

IBM Tivoli Storage Manager for Databases
Version 7.1

*Data Protection for Oracle
for Windows
Installation and User's Guide*



IBM Tivoli Storage Manager for Databases
Version 7.1

*Data Protection for Oracle
for Windows
Installation and User's Guide*



Note:

Before using this information and the product it supports, read the information in “Notices” on page 59.

First edition (December 2013)

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

© **Copyright IBM Corporation 1998, 2013.**

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

Contents

About this publication	v
Who should read this publication	v
Publications	v
Conventions used in this book	vi
Typeface conventions	vi
Reading syntax diagrams	vi

New for Data Protection for Oracle Version 7.1	ix
---	-----------

Chapter 1. Data Protection for Oracle	1
Tivoli Storage Manager overview	1
Overview of Data Protection for Oracle	2
RMAN and Data Protection for Oracle	2
LAN-free data transfer with Data Protection for Oracle	3
Migration and coexistence with Data Protection for Oracle	3
Automated failover for data recovery	3

Chapter 2. Data Protection for Oracle installation	5
Installing Data Protection for Oracle	5
Installation prerequisites	5
Installing Data Protection for Oracle for Windows	6
Installing Data Protection for Oracle in silent mode	7

Chapter 3. Configuring Data Protection for Oracle	13
Configuration with default settings	13
Configuring Data Protection for Oracle	14
Define Data Protection for Oracle options in the tdpo.opt file	15
Register the Data Protection for Oracle node to a Tivoli Storage Manager server	17
Define Tivoli Storage Manager options in the client options file	18
Define Tivoli Storage Manager policy requirements	22
Initialize the password with a Tivoli Storage Manager server	23

Chapter 4. Protecting Oracle Server data	25
RMAN and Data Protection for Oracle	25
Starting RMAN	25
Editing RMAN scripts	26
The Duplex Copy function	28
Removing old backups	30
Setting up a schedule example	30
Setting up a schedule on the Tivoli Storage Manager server	31

Setting up a schedule on the workstation with the Oracle Server	31
Querying backup objects	35
Data deduplication with Data Protection for Oracle	35
Overview of data deduplication	35
Setting up Data Protection for Oracle for client-side data deduplication	36
Determining total data reduction	37

Chapter 5. Commands and utilities for Data Protection for Oracle	39
tdpoconf and tdposync utilities	39
Command line syntax and characteristics	39
tdpoconf utility	40
tdposync utility	42

Appendix A. Tivoli support information	51
Communities and other learning resources	51
Searching knowledge bases	53
Searching the Internet	53
Using IBM Support Assistant	53
Finding product fixes	54
Receiving notification of product fixes	54
Contacting IBM Software Support	54
Setting up and managing support contracts	55
Determining the business impact	55
Describing the problem and gathering background information	55
Submitting the problem to IBM Software Support	56

Appendix B. Accessibility features for the Tivoli Storage Manager product family	57
---	-----------

Notices	59
Trademarks	61
Privacy policy considerations	61

Glossary	63
A.	63
B.	65
C.	66
D.	67
E.	69
F.	70
G.	70
H.	71
I.	72
J.	72
K.	72
L.	73
M.	74
N.	75
O.	76

P.	76
Q.	77
R.	78
S.	79
T.	82
U.	82

V.	83
W	84

Index	85
------------------------	-----------

About this publication

This publication contains information about installing, configuring, administering, and using IBM® Tivoli® Storage Manager for Databases: Data Protection for Oracle.

Data Protection for Oracle runs online or offline backups of Oracle 11g databases to Tivoli Storage Manager storage. This integration with the RMAN Media Management API maximizes the protection of data, and provides a comprehensive storage management solution.

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

Who should read this publication

The target audience for this publication includes system installers, system users, Oracle database administrators, Tivoli Storage Manager administrators, and system administrators.

It is assumed that you have an understanding of the following applications:

- Oracle server
- Tivoli Storage Manager server
- Tivoli Storage Manager backup-archive client
- Tivoli Storage Manager application programming interface

It is assumed that you have an understanding of the Windows Server operating system on which your Oracle database is hosted.

Publications

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

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

You also can find the Tivoli Storage Manager product family information centers and other information centers that contain official product documentation for current and previous versions of Tivoli products at Tivoli Documentation Central. Tivoli Documentation Central is available at [http://www.ibm.com/developerworks/community/wikis/home/wiki/Tivoli Documentation Central](http://www.ibm.com/developerworks/community/wikis/home/wiki/Tivoli%20Documentation%20Central).

Conventions used in this book

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

Typeface conventions

Typeface conventions that are used are presented.

