tdpsqlc query tdp
```

Output:

```
IBM Tivoli Storage Manager for Databases:
Data Protection for Microsoft SQL Server
Version 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011.
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
MOUNTWAITFORDATA ..... Yes
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 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011.
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
MOUNTwaitfordata ..... Yes
NUMBERformat ..... 1
REMOVEDSMAgentnode .....
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 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011.
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 09/23/2011 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 20110923062314
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 09/23/2011 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 20110923063931
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 09/23/2011 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 20110923064114
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 09/23/2011 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 20110923064557
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 09/23/2011 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 20110923063104\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 09/23/2011 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 20110923062314
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 ..... 09/23/2011 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 ..... 20110923064311
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 ..... 09/23/2011 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 ..... 20110923064558
Assigned Management Class ..... DEFAULT

Backup Object Information
-----
SQL Server Name ..... STRINGVM1\STRINGVM1
SQL Database Name ..... model
Backup Method ..... Lgcy
Backup Location ..... Srv
Backup Object Type ..... Full
Backup Object State ..... Active
Backup Creation Date / Time ..... 09/23/2011 06:31:05
Backup Size ..... 2.08 MB
SQL Compressed ..... No
Backup Compressed ..... No
Backup Encryption Type ..... None
Backup Client-deduplicated ..... No
Database Object Name ..... 20110923063105\00001AC4
Number of stripes in backup object ..... 1
Assigned Management Class ..... DEFAULT

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 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011.
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 ..... 09/23/2011 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 ..... 20110923064311  
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 ..... 09/23/2011 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 ..... 20110923064558  
Assigned Management Class ..... DEFAULT
```

Backup Object Information

```
-----  
SQL Server Name ..... STRINGVM1\STRINGVM1  
SQL Database Name ..... model  
Backup Method ..... Lgcy  
Backup Location ..... Srv  
Backup Object Type ..... Full  
Backup Object State ..... Active  
Backup Creation Date / Time ..... 09/23/2011 06:31:05  
Backup Size ..... 2.08 MB  
SQL Compressed ..... No  
Backup Compressed ..... No  
Backup Encryption Type ..... None  
Backup Client-deduplicated ..... No  
Database Object Name ..... 20110923063105\00001AC4  
Number of stripes in backup object ..... 1  
Assigned Management Class ..... DEFAULT
```

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 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011.
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 ..... 09/27/2011 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 ..... 20110927082358\00001A4C
Number of stripes in backup object ..... 1
Assigned Management Class ..... DEFAULT
```

Query 8 –TSM 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 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011.
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 09/23/2011 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 20110923062314
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 09/23/2011 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 20110923063931
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 09/23/2011 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 20110923064114
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 09/23/2011 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 20110923064557
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 09/23/2011 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 20110923063104\00001AC4
Number of stripes in backup object 1
Assigned Management Class DEFAULT
SQL Server Version 10.0.2573 (SQL Server 2008)
MSCS Cluster No
DP Version 6.3.0.0
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

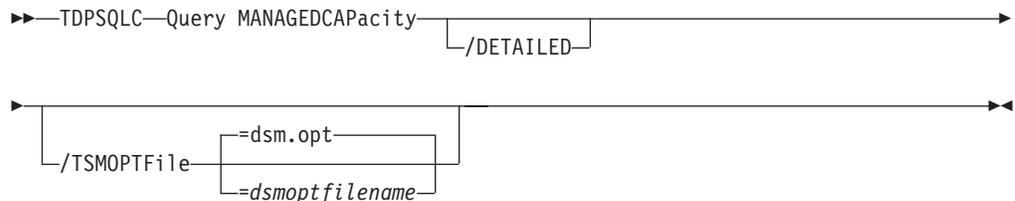
Completed

Query 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 represented in local inventory managed by Data Protection for SQL. This command is valid for all Windows operating systems that are supported by Data Protection for SQL.



Parameters

/DETAILED

Results in a detailed listing of snapped volumes. If this option is not specified then only the total capacity is displayed.

/TSMOPTfile=dsmoptfilename

The `/tsmoptfile` parameter specifies the Data Protection for SQL options file to use.

Considerations:

- The `dsmoptfilename` variable can include a fully qualified path. If you do not include a path, the Data Protection for SQL installation directory is used.
- If the `dsmoptfilename` variable spaces, enclose it in double quotation marks.
- If you do not specify `/tsmoptfile`, the default value is `dsm.opt`.
- If you specify `/tsmoptfile` but not `dsmoptfilename`, the default is also `dsm.opt`.

SQL Server 2005 example

Query the total managed capacity of SQL Server 2005 data represented in the local inventory with a detailed listing of snapped volumes. In this example there is a total of 1.25 TB of managed capacity. The detailed output shows that there are two snapshots of `volumeid_1` and one snapshot of `volumeid_2`. Each unique volume is only counted once so the total adds up to 1.25 TB. “`tdpsqlc query managedcapacity /detailed`”:

```
Managed Capacity: 1.25 TB

Volume: volumeid_1
Snapshot: snapshotid_1
Managed Capacity: 0.50 TB
Volume: volumeid_2
Snapshot: snapshotid_1
Managed Capacity: 0.75 TB

Volume: volumeid_1
Snapshot: snapshotid_2
Managed Capacity: 0.50 TB
```

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)
```

Restore command

Use the **restore** command to restore all or part of one or more SQL databases.

Use this command to restore all or part of one or more SQL databases from Tivoli Storage Manager storage to a SQL server.

Considerations:

- Make sure to review “VSS Restore command-line considerations” on page 169 before attempting any type of VSS Restore operation.
- You cannot restore SQL databases currently in use. By placing SQL databases to be restored in single-user mode, you can avoid attempting such restores. If you are restoring the master database, you *must* start the SQL server in single-user mode by using the -m SQL SERVER startup option. For user mode and master database details, refer to “Setting user mode” on page 98 and “Restoring the master database” on page 87.

Note:

1. The single user of the SQL databases or server must be the same user that Data Protection for SQL uses to log on to the SQL server for the restore.
 2. SQL Enterprise Manager, SQL Server Application Client, and other SQL Server services can be users of databases and the SQL server.
- The user used by Data Protection for SQL to log on to the SQL server must have the SQL Server SYSADMIN fixed server role.
 - You can use the TRANSACT-SQL database consistency checker statement DBCC CHECKDB ('DBNAME') to verify the integrity of the restored SQL databases.

Note: During SQL database restore processing, the SQL Server prepares the database files after first restoring a minimal amount of metadata. For large SQL databases, the preparation of the database files can be time consuming. To prevent a restore operation from ending prematurely, specify a value of at least 10000 in the *commtimeout* option. If the restore operation is performed in a LAN free environment, this value must be specified for the Storage Agent.

Date and time recovery (Legacy only)

The *restoredate* and *restorettime* 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 *restorettime* is restored.
 - All logs up to the first log backed up after the specified *restoredate* and *restorettime* is restored.
 - Recovery up to the specified *restoredate* and *restorettime* (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 *restorettime* is restored.
 - The most recent differential backup (if any exists) prior to the specified *restoredate* and *restorettime* 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 *restorettime* is restored.
 - The most recent differential backup prior to the specified *restoredate* and *restorettime* is restored.
 - All log backups after the differential and up to the first log backed up after the *restoredate* and *restorettime* is restored.
 - Recovery up to the specified *restoredate* and *restorettime* (using *stopat*) is completed.

VSS Restore command-line considerations

Be aware of 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):

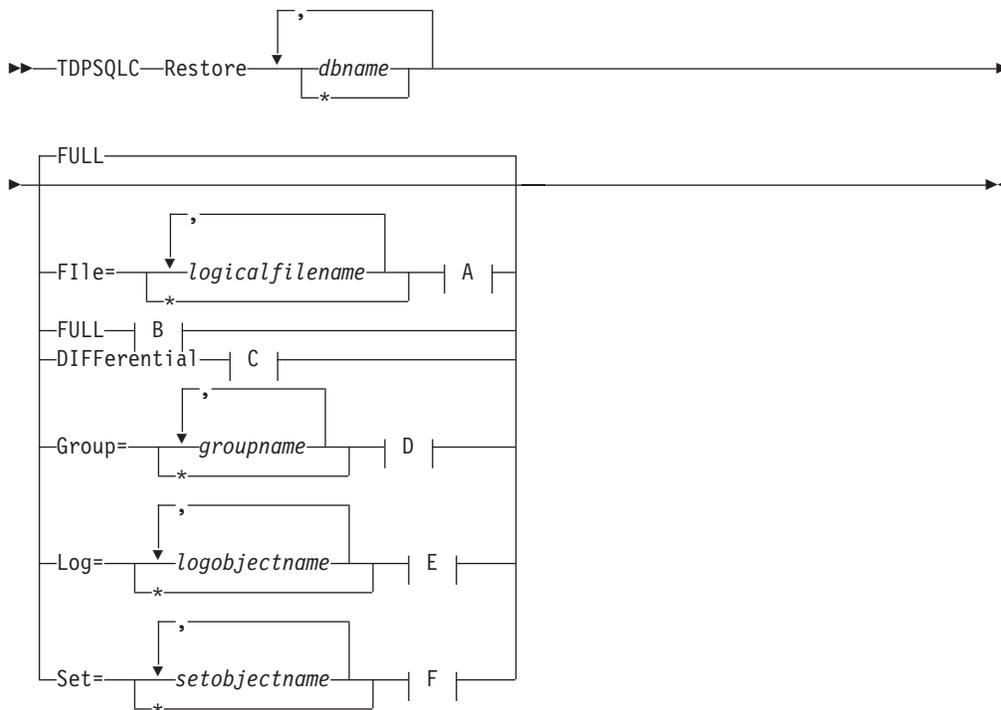
- If you plan to perform a VSS Restore of the master database, see "Restoring the master database" on page 87 for step-by-step instructions.
- A VSS Instant Restore overwrites the entire contents of the source volumes. However, you can avoid overwriting the source volumes by specifying */instantrestore=no*. 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. It is recommended that the source volume contain only the SQL database.
- Be aware that when a VSS restore from local shadow volumes is performed, the bytes transferred will display "0". That is because no data ("0") is restored from the Tivoli Storage Manager server.
- In order to perform a VSS Instant Restore with versions of the IBM Tivoli Storage Manager Client earlier than 6.1.0, the IBM 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.

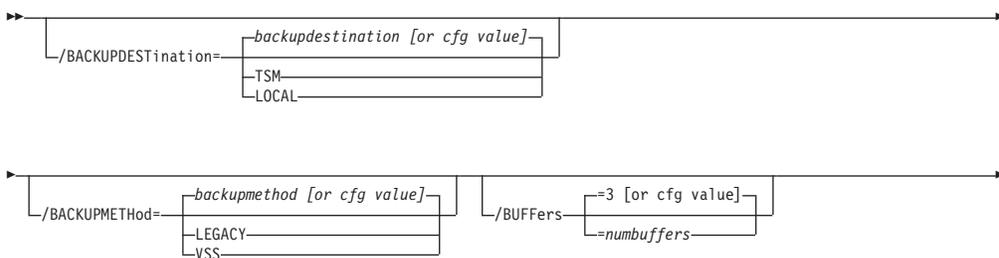
Syntax

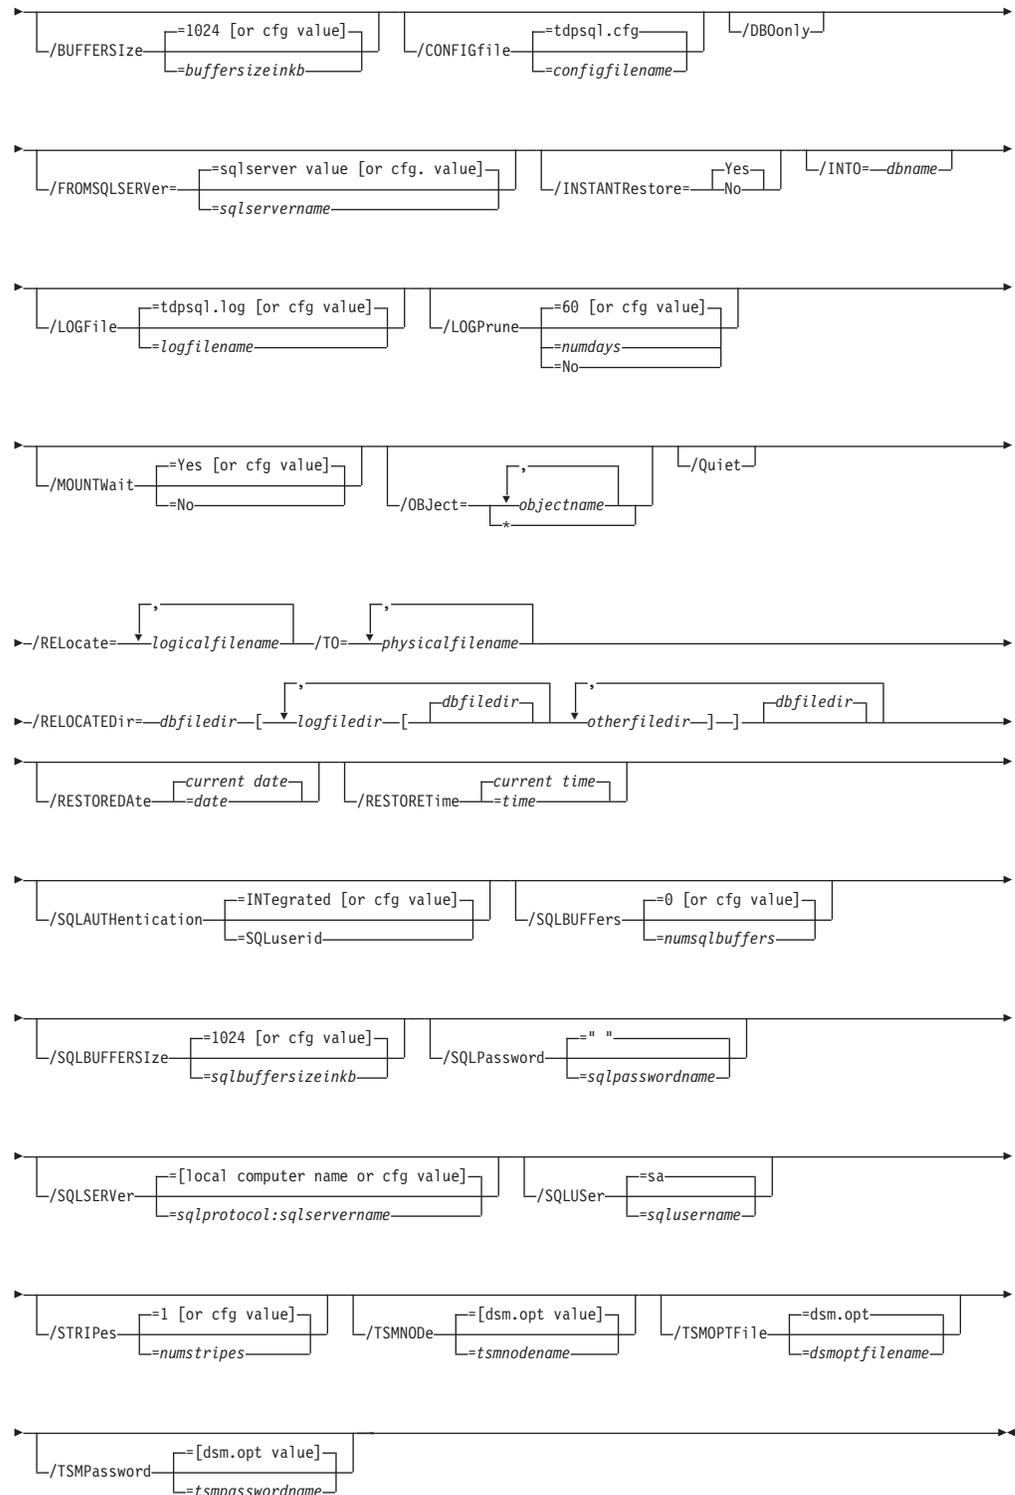


The syntax diagrams of the backup object type options corresponding to the letters above (A,B,C,D,E,F) follow the Optional Parameters below.

For a description of the **restore** positional parameters, see “Restore positional parameters” on page 173.

Restore Optional Parameters:





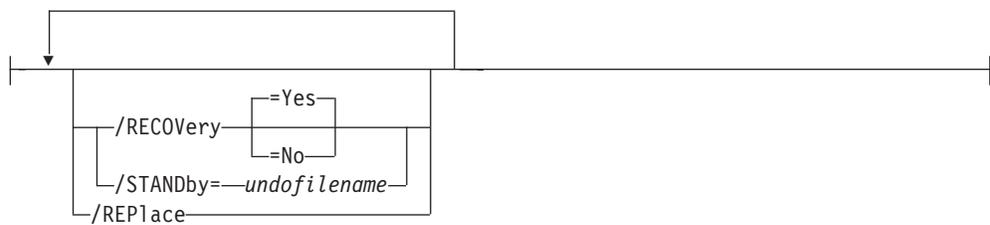
A Restore File Options:



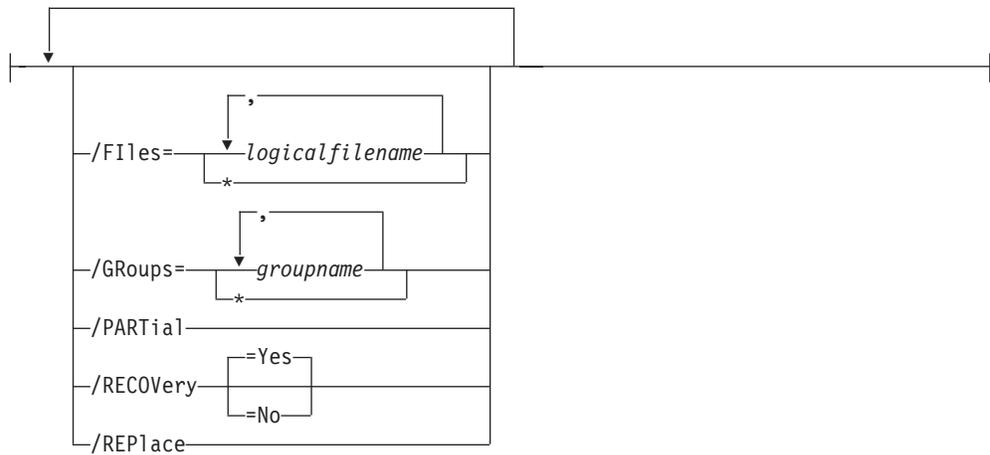
B Restore Full Options:



B1 Restore Full Options 1:



B2 Restore Full Options 2:



C Restore Diff Options:



- Use * as a wildcard character in *logicalfilename* to replace zero or more characters for each occurrence. Specifying only the wildcard character indicates all logical files in the SQL server database. Since each logical file backed up creates a separate backup object on the Tivoli Storage Manager server, specifying only the wildcard character results in a separate backup object for each logical file in the SQL server database.
- If *logicalfilename* includes spaces or special characters, enclose it in double quotes.
- The *logicalfilename* variable is case-sensitive.
- You cannot specify the */recovery* parameter with **restore file** operations.

FULL This option restores all full database backup objects for the SQL databases you specify.

Differential

A **differential** database backup contains only the parts of a SQL server database changed since the latest full backup plus enough of the SQL database's transaction log to make a restore consistent. As such, a differential backup usually takes up less space than a full backup. Use this option so that all individual log backups since the last full database backup do not need to be applied. This option saves time during a restore by replacing the restore of a number of transaction log backups.

Group=* | groupname,...

This option restores all group database backup objects for the SQL databases you specify. The *groupname* variable specifies the names of the SQL server database file groups 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 file groups in the SQL server database.
- If the *groupname* variable includes spaces or special characters, enclose it in double quotes.
- The *groupname* variable is case-sensitive.
- You cannot specify the */recovery* parameter with **restore group** operations.

Log or Log=* | logobjectname,...

This option restores all log database backup objects for the SQL databases you specify. The **log** parameter takes the wildcard or *logobjectname* value. The *logobjectname* variable specifies the log backup objects to restore. Use * as a wildcard character in *logobjectname* to replace zero or more characters for each occurrence. Specifying only the wildcard character indicates all log backup objects for the SQL databases. You can specify this parameter more than once per command invocation.

Set or Set=* | setobjectname,...

This option restores all set database backup objects for the SQL databases you specify. The **set** parameter takes the wildcard or *setobjectname* value. The *setobjectname* variable specifies the set backup objects to restore. Use * as a wildcard character in *setobjectname* to replace zero or more characters for each occurrence. Specifying only the wildcard character indicates all set backup objects for the SQL databases.

Considerations:

- You can specify this parameter more than once per command invocation.
- You cannot specify the */recovery* parameter with **restore set** operations.

Restore optional parameters

Optional parameters follow the **restore** command and positional parameters.

The following are detailed descriptions of each of the optional parameters:

/BACKUPDESTination=TSM|LOCAL

Use the */backupdestination* parameter to specify the location from where the backup is to be restored. The default is the value (if present) specified in the Data Protection for SQL preferences file (tdpsql.cfg). If no value is present, the backup is restored from Tivoli Storage Manager server storage.

You can specify:

TSM The backup is restored from Tivoli Storage Manager server storage. This is the default if no value is specified in the Data Protection for SQL preferences file (tdpsql.cfg).

LOCAL

The backup is restored from the local shadow volumes.

/BACKUPMETHod=LEGACY|VSS

Use the */backupmethod* parameter to specify the manner in which the restore is performed. The default is the value (if present) specified in the Data Protection for SQL preferences file (tdpsql.cfg). If no value is present, the backup is restored with the legacy API.

You can specify:

LEGACY

The restore is performed with the legacy API. This is the default if no value is specified in the Data Protection for SQL preferences file (tdpsql.cfg).

VSS The restore is performed with VSS.

/BUFFers=numbuffers

The */buffers* parameter specifies the number of data buffers used for each data stripe to transfer data between Data Protection for SQL 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 configuration file. This is initially 3.
- If you specify */buffers*, its value is used instead of the value stored in the Data Protection for SQL configuration file. Specifying this parameter does not change the value in the configuration file.
- 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 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 configuration file. This is initially 1024.
- If you specify **/bufferize**, its value is used instead of the value stored in the Data Protection for SQL configuration file. Specifying this parameter does not change the value in the configuration file.
- If you specify **/bufferize** but not *bufferizeinkb*, the default value 1024 is used.

/CONFIGfile=*configfilename*

The **/configfile** parameter specifies the name of the Data Protection for SQL configuration file, which contains the values for the Data Protection for SQL configurable options. See “Set command” on page 222 for details on the contents of the file.

Considerations:

- *configfilename* can include a fully qualified path. If *configfilename* does not include a path, it uses the directory where Data Protection for SQL 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.

/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 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 configuration file.

Considerations:

- If the two SQL server names are different, you must use this parameter even if **/fromsqlserver** was a non-clustered default instance.
- After you restore a SQL database to a different SQL server, the logins of the SQL database may not match the logins for the different SQL server. If appropriate, you can use the SQL stored procedure `SP_CHANGE_USERS_LOGIN` to find and correct such SQL login mismatches.

/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 resides on local shadow volumes. Note that a SAN Volume Controller, Storwize V7000, DS6000, DS8000, or XIV storage subsystem is required to perform VSS Instant Restores.

You can specify:

- Yes** Use volume level snapshot restore for a VSS Backup that resides on local shadow volumes if the backup exists on volumes that support it. This is the default.
- No** Use file-level copy to restore the files from a VSS Backup that resides on local shadow volumes. Note that bypassing volume-level copy means that SQL database files, log files, and the checkpoint file are the only data overwritten on the source volumes.

When performing VSS Instant Restores, a best practice is to make sure that any previous background copies (that involve the volumes being restored) are completed prior to initiating the VSS Instant Restore. However, this check is not necessary for XIV, SAN Volume Controller, or Storwize V7000 with space-efficient target volumes.

/INTO=dbname

For **restore** operations, **/into** specifies the SQL server database that you want a backup object 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.

Important: The **/into** parameter is only available when the restored database is relocated with the **/relocate /to** or **/relocatedir** parameter.

Considerations:

- When you specify **/into**, wildcards (*) may not appear in either the command **dbname** variable or the **/into dbname** variable.
- There must be exactly one item in the **/into dbname** variable list as well as in the command **dbname** list.
- The SQL server database that you want a backup object restored into must exist for the restore to be successful.

/LOGFile=logfilename

The **/logfile** parameter specifies the name of the activity log that is generated by Data Protection for SQL. This activity log records significant events such as completed commands and error messages. The Data Protection for SQL activity log is distinct from the SQL Server error log. The **logfilename** variable identifies the name to be used for the activity log generated by Data Protection for SQL.

Considerations:

- If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file.
- The file name can include a fully-qualified path; however, if you specify no path, the file is written to the directory where Data Protection for SQL is installed.

- You cannot turn Data Protection for SQL activity logging off. If you do not specify **/logfile**, log records are written to the default log file. The default log file is *tdpsql.log*.
- When using multiple simultaneous instances of Data Protection for SQL to perform operations, use the **/logfile** parameter to specify a different log file for each instance used. This directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPrune=numdays | No

The **/logprune** parameter prunes the Data Protection for SQL activity log and specifies how many days of entries are saved. By default, log pruning is enabled and performed once each day Data Protection for SQL is executed; however, this option allows you to disable log pruning or explicitly request a prune of the log for one command run even if the log file has already been pruned for the day. The *numdays* variable represents the number of days to save log entries. By default, 60 days of log entries are saved in the prune process.

Considerations:

- If you specify *numdays*, it can range from 0 to 9999. A value of 0 deletes all entries in the Data Protection for SQL activity log file except for the current command entries.
- If you specify **no**, the log file is not pruned during this command.
- If you do not specify **/logprune**, the default value is that specified by the logprune configurable option in the Data Protection for SQL configuration file. This is initially 60.
- If you specify **/logprune**, its value is used instead of the value stored in the Data Protection for SQL configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/logprune** without specifying *numdays* or **no**; in this case, the default 60 is used.
- Changes to the value of the **timeformat** or **dateformat** parameter can result in an undesired pruning of the &agentname; log file. If you are running a command that may prune the log file and the value of the **timeformat** or **dateformat** parameter has changed, perform one of the following to prevent undesired pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/logfile** parameter or **logfile** setting.

/MOUNTWait=Yes | No

If the Tivoli Storage Manager server is configured to store backup data on removable media such as tapes or optical media, it is possible that the Tivoli Storage Manager server may indicate to Data Protection for SQL that it is waiting for a required storage volume to be mounted. If that occurs, this option allows you to specify whether Data Protection for SQL **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 cannot complete waiting until the initial media for all stripes are available, although Data Protection for SQL 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 command from a prolonged wait. This can be done *only* by terminating the Data Protection for SQL 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 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 terminates the command with an error message. This is also true if the management class for meta objects requires removable media. For **backup**, since the meta objects are not created until after the data objects are complete, the command termination does not occur until after all of the database data is transferred.
- If you do not specify **/mountwait** with **backup** or **restore**, the default value is that specified in the mountwait configurable option in the Data Protection for SQL configuration file. This is initially **yes**. Specifying this parameter does not change the value in the configuration file.
- If you specify **/mountwait** but neither **yes** nor **no**, the default **yes** is used.
- If you do not specify **/mountwait** with a **query TSM /fileinfo** request, the default value **no** is used.

/OBJECT=* *objectname*,...

For **restore** and **inactivate** operations, **/object** specifies that only particular backup objects for the specified SQL databases and backup object type (if specified) be restored or inactivated. For **query** operations, **/object** includes particular objects and object types in the display. The *objectname* variable specifies the names of the backup objects you want to restore or inactivate. The object name uniquely identifies each backup object and is created by Data Protection for SQL. Use **query** to view the names of backup objects.

Considerations:

- If you do not specify restore, only the *active* backup object is included in the restore.
- You can use * as a wildcard character in *objectname* to replace zero or more characters for each occurrence. Specifying only the wildcard character indicates all backup objects of the specified SQL databases and backup object type.

/PARTIAL

The **/partial** parameter restores only part of a SQL database. You can perform partial restores only on **full** database backup objects. The primary purpose of a partial restore is to retrieve lost or damaged data. A partial restore creates a subset of the SQL database. After the partial restore, differential database restores and transaction log restores can return the subset SQL database to a point where the required data exists or is undamaged. You can then copy the required data from the subset SQL

database to the original SQL database. You can also use partial restores whenever you need a subset of a SQL database, such as for development or reporting purposes.

Note: A partial restore always restores the entire backup object from the Tivoli Storage Manager server although only a portion of the restored object may be used to complete a recovery. The statistics displayed reflect the amount of data restored from the Tivoli Storage Manager server only, not the amount of data used by the SQL Server for database recovery.

Considerations:

- You can specify the content of a partial restore with the **files** or **groups** parameters.
 - You can restore only complete SQL groups, even if you did not specify all SQL files in a SQL group with the **files** option.
 - The primary group is always included.
 - SQL groups not restored are marked offline and are not accessible.
- If you are restoring the subset SQL database to a location where it was backed up, you must use the **/relocate** and **/to** parameters.
- The MMC GUI does *not* support the **/relocate** and **/to** parameters. You must use the command line interface when performing a partial restore that requires these parameters.
- You can specify the **/recovery** parameter with **/partial**.

/Quiet The **/quiet** parameter omits displaying status information from the command. However, the information is appended to the Data Protection for SQL activity log.

/RECOVery=Yes | No

For **restore** operations, **/recovery** specifies whether or not you want to make additional restores to a SQL database that is not on a standby SQL server. A restored database cannot be used until the **/recovery=yes** parameter is administered to the database. You can specify:

Yes (default)

Whenever you make a sequence of restores to a SQL database and the current restore is the final restore in the sequence, or is the only restore to a SQL database. This informs the SQL server the restore is complete and ready for uncompleted transactions to be rolled back.

No Whenever you make a sequence of restores to a SQL database and the current restore is not the final restore in the sequence. Issue **/recovery=no** for all **restore** commands except the last one.

Considerations:

- Once 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 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 forces such restores to **/recovery=no**.
- Not specifying this option automatically rolls back uncompleted transactions for the database.

- When you specify **yes** and you are restoring several restore objects for the same database, only the final restore object for the database uses **recovery=yes**; all others use **recovery=no**. This allows you to specify a list of logs without having to specify the final log in a separate command.

The following is a sample scenario:

Sequence of Restores	Specify
full database	no
differential database	no
transaction log backup object	no
transaction log backup object	yes

Note:

1. Data Protection for SQL sorts the restore objects by database name, and, within database name, by backup time stamp from earliest to latest. A **query TSM** command will also display this order.
2. If a restore object fails, then all subsequent restore objects for that database in a single restore command are skipped. This is true no matter what the **/recovery** or **/standby** settings are.

/RELocate=*logicalfilename*,... **/TO**=*physicalfilename*,...

For **restore** operations, the **/relocate** and **/to** parameters as a pair specify the new location of a SQL database file. You must use this parameter for every SQL database file that you are not restoring to its original drive, complete path, and file name. The *logicalfilename* variable specifies the logical file name of the SQL database file you want to relocate. The *physicalfilename* variable specifies the new physical Windows file name where you want to relocate the SQL database file. This parameter is available when restoring Legacy backups only.

Considerations:

- You cannot specify more than one database name as the value for the restore command when specifying **/relocate**.
- **/relocate** and **/to** can each take a list of values and can be specified more than once. However, as a pair, **/relocate** and **/to** must take the same number of values, and the values must be paired in order of appearance. For example,
`/relocate=a,b,c /to=a1,b1,c1`

is valid, but not
`/relocate=a,b,c /to=b1,a1`
- The MMC GUI does *not* support the **/relocate** and **/to** parameters. You must use the command line interface when performing a partial restore that requires these parameters.
- You can use the **query** command with the **/fileinfo** parameter to determine the logical file names and physical file names in the backup object.
- If either *logicalfilename* or *physicalfilename* includes spaces, you must enclose it in double quotes.
- For *physicalfilename*, include the complete drive, path, and file name of the new file.

- The drive and path of the new physical file name must exist, but if the file does not yet exist, SQL Server will create it. Additionally, if the file does exist, you may be required to use the **/replace** parameter.
- The wildcard (*) is not allowed in the values for either **/relocate** or **/to**.

/RELOCATEDir=*dbfiledir* [*logfiledir* [*otherfiledir*]]

The **/relocatedir** parameter specifies the new destination locations in which to restore the backed up SQL databases, logs, and SQL Server full-text index files. FILESTREAM files are included for SQL Server 2008 and SQL Server 2008 R2. This parameter is available when restoring VSS Backups or Legacy backups.

The *dbfiledir* variable specifies the directory location of the SQL database you want to relocate. Note that if the *logfiledir* and/or *otherfiledir* variables are not specified, the logs and SQL Server full-text index files are restored to the directory specified by *dbfiledir*.

The *logfiledir* variable specifies the directory location of the SQL log files you want to relocate. Note that if the *logfiledir* variable is not specified, the SQL log files are restored to the directory specified by *dbfiledir*.

The *otherfiledir* variable specifies the directory location of the SQL Server full-text index files and FILESTREAM files (SQL Server 2008 and SQL Server 2008 R2) are included for you want to relocate. Note that if the *otherfiledir* variable is not specified, the SQL Server full-text index files and FILESTREAM files (SQL Server 2008 and SQL Server 2008 R2) are restored to the directory specified by *dbfiledir*.

/REPlace

For **restore** operations, the **/replace** parameter specifies that you want existing SQL files to be overwritten when they otherwise would not be. You may have to use this parameter in the following instances:

- You are performing a **full** database restore, and one of the following is true:
 - You are using the **/into** parameter, and the **/into** database already exists on the SQL server.
 - The database already exists on the SQL server, and one of the following is also true:
 - The number of SQL files in the existing database differs from the number of SQL files in the full database backup object.
 - The names of one or more SQL files in the existing database are not the names of any of the SQL files in the full database backup object.
- You are performing a **file**, **group**, or **set** restore, and one or more of the SQL files already exist.

/RESTOREDate=*date*

The **/restoredate** parameter specifies a date to which the database identified by *dbname* is to be recovered. The date value must be specified in the same date format defined in the Data Protection for SQL preferences file. If **/restoredate** is not specified but **/restoretime** is specified, the **/restoredate** value is the current date. The **/restoredate** parameter is only available for Legacy restore operations. It can only be specified when restoring a full database backup. The **/restoretime** parameter cannot be used to restore file, group, and set backups.

/RESTORETime=*time*

The **/restoretime** parameter specifies the time of day to which the database

identified by *dbname* is to be recovered. The time value must be specified in the same time format defined in the Data Protection for SQL preferences file. If **/restoretime** is not specified but **/restoredate** is specified, the **/restoretime** is the current time. The **/restoretime** parameter is only available for Legacy restore operations. It can only be specified when restoring a full database backup. The **/restoretime** parameter cannot be used to restore file, group, and set backups.

/SQLAUTHentication=INTEgrated | SQLuserid

This parameter specifies the authorization mode used when logging on to the SQL server. The **integrated** value specifies Windows authentication. The user id you use to log on to Windows is the same id you will use to log on to the SQL server. This is the default value. Use the **sqluserid** value to specify SQL Server user id authorization. The user id specified by the **/sqluserid** parameter is the id you will use to log on to the SQL server. Any SQL user id must have the SQL Server SYSADMIN fixed server role.

/SQLBUFFers=numsqbuffers

The **/sqlbuffers** parameter specifies the total number of data buffers SQL Server uses to transfer data between SQL Server and Data Protection for SQL. The *numsqbuffers* variable refers to the number of data buffers to use. The number can range from 0 to 999. The initial value is 0. When **/sqlbuffers** is set to 0, SQL determines how many buffers should be used.

Considerations:

- The default value is the value specified by the SQL buffers configurable option in the Data Protection for SQL 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 configuration file. Specifying this parameter does not change the value in the configuration file.
- If you specify **/sqlbuffers** but not *numsqbuffers*, 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. The *sqlbuffersizeinkb* variable refers to the size of data buffers in kilobytes. The number can range from 64 to 4096. The default is 1024.

Considerations:

- The default value is the value specified by the SQL buffers configurable option in the Data Protection for SQL 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 configuration file. Specifying this parameter does not change the value in the configuration file.
- If you specify **/sqlbuffersize** but not *sqlbuffersizeinkb*, the default value 1024 is used.

/SQLPassword=sqlpasswordname

This parameter specifies the SQL password that Data Protection for SQL uses to log on to the SQL server that objects are backed up from or restored to.

Considerations:

- Using this parameter means that you are using SQL Server authentication. The SQL Server and the SQL user id for this password must both be configured for SQL Server authentication.
- If you do not specify **/sqlpassword**, the default value is blank (" ").
- If you specify **/sqlpassword** but not *sqlpasswordname*, the default is also blank (" ").

Note: This parameter is ignored if you use the **/sqlauth=integrated** parameter with it.

/SQLSERVER=sqlprotocol:sqlservername

The **/sqlserver** parameter specifies the SQL server that Data Protection for SQL logs on to. For **restore** operations, this is the SQL server that backup objects are restored to. However, if the backup objects were created from a different SQL server name, you must use the **/fromsqlserver** parameter. Use **/sqlserver** for the **query SQL** and **backup** commands, but use **/fromsqlserver** for the **query TSM** and **inactivate** commands. The *sqlprotocol* variable specifies the communication protocol to use. You can specify one of the following protocols:

- *lpc*: Use Shared Memory protocol.
- *np*: Use Named Pipes protocol.
- *tcp*: Use Transmission Control protocol.
- *via*: Use Virtual Interface Architecture protocol.

If no protocol is specified, Data Protection for SQL logs on to the SQL server according to the first protocol that becomes available.

Considerations:

- The default value is the value specified by the SQL server configurable option in the Data Protection for SQL configuration file. This is initially the local computer name.
- If you specify **/sqlserver** but not *sqlservername*, the local computer name is used.
- The following two shortcuts are accepted as the local computer name: . (local) These are a period or the word *local* within parentheses.
- 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 reside on. The TCP/IP host name may not always be the same.

instancename

The name given to the named instance of SQL Server specified during installation of the instance.

virtualservername

The name given to the clustered SQL Server specified during clustering service setup. This is not the cluster or node name.

/SQLUSER=sqlusername

The **/sqluser** parameter specifies the name that Data Protection for SQL uses to log on to the SQL server.

Considerations:

- Using this parameter means that you are using SQL Server authentication. The SQL Server and the SQL user id for this password must both be configured for SQL Server authentication.
- The SQL user id must have the SQL server SYSADMIN fixed server role.
- If you do not specify **/sqluser**, the default is **sa**.
- If you specify **/sqluser** but not *sqlusername*, the default is also **sa**.

Note: This parameter is ignored if you use the **/sqlauth=integrated** parameter with it.

/STANDBY=undofilename

Specifies that the restore is to a standby SQL server, and specifies the name of an undo file.

Considerations:

- You cannot specify more than one database name as the restore command value.
- A standby SQL server can be in read-only mode between restores and can accept additional restores to its databases.
- You can use the same undo file for a database for each restore to the database, but you cannot use a single undo file for more than one database.
- The *undofilename* variable can include a fully qualified path. However, if a fully qualified path is not specified, the undo file is created in the directory specified by the %TEMP% environment variable.
- If *undofilename* includes spaces, you must enclose it in double quotes.
- If the specified undo file does not exist, SQL server creates it. If the file exists but was not used for the same SQL database, SQL Server overwrites it.
- If you specify neither **/recovery** nor **/standby**, the default is **/recovery=yes**.

/STOPAT=datetime

For **restore** operations, **/stopat** specifies the point in time that you restore a SQL database to. Only transaction logs written before the point in time are applied to the SQL database. The *datetime* variable specifies both the date and time separated by a space. Use any valid date and time format accepted by SQL Server.

Considerations:

- This parameter applies only to transaction log restores, but the base restore that the transaction logs apply to must have been a **full** database restore. You cannot restore **file**, **group**, and **set** restores to a point in time.
- You cannot also specify **/recovery=no** or **/standby** with the **/stopat** parameter.
- Because *datetime* includes a space, you must enclose it in double quotes.
- If the restore operation with the **/stopat** parameter does not encounter a transaction in the restored transaction log that has a time stamp equal to or greater than the specified point in time, the SQL database is left in an unrecovered state, even if you also specify **/recovery=yes**.

/STOPATMark=markname [/AFTER=datetime]

The **/stopatmark** parameter specifies a named point in time to restore a database to. This can be after a specified point in time if you specify the **/after** option. Only transaction log records written *up to* and including the named transaction (which may be found at or after the specified point in time) are applied to the SQL database. The *markname* variable specifies the name of a SQL transaction. The SQL transaction may be a local transaction or a distributed transaction. If it is a distributed transaction name, the named mark exists in the transaction log of each SQL database participating in the distributed transaction.

Note: *markname* is the transaction name, not the description that follows the MARK keyword in a SQL BEGIN TRANSACTION or BEGIN DISTRIBUTED TRANSACTION statement.

The *datetime* variable specifies both the date and time separated by a space. Use any valid date and time format accepted by SQL Server.

Considerations:

- This parameter applies only to transaction log restores. The base restore that the transaction logs apply to must have been a **full** database restore. You cannot restore **file**, **group**, and **set** restores to a mark.
- You can use the same named mark for several SQL transactions.
- If you do not specify **/after**, the restore stops *at* the first mark it encounters with the specified name.
- If you specify **/after**, the restore stops *at* the first mark it encounters with the specified name *after* the specified date and time.
- If *markname* includes spaces, you must enclose it in double quotes.
- You can not use a Data Protection for SQLrestore command with **/stopatmark** and also specify **/recovery=no** or **/standby**.
- If the restore operation with **/stopatmark** does not encounter a transaction in the restored transaction log to stop at, the SQL database is left in an unrecovered state, even if you also specify **/recovery=yes**.

/STOPBEFOREMark=markname [/AFTER=datetime]

This parameter specifies a named point in time to restore a database to. This can be after a specified point in time if you specify the **/after** option. Only transaction log records written *before* and not including the named transaction (which may be found at or after the specified point in time) are applied to the SQL database. The *markname* variable specifies the name of a SQL transaction. The SQL transaction may be a local transaction or a

distributed transaction. If it is a distributed transaction name, the named mark exists in the transaction log of each SQL database participating in the distributed transaction.

Note: *markname* is the transaction name, not the description that follows the MARK keyword in a SQL BEGIN TRANSACTION or BEGIN DISTRIBUTED TRANSACTION statement.

The *datetime* variable specifies both the date and time separated by a space. Use any valid date and time format accepted by SQL Server.

Considerations:

- This parameter applies only to transaction log restores. The base restore that the transaction logs apply to must have been a **full** database restore. You cannot restore **file**, **group**, and **set** restores to a mark.
- You can use the same named mark for several SQL transactions.
- If you do not specify **/after**, the restore stops *before* the first mark it encounters with the specified name.
- If you specify **/after**, the restore stops *before* the first mark it encounters with the specified name, or *after* the specified date and time.
- If *markname* includes spaces, you must enclose it in double quotes.
- You can not use a Data Protection for SQL **restore** command with **/stopbeforemark** and also specify **/recovery=no** or **/standby**.
- If the restore operation with **/stopbeforemark** does not encounter a transaction in the restored transaction log to stop before, the SQL database is left in an unrecovered state, even if you also specify **/recovery=yes**.

/STRIPes=numstripes

The **/stripes** parameter specifies the number of data stripes to use in a backup or restore operation. The *numstripes* variable can range from 1 to 64.

Considerations:

- If you do not specify **/stripes**, the default value is that specified in the Data Protection for SQL configuration file. The initial value is 1. For **restore**, the value is the same as that used in the backup operation.
- If you specify **/stripes** but not *numstripes*, the stored value is used.
- You may use *up to* the number used to create the backup. You can determine the number of data stripes used to create a backup object with the Data Protection for SQL command: `query tsm dbname backup_object`
- You must use the MAXNUMMP parameter on a Tivoli Storage Manager REGISTER NODE or UPDATE NODE command to allow a node to use multiple sessions to store data on removable media (which requires you to allocate multiple mount points to that node). The MAXNUMMP value must be equal to or less than the maximum number of stripes you desire.
- When you use data striping, you should use Tivoli Storage Manager server file space collocation to try to keep each stripe on a different storage volume.
- The maximum number of data stripes you can use is one less than the value of 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 uses to log on to the Tivoli Storage Manager server. This identifies which Tivoli Storage Manager client is requesting services. You can also store the node name in the options file. The command line parameter overrides the value in the options file.

Considerations:

- You cannot use the **/tsmnode** parameter if **PASSWORDACCESS GENERATE** is specified in the Tivoli Storage Manager options file. You must specify the **nodename** in the options file. Otherwise, you can change **PASSWORDACCESS** to **PROMPT** to utilize the **/tsmnode** parameter. For details about the Tivoli Storage Manager options file, see the reference manual *IBM Tivoli Storage Manager for Windows Backup-Archive Client Installation and User's Guide*.
- If you do not specify **/tsmnode**, the default value is that specified by the **nodename** option in the Tivoli Storage Manager options file. Specifying this parameter does not change the value in the options file.

/TSMOPTFile=*dsmoptfilename*

The **/tsmoptfile** parameter specifies the Tivoli Storage Manager options file to use. This is similar to selecting a Tivoli Storage Manager server from the server list in the GUI. The Tivoli Storage Manager options file contains the configuration values for the Tivoli Storage Manager API. For details about the Tivoli Storage Manager options file, see the reference manual *IBM Tivoli Storage Manager for Windows Backup-Archive Client Installation and User's Guide*.

Considerations:

- The *tsmoptfilename* variable can include a fully qualified path. If you do not include a path, the directory where Data Protection for SQL is installed is used.
- If *tsmoptfilename* includes spaces, you must enclose it in double quotes.
- If you do not specify **/tsmoptfile**, the default value is *dsm.opt*.
- If you specify **/tsmoptfile** but not *tsmoptfilename*, the default is also *dsm.opt*.

/TSMPassword=*tsmpasswordname*

The **/tsmpassword** parameter specifies the Tivoli Storage Manager password that Data Protection for SQL uses to log on to the Tivoli Storage Manager server. This parameter and the option **PASSWORDACCESS** in the Tivoli Storage Manager options file interact in the following ways:

/tsmpassword	PASSWORDACCESS in Tivoli Storage Manager options file	Password already stored in registry?	Result
specified	<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

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 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011.
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 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011.
All rights reserved.

Starting Sql database restore...

Querying Tivoli Storage Manager server for a list of database backups,
please wait...

Beginning diff full 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 file group 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 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011.
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 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011.
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\20100718141546\00000700,
1 of 1, to database Test1
Full: 0   Read: 88489472   Written: 88489472   Rate: 8,125.58 Kb/Sec
Restore of Test1\20100718141546\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="07/01/2010 13:56:00"
```

Output:

```

IBM Tivoli Storage Manager for Databases
Data Protection for Microsoft SQL Server
Version 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011.
All rights reserved.

Starting Sql database restore...

Querying Tivoli Storage Manager server for a list of
database backups, please wait...

Beginning log restore of backup object Test1\20100701135511\
00000700,
1 of 4,to database Test1
Full: 0 Read: 214528 Written: 214528 Rate: 59.75 Kb/Sec
Restore of Test1\20100701135511\00000700 completed successfully.

Beginning log restore of backup object Test1\20100701135605\
00000700,
2 of 4,to database Test1
Full: 0 Read: 147968 Written: 147968 Rate: 32.15 Kb/Sec
Restore of Test1\20100701135605\00000700 completed successfully.

Beginning log restore of backup object Test1\20100701135712\
00000700,
3 of 4,to database Test1
Full: 0 Read: 0 Written: 0 Rate: 0.00 Kb/Sec
Restore of Test1\20100701135712\00000700 completed successfully.

Skipping Test1\20100701135817\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 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011.
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\20100701102947\
0000065C, 1 of 4, to database Testmark
Full: 0   Read: 159232   Written: 159232   Rate: 61.68 Kb/Sec
Restore of Testmark\20100701102947\0000065C completed successfully.

Beginning log restore of backup object Testmark\20100701103127\
000001DC, 2 of 4, to database Testmark
Full: 0   Read: 159232   Written: 159232   Rate: 34.51 Kb/Sec
Restore of Testmark\20100701103127\000001DC completed successfully.

Beginning log restore of backup object Testmark\20100701103325\
00000680, 3 of 4, to database Testmark
Full: 0   Read: 0   Written: 0   Rate: 0.00 Kb/Sec
Restore of Testmark\20100701103325\00000680 completed successfully.

Skipping Testmark\20100701103556\00000694
because of the preceding failure or point-in-time recovery.

Total database backups inspected:          4
Total database backups requested for restore: 4
Total database backups restored:          3
Total database skipped:                   4

Throughput rate:                           38.60 Kb/Sec
Total bytes transferred:                   318,464
LanFree bytes transferred:                 0
Elapsed processing time:                   8.06 Secs

```

Legacy Restore 7–Log (inactive object)

Legacy Restore 7 begins with a query to display both active and inactive log backup objects for database *Test1*.

Command:

```
tdpsqlc q tsm netapp_db2 log=* /all
```

Output:

```

IBM Tivoli Storage Manager for Databases
Data Protection for Microsoft SQL Server
Version 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011.
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 ..... 09/27/2011 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 ..... 20110927083628\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=20100622135511\00000700
```

Output:

```

IBM Tivoli Storage Manager for Databases
Data Protection for Microsoft SQL Server
Version 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011.
All rights reserved.

Starting Sql database restore...

Querying Tivoli Storage Manager server for a list of database
backups,please wait...

Beginning log restore of backup object Test1\20100622135511\
00000700,1 of 1,to database Test1
Full: 0 Read: 214528 Written: 214528 Rate: 29.47 Kb/Sec
Restore of Test1\20100622135511\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, file group *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 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011.
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 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011.
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 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011.
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 6, Release 3, Level 0.0  
(C) Copyright IBM Corporation 1998, 2011.  
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 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011.
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 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011.
All rights reserved.

Connecting to SQL Server, please wait...

Querying TSM Server for Backups ....

Starting Sql database restore...

Beginning VSS restore of 'svtodb'...

Preparing for restore of 'svtodb' 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

```

In order to restore and relocate the database *svtodb*, its logs, and its full-text index files into their own respective locations, the following command is issued:

```

tdpsqlc restore svtodb full /relocatedir=m:\svtodb,e:\svtodb,f:\svtodb
/backuperdestination=tsm /backupmethod=vss

```

The */relocatedir* values in this command are as follows:

- *m:\svtodb*: The directory where only the *svtodb* database is relocated.
- *e:\svtodb*: The directory where only the *svtodb* logs are relocated.
- *f:\svtodb*: The directory where only the *svtodb* full-text index files are relocated.

Restorefiles command

Use the **restorefiles** command to restore VSS-based backups on the Tivoli Storage Manager server (/BACKUPDESTINATION=TSM), or stored locally (/BACKUPDESTINATION=LOCAL).

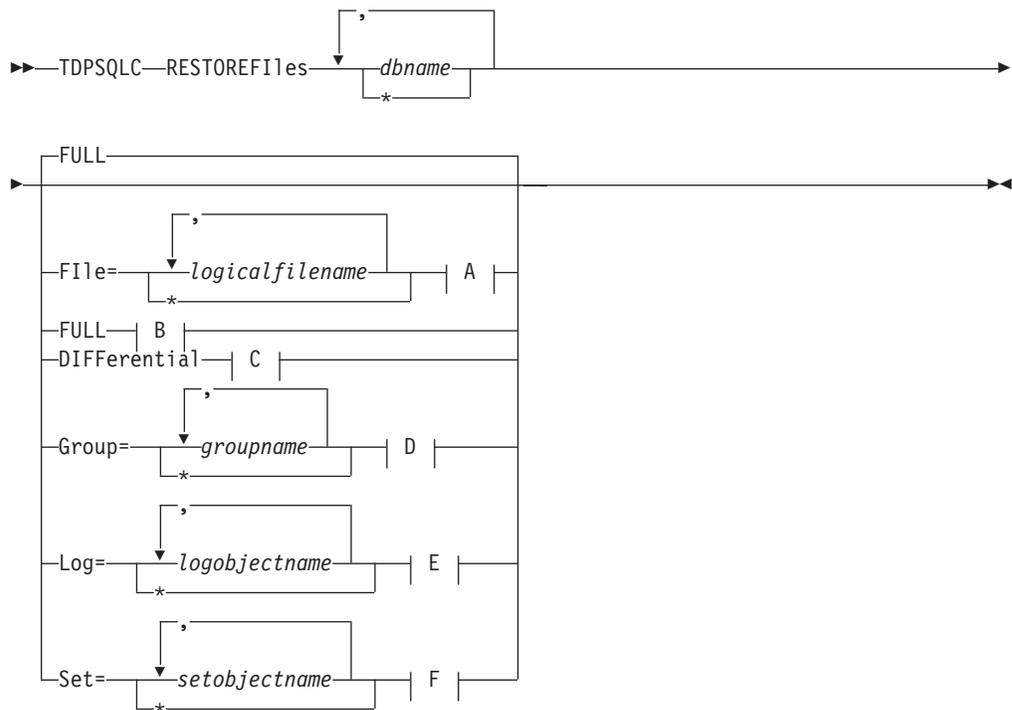
Consider the following information before using the **restorefiles** command.