Bold text

- Commands, keywords, authorization roles, or other information that you must use.
- As an example, use the **query** command to query the Tivoli Storage Manager server for information about objects that have been backed up.

Text in italics

- Values or variables that you must provide.
- Emphasized words and phrases.
- As an example, the node name of the *production node* and *backup node* must not be the same.

Text in bold and italics

- Options and parameters.
- As an example, specify the value for the ***compression*** option.

monospace

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

UPPERCASE

- Environment variables that are associated with Tivoli Storage Manager, operating systems, or Oracle Server.
- As an example, make sure the DSM_DIR environment variable is set correctly.

Reading syntax diagrams






Information about how to read the syntax diagrams that are provided is presented. To read a syntax diagram, follow the path of the line. Read from left to right, and from top to bottom.

- The ►— symbol indicates the beginning of a syntax diagram.
- The —→ symbol at the end of a line indicates that the syntax diagram continues on the next line.
- The ►— symbol at the beginning of a line indicates that 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																		
<p>Abbreviations:</p> <p>Uppercase letters denote the shortest acceptable truncation. If an item displays entirely in uppercase letters, it cannot be truncated.</p> <p>You can type the item in any combination of uppercase or lowercase letters.</p> <p>In this example, you can enter KEYW0, KEYWORD, or KEYWOrd.</p>	<p>The diagram shows the text "KEYWOrd" on a line. Above the "W" is a horizontal arrow pointing to the right, and below the "d" is a horizontal arrow pointing to the left, indicating truncation points.</p>																		
<p>Symbols:</p> <p>Enter these symbols exactly as they display in the syntax diagram.</p>	<table> <tr><td>*</td><td>Asterisk</td></tr> <tr><td>{ }</td><td>Braces</td></tr> <tr><td>:</td><td>Colon</td></tr> <tr><td>,</td><td>Comma</td></tr> <tr><td>=</td><td>Equal Sign</td></tr> <tr><td>-</td><td>Hyphen</td></tr> <tr><td>()</td><td>Parentheses</td></tr> <tr><td>.</td><td>Period</td></tr> <tr><td></td><td>Space</td></tr> </table>	*	Asterisk	{ }	Braces	:	Colon	,	Comma	=	Equal Sign	-	Hyphen	()	Parentheses	.	Period		Space
*	Asterisk																		
{ }	Braces																		
:	Colon																		
,	Comma																		
=	Equal Sign																		
-	Hyphen																		
()	Parentheses																		
.	Period																		
	Space																		
<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 the text "KEYWOrd-var_name" on a line. Above the "W" is a horizontal arrow pointing to the right, and below the "d" is a horizontal arrow pointing to the left, indicating truncation points.</p>																		
<p>Repetition:</p> <p>An arrow that returns to the left means that 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 the text "repeat" on a line. Above the "t" is a horizontal arrow pointing to the left, indicating repetition.</p> <p>The second diagram shows the text "repeat'" on a line. Above the "t" is a horizontal arrow pointing to the left, indicating repetition with a comma separator.</p> <p>The third diagram shows the text "repeat" on a line. Above the "t" is a horizontal arrow pointing to the left, with a footnote "(1)" next to it, indicating a specific number of repetitions.</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 must specify one item.</p> <p>In this example, you must choose A, B, or C.</p>	<p>The diagram shows the text "A", "B", and "C" stacked vertically on a line. A horizontal arrow points to the right from the "A", and a horizontal arrow points to the left from the "C", indicating that one of the choices must be specified.</p>																		

Syntax Diagram Description	Example
<p>Optional Choice:</p> <p>When an item is below the line, that item is optional. In the first example, you can choose A or nothing at all.</p> <p>When two or more items are in a stack below the line, all of them are optional. In the second example, you can choose A, B, C, or nothing at all.</p>	 
<p>Defaults:</p> <p>Defaults are above the line. The default is selected unless you override it. You can override the default by including an option from the stack below the line.</p> <p>In this example, A is the default. You can override A by choosing B or C. You can also specify the default explicitly.</p>	
<p>Repeatable Choices:</p> <p>A stack of items followed by an arrow that returns 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>Syntax Fragments:</p> <p>Some diagrams because of their length, must fragment the syntax. The fragment name displays between vertical bars in the diagram. The expanded fragment displays 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 for Data Protection for Oracle Version 7.1

Read about the new features and other changes in IBM Tivoli Storage Manager for Databases: Data Protection for Oracle Version 7.1.

In V7.1, Data Protection for Oracle can fail over to a secondary Tivoli Storage Manager server for data recovery.