- The **restorefiles** command restores .mdf, ldf, and other flat files from a specified Data Protection for SQL server, VSS-based backup into a specified directory.
- A destination directory can be specified as a directory on a fixed file system (for example C:\temp), or on a network share (for example \\server\dest) that is accessible to the Tivoli Storage Manager server Remote Agent (VSS Requestor)
- The **restorefiles** command does not restore the data to the SQL server.
- This command does not require the SQL Server to be installed on the machine where the **restorefiles** command is run. Files can be restored to another machine or directory on the same machine as the SQL Server.
- A restore continues until it completed, unless the destination volume does not have enough space to fulfill the restore operation.

- VSS-based backups that are located on the Tivoli Storage Manager server (/BACKUPDESTINATION=TSM) can be restored by using **restorefiles** on the same machine that performed the VSS-based backup, or by running the command on a machine that has the Data Protection for SQL client installed and configured for VSS operations.
- The directory specified in the **restorefiles** command has the VSS component name appended so that multiple databases can be restored to the same target directory.
- VSS-based backups that are stored on the local machine by using a persistent snapshot (/BACKUPDESTINATION=LOCAL), can be restored only by running the **restorefiles** command on the same machine that performed the VSS-based backup, and has access to the persistent snapshot.
- To run a full restore: `tdpsqlc restorefiles DB1 FULL relocatedir=d:\tempstore`
- Use /RELOCATEDIR to restore a database that currently exists to a different directory, even if your backup contains files that are located in different directories. Run the **restorefiles** command and specify just one restore destination directory. For example, issue `restorefiles db1 full /relocatedir=d:\temp` to place the files into the `d:\temp\db1*` directory.
- If you are in a non-clustered environment, you can restore only a local snapshot to the machine that generated the snapshot.
- If you are in a clustered environment, you can run a **restorefiles** command from any of the machines in the cluster.

Restorefiles syntax

Use the **restorefiles** command syntax diagram as a reference for available options and truncation requirements.



The syntax diagrams of the backup object type options corresponding to the letters above (A,B,C,D,E,F) follow the Optional Parameters for the **restorefiles** command "Restorefiles optional parameters."

Restorefiles positional parameters

Positional parameters immediately follow the **restorefiles** command and precede the optional parameters.

The following positional parameters specify the object to restore:

tdpsqlc restorefiles

* | **componentname1, ..., 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 **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 configuration file that contains the values for the Data Protection for SQL 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 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.

The *logfilename* variable identifies the name of the activity log file.

If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *logfile* variable can include a fully-qualified path. However, if no path is specified, the log file is written to the Data Protection for SQL installation directory.

If the *logfile* 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 to perform operations, use the **/logfile** parameter to specify a different log file for each instance used. This directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPrune=numdays | No

Use the **/logprune** parameter to disable log pruning or to explicitly request that the log be pruned for one command run. By default, log pruning is enabled and performed once per day. The *numdays* variable represents the number of days to save log entries. By default, **60** days of log entries are saved in the pruning process. You can use the MMC GUI or the **set** command to change the defaults so that log pruning is disabled, or so that more or less days of log entries are saved. If you use the command line, you can use the **/logprune** parameter to override these defaults. When the value of the **/logprune** variable *numdays* is a number in the range 0 to 9999, the log is pruned even if log pruning has already been performed for the day.

Changes to the value of the **timeformat** or **dateformat** parameter can result in the log file being pruned unintentionally. If the value of the **timeformat** or **dateformat** parameter has changed, prior to issuing a Data Protection for SQL 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

Use the **/mountwait** parameter to specify whether Data Protection for SQL 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.

Use the Data Protection for SQL **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 SQL Server full-text index files (SQL Server 2005).

The *dbfiledir* variable specifies the directory location of the SQL database you want to relocate. Note that if the *logfiledir* or *otherfiledir* variables are not specified, the logs and SQL Server full-text index files are restored to the directory specified by *dbfiledir*.

The *logfiledir* variable specifies the directory location of the SQL log files you want to relocate. Note that if the *logfiledir* variable is not specified, the SQL log files are restored to the directory specified by *dbfiledir*.

The *otherfiledir* variable specifies the directory location of the SQL Server full-text index files you want to relocate. Note that if the *otherfiledir* variable is not specified, the SQL Server full-text index files are restored to the directory specified by *dbfiledir*. The **restorefiles** operation creates a subdirectory under the root directory that contains the name of the database name. Restored files are placed in that subdirectory. If the **/relocatedir** parameter is not specified, the files will be restored into the directory where the **restorefiles** command is issued. For example, if Data Protection for SQL 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=tsmnode name

Use the *tsmnode name* variable to refer to the Tivoli Storage Manager node name that Data Protection for SQL 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 options file.

The file name can include a fully qualified path name. If no path is specified, the directory where Data Protection for SQL 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 uses to log on to the Tivoli Storage Manager server. If you specified **PASSWORDACCESS GENERATE** in the Data Protection for SQL 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 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 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
```

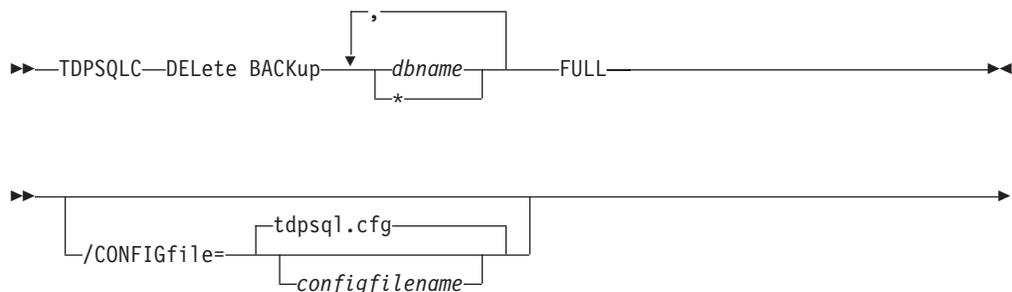
Delete Backup command

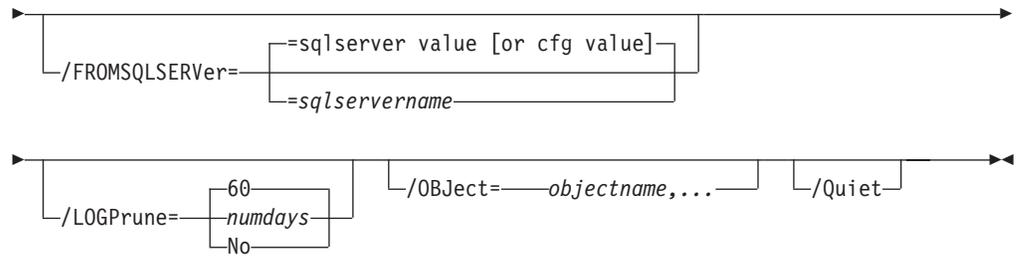
Use the **delete backup** command to delete a VSS Backup of a SQL Server database.

You must have local registry rights (for all versions of SQL Server) to perform a Data Protection for SQL delete backup.

Delete Backup syntax

Use the **delete backup** command syntax diagrams as a reference to view available options and truncation requirements.





Delete Backup positional parameters

Positional parameters immediately follow the **delete backup** command and precede the optional parameters.

The following positional parameters specify the backup to delete:

* | *dbname*

* Delete the active backups of all databases.

dbname

Delete a backup of the specified database. The active backup is deleted unless you specify a different backup with the **/object** optional parameter.

Multiple entries are separated by commas. If separated by commas, make sure there is no space between the comma and the database name. If any database name contains commas or blanks, enclose the database name in double quotation marks.

The following positional parameter specifies the type of delete backup to perform:

FULL Delete full type backups.

Attention: Be careful to delete only the backups that you want.

Delete Backup optional parameters

Optional parameters follow the **delete backup** command and positional parameters.

/BACKUPDESTINATION=TSM|LOCAL

Use the **/backupdestination** parameter to specify the location from where the backup is to be deleted. The default is the value (if present) specified in the Data Protection for SQL preferences file (`tdpsql.cfg`). If no value is present, the backup is deleted from Tivoli Storage Manager server storage.

You can specify:

TSM The backup is deleted from Tivoli Storage Manager server storage. This is the default if no value is specified in the Data Protection for SQL 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 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 installation directory is used. If the **/configfile** parameter is not specified, or if the *configfilename* variable is not specified, the default value is *tdpsql.cfg*.

If the *configfilename* variable includes spaces, enclose the entire **/configfile** parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\file.cfg"
```

See "Set positional parameters" on page 223 for descriptions of available configuration parameters.

/FROMSQLSERVER=server-name

Use the **/fromsqlserver** parameter to specify the name of the SQL Server where the original backup was performed. This parameter is necessary only when the name of the SQL server to delete from, as determined by the **/sqlserver** parameter, is different from the name of the SQL server that the backup objects were created from. The default value is the **/sqlserver** value or the value set in the Data Protection for SQL configuration file.

Considerations:

- If the two SQL server names are different, you must use this parameter even if **/fromsqlserver** was a non-clustered default instance.

/LOGFile=logfilename

Use the **/logfile** parameter to specify the name of the activity log file that is generated by Data Protection for SQL.

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 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 using multiple simultaneous instances of Data Protection for SQL to perform operations, use the **/logfile** parameter to specify a different log file for each instance used. This directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPrune=numdays | No

Use the **/logprune** parameter to disable log pruning or to explicitly request that the log be pruned for one command run. By default, log pruning is enabled and performed once per day. The *numdays* variable represents the number of days to save log entries. By default, **60** days of log entries are saved in the pruning process. You can use the MMC GUI or the **set** command to change the defaults so that log pruning is disabled, or so that more or less days of log entries are saved. If you use the command line, you can use the **/logprune** parameter to override these defaults. When the

value of the **/logprune** variable *numdays* is a number in the range 0 to 9999, the log is pruned even if log pruning has already been performed for the day.

Changes to the value of the **timeformat** or **dateformat** parameter can result in the log file being pruned unintentionally. If the value of the **timeformat** or **dateformat** parameter has changed, prior to issuing a Data Protection for SQL command that might prune the log file, perform one of the following actions to prevent the log file from being pruned:

- Make a copy of the existing log file.
- Specify a new log file with the **/logfile** parameter or **logfile** setting.

/OBJECT=objectname,...

Use the **/object** parameter to specify the names of backup objects you want to delete. The object name uniquely identifies each backup object and is created by Data Protection for SQL.

Use the Data Protection for SQL **query tsm * /all** command to view the names of all available backup objects. This parameter specifies that only particular backup objects for the specified SQL databases and backup object type be deleted. The *objectname* variable specifies the names of the backup objects you want to delete. The object name uniquely identifies each backup object and is created by Data Protection for SQL.

/Quiet This parameter prevents status information from being displayed. This does not affect the level of information written to the activity log.

/TSMNODE=tsmnode

Use the *tsmnode* variable to refer to the Tivoli Storage Manager node name that Data Protection for SQL 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 options file.

The file name can include a fully qualified path name. If no path is specified, the directory where Data Protection for SQL 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 uses to log on to the Tivoli Storage Manager server. If you specified **PASSWORDACCESS GENERATE** in the Data Protection for SQL 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 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 uses to log on to the Tivoli Storage Manager server can be up to 63 characters in length.

Delete Backup example

This output example provides a sample of the text, messages, and process status that displays when using the **delete backup** command.

In this example, the **tdpsqlc delete backup xivdb1 full** command deletes a full backup of database xivdb1. The following output is displayed:

```
Connecting to SQL Server, please wait...
Querying for Backups ....
Backup(s) to be deleted:
<xivdb1 : VSS : full : 02/10/2011 10:03:29>
VSS Delete backup operation completed with rc = 0
  Files Examined   : 1
  Files Completed  : 1
  Files Failed     : 0
  Total Bytes      : 0
```

Inactivate command (Legacy only)

Use the **inactivate** command to inactivate one or more active Legacy backup objects on the Tivoli Storage Manager server.

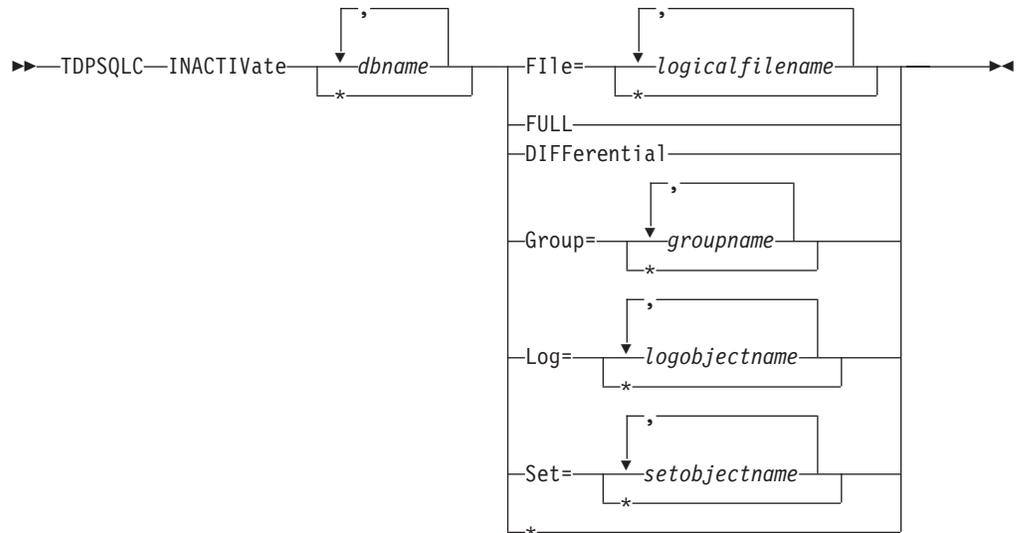
Most backup objects are automatically inactivated as part of the normally scheduled backup processing. For those occasions when that processing is not sufficient, you can use the **inactivate** command.

Tivoli Storage Manager server does not delete *active* backup objects from Tivoli Storage Manager managed storage; it will delete only *inactive* backup objects. Once a backup object becomes inactive, the expiration processing defined in the object's management class determines exactly when the backup object is deleted.

Inactivate syntax

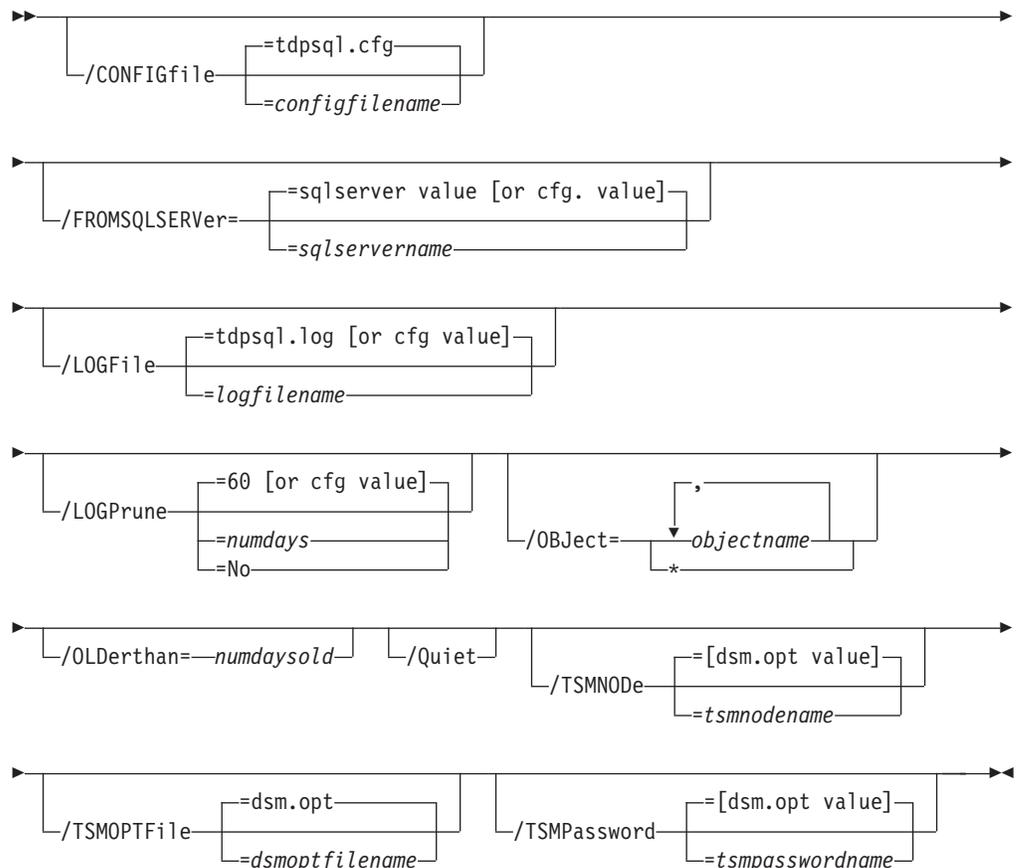
Use the **inactivate** command syntax diagrams as a reference to view available options and truncation requirements.

Syntax



For a description of the **inactivate** positional parameters, see “Inactivate positional parameters” on page 210.

Inactivate Optional Parameters:



Inactivate positional parameters

Positional parameters immediately follow the **inactivate** command and precede the optional parameters.

File=* | *logicalfilename*,...

This option inactivates only the active file backup objects for the SQL databases you specify. The *logicalfilename* variable specifies the names of the SQL server database logical files you want to inactivate.

Considerations:

- You can specify this parameter more than once per command invocation.
- Use * as a wildcard character in *logicalfilename* to replace zero or more characters for each occurrence. Specifying only the wildcard character indicates all logical files in the SQL server database.
- If *logicalfilename* includes spaces or special characters, enclose it in double quotes.
- The *logicalfilename* variable is case-sensitive.

FULL This option inactivates only the active full database backup objects for the SQL databases you specify. Each SQL database backed up creates a separate backup object on the Tivoli Storage Manager server. A new full database backup object inactivates all prior active backup objects for the same SQL database. This inactivation includes any active full backup object as well as any active file, group, set, differential, and log backup objects.

Differential

This option inactivates only the active differential database backup object. Because each SQL database backup creates a separate backup object on the Tivoli Storage Manager server, a new differential database backup object inactivates any active differential backup object for the same SQL database. Use this option so that all individual log backups since the last full database backup do not need to be applied.

Group=* | *groupname*,...

This option inactivates only the active group database backup object for the SQL database you specify. The *groupname* variable specifies the names of the SQL server database file groups you want to inactivate.

Considerations:

- You can specify this parameter more than once per command invocation.
- Use * as a wildcard character in the *groupname* variable to replace zero or more characters for each occurrence.
- Specifying only the wildcard character indicates all file groups in the SQL server database.
- If the *groupname* variable includes spaces or special characters, enclose it in double quotes.
- The *groupname* variable is case-sensitive.

Log or Log=* | *logobjectname*,...

This option inactivates only the active log database backup object for the SQL database you specify. This parameter takes the wildcard or *logobjectname* value. The *logobjectname* variable specifies the log backup objects to inactivate. Use * as a wildcard character in *logobjectname* to replace zero or more characters for each occurrence. Specifying only the wildcard character indicates all log backup objects for the SQL databases. You can specify this parameter more than once per command invocation.

Set or Set=* | setobjectname,...

This option inactivates only the active set database backup object for the SQL database you specify. This parameter takes the wildcard or *setobjectname* value. The *setobjectname* variable specifies the set backup objects to inactivate. Use * as a wildcard character in *setobjectname* to replace zero or more characters for each occurrence. Specifying only the wildcard character indicates all set backup objects for the SQL databases. You can specify this parameter more than once per command invocation.

Inactivate optional parameters

Optional parameters follow the **inactivate** command and positional parameters.

The following are detailed descriptions of each of the optional parameters:

/CONFIGfile=configfilename

The **/configfile** parameter specifies the name of the Data Protection for SQL configuration file, which contains the values for the Data Protection for SQL configurable options. See “Set command” on page 222 for details on the file's contents.

Considerations:

- *configfilename* can include a fully qualified path. If *configfilename* does not include a path, it uses the directory where Data Protection for SQL is installed.
- If *configfilename* includes spaces, place it in double quotes.
- If you do not specify **/configfile**, the default value is *tdpsql.cfg*.
- If you specify **/configfile** but not *configfilename*, the default value *tdpsql.cfg* is used.

/FROMSQLSERVER=sqlservername

The **/fromsqlserver** parameter specifies the SQL server that backup objects were backed up from. This parameter is necessary only when the name of the SQL server to inactivate from, as determined by the **/sqlserver** parameter, is different from the name of the SQL server that the backup objects were created from. The default value is the **/sqlserver** value or the value set in the Data Protection for SQL configuration file. If the two SQL server names are different, you must use this parameter even if **/fromsqlserver** was a non-clustered default instance.

/LOGFile=logfilename

The **/logfile** parameter specifies the name of the activity log that is generated by Data Protection for SQL. This activity log records significant events such as completed commands and error messages. The Data Protection for SQL activity log is distinct from the SQL Server error log. The *logfilename* variable identifies the name to be used for the activity log generated by Data Protection for SQL.

Considerations:

- If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file.
- The file name can include a fully-qualified path; however, if you specify no path, the file is written to the directory where Data Protection for SQL is installed.
- You cannot turn Data Protection for SQL activity logging off. If you do not specify **/logfile**, log records are written to the default log file. The default log file is *tdpsql.log*.

- When using multiple simultaneous instances of Data Protection for SQL to perform operations, use the **/logfile** parameter to specify a different log file for each instance used. This directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPrune=numdays | No

The **/logprune** parameter prunes the Data Protection for SQL activity log and specifies how many days of entries are saved. By default, log pruning is enabled and performed once each day Data Protection for SQL is executed; however, this option allows you to disable log pruning or explicitly request a prune of the log for one command run even if the log file has already been pruned for the day. The *numdays* variable represents the number of days to save log entries. By default, 60 days of log entries are saved in the prune process.

Considerations:

- If you specify *numdays*, it can range from 0 to 9999. A value of 0 deletes all entries in the Data Protection for SQL activity log file except for the current command entries.
- If you specify **no**, the log file is not pruned during this command.
- If you do not specify **/logprune**, the default value is that specified by the **logprune** configurable option in the Data Protection for SQL configuration file. This is initially 60.
- If you specify **/logprune**, its value is used instead of the value stored in the Data Protection for SQL configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/logprune** without specifying *numdays* or **no**; in this case, the default 60 is used.
- Changes to the value of the **timeformat** or **dateformat** parameter can result in an undesired pruning of the &agentname; log file. If you are running a command that may prune the log file and the value of the **timeformat** or **dateformat** parameter has changed, perform one of the following to prevent undesired pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/logfile** parameter or **logfile** setting.

/OBJECT=* | objectname,...

This parameter specifies that only particular backup objects for the specified SQL databases and backup object type (if specified) be inactivated. The *objectname* variable specifies the names of the backup objects you want to inactivate. The object name uniquely identifies each backup object and is created by Data Protection for SQL. Use **query** to view the names of backup objects. You can use * as a wildcard character in *objectname* to replace zero or more characters for each occurrence. Specifying only the wildcard character indicates all backup objects of the specified SQL databases and backup object type.

/OLDERthan=numdaysold

This parameter specifies how old a backup object must be before the command can inactivate it.

Considerations:

- The *numdaysold* variable can range from 0 to 9999.
- If you specify 0, you inactivate all selected backup objects.

- If you specify 1, you inactivate all selected backup objects created prior to the current date. Any part of a day counts as a whole day.
- There is no default value for **/olderthan**.

/Quiet The **/quiet** parameter omits displaying status information from the command. However, the information is appended to the Data Protection for SQL activity log.

/TSMNODE=tsmnodename

The **/tsmnode** parameter specifies the Tivoli Storage Manager node name that Data Protection for SQL uses to log on to the Tivoli Storage Manager server. This identifies which Tivoli Storage Manager client is requesting services. You can also store the node name in the options file. The command line parameter overrides the value in the options file.

Considerations:

- You cannot use the **/tsmnode** parameter if **PASSWORDACCESS GENERATE** is specified in the Tivoli Storage Manager options file. You must specify the nodename in the options file. Otherwise, you can change **PASSWORDACCESS** to **PROMPT** to utilize the **/tsmnode** parameter. For details about the Tivoli Storage Manager options file, see the reference manual *IBM Tivoli Storage Manager for Windows Backup-Archive Client Installation and User's Guide*.
- If you do not specify **/tsmnode**, the default value is that specified by the nodename option in the Tivoli Storage Manager options file. Specifying this parameter does not change the value in the options file.

/TSMOPTFile=dsmoptfilename

The **/tsmoptfile** parameter specifies the Tivoli Storage Manager options file to use. This is similar to selecting a Tivoli Storage Manager server from the server list in the GUI. The Tivoli Storage Manager options file contains the configuration values for the Tivoli Storage Manager API. For details about the Tivoli Storage Manager options file, see the reference manual *IBM Tivoli Storage Manager for Windows Backup-Archive Client Installation and User's Guide*.

Considerations:

- The *tsmoptfilename* variable can include a fully qualified path. If you do not include a path, the directory where Data Protection for SQL is installed is used.
- If *tsmoptfilename* includes spaces, you must enclose it in double quotes.
- If you do not specify **/tsmoptfile**, the default value is *dsm.opt*.
- If you specify **/tsmoptfile** but not *tsmoptfilename*, the default is also *dsm.opt*.

/TSMPassword=tsmpasswordname

The **/tsmpassword** parameter specifies the Tivoli Storage Manager password that Data Protection for SQL uses to log on to the Tivoli Storage Manager server. This parameter and the option **PASSWORDACCESS** in the Tivoli Storage Manager options file interact in the following ways:

/tsmpassword	PASSWORDACCESS in Tivoli Storage Manager options file	Password already stored in registry?	Result
specified	<i>generate</i>	yes	<i>/tsmpassword</i> ignored
specified	<i>generate</i>	no	<i>/tsmpassword</i> used and stored

<i>/tsmpassword</i>	PASSWORDACCESS in Tivoli Storage Manager options file	Password already stored in registry?	Result
specified	<i>prompt</i>	—	<i>/tsmpassword</i> used
not specified	<i>prompt</i>	—	user is prompted

Inactivate output examples

These output examples provide a sample of the text, messages, and process status that displays when using the **inactivate** command.

The following operation explicitly inactivates database backup objects. Once a backup object is inactivated, it will expire automatically according to retention policy. In this case, the objects were backed up from a different SQL server. First, a query is performed to display status information such as active state and backup date.

Command:

```
tdpsqlc query tsm DB1_XIVmini_G_BAS,model * /fromsqlserv=STRINGVM1\STRINGVM1
```

Output:

IBM Tivoli Storage Manager for Databases
Data Protection for Microsoft SQL Server
Version 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011.
All rights reserved.

Connecting to TSM Server as node 'STRINGVM1_SQL'...

Querying TSM Server for Backups

Backup Object Information

```
-----  
SQL Server Name ..... STRINGVM1\STRINGVM1  
SQL Database Name ..... DB1_XIVmini_G_BAS  
Backup Method ..... Lgcy  
Backup Location ..... Srv  
Backup Object Type ..... Full  
Backup Object State ..... Active  
Backup Creation Date / Time ..... 09/23/2011 06:31:04  
Backup Size ..... 3.35 MB  
SQL Compressed ..... No  
Backup Compressed ..... No  
Backup Encryption Type ..... None  
Backup Client-deduplicated ..... No  
Database Object Name ..... 20110923063104\00001AC4  
Number of stripes in backup object ..... 1  
Assigned Management Class ..... DEFAULT
```

Backup Object Information

```
-----  
SQL Server Name ..... STRINGVM1\STRINGVM1  
SQL Database Name ..... DB1_XIVmini_G_BAS  
Backup Method ..... Lgcy  
Backup Location ..... Srv  
Backup Object Type ..... Full  
Backup Object State ..... Active  
Backup Creation Date / Time ..... 09/20/2011 05:35:14  
Backup Size ..... 3.35 MB  
SQL Compressed ..... No  
Backup Compressed ..... No  
Backup Encryption Type ..... None  
Backup Client-deduplicated ..... No  
Database Object Name ..... 20110920053514\00001AC4  
Number of stripes in backup object ..... 1  
Assigned Management Class ..... DEFAULT
```

Backup Object Information

```
-----  
SQL Server Name ..... STRINGVM1\STRINGVM1  
SQL Database Name ..... DB1_XIVmini_G_BAS  
Backup Method ..... Lgcy  
Backup Location ..... Srv  
Backup Object Type ..... Full  
Backup Object State ..... Active  
Backup Creation Date / Time ..... 09/19/2011 07:01:39  
Backup Size ..... 3.35 MB  
SQL Compressed ..... No  
Backup Compressed ..... No  
Backup Encryption Type ..... None  
Backup Client-deduplicated ..... No  
Database Object Name ..... 20110919070139\00001AC4  
Number of stripes in backup object ..... 1  
Assigned Management Class ..... DEFAULT
```

Backup Object Information

```
-----  
SQL Server Name ..... STRINGVM1\STRINGVM1  
SQL Database Name ..... model  
Backup Method ..... Lgcy  
Backup Location ..... Srv  
Backup Object Type ..... Full  
Backup Object State ..... Active  
Backup Creation Date / Time ..... 09/23/2011 06:31:05  
Backup Size ..... 2.08 MB  
SQL Compressed ..... No  
Backup Compressed ..... No  
Backup Encryption Type ..... None  
Backup Client-deduplicated ..... No  
Database Object Name ..... 20110923063105\00001AC4  
Number of stripes in backup object ..... 1  
Assigned Management Class ..... DEFAULT
```

Backup Object Information

```
-----  
SQL Server Name ..... STRINGVM1\STRINGVM1  
SQL Database Name ..... model  
Backup Method ..... Lgcy  
Backup Location ..... Srv  
Backup Object Type ..... Full  
Backup Object State ..... Active  
Backup Creation Date / Time ..... 09/19/2011 11:26:15  
Backup Size ..... 2.08 MB  
SQL Compressed ..... No  
Backup Compressed ..... No  
Backup Encryption Type ..... None  
Backup Client-deduplicated ..... No  
Database Object Name ..... 20110919112615\00001AC4  
Number of stripes in backup object ..... 1  
Assigned Management Class ..... DEFAULT
```

Backup Object Information

```
-----  
SQL Server Name ..... STRINGVM1\STRINGVM1  
SQL Database Name ..... model  
Backup Method ..... Lgcy  
Backup Location ..... Srv  
Backup Object Type ..... Full  
Backup Object State ..... Active  
Backup Creation Date / Time ..... 09/17/2011 01:15:48  
Backup Size ..... 2.08 MB  
SQL Compressed ..... No  
Backup Compressed ..... No  
Backup Encryption Type ..... None  
Backup Client-deduplicated ..... No  
Database Object Name ..... 20110917011548\00001AC4  
Number of stripes in backup object ..... 1  
Assigned Management Class ..... DEFAULT
```

Completed

The user then decides to inactivate all *DB1_XIVmini_G_BAS* database objects older than two days (older than September 23), of which there are two.

Command:

```
tdpsqlc inactivate DB1_XIVmini_G_BAS * /fromsqlserv=STRINGVM1 /olderthan=2
```

Output:

```

IBM Tivoli Storage Manager for Databases
Data Protection for Microsoft SQL Server
Version 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011.
All rights reserved.

Starting Sql database backup inactivation...
Querying Tivoli Storage Manager server for a list of database
backups,please wait...

Inactivating full backup DB1_XIVmini_G_BAS
Inactivating log backup DB1_XIVmini_G_BAS\20110920053514\00001AC4

Inactivating full backup DB1_XIVmini_G_BAS
Inactivating log backup DB1_XIVmini_G_BAS\20110919070139\00001AC4

Total database backups inspected:           2
Total database backups requested for inactivation: 2
Total database backups inactivated:        2
Total database skipped:                    0

Elapsed processing time:                    2.18 Secs

```

Another Tivoli Storage Manager query displays the current status of these backup objects using the */all* parameter; a full and a log backup of *test1* are now both inactive.

Command:

```
tdpsqlc query tsm test1 /fromsqlserv=STRINGVM1 /all
```

Output:

```

IBM Tivoli Storage Manager for Databases
Data Protection for Microsoft SQL Server
Version 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011.
All rights reserved.

Backup Object Information
-----

SQL Server Name ..... STRINGVM1
SQL Database Name ..... DB1_XIVmini_G_BAS
Backup Object Type ..... Log
Backup Object State ..... Inactive
Backup Creation Date / Time ..... 09/20/2011 05:35:14
Backup Size ..... 3,349
Database Object Name ..... 20110920053514\00001AC4
Number of stripes in backup object ..... 1

SQL Server Name ..... STRINGVM1
SQL Database Name ..... DB1_XIVmini_G_BAS
Backup Object Type ..... Full
Backup Object State ..... Inactive
Backup Creation Date / Time ..... 09/19/2011 07:01:39
Backup Size ..... 3,349
Database Object Name ..... 20110920053514\00001AC4
Number of stripes in backup object ..... 1

```

Help command

Use the **help** command to display the syntax of all or selected Data Protection for SQL commands using a textual notation.

Help uses the following notation:

- [*a*] *a* is optional; *a* may occur zero or one time
- {*a* | *b*} select either *a* or *b*, but not both
- {*a* } + *a* must occur at least one time
- {*a* } * *a* may occur zero or more times
- (*a*) comments that are not part of the command

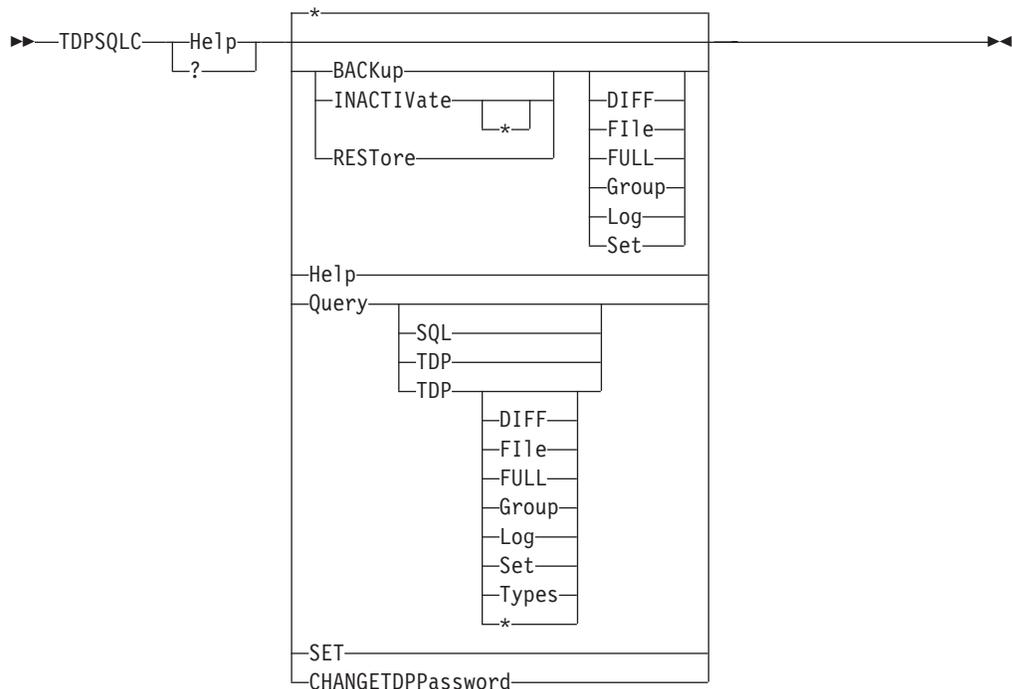
UPPERCASE

minimum abbreviation (which you can also enter in lowercase)

Note: When using languages other than English, you might need to set the width of your screen display to a value greater than 80 characters in order to view the entire help description in one screen. For example, set the screen width to 100 characters.

Help syntax

Use the **help** command syntax diagrams as a reference to view available options and truncation requirements.



Help positional parameters

Positional parameters immediately follow the **help** command. There are no optional parameters with this command.

Use the help command to display the syntax of all or selected Data Protection for SQL commands using a textual notation.

Help uses the following notation:

[*a*] *a* is optional; *a* may occur zero or one time

{*a* | *b*} select either *a* or *b*, but not both

{*a* } + *a* must occur at least one time

{*a* } * *a* may occur zero or more times

(*a*) comments that are not part of the command

UPPERCASE

minimum abbreviation (which you can also enter in lowercase)

Help output examples

These output examples provide a sample of the text, messages, and process status that displays when using the **help** command.

Help 1-Query TSM

Command:

```
tdpsqlc help query tsm *
```

Output:

```
IBM Tivoli Storage Manager for Databases:
Data Protection for Microsoft SQL Server
Version 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011.
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=tsmnode]              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 6, Release 3, Level 0.0
 (C) Copyright IBM Corporation 1998, 2011.
 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
[/DBOonly]
[/FILES=*|logicalname[,logicalname,...] ]
[/FROMSQLserver=sqlservername] default: sqlserver value (or cfg value)
[/GROups=*|groupname[,groupname,...] ]
[/INSTANTRestore=Yes|No] default: Yes
[/INTO=dbname]
[/LOGFile=logfilename] default: tdpsql.log (or cfg value)
[/LOGPrune=numdays|No] default: 60 (or cfg value)
[/MOUNTWait=Yes|No] default: Yes (or cfg value)
[/OBJECT=*|objectname[,objectname,...] ]
[/PARTial]
[/Quiet]
[/RECOVery=Yes|No] default: Yes
[/RELocate=lname /TO=pname [/RELocate=lname /TO=pname ...] ]
[/RELOCATEDir=directory[,logfiledirectory[,otherfiledirectory]] ]
[/REPlace]
[/SQLAUTHentication=INTEgrated|SQLuserid] default: INTEgrated (or cfg value)
[/SQLBUFFers=numsqlbuffers] default: 0 (or cfg value)
[/SQLBUFFERSize=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=undofilename]
[/STRIPes=numstripes] default: 1 (or cfg value)
[/TSMNODE=tsmnode name] default: dsm.opt value
[/TSMOPTFile=dsmoptfilename] default: dsm.opt
[/TSMPassword=tsmpassword] default: dsm.opt value

```

Help 3-Restore Log

Command:

```
tdpsqlc help rest log
```

Output:

IBM Tivoli Storage Manager for Databases:
 Data Protection for Microsoft SQL Server
 Version 6, Release 3, Level 0.0
 (C) Copyright IBM Corporation 1998, 2011.
 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)
[/MOUNTWait=Yes|No] default: Yes (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=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)

[/SQLUSer=sqlusername] default: sa
[/STANDBy=undofilename]
[/STOPAT=datetime]
[/STOPATMark=markname [/AFTER=datetime] ]
[/STOPBEFOREMark=markname [/AFTER=datetime] ]
[/STRIPes=numstripes] default: 1 (or cfg value)
[/TSMNODE=tsmnodename] default: dsm.opt value
[/TSMOPTFile=dsmoptfilename] default: dsm.opt
[/TSMPassword=tsmpassword] default: dsm.opt value
  
```

Help 4-Set

Command:

```
tdpsqlc help set
```

Output:

IBM Tivoli Storage Manager for Databases:
 Data Protection for Microsoft SQL Server
 Version 6, Release 3, Level 0.0
 (C) Copyright IBM Corporation 1998, 2011.
 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=buffer size (64..8192)
 DATEformat=dateformatnum
 1 MM/DD/YYYY
 2 DD-MM-YYYY
 3 YYYY-MM-DD
 4 DD.MM.YYYY
 5 YYYY.MM.DD
 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
 MOUNTWaitfordata=[Yes|No]
 NUMBERformat=numberformatnum
 1 n,nnn.dd
 2 n,nnn,dd
 3 n nnn,dd
 4 n nnn.dd
 5 n.nnn,dd
 6 n'nnn,dd
 REMOTEDSMAgentnode=nodename
 SQLAUTHentication=[INTEgrated|SQLuserid]
 SQLBUFFers=numsqllibuffers (0..999)
 SQLBUFFERSize=sqlbuffer size (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

Set command

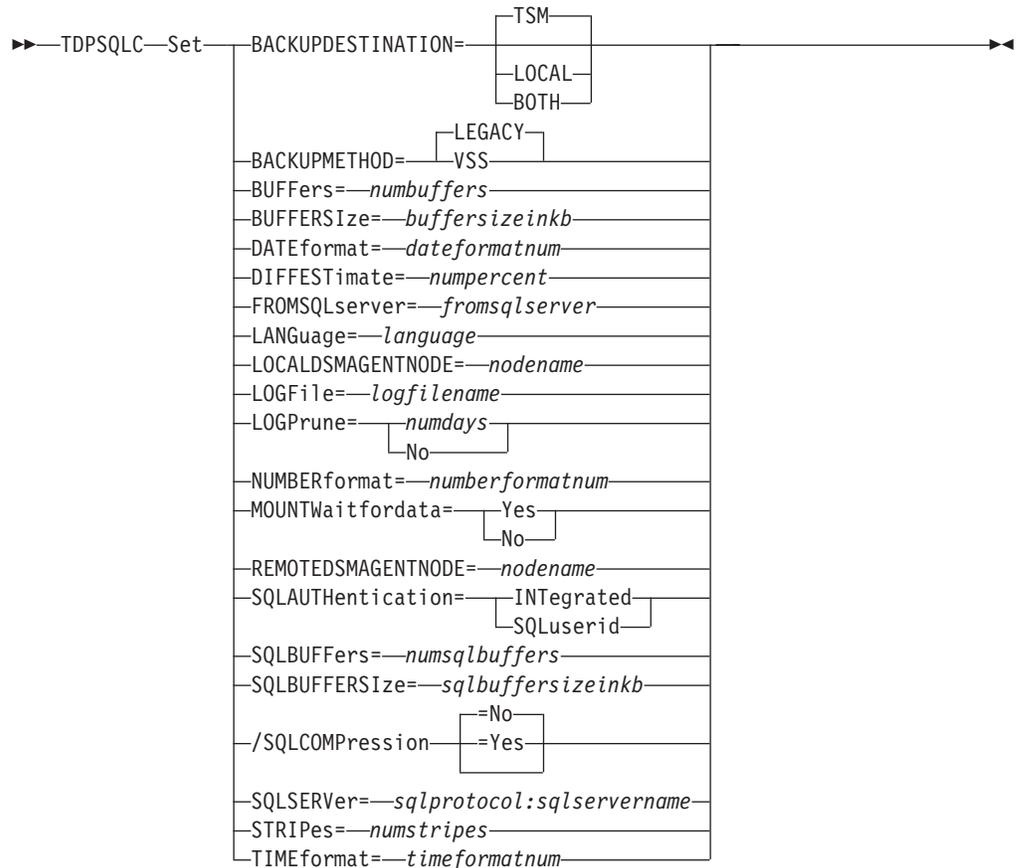
Use the **set** command to change the values for the Data Protection for SQL configurable parameters and options.

The values are saved in a configuration file. The default file is `tdpsql.cfg`. Configuration values can also be set in the GUI **Edit** menu bar item.

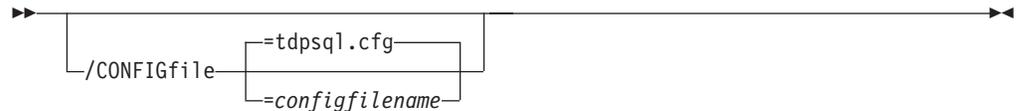
Note: If a configuration file is not specified, the `tdpsql.cfg` values are used, and a default configuration file is created with just the *lastprunedate* value. If an invalid or non-existent file is specified, the default values are used.

Set syntax

Use the **set** command syntax diagrams as a reference to view available options and truncation requirements.



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 configuration file, specify one of the following when issuing a **set** command.

BACKUPDESTINATION=TSM | LOCAL | BOTH

Use the **BACKUPDESTINATION** positional parameter to specify the storage location for your backup. You can specify:

TSM The backup is stored on Tivoli Storage Manager server storage only. This is the default.

LOCAL

The backup is stored on local shadow volumes only.

BOTH The backup is stored on both Tivoli Storage Manager server storage and local shadow volumes.

BACKUPMETHod=Legacy | VSS

Use the BACKUPMETHOD positional parameter to specify the method for your backup. You can specify:

LEGACY

Data Protection for SQL uses the legacy API to perform the backup. This is the default.

VSS Data Protection for SQL uses VSS to perform the backup.

BUFFers=*numbuffers*

The **buffers** parameter specifies the number of data buffers used for each data stripe to transfer data between Data Protection for SQL and the Tivoli Storage Manager API. You can improve throughput by increasing the number of buffers, but you will also increase storage use. Each buffer is the size specified by the **/buffersize** parameter. The *numbuffers* variable refers to the number of data buffers to use. The number can range from 2 to 8. The initial value is 3.

BUFFERSize=*buffersizein kb*

The **buffersize** parameter specifies the size of each Data Protection for SQL buffer specified by the **buffers** parameter. The *buffersizein kb* variable refers to the size of data buffers in kilobytes. The number can range from 64 to 8192. The default is initially 1024.

DATEformat=*dateformatnum*

The **/dateformat** parameter selects the format you want to use to display dates.

The *dateformatnum* variable can range from 1 to 7. The initial value is 1. The number values specify the following formats:

- 1 MM/DD/YYYY.
- 2 DD-MM-YYYY.
- 3 YYYY-MM-DD.
- 4 DD.MM.YYYY.
- 5 YYYY.MM.DD.
- 6 YYYY/MM/DD.
- 7 DD/MM/YYYY.

Changes to the value of the **dateformat** parameter can result in an undesired pruning of the Data Protection for SQL log file (tdpsql.log by default). You can avoid losing existing log file data by performing one of the following:

- After changing the value of the **dateformat** parameter, make a copy of the existing log file before running Data Protection for SQL.
- Specify a new log file with the **/logfile** parameter.

DIFFESTimate=*numpercent*

For **differential** database backups using the Data Protection for SQL **backup** command, **diffestimate** specifies the estimated fraction of an entire

SQL database that has changed since its last full database backup. This estimate is needed because SQL Server does not provide a way to determine the size of a differential backup, and because the Tivoli Storage Manager server requires an accurate size estimate to efficiently allocate space and place objects. The Tivoli Storage Manager server uses this value to determine if there is enough space in the primary storage pool to contain the SQL database backup. Because a separate backup object is created for each specified SQL database, this estimate applies to each specified SQL database individually. The *numpercent* variable can range from 1 to 99. Because a differential backup backs up database pages, this number is the percent of database pages changed since the last full database backup. The initial value is 20.

Considerations:

- If the estimate is significantly smaller than the actual quantity of changes, the Tivoli Storage Manager server may be forced to abnormally end the backup because the backup size is larger than the space the Tivoli Storage Manager server allocated for it.
- If the estimate is significantly larger than the actual quantity of changes, the server may be forced to place the backup object higher in the storage pool hierarchy than otherwise necessary, possibly on removable media.

FROMSQLSERVER=sqlservername

The **fromsqlserver** parameter specifies the SQL server that backup objects were backed up from. This parameter is necessary only when the name of the SQL server to restore to, as determined by the **sqlserver** parameter, is different from the name of the SQL server that the backup objects were created from. Use **fromsqlserver** for **query TSM** and **inactivate** commands, but use **sqlserver** for **query SQL** commands. The default value is the **sqlserver** value or the value set in the Data Protection for SQL configuration file.

LANGUage=language

Specify the three-character code of the language you want to use to display messages:

- CHS** Simplified Chinese
- CHT** Traditional Chinese
- DEU** Standard German
- ENU** American English (This is the default.)
- ESP** Standard Spanish
- FRA** Standard French
- ITA** Standard Italian
- JPN** Japanese
- KOR** Korean
- PTB** Brazilian Portuguese

LOCALDSMAgentnode=nodename

Specify the node name of the local machine that performs the VSS backups. This positional parameter must be specified for VSS operations to be performed.

LOGFile=logfilename

The **logfile** parameter specifies the name of the activity log that is

generated by Data Protection for SQL. The activity log records significant events such as completed commands and error messages. This 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.

Considerations:

- If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file.
- The file name can include a fully-qualified path; however, if you specify no path, the file is written to the directory where Data Protection for SQL is installed.
- You cannot turn Data Protection for SQL activity logging off. If you do not specify **/logfile**, log records are written to the default log file. The default log file is *tdpsql.log*.

LOGPrune=*numdays* | No

The **logprune** parameter prunes the Data Protection for SQL activity log and specifies how many days of entries to save. By default, log pruning is enabled and performed once each day Data Protection for SQL is executed; however, this option allows you to disable log pruning. The *numdays* variable represents the number of days to save log entries.

Considerations:

- If you specify *numdays*, it can range from 0 to 9999. The initial value is 60. A value of 0 deletes all entries in the Data Protection for SQL activity log file except for the current command entries.
- If you specify **no**, the log file is not pruned.

NUMBERformat=*numberformatnum*

The **numberformat** parameter specifies the format of the numbers displayed by Data Protection for SQL. The *numberformatnum* variable can range from 1 to 6. The initial value is 1. The number values specify the following formats:

1	1,000.00
2	1,000,00
3	1 000,00
4	1 000.00
5	1.000,00
6	1'000,00

MOUNTWaitfordata=Yes | No

If the Tivoli Storage Manager server is configured to store backup data on removable media such as tapes or optical media, it is possible that the Tivoli Storage Manager server may indicate to Data Protection for SQL that it is waiting for a required storage volume to be mounted. If that occurs, this option allows you to specify whether Data Protection for SQL **backup** and **restore** commands wait for the media mount or stop the current operation. The initial value is **yes**.

Considerations:

- If you use data striping, Data Protection for SQL cannot complete waiting until the initial media for all stripes are available, although Data Protection for SQL 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 command from a prolonged wait. This can be done *only* by terminating the Data Protection for SQL 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 waits for that volume. During backup operations, the wait occurs *after* all of the data is transferred because meta objects are not created until after the data objects are complete. During restore operations, if the metadata is required, the wait occurs *before* any of the data is transferred.
- If you specify **no** and any removable media are required, Data Protection for SQL terminates the command with an error message. This is also true if the management class for meta objects requires removable media, but, during backups, the command termination does not occur until after all of the data is transferred.

REMOTEDSMAgentnode=nodename

Specify the node name of the machine that moves the VSS data to Tivoli Storage Manager server storage during off-loaded backups.

SQLAUTHentication=INTEgrated | SQLuserid

This parameter specifies the authorization mode used when logging on to the SQL server. The **integrated** value specifies Windows NT or Windows 2000 authentication. The user id you use to log on to Windows is the same id you will use to log on to the SQL server. This is the default value. Use the **sqluserid** value to specify SQL Server user id authorization. The user id specified by the **sqluserid** parameter is the id you will use to log on to the SQL server. That user id must have the SQL Server SYSADMIN fixed server role.

SQLBUFFers=numsqlbuffers

The **sqlbuffers** parameter specifies the total number of data buffers SQL Server uses to transfer data between SQL Server and Data Protection for SQL. The *numsqlbuffers* variable refers to the number of data buffers to use. The number can range from 0 to 999. The default value is 0. When **sqlbuffers** is set to 0, SQL determines how many buffers should be used. The *numsqlbuffers* variable is limited by storage restrictions. If you specify a value other than 0, the number you specify must be equal to or greater than the number of data stripes that you use. Up to 64 stripes may be used. If you specify a value other than 0 and receive errors during a backup, specify a value of 0 and try the backup again.

SQLBUFFERSize=sqlbuffersizeinkb

The **sqlbuffersize** parameter specifies the size of each buffer (specified by the **sqlbuffers** parameter) SQL Server uses to transfer

data to Data Protection for SQL. The *sqlbuffersizeinkb* variable refers to the size of data buffers in kilobytes. The number can range from 64 to 4096. The default is initially 1024.

/SQLCOMPRESSION=Yes | No

The **/sqlcompression** parameter specifies whether SQL compression is applied. If you do not specify **sqlcompression**, the default value *No* is used.

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

SQL Server 2008 backup compression is generally faster and more effective than using it together with Tivoli Storage Manager compression. Tivoli Storage Manager compression is also effective at reducing the backup window needed to protect SQL Server and is a good choice for SQL Server 2005 (or earlier).

SQLSERVER=sqlprotocol:sqlservername

The **sqlserver** parameter specifies the SQL server that Data Protection for SQL logs on to. This is the SQL server that backup objects are restored to. However, if the backup objects were created from a different SQL server name, you must use the **fromsqlserver** parameter. Use **sqlserver** for the **query SQL** command, but use **fromsqlserver** for the **query TSM** and **inactivate** commands. The *sqlprotocol* variable specifies the communication protocol to use. You can specify one of the following protocols:

- *lpc*: Use Shared Memory protocol.
- *np*: Use Named Pipes protocol.
- *tcp*: Use Transmission Control protocol.
- *via*: Use Virtual Interface Architecture protocol.

If no protocol is specified, Data Protection for SQL logs on to the SQL server according to the first protocol that becomes available.

STRIPES=numstripes

The **stripes** parameter specifies the number of data stripes to use in a backup or restore operation. The *numstripes* variable can range from 1 to 64. The default is initially 1. Note that stripes are not available for VSS operations.

TIMEFORMAT=timeformatnum

The **timeformat** parameter specifies the format of the times displayed by Data Protection for SQL. The *timeformatnum* variable can range from 1 to 4. The initial value is 1. The number values specify the following formats:

- | | |
|---|-------------|
| 1 | 23:00:00 |
| 2 | 23,00,00 |
| 3 | 23.00.00 |
| 4 | 11:00:00A/P |

Changes to the value of the **timeformat** parameter can result in an undesired pruning of the Data Protection for SQL log file (tdpsql.log by default). You can avoid losing existing log file data by performing one of the following:

- After changing the value of the **timeformat** parameter, make a copy of the existing log file before running Data Protection for SQL.
- Specify a new log file with the **/logfile** parameter.

Set optional parameters

Optional parameters follow the **set** command and positional parameters.

/CONFIGfile=*configfilename*

The **/configfile** parameter specifies the name of the Data Protection for SQL configuration file, which contains the values for the Data Protection for SQL configurable options.

Considerations:

- *configfilename* can include a fully qualified path. If *configfilename* does not include a path, it uses the directory where Data Protection for SQL is installed.
- If *configfilename* includes spaces, place it in double quotes.
- If you do not specify **/configfile**, the default value is *tdpsql.cfg*.
- If you specify **/configfile** but not *configfilename*, the default value *tdpsql.cfg* is used.

Set output examples

These output examples provide a sample of the text, messages, and process status that displays when using the **set** command.

The following specifies the *STRINGVM1* server as the default SQL server in the configuration file.

Command:

```
tdpsqlc set sqlserver=STRINGVM1
```

Output:

```
IBM Tivoli Storage Manager for Databases
Data Protection for Microsoft SQL Server
Version 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011. All rights reserved.
AC05054I The configuration option was set successfully.
```

Changetsmpassword command

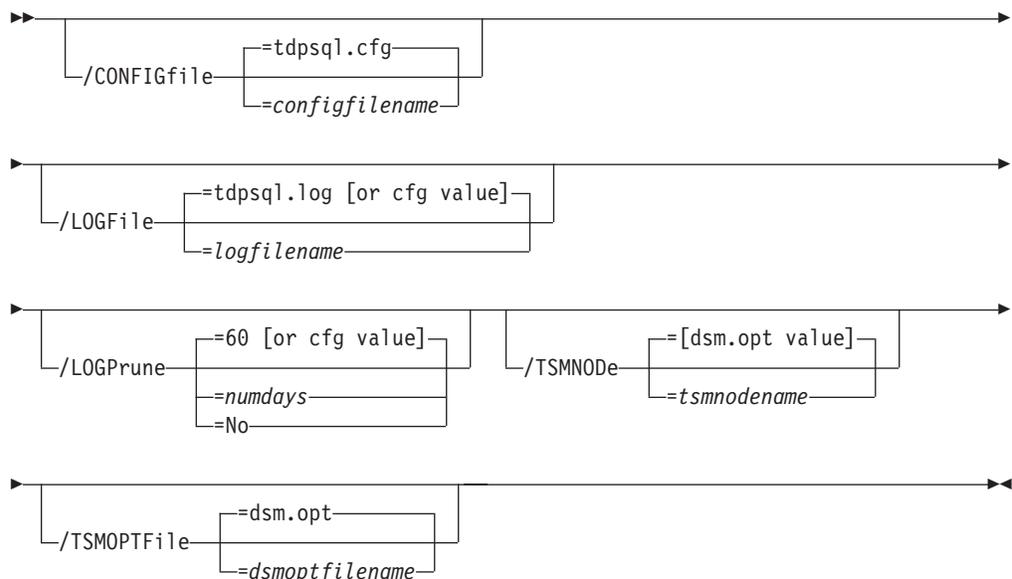
Use the **changetsmpassword** command to change the Tivoli Storage Manager password used by Data Protection for SQL.

Changetsmpassword

Use the **changetsmpassword** command syntax diagrams as a reference to view available options and truncation requirements.



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 configuration file, which contains the values for the Data Protection for SQL configurable options. See “Set positional parameters” on page 223 for details on the file's contents.

Considerations:

- *configfilename* can include a fully qualified path. If *configfilename* does not include a path, it uses the directory where Data Protection for SQL is installed.
- If *configfilename* includes spaces, place it in double quotes.
- If you do not specify **/configfile**, the default value is *tdpsql.cfg*.
- If you specify **/configfile** but not *configfilename*, the default value *tdpsql.cfg* is used.

/LOGFile=*logfile*

The **/logfile** parameter specifies the name of the activity log that is generated by Data Protection for SQL. This activity log records significant events such as completed commands and error messages. The Data Protection for SQL activity log is distinct from the SQL Server error log. The *logfile* variable identifies the name to be used for the activity log generated by Data Protection for SQL.

Considerations:

- If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file.
- The file name can include a fully-qualified path; however, if you specify no path, the file is written to the directory where Data Protection for SQL is installed.
- You cannot turn Data Protection for SQL activity logging off. If you do not specify **/logfile**, log records are written to the default log file. The default log file is *tdpsql.log*.
- When using multiple simultaneous instances of Data Protection for SQL to perform operations, use the **/logfile** parameter to specify a different log file for each instance used. This directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPrune=*numdays* | **No**

The **/logprune** parameter prunes the Data Protection for SQL activity log and specifies how many days of entries are saved. By default, log pruning is enabled and performed once each day Data Protection for SQL is executed; however, this option allows you to disable log pruning or explicitly request a prune of the log for one command run even if the log file has already been pruned for the day. The *numdays* variable represents the number of days to save log entries. By default, 60 days of log entries are saved in the prune process.

Considerations:

- If you specify *numdays*, it can range from 0 to 9999. A value of 0 deletes all entries in the Data Protection for SQL activity log file except for the current command entries.
- If you specify **no**, the log file is not pruned during this command.
- If you do not specify **/logprune**, the default value is that specified by the **logprune** configurable option in the Data Protection for SQL configuration file. This is initially 60.
- If you specify **/logprune**, its value is used instead of the value stored in the Data Protection for SQL configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/logprune** without specifying *numdays* or **no**; in this case, the default 60 is used.
- Changes to the value of the **timeformat** or **dateformat** parameter can result in an undesired pruning of the &agentname; log file. If you are running a command that may prune the log file and the value of the **timeformat** or **dateformat** parameter has changed, perform one of the following to prevent undesired pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/logfile** parameter or **logfile** setting.

/TSMNODE=tsmnodename

The **/tsmnode** parameter specifies the Tivoli Storage Manager node name that Data Protection for SQL uses to log on to the Tivoli Storage Manager server. This identifies which Tivoli Storage Manager client is requesting services. You can also store the node name in the options file. The command line parameter overrides the value in the options file.

Considerations:

- You cannot use the **/tsmnode** parameter if **PASSWORDACCESS GENERATE** is specified in the Tivoli Storage Manager options file. You must specify the nodename in the options file. Otherwise, you can change **PASSWORDACCESS** to **PROMPT** to utilize the **/tsmnode** parameter. For details about the Tivoli Storage Manager options file, see the reference manual *IBM Tivoli Storage Manager for Windows Backup-Archive Client Installation and User's Guide*.
- If you do not specify **/tsmnode**, the default value is that specified by the nodename option in the Tivoli Storage Manager options file. Specifying this parameter does not change the value in the options file.

/TSMOPTFile=dsmoptfilename

The **/tsmoptfile** parameter specifies the Tivoli Storage Manager options file to use. This is similar to selecting a Tivoli Storage Manager server from the server list in the GUI. The Tivoli Storage Manager options file contains the configuration values for the Tivoli Storage Manager API. For details about the Tivoli Storage Manager options file, see the reference manual *IBM Tivoli Storage Manager for Windows Backup-Archive Client Installation and User's Guide*.

Considerations:

- The *tsmoptfilename* variable can include a fully qualified path. If you do not include a path, the directory where Data Protection for SQL is installed is used.
- If *tsmoptfilename* includes spaces, you must enclose it in double quotes.
- If you do not specify **/tsmoptfile**, the default value is *dsm.opt*.

- If you specify `/tsmoptfile` but not `tsmoptfilename`, the default is also `dsm.opt`.

Changetsmpassword output examples

This output example provides a sample of the text, messages, and process status that displays when using the `changetsmpassword` command.

The following displays changing the Tivoli Storage Manager password.

Command:

```
tdpsqlc changetsm ssqlv2old ssqlv2new ssqlv2new
```

Output:

```
IBM Tivoli Storage Manager for Databases:
Data Protection for Microsoft SQL Server
Version 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011.
All rights reserved.

AC00260I Password successfully changed.
```

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
```

Examples of Tivoli Storage Manager policy binding using include/exclude and VSSPOLICY statements

To exploit automatic version control and expiration, you are able to set policy for each type of backup data. The method of setting policy is different for Legacy backups and VSS Backups:

- **Legacy backups:** Use INCLUDE/EXCLUDE statements in the Data Protection for SQL options file (`dsm.opt` by default).
- **VSS Backups:** Use the VSSPOLICY statement in the Data Protection for SQL configuration file (`tdpsql.cfg` by default).

VSS examples

VSS Backups use the VSSPOLICY statement in the Data Protection for SQL configuration file (see “Setting automatic expiration (VSS and Legacy)” on page 41 for the general syntax):

```
VSSPOLICY * * FULL TSM MC1
VSSPOLICY * * FULL LOCAL MC2
VSSPOLICY * largdb1 FULL TSM MC3
VSSPOLICY * largdb1 FULL LOCAL MC4
VSSPOLICY SERVER1 * FULL TSM MC5
VSSPOLICY SERVER1 * FULL LOCAL MC6
```

Legacy examples

Legacy backups use INCLUDE/EXCLUDE statements in the Data Protection for SQL options file (see “Setting automatic expiration (VSS and Legacy)” on page 41 for the general syntax):

BackupType Object Matches	Specification
All objects	\...*
All <i>backupType</i> objects (full, diff, log, group, file, set)	\...\full* \...\diff* etc.
All group or file object names (g1, f1)	\...g1\group* \...f1\file*
All group or file object names beginning with g or f	\...g*\group* \...f*\file*
Same as \...\group* or \...\file*	\...*\group* \...*\file*
BackupType Object with Database Matches	Specification
All objects with database name Db1	\...\Db1\...*
All objects with database name Db1 beginning with Db	\...\Db*\...*
Ambiguous	\...*\...*
All full or diff objects with database name Db1	\...\Db1\full* \...\Db1\diff*
All log, group, file, or set objects with database name Db1	\...\Db1\...\log* \...\Db1\...\group* etc.
All group or file object names (g1, f1) with database name Db1	\...\Db1g1\group* \...\Db1f1\file*

BackupType Object with Database Matches	Specification
All group or file object names beginning with g or f with database name Db1	\...\Db1\g*\group* \...\Db1\f*\file*
Same as \...\Db1\...\group* or file*	\...\Db1*\group* \...\Db1*\file*
Same as \...\Db1\full*	\...\Db1\...\full*
Same as \...\full*	\...*\full*
Same as \...\group*	\...**\group*
Same as \...\g1\group*	\...*\g1\group*
Ambiguous	\...*\...\log*
Nothing (<i>typeInfo</i> missing)	\...\Db1\set*
Meta & Data Object Matches	Specification
All meta or data objects	\...\meta\...* \...\data\...*
All meta or data full, log, or group objects	\...\meta\...\full* etc.
All meta or data group object names (g1)	\...\meta\...\g1\group* etc.
All meta or data group object names beginning with g	\...\meta\...\g*\group* etc.
Same as \...\meta\data\...\group*	\...\meta\...*\group* etc.
Nothing (qualifiers missing)	\...\meta*\...\data*
Meta & Data Object with Database Matches	Specification
All meta or data objects with database name Db1	\...\meta\...\Db1\...* \...\data\...\Db1\...*
All meta or data full objects with database name Db1	\...\meta\...\Db1\full* etc.
All meta or data log or group objects with database name Db1	\...\meta\...\Db1\...\log* etc.
All meta or data group object names (g1) with database name Db1	\...\meta\...\Db1\g1\group* etc.

Meta & Data Object with Database Matches	Specification
All meta or data group object names beginning with g with database name Db1	\...\meta\...\Db1\g*\group* etc.
Same as \...\meta\data\...\Db1\...\group*	\...\meta\...\Db1*\group* etc.
Same as \...\meta\data\...\full*	\...\meta\...*\full* etc.
Same as \...\meta\data\...\group*	\...\meta\...*\group* etc.
Same as \...\meta\data\...\g1\group*	\...\meta\...*g1\group* etc.
Ambiguous	\...\meta\...*\...\log* \...\data\...*\...\log*
Nothing (qualifiers missing)	\...\meta*\...\data*
Server Matches	Specification
All objects from all servers beginning with SQL	SQL*\...*
All objects from all server instances with host SQL2000	SQL2000\...*
All objects from server SQL2000\INST1	SQL2000\INST1\...*
All objects from all servers beginning with SQL2000\INST	SQL2000\INST*\...*
Same as SQL2000\...*	SQL2000*\...*
All meta or data objects from server SQL2000\INST1	SQL2000\INST1\meta\...** SQL2000\INST1\data\...**
All meta or data objects from all <i>named</i> server instances with host SQL2000	SQL2000*\meta\...** etc.
All meta or data objects from all server instances with host SQL2000	SQL2000\...\meta\...** etc.
All objects from server default instance (if no instance name matches ??ta)	SQL2000\??ta\meta\...** SQL2000\??ta\data\...**

Transitioning SQL Server backups from Tivoli Storage FlashCopy Manager to Tivoli Storage Manager

Configure Tivoli Storage FlashCopy Manager so that you can access both a local and Tivoli Storage Manager server at the same time. This might be useful if you decide to move to a Tivoli Storage Manager environment and want to continue to interact with the locally managed snapshots until policy marks them for expiration.

Tivoli Storage FlashCopy Manager works when connected to the local Tivoli Storage FlashCopy Manager server or a Tivoli Storage Manager server. The Tivoli Storage Manager server can be located anywhere on your network. The Tivoli Storage FlashCopy Manager Snap-in includes two configuration wizards. These enable you to do a local configuration and a Tivoli Storage Manager configuration. You can move from one type of server to another by running the corresponding configuration wizard at any time.

Using the Tivoli Storage Manager server wizard

You can use the Tivoli Storage Manager server wizard to transition SQL Server backups from Tivoli Storage FlashCopy Manager to Tivoli Storage Manager. The wizard leads you through all the steps necessary to perform the configuration. You do not then need to perform the manual steps listed below.

Implement these tasks on the Tivoli Storage Manager server

Coordinate efforts with your Tivoli Storage Manager server administrator to get these tasks completed:

1. Select or create the policy definitions that will be used for each type of backup you plan to use. You can provide the administrator with the existing locally-defined policy settings in your Tivoli Storage FlashCopy Manager stand-alone environment. Use the GUI or the command-line interface of Data Protection for SQL to retrieve this information.
2. Register your Data Protection for SQL node name and password with the Tivoli Storage Manager **register node** command. For example:
`register node DPnodename DPpassword`
3. If not already defined in the Tivoli Storage Manager server, register the Tivoli Storage Manager backup-archive client node name and password for the workstation where the SQL server is installed. For example:
`register node BAnodename 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`

Implement these tasks on the workstation running the SQL Server

1. In the directory where the Data Protection for SQL is installed, make a copy of the options file named `dsm.opt`. After you begin using the Tivoli Storage Manager server, the copy is used for access to the Tivoli Storage FlashCopy Manager stand-alone environment. One method of making the copy is to start the SQL command line prompt from the Tivoli Storage FlashCopy Manager Snapin: In the Tivoli Storage FlashCopy Manager Snapin Tree view, an SQL server node is displayed for each SQL server instance on the computer.
 - a. Select an SQL server instance in the tree view. The integrated command line and an Actions pane is displayed.
 - b. Launch the Data Protection for SQL 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 configuration file. For example:
`copy tdpsql.cfg tdpsql_local.cfg`

Preserve the contents of the local configuration file if:

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

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

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

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

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

For more information on using the **dsmcutil** command, refer to the information on using the client service configuration utility in the Tivoli Storage Manager Windows Backup-Archive Clients Installation and User's Guide.

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

```
dsmcutil install cad /name:tsmcad4local
/node:my_backup-archive_client_node
/password:my_TSM_server_password
/optfile:"C:\Program Files\Tivoli\TSM\baclient\dsm_local.opt"
/httpport:1583
```
7. Edit the `dsm_local.opt` file in the Data Protection for SQL installation directory. Add this line:

```
HTTPPORT 1583
```
8. Start the alternate Tivoli Storage Manager Client Acceptor service:

```
dsmcutil start /name:tsmcad4local
```
9. Stop and restart the original Tivoli Storage Manager Client Acceptor service so that the new values in the `dsm.opt` file are activated. You can do this through the Windows Services GUI or by using the **dsmcutil** command:

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

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

You can also make this change using the Local Policy Management dialog that is accessed from the Utilities menu of the Data Protection for SQL Backup/Restore GUI. Information on how to start the GUI is located below in the section describing how to access the Tivoli Storage FlashCopy Manager stand-alone environment.

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

```
tdpsqlc query tsm * /all
/tsmoptfile=dsm_local.opt
/configfile=tdpsql_local.cfg
```

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

The very last backup, that is created while running the stand-alone environment, will not be automatically deleted by the process of expiring the backups. For that, you will need to use the explicit delete operation, as described next.

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

```
tdpsqlc query tsm * /all
/tsmoptfile=dsm_local.opt
/configfile=tdpsql_local.cfg
```
 - c. Start the GUI (from the Command Line prompt) by issuing the GUI invocation command, for example:

```
tdpsql /tsmoptfile=dsm_local.opt
/configfile=tdpsql_local.cfg
```
12. If necessary, start the Tivoli Storage FlashCopy Manager stand-alone environment to restore from a backup that was created in that environment.
13. When the transition is complete and you no longer need to access the Tivoli Storage FlashCopy Manager stand-alone environment, you can remove the alternate services. To do this, use the Windows Services GUI or the **dsmcutil** command:

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

Appendix A. Frequently asked questions

Answers related to frequently asked questions about Data Protection for SQL are provided.

How can I compress my Data Protection for SQL backups?

You can use the following methods to compress your Data Protection for SQL 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 options file.
- For VSS Backups, specify the **compression** option in the backup-archive client options file that is used as the Local DSMAGENT Node. If the environment is configured for VSS offloaded backups, you must also specify the compression option in the backup-archive client options file that is used as the Remote DSMAGENT Node. Review the compression information available in the client documentation before attempting to compress your data.

See “Specifying Data Protection for SQL options” on page 37 for more information about the **compression** option.

- For Microsoft SQL Server 2008, you can specify SQL backup compression from the SQL Properties windows in the MMC GUI, or you can use the **sqlcompression** option from the command line to set SQL native backup compression for Data Protection for SQL backups.

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

How do I encrypt my Data Protection for SQL 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 options file.
- For VSS Backups, specify the encryption options in the backup-archive client options file that is used as the Local DSMAGENT Node. If the environment is configured for VSS offloaded backups, you must also specify the encryption options in the backup-archive client options file that is used as the Remote DSMAGENT Node. Review the encryption information available in the client documentation before attempting to encrypt your databases.

See “Specifying Data Protection for SQL options” on page 37 for more information about the **enableclientencryptkey** and **encryptiontype** options.

How do I deduplicate my Data Protection for SQL 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 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.

See “Specifying Data Protection for SQL options” on page 37 for more information about the **deduplication** option.

Can I restore an individual table from a SQL Server backup?

Yes, but only for Legacy backups. You cannot restore an individual table from a VSS Backup. To restore an individual table from a Legacy SQL Server backup, place the tables that require individual restore granularity into their own file group. Then, use Data Protection for SQL to restore a single file group from a full backup.

How can I restore a SQL database backup to an alternate SQL Server machine or database?

For VSS Backups, you cannot restore VSS Backups to an alternate SQL Server. This feature is not supported by Microsoft.

For Legacy backups, see “Restoring to an alternate machine” on page 86 for information about performing this procedure.

Can I restore VSS Backups to alternate locations?

Yes, this feature is supported by Data Protection for SQL.

- On the command-line interface, use the *relocatedir* parameter. See “Restore optional parameters” on page 175 for more information about this parameter.
- In the GUI, use the **Relocate** option in the Restore Databases window. See “Restore options” on page 88 for more information about this option.

Can I restore VSS Backups to alternate database names?

Yes, this feature is supported by Data Protection for SQL.

- On the command-line interface, use the *into* parameter. See “Restore optional parameters” on page 175 for more information about this parameter.
- In the GUI, use the **Restore Into** option in the Restore Databases window. See “Restore options” on page 88 for more information about this option.

Can I use Data Protection for SQL to back up SQL databases, logs, and then also shrink the transaction log file?

Modify the command file that is used for scheduled backups with an entry that calls a T-SQL command file that shrinks the transaction log file. For example, in the following command file that is used for scheduled backups:

```
tdpsqlc backup * full
tdpsqlc backup * log
osql -E -i shrinkjob.sql
```

The file *shrinkjob.sql* is a T-SQL command file that shrinks the transaction log file. See the following sections for more information:

- “Using the Tivoli Storage Manager scheduler” on page 91
- “Backup command” on page 121

Should I create a separate node name in order to create an archive backup of a SQL database?

First, use the same node name as the primary SQL node but add an extension for the archive node. For example:

```
Primary: SQLSRV550_SQL
Archive: SQLSRV550_SQL_ARCH
```

Second, use a separate Data Protection for SQL options file (*dsmarchive.opt*) that contains the archive node with the archive settings that you want. See the following sections for more information about nodes and options:

- “Data Protection for SQL node name: Recommended settings” on page 34
- “Specifying Data Protection for SQL options” on page 37

How do I back up a SQL Server 2005 or later database with database mirroring?

See the Microsoft TechNet article “Database Mirroring in SQL Server 2005” at the following URL:<http://www.microsoft.com/technet/prodtechnol/sql/2005/dbmirror.mspx>

Can I perform VSS operations in a clustered SQL Server environment?

Yes, Data Protection for SQL supports VSS operations in a clustered SQL Server environment. See “Using VSS operations in a SQL Server Failover Cluster environment” on page 25 for detailed information.

Why can I not perform VSS offloaded backups, or manage local snapshots?

You must install IBM Tivoli Storage FlashCopy Manager to perform VSS offloaded backups, or back up and restore local snapshots. See “Minimum software and operating system requirements” on page 48 for detailed information.

Why can I not perform VSS Instant Restore even though I have SAN Volume Controller or a DS storage subsystem installed?

If your Tivoli Storage Manager Client version is earlier than 6.1.0, you must install the IBM Tivoli Storage Manager for Copy Services Hardware Devices Snapshot Integration Module. See “Minimum software and operating system requirements” on page 48 for detailed information.

How can I use VSS and Legacy backups together in a common backup strategy?

See “Using VSS and Legacy Backups together” on page 24 and “Back up to Tivoli Storage Manager storage versus back up to local shadow volumes” on page 28 for detailed information.

Can I use Legacy backups and VSS Backups together?

Yes, you can apply Legacy differential and Legacy log backups after a full VSS Backup has been restored. In order to do this, you must leave the database in a recovering state by specifying `/recovery=no` on the command-line interface or by making sure that the **Recovery** option in the GUI Restore Databases or Restore Groups/Files is not selected when restoring the VSS Backup. VSS supports only full backups. Log, differential, individual file groups, individual files, and set backups are not supported by VSS. See “Using VSS and Legacy Backups together” on page 24 for more information.

When restoring very large SQL databases, how can I prevent the restore operation from failing due to a timeout error?

SQL Server rebuilds and formats new physical files into which the backup data is restored. Because this process can continue for more than an hour for large databases, the Tivoli Storage Manager session might timeout and cause the restore process to fail. To prevent such a failure, set the value of the Tivoli Storage Manager `COMMTIMEOUT` option to 3600 (or higher). Set the value to 10000 (or higher) for databases larger than 100 GB. For a LANFREE restore operation, increase the value of both the `COMMTIMEOUT` and `IDLETIMEOUT` options for the Storage Agent.

How does VSS Instant Restore work?

VSS Instant Restore is a volume-level hardware-assisted copy where target volumes (that contain the snapshot) are copied back to the original source volumes. A SAN Volume Controller, Storwize V7000, DS6000, DS8000, or XIV storage subsystem is required to perform VSS Instant Restores. See “VSS Instant Restore” on page 8 for more information.

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. See “How Tivoli Storage Manager server policy affects Data Protection for SQL” on page 19 and “Back up to Tivoli Storage Manager storage versus back up to local shadow volumes” on page 28 for more information.

Why do I receive a TCP/IP timeout failure when I have Windows internal VSS tracing turned on?

Data Protection for SQL 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 should my Data Protection for SQL performance settings be?

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 performance and should be considered when determining your settings. Note that the `buffers` and `buffersize` parameters apply to Legacy backups only. See the following sections for more information:

- Chapter 7, “Performance,” on page 115
- “Specifying Data Protection for SQL options” on page 37

- "*/buffers* and */buffersize* parameters" (with the **backup** command) on "Backup optional parameters" on page 127.
- "*/buffers* and */buffersize* parameters" (with the **restore** command) on "Restore optional parameters" on page 175.
- "*/buffers* and */buffersize* parameter" (with the **set** command) on "Set optional parameters" on page 229.