Support for Data Protection for Oracle on Windows 2012 is added in Version 7.1.

Chapter 1. Data Protection for Oracle

A brief overview of IBM Tivoli Storage Manager and IBM Tivoli Storage Manager for Databases: Data Protection for Oracle is provided.

Tivoli Storage Manager overview

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

Tivoli Storage Manager provides these functions:

- Reduces network complexity
Tivoli Storage Manager reduces network complexity with interfaces and functions that span network environments. Consistency across different operating systems and hardware is provided.
- Increases administrator productivity
Tivoli Storage Manager can reduce the cost of network administration by allowing administrators to:
 - Automate repetitive processes.
 - Schedule unattended processes.
 - Administer Tivoli Storage Manager from anywhere in the network.
- Reduces the risk of data loss
Many users do not back up their data. Other users apply stand alone backup techniques with diskettes and tapes as the only protection for business data. These backup systems often produce disappointing results during recovery operations. Tivoli Storage Manager schedules routine backups that enable users to recover from accidental data deletion without administrator involvement.
- Optimizes existing storage resources
Tivoli Storage Manager allows users to move files from client file systems to Tivoli Storage Manager storage. This optimization saves space on client file systems and can eliminate the expense of upgrading client storage hardware. Tivoli Storage Manager monitors client storage space and moves files from client file systems to Tivoli Storage Manager storage if an out-of-space condition threatens. This function can also eliminate the expense of client hardware upgrades.

Tivoli Storage Manager provides these services:

- Backup and restore services
These services generate backup copies of data at specified intervals, and restore the data from these copies when required. The services protect against workstation or file server media failure, accidental file deletion, data corruption, data vandalism, or site disasters.
- Archive and retrieve services
These services provide backup-archive clients with point-in-time copies of data for long-term storage.
- Server hierarchical storage management services

These services migrate client files from expensive storage media to less expensive storage media, for example from disk to tape. Administrator-defined thresholds determine file migration for each storage pool. Migration applies to all backup and archive client files.

- Automation services

Tivoli Storage Manager administrators can increase productivity by automating common storage administration tasks.

- Administration services

Tivoli Storage Manager administration services provide support for routine monitoring, administration, and accounting. Administrators can manage the server from another system or the same system. The Tivoli Storage Manager utilities allow the administrator to:

- Set client and server options.
- Define devices.
- Format storage volumes.
- Add more clients.
- Label tape volumes.

Tivoli Storage Manager monitors scheduled operations and maintains status information in the database. An administrator can export data to removable media. This data can be imported by another server, making the export and import features a convenient utility for moving server data. The administrator can specify the accounting option that is generated at the end of each client session.

- Security services

Security services control user access to Tivoli Storage Manager data, storage, policy definitions, and administrative commands.

- Disaster recovery management

Disaster recovery management helps the administrator implement a comprehensive backup and recovery procedure for important business applications, data, and records.

Overview of Data Protection for Oracle

Data Protection for Oracle interfaces with the Oracle Recovery Manager (RMAN) to send backup versions of Oracle databases to the Tivoli Storage Manager server.

Data Protection for Oracle currently supports Oracle 11g databases with the Oracle Recovery Manager. See Chapter 2, “Data Protection for Oracle installation,” on page 5 for specific levels of supported Oracle databases.

RMAN and Data Protection for Oracle

Oracle Recovery Manager (RMAN) provides consistent and secure backup, restore, and recovery performance for Oracle databases. While the Oracle RMAN initiates a backup or restore, Data Protection for Oracle acts as the interface to the Tivoli Storage Manager server. The Tivoli Storage Manager server then applies administrator-defined storage management policies to the data. Data Protection for Oracle implements the Oracle defined Media Management application programming interface (SBTAPI) 2.0. This SBTAPI communicates with RMAN and translates Oracle commands into Tivoli Storage Manager API calls to the Tivoli Storage Manager server.

You can use RMAN Data Protection for Oracle to run backup and restore functions that are listed.

- Full and incremental backup functions online or offline for:
 - Databases
 - Table spaces
 - Data files
 - Archive log files
 - Control files
- Full database restores while offline.
- Table space and data file restore online or offline.

LAN-free data transfer with Data Protection for Oracle

Data Protection for Oracle supports backup and restore operations in a LAN-free environment. This environment shifts the movement of data from the communications network to a storage area network (SAN). Data moves over the SAN to a SAN-attached storage device by the Tivoli Storage Manager Storage Agent. Running Data Protection for Oracle in a LAN-free environment avoids constraints of the network. The load on the Tivoli Storage Manager server is decreased, allowing the server to support a greater number of simultaneous connections.