How do I schedule Data Protection for SQL backups?

You can schedule Data Protection for SQL backups by using the Tivoli Storage Manager backup-archive client scheduler or the MMC GUI scheduler. See the following sections for more information:

- "Using the Tivoli Storage Manager scheduler" on page 91
- "Backup types" on page 5
- "Backup strategies" on page 21
- "Automating tasks" on page 99

How do I set up Data Protection for SQL to run in a cluster?

The following sections contain information about using Data Protection for SQL in a cluster environment:

- "Using Data Protection for SQL in a Microsoft Windows Failover Clustering environment" on page 27
- "Using Data Protection for SQL in a Veritas Cluster Server environment" on page 28
- "*clusternode* option" in "Specifying Data Protection for SQL options" on page 37

How do I know if my backup ran successfully?

A message displays that states the backup completed successfully. In addition, the Task Manager in the MMC GUI provides centralized information about the status of your tasks. Processing information is also available in the following files:

- Data Protection for SQL 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 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.

See "Specifying Data Protection for SQL options" on page 37 for more information.

How do I set up LAN Free to back up Data Protection for SQL over my SAN?

See the LAN-free section in Chapter 7, "Performance," on page 115.

For more information, go to <http://www.redbooks.ibm.com/abstracts/sg246148.html>.

Appendix B. Accessibility features for Data Protection for SQL

Accessibility features help users who have a disability, such as restricted mobility or limited vision, to use information technology products successfully.

Accessibility features

The following list includes the major accessibility features in Data Protection for SQL:

- Keyboard-only operation
- Interfaces that are commonly used by screen readers
- Keys that are discernible by touch but do not activate just by touching them
- Industry-standard devices for ports and connectors
- The attachment of alternative input and output devices

The Tivoli Information Center, and its related publications, are accessibility-enabled. The accessibility features of the information center are described at Accessibility and keyboard shortcuts in the information center.

Keyboard navigation

Standard shortcut and accelerator keys are used by the product and are documented by the operating system. Refer to the documentation provided by your operating system for more information.

Interface information

You can enlarge information on the product windows using facilities provided by the operating systems on which the product is run. For example, in a Microsoft Windows environment, you can lower the resolution of the screen to enlarge the font sizes of the text on the screen. Refer to the documentation provided by your operating system for more information.

IBM and accessibility

See the IBM Human Ability and Accessibility Center for more 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 give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

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 Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Information Enabling Requests
Dept. M13
5600 Cottle Road
San Jose CA 95193-0001
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this 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.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com)[®] are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at <http://www.ibm.com/legal/copytrade.shtml>.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Other product and service names might be trademarks of IBM or other companies.

Glossary

This glossary includes terms and definitions for IBM Tivoli Storage Manager and IBM Tivoli Storage FlashCopy Manager products.

To view glossaries for other IBM products, go to <http://www.ibm.com/software/globalization/terminology/>.

The following cross-references are used in this glossary:

- *See* refers the reader from a term to a preferred synonym, or from an acronym or abbreviation to the defined full form.
- *See also* refers the reader to a related or contrasting term.

A

absolute mode

In storage management, a backup copy-group mode that specifies that a file is considered for incremental backup even if the file has not changed since the last backup. See also *modified mode*.

access control list (ACL)

In computer security, a list associated with an object that identifies all the subjects that can access the object and their access rights. For example, an access control list is associated with a file that identifies the users who can access that file and their access rights.

access mode

An attribute of a storage pool or a storage volume that specifies whether the server can write to or read from the storage pool or storage volume. The access mode can be read/write, read-only, or unavailable. Volumes in primary storage pools can also have an access mode of destroyed. Volumes in copy storage pools can also have an access mode of offsite.

acknowledgment

The transmission of acknowledgment characters as a positive response to a data transmission.

ACL See *access control list*.

activate

To validate the contents of a policy set and then make it the active policy set.

active-data pool

A named set of storage pool volumes that contain only active versions of client backup data.

active file system

A file system to which space management has been added. With space management, tasks for an active file system include automatic migration, reconciliation, selective migration, and recall. Contrast with *inactive file system*.

active policy set

The activated policy set that contains the policy rules in use by all client nodes that are assigned to the policy domain. See also *policy domain* and *policy set*.

active version

The most recent backup copy of a file stored. The active version of a file cannot be deleted until a backup process detects that the user has either replaced the file with a newer version or has deleted the file from the file server or workstation. Contrast with *inactive version*.

activity log

A log that records normal activity messages that are generated by the server. These messages include information about server and client operations, such as the start time of sessions or device I/O errors.

adaptive subfile backup

A type of backup that sends only changed portions of a file to the server, instead of sending the entire file. Adaptive subfile backup reduces network traffic and increases the speed of the backup.

administrative client

A program that runs on a file server, workstation, or mainframe that administrators use to control and monitor the Tivoli Storage Manager server. Contrast with *backup-archive client*.

administrative command schedule

A database record that describes the

planned processing of an administrative command during a specific time period. See also *client schedule*.

administrative privilege class

See *privilege class*.

administrative session

A period of time during which an administrator user ID communicates with a server to perform administrative tasks. Contrast with *client node session*.

administrator

A user who is registered to the server as an administrator, and who is authorized to perform tasks and issue commands through the assignment of an administrative privilege class.

Advanced Program-to-Program Communication (APPC)

An implementation of the SNA LU 6.2 protocol that allows interconnected systems to communicate and share the processing of programs.

agent node

A client node that has been granted proxy authority to perform operations on behalf of another client node, which is the target node.

aggregate

An object, stored in one or more storage pools, consisting of a group of logical files that are packaged together. See also *logical file* and *physical file*.

aggregate data transfer rate

A performance statistic that indicates the average number of bytes that were transferred per second while processing a given operation.

APPC See *Advanced Program-to-Program Communication*.

application client

A program that is installed on a system to protect an application. The Tivoli Storage Manager server provides backup services to an application client.

archive

To copy programs, data, or files to other storage media, usually for long-term storage or security. Contrast with *retrieve*.

archive copy

A file or group of files that was archived to server storage.

archive copy group

A policy object containing attributes that control the generation, destination, and expiration of archived files.

archive-retention grace period

The number of days that the storage manager retains an archived file when the server is unable to rebind the file to an appropriate management class. See also *bind*.

association

(1) The defined relationship between a client node and a client schedule. An association identifies the name of a schedule, the name of the policy domain to which the schedule belongs, and the name of a client node that performs scheduled operations.

(2) On a configuration manager, the defined relationship between a profile and an object such as a policy domain. Profile associations define the configuration information that is distributed to a managed server when it subscribes to the profile.

audit

To check for logical inconsistencies between information that the server has and the actual condition of the system. The storage manager can audit information about items such as volumes, libraries, and licenses. For example, when a storage manager audits a volume, the server checks for inconsistencies between information about backed-up or archived files that are stored in the database and the actual data that are associated with each backup version or archive copy in server storage.

authentication

The process of checking a user's password before permitting user access to the Tivoli Storage Manager server. Authentication can be turned on or off by an administrator with system privilege.

authentication rule

A specification that another user can use to either restore or retrieve files from storage.

authority

The right to access objects, resources, or functions. See also *privilege class*.

authorization rule

A specification that permits another user to either restore or retrieve a user's files from storage.

authorized user

A user who has administrative authority for the Tivoli Storage Manager client on a workstation. This user changes passwords, performs open registrations, and deletes file spaces.

AutoFS

See *automounted file system*.

automatic detection

A feature that detects, reports, and updates the serial number of a drive or library in the database when the path from the local server is defined.

automatic migration

The process that is used to automatically move files from a local file system to storage, based on options and settings that are chosen by a root user on a workstation. See also *threshold migration* and *demand migration*.

automatic reconciliation

The process that is used to reconcile file systems at regular intervals. The intervals are set by a user with root user authority. See also *reconciliation*.

automounted file system (AutoFS)

A file system that is managed by an automounter daemon. The automounter daemon monitors a specified directory path, and automatically mounts the file system to access data.

B**backup-archive client**

A program that runs on a workstation or file server and provides a means for users to back up, archive, restore, and retrieve files. Contrast with *administrative client*.

backup copy group

A policy object containing attributes that control the generation, destination, and expiration of backup versions of files. A backup copy group belongs to a management class.

backup-retention grace period

The number of days the storage manager retains a backup version after the server is unable to rebind the file to an appropriate management class.

backup set

A portable, consolidated group of active versions of backup files that are generated for a backup-archive client.

backup set collection

A group of backup sets that are created at the same time and which have the same backup set name, volume names, description, and device classes. The server identifies each backup set in the collection by its node name, backup set name, and file type.

backup version

A file or directory that a client node backed up to server storage. More than one backup version can exist in server storage, but only one backup version is the active version. See also *active version* and *inactive version*.

bind To associate all versions of a file with a management class name. See *rebind*.

bindery

A database that consists of three system files for a NetWare server. The files contain user IDs and user restrictions.

C

cache To place a duplicate copy of a file on random access media when the server migrates a file to another storage pool in the hierarchy.

cache file

A snapshot of a logical volume created by Logical Volume Snapshot Agent. Blocks are saved immediately before they are modified during the image backup and their logical extents are saved in the cache files.

CAD See *client acceptor*.

central scheduler

A function that permits an administrator to schedule client operations and administrative commands. The operations can be scheduled to occur periodically or on a specific date. See *client schedule* and *administrative command schedule*.

client A software program or computer that requests services from a server.

client acceptor

An HTTP service that serves the applet for the web client to web browsers. On Windows systems, the client acceptor is installed and run as a service. On AIX®, UNIX, and Linux systems, the client acceptor is run as a daemon, and is also called the *client acceptor daemon (CAD)*.

client acceptor daemon (CAD)

See *client acceptor*.

client domain

The set of drives, file systems, or volumes that the user selects to back up or archive data, using the backup-archive client.

client node

A file server or workstation on which the backup-archive client program has been installed, and which has been registered to the server.

client node session

A session in which a client node communicates with a server to perform backup, restore, archive, retrieve, migrate, or recall requests. Contrast with *administrative session*.

client options file

An editable file that identifies the server and communication method, and provides the configuration for backup, archive, hierarchical storage management, and scheduling.

client option set

A group of options that are defined on the server and used on client nodes in conjunction with client options files.

client-polling scheduling mode

A method of operation in which the client queries the server for work. Contrast with *server-prompted scheduling mode*.

client schedule

A database record that describes the planned processing of a client operation during a specific time period. The client operation can be a backup, archive, restore, or retrieve operation, a client operating system command, or a macro. See also *administrative command schedule*.

client/server

Pertaining to the model of interaction in

distributed data processing in which a program on one computer sends a request to a program on another computer and awaits a response. The requesting program is called a client; the answering program is called a server.

client system-options file

A file, used on AIX, UNIX, or Linux system clients, containing a set of processing options that identify the servers to be contacted for services. This file also specifies communication methods and options for backup, archive, hierarchical storage management, and scheduling. This file is also called the *dsm.sys* file. See also *client user-options file*.

client user-options file

A file that contains the set of processing options that the clients on the system use. The set can include options that determine the server that the client contacts, and options that affect backup operations, archive operations, hierarchical storage management operations, and scheduled operations. This file is also called the *dsm.opt* file. For AIX, UNIX, or Linux systems, see also *client system-options file*.

closed registration

A registration process in which only an administrator can register workstations as client nodes with the server. Contrast with *open registration*.

collocation

The process of keeping all data belonging to a single-client file space, a single client node, or a group of client nodes on a minimal number of sequential-access volumes within a storage pool. Collocation can reduce the number of volumes that must be accessed when a large amount of data must be restored.

collocation group

A user-defined group of client nodes whose data is stored on a minimal number of volumes through the process of collocation.

commit point

A point in time when data is considered consistent.

Common Programming Interface for Communications (CPI-C)

A call-level interface that provides a consistent application programming interface (API) for applications that use program-to-program communications. CPI-C uses LU 6.2 architecture to create a set of interprogram services that can establish and end a conversation, send and receive data, exchange control information, and notify a partner program of errors.

communication method

The method by which a client and server exchange information. See also *Transmission Control Protocol/Internet Protocol*.

communication protocol

A set of defined interfaces that permit computers to communicate with each other.

compression

A function that removes repetitive characters, spaces, or strings of characters from the data being processed and replaces the repetitive characters with control characters. Compression reduces the amount of storage space that is required for the data.

configuration manager

A server that distributes configuration information, such as policies and schedules, to managed servers according to their profiles. Configuration information can include policy and schedules. See also *managed server* and *profile*.

conversation

A connection between two programs over a session that allows them to communicate with each other while processing a transaction.

copy backup

A full backup in which the transaction log files are not deleted so that backup procedures that use incremental or differential backups are not disrupted.

copy group

A policy object containing attributes that control how backup versions or archive copies are generated, where backup versions or archive copies are initially

located, and when backup versions or archive copies expire. A copy group belongs to a management class. See also *archive copy group*, *backup copy group*, *backup version*, and *management class*.

copy storage pool

A named set of volumes that contain copies of files that reside in primary storage pools. Copy storage pools are used only to back up the data that is stored in primary storage pools. A copy storage pool cannot be a destination for a backup copy group, an archive copy group, or a management class (for space-managed files). See also *primary storage pool* and *destination*.

CPI-C See *Common Programming Interface for Communications*.

D**daemon**

A program that runs unattended to perform continuous or periodic functions, such as network control.

damaged file

A physical file in which Tivoli Storage Manager has detected read errors.

data access control mode

A mode that controls whether a command can access a migrated file, see a migrated file as zero-length, or receive an input/output error if it attempts to access a migrated file. See also *execution mode*.

database backup series

One full backup of the database, plus up to 32 incremental backups made since that full backup. Each full backup that is run starts a new database backup series. A number identifies each backup series.

database snapshot

A complete backup of the entire database to media that can be taken off-site. When a database snapshot is created, the current database backup series is not interrupted. A database snapshot cannot have incremental database backups associated with it. See also *database backup series*. Contrast with *full backup*.

data deduplication

A method of reducing storage needs by eliminating redundant data. Only one instance of the data is retained on storage

media. Other instances of the same data are replaced with a pointer to the retained instance.

data manager server

A server that collects metadata information for client inventory and manages transactions for the storage agent over the local area network. The data manager server informs the storage agent with applicable library attributes and the target volume identifier.

data mover

A device that moves data on behalf of the server. A network-attached storage (NAS) file server is a data mover.

data storage-management application-programming interface (DSMAPI)

A set of functions and semantics that can monitor events on files, and manage and maintain the data in a file. In an HSM environment, a DSMAPI uses events to notify data management applications about operations on files, stores arbitrary attribute information with a file, supports managed regions in a file, and uses DSMAPI access rights to control access to a file object.

default management class

A management class that is assigned to a policy set. This class is used to govern backed up or archived files when a file is not explicitly associated with a specific management class through the include-exclude list.

deduplication

See *data deduplication*.

demand migration

The process that is used to respond to an out-of-space condition on a file system for which hierarchical storage management (HSM) is active. Files are migrated to server storage until space usage drops to the low threshold that was set for the file system. If the high threshold and low threshold are the same, one file is migrated.

desktop client

The group of backup-archive clients that includes clients on Microsoft Windows, Apple, and Novell NetWare operating systems.

destination

A copy group or management class attribute that specifies the primary storage pool to which a client file will be backed up, archived, or migrated.

device class

A named set of characteristics that are applied to a group of storage devices. Each device class has a unique name and represents a device type of disk, file, optical disk, or tape.

device configuration file

(1) For a server, a file that contains information about defined device classes, and, on some servers, defined libraries and drives. The information is a copy of the device configuration information in the database.

(2) For a storage agent, a file that contains the name and password of the storage agent, and information about the server that is managing the SAN-attached libraries and drives that the storage agent uses.

device driver

A program that provides an interface between a specific device and the application program that uses the device.

disaster recovery manager (DRM)

A function that assists in preparing and using a disaster recovery plan file for the server.

disaster recovery plan

A file that is created by the disaster recovery manager (DRM) that contains information about how to recover computer systems if a disaster occurs and scripts that can be run to perform some recovery tasks. The file includes information about the software and hardware that is used by the server, and the location of recovery media.

domain

A grouping of client nodes with one or more policy sets, which manage data or storage resources for the client nodes. See *policy domain* or *client domain*.

DRM See *disaster recovery manager*.

DSMAPI

See *data storage-management application-programming interface*.

dynamic serialization

A type of copy serialization in which a file or folder is backed up or archived on the first attempt regardless of whether it changes during a backup or archive.

E

EA See *extended attribute*.

EB See *exabyte*.

EFS See *Encrypted File System*.

Encrypted File System (EFS)

A file system that uses file system-level encryption.

enterprise configuration

A method of setting up servers so that the administrator can distribute the configuration of one of the servers to the other servers, using server-to-server communication. See also *configuration manager*, *managed server*, *profile*, and *subscription*.

enterprise logging

The process of sending events from a Tivoli Storage Manager server to a designated event server. The event server routes the events to designated receivers, such as to a user exit. See also *event*.

error log

A data set or file that is used to record error information about a product or system.

estimated capacity

The available space, in megabytes, of a storage pool.

event (1) An administrative command or a client operation that is scheduled to be run using Tivoli Storage Manager scheduling.

(2) A message that an Tivoli Storage Manager server or client issues. Messages can be logged using Tivoli Storage Manager event logging.

event record

A database record that describes actual status and results for events.

event server

A server to which other servers can send events for logging. The event server routes the events to any receivers that are enabled for the sending server's events.

exabyte (EB)

For processor storage, real and virtual storage, and channel volume, 1 152 921 504 606 846 976 bytes. For disk storage capacity and communications volume, 1 000 000 000 000 000 000 bytes.

exclude

The process of identifying files in an include-exclude list. This process prevents the files from being backed up or migrated whenever a user or schedule enters an incremental or selective backup operation. A file can be excluded from backup and space management, backup only, or space management only.

exclude-include list

See *include-exclude list*.

execution mode

A mode that controls the space-management related behavior of commands that run under the **dsmmode** command.

expiration

The process by which files, data sets, or objects are identified for deletion because their expiration date or retention period has passed.

expiring file

A migrated or premigrated file that has been marked for expiration and removal from storage. If a stub file or an original copy of a premigrated file is deleted from a local file system, or if the original copy of a premigrated file is updated, the corresponding migrated or premigrated file is marked for expiration the next time reconciliation is run.

extend

To increase the portion of available space that can be used to store database or recovery log information.

extended attribute (EA)

Names or value pairs that are associated with files or directories. There are three classes of extended attributes: user attributes, system attributes, and trusted attributes.

extent The part of a file that is created during the data-deduplication process. Extents are compared with other file extents to identify duplicates.

external library

A type of library that is provided by Tivoli Storage Manager that permits LAN-free data movement for StorageTek libraries that are managed by Automated Cartridge System Library Software (ACSL). To activate this function, the Tivoli Storage Manager library type must be EXTERNAL.

F**file access time**

On AIX, UNIX, or Linux systems, the time when the file was last accessed.

file age

For migration prioritization purposes, the number of days since a file was last accessed.

file device type

A device type that specifies the use of sequential access files on disk storage as volumes.

file server

A dedicated computer and its peripheral storage devices that are connected to a local area network that stores programs and files that are shared by users on the network.

file space

A logical space in server storage that contains a group of files that have been backed up or archived by a client node, from a single logical partition, file system, or virtual mount point. Client nodes can restore, retrieve, or delete their file spaces from server storage. In server storage, files belonging to a single file space are not necessarily stored together.

file space ID (FSID)

A unique numeric identifier that the server assigns to a file space when it is stored in server storage.

file state

The space management mode of a file that resides in a file system to which space management has been added. A file can be in one of three states: resident, premigrated, or migrated. See also *resident file*, *premigrated file*, and *migrated file*.

file system migrator (FSM)

A kernel extension that intercepts all file system operations and provides any space

management support that is required. If no space management support is required, the operation is passed to the operating system, which performs its normal functions. The file system migrator is mounted over a file system when space management is added to the file system.

file system state

The storage management mode of a file system that resides on a workstation on which the hierarchical storage management (HSM) client is installed. A file system can be in one of these states: native, active, inactive, or global inactive.

frequency

A copy group attribute that specifies the minimum interval, in days, between incremental backups.

FSID See *file space ID*.

FSM See *file system migrator*.

full backup

The process of backing up the entire server database. A full backup begins a new database backup series. See also *database backup series* and *incremental backup*. Contrast with *database snapshot*.

fuzzy backup

A backup version of a file that might not accurately reflect what is currently in the file because the file was backed up at the same time as it was being modified.

fuzzy copy

A backup version or archive copy of a file that might not accurately reflect the original contents of the file because it was backed up or archived the file while the file was being modified. See also *backup version* and *archive copy*.

G**General Parallel File System**

A high-performance shared-disk file system that can provide data access from nodes in a cluster environment.

gigabyte (GB)

In decimal notation, 1 073 741 824 when referring to memory capacity; in all other cases, it is defined as 1 000 000 000.

global inactive state

The state of all file systems to which

space management has been added when space management is globally deactivated for a client node. When space management is globally deactivated, hierarchical storage management (HSM) cannot perform migration, recall, or reconciliation. However, a root user can update space management settings and add space management to additional file systems. Users can access resident and premigrated files.

Globally Unique Identifier (GUID)

An algorithmically determined number that uniquely identifies an entity within a system.

GPFS™

See *General Parallel File System*.

GPFS node set

A mounted, defined group of GPFS file systems.

group backup

The backup of a group containing a list of files from one or more file space origins.

GUID See *Globally Unique Identifier*.

H

hierarchical storage management (HSM)

A function that automatically distributes and manages data on disk, tape, or both by regarding devices of these types and potentially others as levels in a storage hierarchy that range from fast, expensive devices to slower, cheaper, and possibly removable devices. The objectives are to minimize access time to data and maximize available media capacity.

hierarchical storage management (HSM) client

A client program that works with the Tivoli Storage Manager server to provide hierarchical storage management (HSM) for a system. See also *hierarchical storage management* and *space manager client*.

HSM See *hierarchical storage management*.

HSM client

See *hierarchical storage management client*.

I

ILM See *information lifecycle management*.

image A file system or raw logical volume that is backed up as a single object.

image backup

A backup of a full file system or raw logical volume as a single object.

inactive file system

A file system for which space management has been deactivated. Contrast with *active file system*.

inactive version

A backup version of a file that is either not the most recent backup version, or that is a backup version of a file that no longer exists on the client system. Inactive backup versions are eligible for expiration processing according to the management class assigned to the file. Contrast with *active version*.

include-exclude file

A file containing statements to determine the files to back up and the associated management classes to use for backup or archive. See also *include-exclude list*.

include-exclude list

A list of options that include or exclude selected files for backup. An exclude option identifies files that should not be backed up. An include option identifies files that are exempt from the exclusion rules or assigns a management class to a file or a group of files for backup or archive services.

incremental backup

(1) A copy of all database data that has changed since the most recent successful full backup operation. An incremental backup is also known as a *cumulative backup image* because each incremental backup includes the contents of the previous incremental backup.

(2) The process of backing up information in the database that is new or changed since the last full backup. Contrast with *full backup*. See also *database backup series*.

(3) For Data Protection for Microsoft Exchange Server, a backup in which the transaction logs are backed up and then cleared.

individual mailbox restore

See *mailbox restore*.

information lifecycle management (ILM)

GPFS policy-based file management for storage pools and file sets.

inode The internal structure that describes the individual files on AIX, UNIX, or Linux systems. An inode contains the node, type, owner, and location of a file.

inode number
A number specifying a particular inode file in the file system.

IP address
A unique address for a device or logical unit on a network that uses the IP standard.

J

job file
A generated file that contains configuration information for a migration job. The file is XML format and can be created and edited in the hierarchical storage management (HSM) client for Windows client graphical user interface.

journal-based backup
A method for backing up Windows clients and AIX clients that exploits the change notification mechanism in a file to improve incremental backup performance by reducing the need to fully scan the file system.

journal daemon
On AIX, UNIX, or Linux systems, a program that tracks change activity for files residing in file systems.

journal service
In Microsoft Windows, a program that tracks change activity for files residing in file systems.

K

kilobyte (KB)
For processor storage, real and virtual storage, and channel volume, 210 or 1 024 bytes. For disk storage capacity and communications volume, 1 000 bytes.

L

LAN See *local area network*.

LAN-free data movement
The movement of client data between a client system and a storage device on a storage area network (SAN), bypassing the local area network. This process is also referred to as *LAN-free data transfer*.

LAN-free data transfer
See *LAN-free data movement*.

leader data
Bytes of data, from the beginning of a migrated file, that are stored in the file's corresponding stub file on the local file system. The amount of leader data that is stored in a stub file depends on the stub size that is specified.

library
(1) A repository for demountable recorded media, such as magnetic disks and magnetic tapes.
(2) A collection of one or more drives, and possibly robotic devices (depending on the library type), which can be used to access storage volumes.

library client
A server that uses server-to-server communication to access a library that is managed by another storage management server. See also *library manager*.

library manager
A server that controls device operations when multiple storage management servers share a storage device. See also *library client*.

local
(1) Pertaining to a device, file, or system that is accessed directly from a user system, without the use of a communication line.
(2) For HSM products, pertaining to the destination of migrated files that are being moved.

local area network (LAN)
A network that connects several devices in a limited area (such as a single building or campus) and that can be connected to a larger network.

local shadow volumes
Data that is stored on shadow volumes localized to a disk storage subsystem.

LOFS See *loopback virtual file system*.

logical file
A file that is stored in one or more server storage pools, either by itself or as part of an aggregate. See also *aggregate* and *physical file*.

logical occupancy
The space that is used by logical files in a

storage pool. This space does not include the unused space created when logical files are deleted from aggregate files, so it might be less than the physical occupancy.

logical unit (LU)

An access point through which a user or application program accesses the Systems Network Architecture (SNA) network to communicate with another user or application program.

logical unit number (LUN)

In the Small Computer System Interface (SCSI) standard, a unique identifier that is used to differentiate devices, each of which is a logical unit (LU).

logical volume

A portion of a physical volume that contains a file system.

logical volume backup

A backup of a file system or logical volume as a single object.

Logical Volume Snapshot Agent (LVSA)

Software that can act as the snapshot provider for creating a snapshot of a logical volume during an online image backup.

loopback virtual file system (LOFS)

A file system that is created by mounting a directory over another local directory, also known as mount-over-mount. A LOFS can also be generated using an automounter.

LU See *logical unit*.

LUN See *logical unit number*.

LVSA See *Logical Volume Snapshot Agent*.

M

macro file

A file that contains one or more storage manager administrative commands, which can be run only from an administrative client using the MACRO command. Contrast with *Tivoli Storage Manager command script*.

mailbox restore

A function that restores Microsoft Exchange Server data (from IBM Data Protection for Microsoft Exchange backups) at the mailbox level or mailbox-item level.

managed object

In Tivoli Storage Manager, a definition in the database of a managed server that was distributed to the managed server by a configuration manager. When a managed server subscribes to a profile, all objects that are associated with that profile become managed objects in the database of the managed server. In general, a managed object cannot be modified locally on the managed server. Objects can include policy, schedules, client option sets, server scripts, administrator registrations, server definitions, and server group definitions.

managed server

A Tivoli Storage Manager server that receives configuration information from a configuration manager using a subscription to one or more profiles. Configuration information can include definitions of objects such as policy and schedules. See also *configuration manager*, *subscription*, and *profile*.

management class

A policy object that users can bind to each file to specify how the server manages the file. The management class can contain a backup copy group, an archive copy group, and space management attributes. See also *copy group*, *space manager client*, *bind*, and *rebind*.

maximum transmission unit

The largest possible unit of data that can be sent on a given physical medium in a single frame. For example, the maximum transmission unit for Ethernet is 1500 bytes.

MB See *megabyte*.

media server

In a z/OS[®] environment, a program that provides access to z/OS disk and tape storage for Tivoli Storage Manager servers that run on operating systems other than z/OS.

megabyte (MB)

(1) 1 048 576 bytes (2 to the 20th power) when used in this publication.

(2) For processor storage, real and virtual storage, and channel volume, 2 to the power of 20 or 1 048 576 bits. For disk

storage capacity and communications volume, 1 000 000 bits.

metadata

Data that describes the characteristics of data; descriptive data.

migrate

To move data from one storage location to another. In Tivoli Storage Manager products, migrating can mean moving data from a client node to server storage, or moving data from one storage pool to the next storage pool defined in the server storage hierarchy. In both cases the movement is controlled by policy, such as thresholds that are set. See also *migration threshold*.

migrated file

A file that has been copied from a local file system to Tivoli Storage Manager storage. For HSM clients on UNIX or Linux systems, the file is replaced with a stub file on the local file system. On Windows systems, creation of the stub file is optional. See also *stub file* and *resident file*. For HSM clients on UNIX or Linux systems, contrast with *premigrated file*.

migrate-on-close recall mode

A mode that causes a migrated file to be recalled back to its originating file system temporarily. Contrast with *normal recall mode* and *read-without-recall recall mode*.

migration job

A specification of files to migrate, and actions to perform on the original files after migration. See also *job file*.

migration threshold

High and low capacities for storage pools or file systems, expressed as percentages, at which migration is set to start and stop.

mirroring

The process of writing the same data to multiple locations at the same time. Mirroring data protects against data loss within the recovery log.

mode

A copy group attribute that specifies whether to back up a file that has not been modified since the last time the file was backed up. See *modified mode* and *absolute mode*.

modified mode

In storage management, a backup copy-group mode that specifies that a file is considered for incremental backup only if it has changed since the last backup. A file is considered a changed file if the date, size, owner, or permissions of the file have changed. See also *absolute mode*.

mount limit

The maximum number of volumes that can be simultaneously accessed from the same device class. The mount limit determines the maximum number of mount points. See also *mount point*.

mount point

On the Tivoli Storage Manager server, a logical drive through which volumes in a sequential access device class are accessed. For removable-media device types, such as tape, a mount point is a logical drive that is associated with a physical drive. For the file device type, a mount point is a logical drive that is associated with an I/O stream. The number of mount points for a device class is defined by the value of the mount limit attribute for that device class. See also *mount limit*.

mount retention period

The maximum number of minutes that the server retains a mounted sequential-access media volume that is not being used before it dismounts the sequential-access media volume.

mount wait period

The maximum number of minutes that the server waits for a sequential-access volume mount request to be satisfied before canceling the request.

MTU See *maximum transmission unit*.

N**Nagle algorithm**

An algorithm that reduces congestion of TCP/IP networks by combining smaller packets and sending them together.

named pipe

A type of interprocess communication that permits message data streams to pass between peer processes, such as between a client and a server.

NAS See *network-attached storage*.

NAS node

A client node that is a network-attached storage (NAS) file server. Data for the NAS node is transferred by a NAS file server that is controlled by the network data management protocol (NDMP). A NAS node is also called a NAS file server node.

native file system

A file system that is locally added to the file server and is not added for space management. The hierarchical storage manager (HSM) client does not provide space management services to the file system.

native format

A format of data that is written to a storage pool directly by the Tivoli Storage Manager server. Contrast with *non-native data format*.

NDMP

See *Network Data Management Protocol*.

NetBIOS

See *Network Basic Input/Output System*.

network-attached storage (NAS) file server

A dedicated storage device with an operating system that is optimized for file-serving functions. A NAS file server can have the characteristics of both a node and a data mover.

Network Basic Input/Output System (NetBIOS)

A standard interface to networks and personal computers that is used on local area networks to provide message, print-server, and file-server functions. Application programs that use NetBIOS do not have to handle the details of LAN data link control (DLC) protocols.

Network Data Management Protocol (NDMP)

A protocol that allows a network storage-management application to control the backup and recovery of an NDMP-compliant file server, without installing vendor-acquired software on that file server.

network data-transfer rate

A rate that is calculated by dividing the total number of bytes that are transferred by the data transfer time. For example, this rate can be the time that is spent transferring data over a network.

node A file server or workstation on which the backup-archive client program has been installed, and which has been registered to the server.

node name

A unique name that is used to identify a workstation, file server, or PC to the server.

node privilege class

A privilege class that gives an administrator the authority to remotely access backup-archive clients for a specific client node or for all clients in a policy domain. See also *privilege class*.

non-native data format

A format of data that is written to a storage pool that differs from the format that the server uses for operations.

normal recall mode

A mode that causes a migrated file to be copied back to its originating file system when it is accessed.

O**offline volume backup**

A backup in which the volume is locked so that no other system applications can access it during the backup operation.

online volume backup

A backup in which the volume is available to other system applications during the backup operation.

open registration

A registration process in which users can register their workstations as client nodes with the server. Contrast with *closed registration*.

operator privilege class

A privilege class that gives an administrator the authority to disable or halt the server, enable the server, cancel server processes, and manage removable media. See also *privilege class*.

options file

A file that contains processing options. On Windows and NetWare systems, the file is called *dsm.opt*. On AIX, UNIX, Linux, and Mac OS X systems, the file is called *dsm.sys*.

originating file system

The file system from which a file was

migrated. When a file is recalled using normal or migrate-on-close recall mode, it is always returned to its originating file system.

orphaned stub file

A file for which no migrated file can be found on the Tivoli Storage Manager server that the client node is contacting for space management services. For example, a stub file can be orphaned when the client system-options file is modified to contact a server that is different than the one to which the file was migrated.

out-of-space protection mode

A mode that controls whether the program intercepts out-of-space conditions. See also *execution mode*.

P

pacing

In SNA, a technique by which the receiving system controls the rate of transmission of the sending system to prevent overrun.

packet In data communication, a sequence of binary digits, including data and control signals, that is transmitted and switched as a composite whole.

page A defined unit of space on a storage medium or within a database volume.

partial-file recall mode

A recall mode that causes the hierarchical storage management (HSM) function to read just a portion of a migrated file from storage, as requested by the application accessing the file.

password generation

A process that creates and stores a new password in an encrypted password file when the old password expires. Automatic generation of a password prevents password prompting. Password generation can be set in the options file (passwordaccess option). See also *options file*.

path An object that defines a one-to-one relationship between a source and a destination. Using the path, the source accesses the destination. Data can flow from the source to the destination, and back. An example of a source is a data

mover (such as a network-attached storage [NAS] file server), and an example of a destination is a tape drive.

pattern-matching character

See *wildcard character*.

physical file

A file that is stored in one or more storage pools, consisting of either a single logical file, or a group of logical files that are packaged together as an aggregate. See also *aggregate* and *logical file*.

physical occupancy

The amount of space that is used by physical files in a storage pool. This space includes the unused space that is created when logical files are deleted from aggregates. See also *physical file*, *logical file*, and *logical occupancy*.

plug-in

A self-contained software component that modifies (adds, or changes) the function in a particular system. When a plug-in is added to a system, the foundation of the original system remains intact.

policy domain

A grouping of policy users with one or more policy sets, which manage data or storage resources for the users. The users are client nodes that are associated with the policy domain.

policy privilege class

A privilege class that gives an administrator the authority to manage policy objects, register client nodes, and schedule client operations for client nodes. Authority can be restricted to certain policy domains. See also *privilege class*.

policy set

A group of rules in a policy domain. The rules specify how data or storage resources are automatically managed for client nodes in the policy domain. Rules can be contained in management classes. See also *active policy set* and *management class*.

premigrated file

A file that has been copied to Tivoli Storage Manager storage, but has not been replaced with a stub file on the local file system. An identical copy of the file resides both on the local file system and

in Tivoli Storage Manager storage. Premigrated files occur on UNIX and Linux file systems to which space management has been added. Contrast with *migrated file* and *resident file*.

premigrated files database

A database that contains information about each file that has been premigrated to Tivoli Storage Manager storage. The database is stored in a hidden directory named `.SpaceMan` in each file system to which space management has been added.

premigration

The process of copying files that are eligible for migration to Tivoli Storage Manager storage, but leaving the original file intact on the local file system.

premigration percentage

A space management setting that controls whether the next eligible candidates in a file system are premigrated following threshold or demand migration.

primary storage pool

A named set of volumes that the server uses to store backup versions of files, archive copies of files, and files migrated from client nodes. See also *destination* and *copy storage pool*.

privilege class

A level of authority that is granted to an administrator. The privilege class determines which administrative tasks the administrator can perform. See also *node privilege class*, *operator privilege class*, *policy privilege class*, *storage privilege class*, and *system privilege class*.

profile

A named group of configuration information that can be distributed from a configuration manager when a managed server subscribes. Configuration information can include registered administrator IDs, policies, client schedules, client option sets, administrative schedules, storage manager command scripts, server definitions, and server group definitions. See also *configuration manager* and *managed server*.

Q

quota (1) For HSM on AIX, UNIX, or Linux systems, the limit (in megabytes) on the

amount of data that can be migrated and premigrated from a file system to server storage.

(2) For HSM on Windows systems, a user-defined limit to the space that is occupied by recalled files.

R

randomization

The process of distributing schedule start times for different clients within a specified percentage of the schedule's startup window.

raw logical volume

A portion of a physical volume that is comprised of unallocated blocks and has no journaled file system (JFS) definition. A logical volume is read/write accessible only through low-level I/O functions.

read-without-recall recall mode

A mode that causes hierarchical storage management (HSM) to read a migrated file from storage without storing it back on the local file system. The last piece of information read from the file is stored in a buffer in memory on the local file system. Contrast with *normal recall mode* and *migrate-on-close recall mode*.

rebind

To associate all backed-up versions of a file with a new management class name. For example, a file that has an active backup version is rebound when a later version of the file is backed up with a different management class association. See also *bind*.

recall In Tivoli Storage Manager, to copy a migrated file from server storage back to its originating file system using the space management client. See also *transparent recall*, *selective recall*, and *recall mode*.

recall mode

A mode that is assigned to a migrated file with the `dsmatrr` command that determines how the file is processed when it is recalled. It determines whether the file is stored on the local file system, is migrated back to Tivoli Storage Manager storage when it is closed, or is read from Tivoli Storage Manager storage without storing it on the local file system.

receiver

A server repository that contains a log of server and client messages as events. For example, a receiver can be a file exit, a user exit, or the Tivoli Storage Manager server console and activity log. See also *event*.

reclamation

The process of consolidating the remaining data from many sequential-access volumes onto fewer, new sequential-access volumes.

reclamation threshold

The percentage of space that a sequential-access media volume must have before the server can reclaim the volume. Space becomes reclaimable when files are expired or are deleted.

reconciliation

The process of synchronizing a file system with the Tivoli Storage Manager server, and then removing old and obsolete objects from the Tivoli Storage Manager server.

recovery log

A log of updates that are about to be written to the database. The log can be used to recover from system and media failures. The recovery log consists of the active log (including the log mirror) and archive logs.

register

To define a client node or administrator ID that can access the server.

registry

A repository that contains access and configuration information for users, systems, and software.

remote

- (1) Pertaining to a system, program, or device that is accessed through a communication line.
- (2) For HSM products, pertaining to the origin of migrated files that are being moved.

resident file

On a Windows system, a complete file on a local file system that might also be a migrated file because a migrated copy can exist in Tivoli Storage Manager storage. On a UNIX or Linux system, a complete

file on a local file system that has not been migrated or premigrated, or that has been recalled from Tivoli Storage Manager storage and modified. Contrast with *stub file* and *premigrated file*. See *migrated file*.

restore

To copy information from its backup location to the active storage location for use. For example, to copy information from server storage to a client workstation.

retention

The amount of time, in days, that inactive backed-up or archived files are kept in the storage pool before they are deleted. Copy group attributes and default retention grace periods for the domain define retention.

retrieve

To copy archived information from the storage pool to the workstation for use. The retrieve operation does not affect the archive version in the storage pool.

roll back

To remove changes that were made to database files since the last commit point.

root user

A system user who operates without restrictions. A root user has the special rights and privileges needed to perform administrative tasks.

S

SAN See *storage area network*.

schedule

A database record that describes client operations or administrative commands to be processed. See *administrative command schedule* and *client schedule*.

scheduling mode

The type of scheduling operation for the server and client node that supports two scheduling modes: client-polling and server-prompted.

scratch volume

A labeled volume that is either blank or contains no valid data, that is not defined, and that is available for use.

script

A series of commands, combined in a file, that carry out a particular function when the file is run. Scripts are interpreted as

they are run. Contrast with *Tivoli Storage Manager command script*.

Secure Sockets Layer (SSL)

A security protocol that provides communication privacy. With SSL, client/server applications can communicate in a way that is designed to prevent eavesdropping, tampering, and message forgery.

selective backup

The process of backing up certain files or directories from a client domain. The files that are backed up are those that are not excluded in the include-exclude list. The files must meet the requirement for serialization in the backup copy group of the management class that is assigned to each file. Contrast with *incremental backup*.

selective migration

The process of copying user-selected files from a local file system to Tivoli Storage Manager storage and replacing the files with stub files on the local file system. Contrast with *threshold migration* and *demand migration*.

selective recall

The process of copying user-selected files from Tivoli Storage Manager storage to a local file system. Contrast with *transparent recall*.

serialization

The process of handling files that are modified during backup or archive processing. See *dynamic serialization*, *static serialization*, *shared static serialization*, and *shared dynamic serialization*.

server A software program or a computer that provides services to other software programs or other computers.

server options file

A file that contains settings that control various server operations. These settings affect such things as communications, devices, and performance.

server-prompted scheduling mode

A client/server communication technique where the server contacts the client node when tasks must be done. Contrast with *client-polling scheduling mode*.

server storage

The primary, copy, and active-data storage

pools that are used by the server to store user files such as backup versions, archive copies, and files migrated from space manager client nodes (space-managed files). See also *active-data pool*, *primary storage pool*, *copy storage pool*, *storage pool volume*, and *volume*.

session

A logical or virtual connection between two stations, software programs, or devices on a network that allows the two elements to communicate and exchange data.

session resource usage

The amount of wait time, processor time, and space that is used or retrieved during a client session.

shared dynamic serialization

A value for serialization that specifies that a file must not be backed up or archived if it is being modified during the operation. Tivoli Storage Manager retries the backup or archive operation a number of times; if the file is being modified during each attempt, Tivoli Storage Manager will back up or archive the file on its last try. See also *serialization*. Contrast with *dynamic serialization*, *shared static serialization*, and *static serialization*.

shared library

A library device that is used by multiple storage manager servers.

shared static serialization

A copy-group serialization value that specifies that a file must not be modified during a backup or archive operation. Tivoli Storage Manager attempts to retry the operation a number of times. If the file is in use during each attempt, the file is not backed up or archived. See also *serialization*. Contrast with *dynamic serialization*, *shared dynamic serialization*, and *static serialization*.

snapshot

An image backup type that consists of a point-in-time view of a volume.

space-managed file

A file that is migrated from a client node by the space manager client. The space manager client recalls the file to the client node on demand.

space management

The process of keeping sufficient free storage space available on a local file system for new data by migrating files to server storage. Synonymous with *hierarchical storage management*.

space manager client

A program that runs on a UNIX or Linux system to manage free space on the local file system by migrating files to server storage. The program can recall the files either automatically or selectively. Also called *hierarchical storage management (HSM) client*.

space monitor daemon

A daemon that checks space usage on all file systems for which space management is active, and automatically starts threshold migration when space usage on a file system equals or exceeds its high threshold.

sparse file

A file that is created with a length greater than the data it contains, leaving empty spaces for the future addition of data.

special file

On AIX, UNIX, or Linux systems, a file that defines devices for the system, or temporary files that are created by processes. There are three basic types of special files: first-in, first-out (FIFO); block; and character.

SSL See *Secure Sockets Layer*.

stabilized file space

A file space that exists on the server but not on the client.

stanza A group of lines in a file that together have a common function or define a part of the system. Each stanza is identified by a name that occurs in the first line of the stanza. Depending on the type of file, a stanza is ended by the next occurrence of a stanza name in the file, or by an explicit end-of-stanza marker. A stanza can also be ended by the end of the file.

startup window

A time period during which a schedule must be initiated.

static serialization

A copy-group serialization value that specifies that a file must not be modified

during a backup or archive operation. If the file is in use during the first attempt, the storage manager cannot back up or archive the file. See also *serialization*. Contrast with *dynamic serialization*, *shared dynamic serialization*, and *shared static serialization*.

storage agent

A program that enables the backup and restoration of client data directly to and from storage attached to a storage area network (SAN).

storage area network (SAN)

A dedicated storage network that is tailored to a specific environment, combining servers, systems, storage products, networking products, software, and services.

storage hierarchy

(1) A logical order of primary storage pools, as defined by an administrator. The order is typically based on the speed and capacity of the devices that the storage pools use. The storage hierarchy is defined by identifying the next storage pool in a storage pool definition. See also *storage pool*.

(2) An arrangement of storage devices with different speeds and capacities. The levels of the storage hierarchy include: main storage, such as memory and direct-access storage device (DASD) cache; primary storage (DASD containing user-accessible data); migration level 1 (DASD containing data in a space-saving format); and migration level 2 (tape cartridges containing data in a space-saving format).

storage pool

A named set of storage volumes that are the destination that is used to store client data. A storage pool contains backup versions, archive copies, and files that are migrated from space manager client nodes. A primary storage pool is backed up to a copy storage pool. See also *primary storage pool*, *copy storage pool*, and *active-data pool*.

storage pool volume

A volume that has been assigned to a storage pool. See also *volume*, *active-data pool*, *copy storage pool*, and *primary storage pool*.

storage privilege class

A privilege class that gives an administrator the authority to control how storage resources for the server are allocated and used, such as monitoring the database, the recovery log, and server storage. See also *privilege class*.

stub A shortcut on the Windows file system that is generated by the hierarchical storage management (HSM) client for a migrated file that allows transparent user access. A stub is the sparse file representation of a migrated file, with a reparse point attached.

stub file

A file that replaces the original file on a local file system when the file is migrated to storage. A stub file contains the information that is necessary to recall a migrated file from Tivoli Storage Manager storage. It also contains additional information that can be used to eliminate the need to recall a migrated file.

stub file size

The size of a file that replaces the original file on a local file system when the file is migrated to Tivoli Storage Manager storage. The size that is specified for stub files determines how much leader data can be stored in the stub file. The default for stub file size is the block size defined for a file system minus 1 byte.

subscription

In a Tivoli environment, the process of identifying the subscribers that the profiles are distributed to. For Tivoli Storage Manager, a subscription is the process by which a managed server receives configuration information associated with a particular profile on a configuration manager. See also *managed server*, *configuration manager*, and *profile*.

system privilege class

A privilege class that gives an administrator the authority to issue all server commands. See also *privilege class*.

Systems Network Architecture (SNA)

The description of the logical structure, formats, protocols, and operational sequences for transmitting information through and controlling the configuration and operation of networks.

T**tape library**

A set of equipment and facilities that support an installation's tape environment. The tape library can include tape storage racks, mechanisms for automatic tape mounting, a set of tape drives, and a set of related tape volumes mounted on those drives.

tape volume prefix

The high-level-qualifier of the file name or the data set name in the standard tape label.

target node

A client node for which other client nodes (called agent nodes) have been granted proxy authority. The proxy authority allows the agent nodes to perform operations such as backup and restore on behalf of the target node, which owns the data.

TCA See *trusted communications agent*.

TCP/IP

See *Transmission Control Protocol/Internet Protocol*.

threshold migration

The process of moving files from a local file system to Tivoli Storage Manager storage based on the high and low thresholds that are defined for the file system. Contrast with *demand migration*, *selective migration*, and *migration job*.

throughput

In storage management, the total bytes in the workload, excluding overhead, that are backed up or restored, divided by elapsed time.

timeout

A time interval that is allotted for an event to occur or complete before operation is interrupted.

timestamp control mode

A mode that determines whether commands preserve the access time for a file or set it to the current time.

Tivoli Storage Manager command script

A sequence of Tivoli Storage Manager administrative commands that are stored in the database of the Tivoli Storage Manager server. The script can run from any interface to the server. The script can

include substitution for command parameters and conditional logic.

tombstone object

A small subset of attributes of a deleted object. The tombstone object is retained for a specified period, and at the end of the specified period, the tombstone object is permanently deleted.

Transmission Control Protocol/Internet Protocol (TCP/IP)

An industry-standard, nonproprietary set of communication protocols that provides reliable end-to-end connections between applications over interconnected networks of different types.

transparent recall

The process that is used to automatically recall a file to a workstation or file server when the file is accessed. See also *recall mode*. Contrast with *selective recall*.

trusted communications agent (TCA)

A program that handles the sign-on password protocol when clients use password generation.

U

UCS-2 A 2-byte (16-bit) encoding scheme based on ISO/IEC specification 10646-1. UCS-2 defines three levels of implementation: Level 1-No combining of encoded elements allowed; Level 2-Combining of encoded elements is allowed only for Thai, Indic, Hebrew, and Arabic; Level 3-Any combination of encoded elements are allowed.

UNC See *Universal Naming Convention name*.

Unicode

A character encoding standard that supports the interchange, processing, and display of text that is written in the common languages around the world, plus some classical and historical texts. The Unicode standard has a 16-bit character set defined by ISO 10646.

Unicode-enabled file space

Unicode file space names provide support for multilingual workstations without regard for the current locale.

Unicode transformation format 8

Unicode Transformation Format (UTF), 8-bit encoding form, which is designed for ease of use with existing ASCII-based

systems. The CCSID value for data in UTF-8 format is 1208.

Universal Naming Convention (UNC) name

A name that is used to access a drive or directory containing files shared across a network. The UNC name includes the system name and a SharePoint name that represents the shared drive or directory.

Universally Unique Identifier (UUID)

The 128-bit numeric identifier that is used to ensure that two components do not have the same identifier.

UTF-8 See *Unicode transformation format 8*.

UUID See *Universally Unique Identifier*.

V**validate**

To check a policy set for conditions that can cause problems if that policy set becomes the active policy set. For example, the validation process checks whether the policy set contains a default management class.

version

A backup copy of a file stored in server storage. The most recent backup copy of a file is the active version. Earlier copies of the same file are inactive versions. The number of versions retained by the server is determined by the copy group attributes in the management class.

virtual file space

A representation of a directory on a network-attached storage (NAS) file system as a path to that directory.

virtual volume

An archive file on a target server that represents a sequential media volume to a source server.

volume

A discrete unit of storage on disk, tape or other data recording medium that supports some form of identifier and parameter list, such as a volume label or input/output control. See also *scratch volume*, and *storage pool volume*.

volume history file

A file that contains information about volumes that have been used by the server for database backups and for export of administrator, node, policy, or

server data. The file also has information about sequential-access storage pool volumes that have been added, reused, or deleted. The information is a copy of volume information that is recorded in the server database.

Volume Shadow Copy Service

A set of Microsoft application-programming interfaces (APIs) that you can use to create shadow copy backups of volumes, exact copies of files, including all open files, and so on.

VSS See *Volume Shadow Copy Service*.

VSS Backup

A backup operation that uses Microsoft Volume Shadow Copy Service (VSS) technology. The backup operation produces an online snapshot (point-in-time consistent copy) of Microsoft Exchange data. This copy can be stored on local shadow volumes or on Tivoli Storage Manager server storage.

VSS Fast Restore

A function that uses a Microsoft Volume Shadow Copy Service (VSS) software provider to restore VSS Backups (IBM Data Protection for Microsoft Exchange database files and log files) that reside on local shadow volumes.

VSS Instant Restore

A volume-level hardware-assisted Microsoft Volume Shadow Copy Service (VSS) function where target volumes that contain the snapshot are copied back to the original source volumes.

VSS offloaded backup

A backup operation that uses a Microsoft Volume Shadow Copy Service (VSS) hardware provider (installed on an alternate system) to move IBM Data Protection for Microsoft Exchange data to the Tivoli Storage Manager server. This type of backup operation shifts the backup load from the production system to another system.

VSS Restore

A function that uses a Microsoft Volume Shadow Copy Service (VSS) software provider to restore VSS Backups (IBM Data Protection for Microsoft Exchange database files and log files) that reside on

Tivoli Storage Manager server storage to their original location.

W**wildcard character**

A special character such as an asterisk (*) or a question mark (?) that can be used to represent one or more characters. Any character or set of characters can replace the wildcard character.

workstation

A terminal or personal computer at which a user can run applications and that is usually connected to a mainframe or a network.

worldwide name

A 64-bit, unsigned name identifier that is unique.

workload partition (WPAR)

A partition within a single operating system instance.

Index

A

- access to databases, restricting 176
- accessibility features for this
 - product 247
- active/inactive state
 - affected by full backup 124, 210
 - in restore operations 179
- adjustkbtmestimate parameter 127
- adjustpercentestimate parameter 127
- API, Tivoli Storage Manager 149
- authorization mode, setting
 - using the CLI 130, 131, 132, 154, 155, 183, 185, 227
- auto select option, GUI 88
- automatic expiration policy, setting 41

B

- backing up SQL databases
 - by using the Legacy method 81
- backup command
 - and /backupdestination parameter 127
 - and /backupmethod parameter 127
 - and /logfile parameter 206
 - and /logprune parameter 206
 - and /offload parameter 130
 - and /quiet parameter 207
- backup object types 5, 79
 - differential 124, 174, 210
 - file 124, 173, 210
 - for query Data Protection for SQL 150
 - full 124, 174, 210
 - group 125, 174, 210
 - log 125, 174, 210
 - set 126, 174, 211
- backup objects
 - compatibility with server 151
 - query of 149
- backup operations
 - overview 1
 - using the GUI
 - backup databases tab 80
 - backup groups/files tab 83
- backup strategy
 - Tivoli Storage Manager versus local shadow volumes 28
 - VSS and Legacy together 24
 - VSS cluster 25, 29, 30
- backupdestination parameter
 - and backup command 127
 - and delete backup command 205
 - and restore command 175
 - and set command 223
- backupmethod parameter
 - and backup command 127
 - and restore command 175
 - and restorefiles command 201
 - and set command 224

- binary sort order 152
- books
 - See publications
- buffering data
 - Data Protection for SQL performance 115
 - for Data Protection for SQL 128, 151, 175, 224
 - for SQL server 131, 183, 227
 - for TDP 128, 151, 175
 - used with stripes 131, 183, 227
- buffers parameter 128, 151, 175, 224
- buffersize parameter 128, 151, 175

C

- capacity
 - determining managed storage 77
- capturing a log of the installation 57
- Central Scheduler Service
 - using 38
- changetmpspassword command
 - positional parameters 230
- cluster
 - VSS 25, 29, 30
- cluster server
 - running Data Protection for SQL on 28
- clustering 27, 28
 - strategy 23
- clustering state
 - querying 149
- clusternode
 - options file value 39
- code page ID
 - querying 150
- collocation, using
 - for restore performance 23
 - with striping 23
- command
 - policy 233
- command line parameters
 - /backupdestination
 - and backup 127
 - and restore 175, 205
 - and set 223
 - /backupmethod
 - and backup 127
 - and restore 175
 - and set 224
 - /instantrestore
 - and restore 177
 - /localdsmagentnode
 - and set 225
 - /offload
 - and backup 130
 - /quiet
 - and backup 207
 - /remotedsmagentnode
 - and set 227
- command-line interface
 - overview 117
- command-line parameters
 - /backupdestination
 - and restorefiles 201
 - /configfile
 - and delete backup 205
 - and restorefiles 201
 - /fromsqlserver
 - and delete backup 206
 - and restorefiles 201
 - /into
 - and restorefiles 203
 - /logfile
 - and backup 206
 - and restorefiles 201
 - /logprune
 - and backup 206
 - and restorefiles 202
 - /mountwait
 - and restorefiles 202
 - /object
 - and delete backup 207
 - and restorefiles 202
 - /quiet
 - and restorefiles 203
 - /tsmnode
 - and restore 207
 - and restorefiles 203
 - /tsmoptfile
 - and restore 207
 - and restorefiles 203
 - /tsmpspassword
 - and restore 207
 - and restorefiles 204
 - and local 201
 - and tsm 201
 - and vss 201
- commands
 - query managedcapacity 167
- commands, Data Protection for SQL
 - tdpsqlc changetmpspassword 230
 - tdpsqlc help 219
 - tdpsqlc set 223
- communication protocol
 - options file values 37
- compatibility level
 - querying 149
- compatibilityinfo
 - query of Data Protection for SQL 150
- compatibilityinfo parameter 151
- compressing data 23
 - options file value 38
- compression
 - SQL Server 2008 44
- configfile parameter 128, 152, 176, 211, 229, 231
 - and delete backup command 205
 - and restorefiles command 201

- configuration
 - manual procedure
 - Off-loaded backups 63
 - SQL Server 62
 - Tivoli Storage Manager server 63
 - verify 64
 - options 37
- configuration file, Data Protection for SQL
 - query TDP 149
 - setting values, CLI 128, 152, 176, 211, 223, 229, 231
- configuration information, listing 151
- configuration preferences
 - Data Protection for SQL 67
- configuring
 - binding
 - policy 76
 - policy 76
 - quick instructions 45
- consistency checker 24
- conventions
 - typeface xviii
- copy group 32
- create index operation 124, 125
- custom settings 74
- customer support
 - contact xvi

D

- data and meta object values 32
- data compression
 - options file value 38
 - time-saving strategy 23
- data objects
 - in object naming 234
- data protection
 - SQL with VSS backup/restore
 - support 101, 102
 - determining the issue 103
 - gathering files before calling IBM 107
 - gathering information before calling IBM 106
 - tracing when using VSS 106
 - troubleshooting 108
- Data Protection for SQL
 - backup types 5, 79
 - buffer size 128, 151, 175, 224
 - commands 117
 - configuration file query 149
 - configuration file, setting
 - using the CLI 128, 152, 176, 211, 229, 231
 - configuration parameters 40, 41
 - configuring options 37
 - creating an installation package 58
 - creating an installation package on a DVD 58
 - features
 - new xxi
 - functions overview
 - backup 1
 - restore 7
 - installation 49
 - hardware requirements 47

- Data Protection for SQL (*continued*)
 - installation (*continued*)
 - software requirements 48
 - installing on a local system 49
 - installing the language packs 50
 - number of buffers 128, 151, 175, 224
 - options file
 - automatic expiration policy 41
 - clusternode 27, 28
 - planning 19
 - policy settings 19
 - protecting data 67
 - quick configuration 45
 - quick installation 45
 - reference 117
 - registering 34
 - running on MS Cluster Server 28
 - silent installation 52
 - silent installation with batch file 55
 - version query 150
- Data Protection for SQL scripts
 - adding 111
 - editing 111
 - viewing 111
- Data Protection for SQL silent installation
 - capturing a log 57
 - playing back the installation 58
 - setup error messages 59
- Data Protection for SQL tasks
 - automating 99
 - Data Protection for SQL tasks 99
- Data Protection for SQL trace and log files
 - viewing 112
- Data Protection for SQL VSS backup
 - policy binding 77
- data stripes
 - buffering 128, 151, 175, 224
- Data Protection for SQL
 - performance 115
 - querying 149
 - recommendations 23
 - used with /sqlbuffers 131, 133, 183, 187, 227
 - using the CLI 133, 187, 228
- data striping
 - buffering 128, 151, 175
- database
 - options set 149
 - querying 149
 - restoring master 87
 - restoring to alternate 86
 - restoring using the GUI 83, 85
- database name
 - restorefiles
 - command line 201
- database owner option, GUI 88
- dateformat parameter 224
- dbcc check options 24
- dboonly parameter 176
- definitions 251
- delete backup
 - storage group
 - command line 205
- delete backup command
 - and /backupdestination parameter 205

- delete backup command (*continued*)
 - and /configfile parameter 205
 - and /fromsqlserver parameter 206
 - and /object parameter 207
 - overview 204
 - syntax diagram 204
- deleting SQL Server VSS backups 91
- described 150
- diagnostics properties 70
- differential backup
 - overview 5, 79
- differential parameter
 - described 124, 150, 174, 210
- differential restore
 - overview 9
 - using the GUI 83, 85
- differential versus log backup
 - strategy 22
- diffestimate parameter 224
- disaster recovery plan 24
- disk-only storage pool
 - for query operations 32
- documentation
 - See publications
- domain 31
- DS6000
 - considerations 6
- DS8000
 - considerations 6
- dsm.opt file
 - configuring 37
 - enableclientencryptkey 39
 - creating multiple 40
 - Tivoli Storage Manager scheduler 93
- dsmcutil.exe utility
 - using 38

E

- education
 - see Tivoli technical training xiii
- enableclientencryptkey
 - options file value 39
- enableclientencryptkey option 39
- enablelanfree
 - options file value 39
- encryption 39
- error log files 101
- example
 - restorefiles command 204
- excludedb parameter 129
- expiration policy, setting 41
- expiring VSS Backup s
 - policy 76

F

- FAQ 241
- file backup
 - overview 5, 80
 - strategy 22
- file parameter
 - described 124, 150, 173, 210
- file restore
 - overview 9
 - using the GUI 83, 85

- files
 - Data Protection for SQL options 204, 207
 - dsm.opt 37, 93
 - sqlfull.log 98
 - sqlsched.log 98
 - tdpexc.cfg
 - and delete backup command 205
 - tdpexc.log
 - and delete backup command 206
 - tdpsql.cfg
 - and restorefiles command 201
 - setting 40
 - vsspolicy option 40, 41
 - tdpsql.log 245
 - and restorefiles command 201
 - tdpsqlc.exe 117
- fixes, obtaining xv
- frequently asked questions 241
- fromsqlserver parameter 152, 176, 211, 225
 - and delete backup command 206
 - and restorefiles command 201
- full backup
 - overview 5, 79
 - strategy 21
- full parameter
 - described 124, 150, 174, 210
- full plus differential plus log backup strategy 22
- full plus log backup strategy 21
- full restore
 - overview 9
 - using the GUI 83, 85

G

- general properties for SQL Server 72
- generate password, setting
 - options file value 37
- glossary 251
- graphical user interface (GUI)
 - backup databases tab 80
 - backup groups/files tab 83
 - inactivating SQL databases 90
 - restore databases tab 83
 - restore groups/files tab 85
 - restore options 88, 89
 - restoring SQL databases 83, 85
 - shortcut menus
 - for inactivate 91
 - for restore 89
- group backup
 - overview 6, 9, 80
 - strategy 22
- group parameter
 - described 125, 150, 174, 210
- group restore
 - using the GUI 83, 85

H

- hardware provider 4
- hardware requirements 47
- help command
 - described 219

- I**
 - IBM Publications Center x, xiii
 - IBM Support Assistant xv
 - inactivate operations
 - using the GUI 90
 - include/exclude
 - sample statements 234
 - syntax 42
 - indexes and tables
 - backing up 22
 - installation
 - configuring options 37
 - hardware requirements 47
 - prerequisites 47
 - registering Data Protection for SQL 34
 - software requirements 48
 - installationVSS Instant Restore VSS support 49
 - installing
 - creating an installation package 58
 - creating an installation package on a DVD 58
 - Data Protection for SQL language packs 50
 - on a local system 49
 - quick instructions 45
 - silently with batch file 55
 - silently with msiexec.exe 55
 - Tivoli Storage FlashCopy Manager 50
 - installing Data Protection for SQL
 - on multiple servers (silent) 52
 - unattended (silent) 52
 - instantrestore parameter
 - and restore command 177
 - integrated user id mode 130, 154, 183, 227
 - Internet, searching for problem resolution xiv, xv
 - into parameter 177
 - and restorefiles command 203

K

- knowledge bases, searching xiv

L

- LAN-free
 - Data Protection for SQL performance 116
- language packs 50
- Legacy backup
 - and VSS 24
 - hardware requirements 47
- Legacy restore
 - hardware requirements 47
- local shadow volumes
 - storage space 20
- localdsmagentnode parameter
 - and set command 225
- log backup
 - overview 5, 79
 - strategy 21, 24

- log files
 - using for problem determination 101
- log parameter
 - described 125, 150, 174, 210
- log restore
 - overview 9
 - using the GUI 83, 85
- log truncation strategy 23
- logfile parameter 129, 152, 177, 211, 225, 231
 - and delete backup command 206
 - and restorefiles command 201
- logging properties 73
- login settings
 - using the CLI 130, 154, 183, 227
- logprune parameter 129, 153, 178, 212, 226, 231
 - and delete backup command 206
 - and restorefiles command 202

M

- managed storage
 - determining capacity 77
- management class 32
- management class
 - include statements 234
 - meta and data objects 234
 - object naming 42
- manuals
 - See publications
- master database, restoring 87
- maxnummp parameter 24
- media migration 41
- media mounts
 - preventing for query 32
 - restore considerations 23
- messages
 - verification 66
- meta and data object values 32
- meta objects
 - in object naming 234
 - storing in storage pool 32
- metadata policy, setting 234
- Microsoft Windows Failover Clustering environment
 - supported 27
- mountwait parameter 130, 178
 - and restorefiles command 202
- mountwaitfordata parameter 226
- MSCS environment
 - converting to 28
- msiexec.exe
 - used for silent installation 55
- multiple mount points 24
- multiple SQL Servers
 - strategy 23
- multiple-user mode 98

N

- named-marks restore
 - overview 9
- node name
 - Data Protection for SQL recommendations 34
 - Off-loaded backup 36

- node name (*continued*)
 - proxy nodes 35
 - VSS 35
- nodename
 - options file value 37
- number of buffers
 - for Data Protection for SQL 128, 151, 175, 224
 - for SQL server 131, 183, 227
 - used with stripes 131, 183, 227
- numberformat parameter 226

O

- object parameter 153, 179, 212
 - and delete backup command 207
 - and restorefiles command 202
- Off-loaded backup
 - configuration procedure 63
 - hardware requirements 47
 - node names 36
- offload parameter
 - and backup command 130
- offloaded backup
 - description 7
- olderthan parameter 212
 - output sample 216
- options file, Data Protection for SQL
 - clustering 27, 28
 - include/exclude statements 234

P

- parameters
 - /backupdestination
 - and backup command 127
 - and delete backup command 205
 - and restore command 175
 - and restorefiles command 201
 - and set command 223
 - /backupmethod
 - and backup command 127
 - and restore command 175
 - and set command 224
 - /configfile
 - and delete backup command 205
 - and restorefiles command 201
 - /fromsqlserver
 - and delete backup command 206
 - and restorefiles command 201
 - /instantrestore
 - and restore command 177
 - /into
 - and restorefiles command 203
 - /localdsmagentnode
 - and set command 225
 - /logfile
 - and delete backup command 206
 - and restorefiles command 201
 - /logprune
 - and delete backup command 206
 - and restorefiles command 202
 - /mountwait
 - and restorefiles command 202
 - /object
 - and delete backup command 207

- parameters (*continued*)
 - /object (*continued*)
 - and restorefiles command 202
 - /offload
 - and backup command 130
 - /quiet
 - and delete backup command 207
 - and restorefiles command 203
 - /remotedsmagentnode
 - and set command 227
 - /tsmnode
 - and restore command 207
 - and restorefiles command 203
 - /tsmoptfile
 - and restore command 207
 - and restorefiles command 203
 - /tsmpassword
 - and restore command 207
 - and restorefiles command 204
- parameters, described
 - /to 181
 - optional
 - /adjustkbtstestimate 127
 - /buffers 128, 151, 175
 - /buffersize 128, 151, 175
 - /compatibilityinfo 151
 - /configfile 128, 152, 176, 211, 229, 231
 - /dboonly 176
 - /excludedb 129
 - /fromsqlserver 152, 176, 211
 - /into 177
 - /logfile 129, 152, 177, 211, 231
 - /logprune 129, 153, 178, 212, 231
 - /mountwait 130, 178
 - /object 153, 179, 212
 - /olderthan 212
 - /partial 179
 - /quiet 180, 213
 - /recovery 180
 - /relocate 181
 - /relocatedir 182
 - /replace 182
 - /restoredate 182
 - /restoretme 182
 - /sqlauthentication 130, 154, 183
 - /sqlbuffers 131, 183
 - /sqlbuffersize 131, 183
 - /sqlpassword 131, 154, 183
 - /sqlserver 131, 154, 184
 - /sqluser 132, 155, 185
 - /standby 185
 - /stopat 185
 - /stopatmark 186
 - /stopbeforemark 186
 - /stripes 133, 187
 - /tsmnode 133, 155, 188, 213, 232
 - /tsmoptfile 133, 156, 188, 213, 232
 - /tsmpassword 134, 156, 188, 213
 - SQLCOMPRESSSION 129
 - positional
 - backup object types 150
 - differential 124, 174, 210
 - file 124, 173, 210
 - for changetsmpassword command 230
 - for set command 223

- parameters, described (*continued*)
 - positional (*continued*)
 - full 124, 174, 210
 - group 125, 174, 210
 - log 125, 174, 210
 - set 126, 174, 211
 - partial parameter 179
 - partial restore option
 - overview 9
 - Passport Advantage xvi
 - password, Tivoli Storage Manager
 - changing
 - using the CLI 230
 - using dsmsutil to manage 38
 - passwordaccess generate, setting 37
 - performance issues 124, 173, 210
 - performance properties 74
 - performance recommendations 38
 - planning
 - Data Protection for SQL 19
 - point in time named marks restore
 - using the CLI 186
 - using the GUI 89
 - point in time restore option, GUI 89
 - point-in-time restore
 - overview 9
 - policy 75
 - binding 76
 - binding Data Protection for SQL VSS backups 77
 - configuring 76
 - expiring VSS Backup s 76
 - policy command
 - overview 233
 - policy set 31
 - policy settings
 - copy group 32
 - Data Protection for SQL and Tivoli Storage Manager 19
 - domain 31
 - management class 32
 - policy set 31
 - recommended 31
 - storage pool 34
 - printing reports 100
 - problem determination 101, 151
 - describing problem for IBM Software Support xvii
 - determining business impact for IBM Software Support xvi
 - submitting a problem to IBM Software xvii
 - properties
 - custom settings 74
 - diagnostics 70
 - general SQL Server 72
 - logging 73
 - performance 74
 - regional settings 73
 - SQL login 75
 - VSS backup 74
 - property pages for Data Protection for SQL 67
 - protecting data
 - Data Protection for SQL 67
 - proxy nodes 35

- publications
 - download x
 - order x
 - search x
 - Tivoli Storage FlashCopy Manager xiii
 - Tivoli Storage Manager x

Q

- query managedcapacity command
 - overview 167
- query operations
 - preventing media mount 32
 - query SQL 149
 - query TDP 149
 - query TSM 149
 - sample output 156
- query TSM options 148, 150
- querying backup objects 149
- quiet parameter 180, 213
 - and delete backup command 207
 - and restorefiles command 203

R

- recovery model
 - querying 149
- recovery option, GUI 88
- recovery parameter 180
- reference
 - Data Protection for SQL 117
- regional properties 73
- register node command 24
- registration 34
- relocate option
 - overview 9
- relocate option, GUI 90
- relocate...to parameters 181
- relocatedir option
 - overview 10
- relocatedir parameter 182
- remotedsmagentnode parameter
 - and set command 227
- removable media 24, 41
 - recommendations 32
- replace option, GUI 88, 89
- replace parameter 182
- reports
 - viewing, printing, and saving 100
- restore 9
- restore command
 - and /backupdestination parameter 175
 - and /backupmethod parameter 175
 - and /instantrestore parameter 177
 - and /tsmnode parameter 207
 - and /tsmoptfile parameter 207
 - and /tsmpassword parameter 207
- restore operations
 - collocation for performance 23
 - master database 87
 - named marks 89, 186
 - of inactive objects 179
 - overview 7
 - point in time 89, 185, 186

- restore operations (*continued*)
 - time-saving strategy 22, 23
 - to a different SQL Server 23
 - to alternate machine 86
 - using the GUI 83, 85
 - auto select option 88
 - database owner option 88
 - instant restore 88
 - point in time restore option 89
 - recovery option 88
 - relocate option 90
 - replace option 88, 89
 - restore into option 89
 - restore options 88
 - smart select option 88
 - standby server undo file option 90
 - stripes option 88
 - tab controls 83
 - wait for tape mounts options 89
- restore types 9
- restoredate parameter 182
- restorefiles command
 - and /backupmethod parameter 201
 - and /configfile parameter 201
 - and /fromsqlserver parameter 201
 - and /into parameter 203
 - and /logfile parameter 201
 - and /logprune parameter 202
 - and /mountwait parameter 202
 - and /object parameter 202
 - and /quiet parameter 203
 - and /tsmnode parameter 203
 - and /tsmoptfile parameter 203
 - and /tsmpassword parameter 204
 - example 204
- restorettime parameter 182
- restricting database access 176

S

- sample output
 - query command 156
 - set command 229
- SAN Volume Controller 6
 - considerations 6
- SAN Volume Controller configurations 13
- saving reports 100
- scheduler
 - example procedure 92
- scheduling backups 28
- scheduling Data Protection for SQL tasks
 - on the Tivoli Storage Manager server 92
- server, SQL
 - querying 149
- server, Tivoli Storage Manager
 - querying 149
 - using multiple 40
- set backup
 - overview 6, 9, 80
 - strategy 24
- set command
 - and /backupdestination parameter 223
 - and /backupmethod parameter 224
- set command (*continued*)
 - and /localdsmagentnode parameter 225
 - and /remotedsmagentnode parameter 227
 - positional parameters 223
 - sample output 229
- set parameter
 - described 126, 174, 211
- set restore
 - using the GUI 83, 85
- setup.exe
 - used for silent installation 54
- shortcut menus, GUI 89, 91
- silent installation of Data Protection for SQL 52
- silent installation
 - with setup.exe 54
- simple recovery model 24
- single-user mode 98
- size of buffers
 - for SQL server 131, 183, 227
 - for TDP 128, 151, 175, 224
- smart select option, GUI 88
- software provider 4
- software requirements 48
- software support
 - describing problem for IBM Software Support xvii
 - determining business impact for IBM Software Support xvi
 - submitting a problem xvii
- Software Support
 - contact xvi
- sort order ID
 - querying 150
- space allocated and used
 - querying 149, 150
- space-saving considerations 124, 173, 174, 210
- space, saving
 - strategy 22
- SQL login properties 75
- SQL server
 - number of buffers 131, 183, 227
 - size of buffers 131, 183, 227
 - user id 130, 154, 183, 227
- SQL Server
 - running in clustered environment 28
- SQL Server 2000
 - clustering recommendation 23
 - differential strategy 22
 - named-marks restore 9
 - partial restore 9, 89
 - point in time restore 89
 - query of 149
 - simple recovery model 24
- SQL Server 2008 compression 44
- SQL Server VSS backup
 - deleting 91
- SQL services 28
- sqlauthentication parameter 130, 154, 183, 227
- sqlbuffers parameter 131, 183
- sqlbuffersize parameter 131, 183, 227
- sqlcompression parameter 228
- SQLCOMPRESSSION parameter 129

- sqlfull.log 98
- sqlpassword parameter 131, 154, 183
- sqlsched.log 98
- sqlserver parameter 131, 154, 184, 228
- sqluser parameter 132, 155, 185
- sqluserid parameter 227
- standby parameter 185
- standby server undo file option, GUI 90
- stop (point in time) options, GUI 89
- stopat parameter 185
- stopatmark parameter 186
- stopbeforemark parameter 186
- storage
 - determining managed capacity 77
- storage group
 - delete backup
 - command line 205
- storage management, policy 75
- storage pool 34
 - for meta objects 32
- Storwize V7000
 - considerations 6
- Storwize® V7000 configurations 13
- stripes option, GUI 88
- stripes parameter 133, 187, 228
- stripes, data
 - Data Protection for SQL
 - performance 115
 - time-saving strategy 23
 - used with /sqlbuffers 131, 133, 183, 187, 227
 - using the CLI 133, 187, 228
 - using with tape media 23
- support contract xvi
- support files
 - sending to IBM with email 110
- support information xiii
- support subscription xvi
- syntax diagrams
 - delete backup command 204
- sysadmin fixed server role 168, 227
- system provider 4

T

- tables and indexes
 - backing up 22
- tape media
 - using with stripes 23
- task manager 78
- tdpexc.cfg file
 - and delete backup command 205
- tdpexc.log file
 - and delete backup command 206
- tdpsql.cfg file
 - and restorefiles command 201
 - parameters
 - setting 40
- tdpsql.cfg, setting values
 - using the CLI 128, 152, 176, 211, 229, 231
- tdpsql.log file
 - and restorefiles command 201
- tdsqlc.exe
 - overview 117
- tempdb 24

- throughput, improving 128, 151, 175, 176
- time-saving considerations 124, 173, 174, 210
- time, saving
 - strategy 22, 23
- timeformat parameter 228
- Tivoli Storage FlashCopy Manager
 - transitioning backups 237
- Tivoli Storage Manager
 - policy settings 19
- Tivoli Storage Manager API 149
- Tivoli Storage Manager backup-archive client
 - disaster recovery 24
- Tivoli Storage Manager scheduler 28
- Tivoli Storage Manager server
 - querying 149
 - txngroupmax value for striping 23
- Tivoli technical training xiii
- training, Tivoli technical xiii
- transact-SQL command 98
- transaction log
 - querying 149
- truncate log
 - strategy 23
- truncate log on checkpoint option 24
 - strategy 23
- tsmnode parameter 133, 155, 188, 213, 232
 - and restore command 207
 - and restorefiles command 203
- tsmoptfile parameter 133, 156, 188, 213, 232
 - and restore command 207
 - and restorefiles command 203
- tsmpassword parameter 134, 156, 188, 213
 - and restore command 207
 - and restorefiles command 204
- txngroupmax option 23
- typeface conventions xviii
- types parameter
 - described 151
 - syntax 148

U

- Unicode information
 - querying 150
- update node command 24
- user mode, setting 98

V

- VCS environment
 - supported 28
- viewing reports 100
- viewing system information for Data Protection for SQL 111
- volume failure 23
- volume mount 41
- VSS
 - and Legacy backups 24
 - cluster 25, 29, 30
 - DS 29

- VSS (*continued*)
 - hardware provider 4
 - N-series and NetApp
 - storage 30
 - node names 35
 - offloaded backup 7
 - proxy nodes 35
 - SAN Volume Controller 29
 - software provider 4
 - Storwize V7000 29
 - system provider 4
- VSS Backup
 - considerations 2
- VSS backup properties 74
- VSS Fast Restore
 - method 8
- VSS Instant Restore
 - hardware requirements 47
 - method 8
- VSS Provider 3
- VSS Requestor 3
- VSS Restore
 - method 7
- VSS Service
 - overview 3
- VSS Writer 3
- vsspolicy option 40, 41

W

- wait for tape mounts options, GUI 89
- Windows authentication mode, setting
 - using the CLI 130, 154, 183, 227

X

- XIV
 - considerations 6



Product Number: 5608-E04

Printed in USA

GC27-4010-00