Before you enable LAN-free support, you must install the Tivoli Storage Manager Managed System for SAN Storage Agent on the same system as Data Protection for Oracle. See the *IBM Tivoli Storage Manager for SAN* for your operating environment for more information about LAN-free requirements.

Migration and coexistence with Data Protection for Oracle

The migration considerations to the new version of Data Protection for Oracle are provided.

- Existing backups that are created with a previous version of Data Protection for Oracle are restorable with Data Protection for Oracle 7.1.
- Backups that are created with Data Protection for Oracle 7.1 cannot be restored with previous versions of Data Protection for Oracle.

Related tasks:

Chapter 3, “Configuring Data Protection for Oracle,” on page 13

“Editing RMAN scripts” on page 26

Automated failover for data recovery

When there is an outage on the Tivoli Storage Manager server, Data Protection for Oracle can fail over to a secondary server for data recovery operations.

The Tivoli Storage Manager server that Data Protection for Oracle connects to for backup operations is called the *primary server*. When the primary server and the Data Protection for Oracle node are set up for node replication on the primary server, the node can be replicated to another Tivoli Storage Manager server, called the *secondary server*.

During normal operations, connection information for the secondary server is automatically sent to Data Protection for Oracle from the primary server. The

secondary server information is saved to the client options file (dsm.opt) on the Data Protection for Oracle node. No manual intervention is required by you to add the information for the secondary server.

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

You can confirm that Data Protection for Oracle has failed over by looking for entries about the secondary server in the `dsierror.log` file.

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

- Data Protection for Oracle must be at the V7.1 level.
- The primary server and secondary server must be at the V7.1 level.
- The primary and secondary servers must be set up for node replication.
- The Data Protection for Oracle node must be configured for replication with the `replstate=enabled` option in the node definition on the server.
- Before the connection information for the secondary server can be sent to Tivoli Storage FlashCopy Manager, the following processes must occur:
 - You must back up data at least one time to the primary server.
 - The Data Protection for Oracle node on the primary server must be replicated at least one time to the secondary server.

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

- Any operation that requires data to be stored on the Tivoli Storage Manager server, such as backup operations, are not available. You can use only data recovery functions, such as restore or query operations.
- Schedules are not replicated to the secondary server. Therefore, schedules are not run while the primary server is unavailable.
- If the primary server goes down before or during node replication, the most recent backup data is not successfully replicated to the secondary server. The replication status of the file space is not current. If you attempt to restore data in failover mode and the replication status is not current, the recovered data might not be usable. You must wait until the primary server comes back online before you can restore the data.
- For more information about the failover capabilities of Tivoli Storage Manager components, see <http://www.ibm.com/support/docview.wss?uid=swg21649484>.

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

Chapter 2. Data Protection for Oracle installation

Install IBM Tivoli Storage Manager for Databases: Data Protection for Oracle to protect your Oracle server databases.

Installing Data Protection for Oracle

Verify installation prerequisites and follow the instructions to install Data Protection for Oracle for Windows.

Before you begin

Hardware, software, and operating system requirements must be met before you attempt to install Data Protection for Oracle.

Installation prerequisites

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

The minimum hardware and software requirements for the Data Protection for Oracle release are available in the hardware and software requirements technote for each particular release. For current requirements, review the Hardware and Software Requirements technote for your version of Data Protection for Oracle. This technote is available in the *TSM for Databases - All Requirements Documents* website at <http://www.ibm.com/support/docview.wss?uid=swg21218747>. From the page, follow the link to the requirements technote for your specific release or update level.

Note:

- You cannot install the 32-bit version of Data Protection for Oracle on a 64-bit operating system.
- You must have administrative authority to install Data Protection for Oracle on a Windows server where the target database is stored.
- If Data Protection for Oracle was not previously installed on your system, you can install it to the Tivoli Storage Manager base directory. For best results, install all Tivoli Storage Manager products and components into the same base directory.
- The installation process does not overwrite the existing `dsm.opt` options file, `tdpo.opt` configuration file, or log files.

Minimum hardware requirements

Your system must meet the minimum hardware requirements for installing or operating Data Protection for Oracle in a Windows environment.

Hardware requirements for the x86 platform

The minimum hardware requirements for the Data Protection for Oracle release are available in the hardware and software requirements technote for each particular release. For current requirements, review the Hardware and Software Requirements technote for your version of Data Protection for Oracle. This technote is available in the *TSM for Databases - All Requirements Documents* website at

<http://www.ibm.com/support/docview.wss?uid=swg21218747>. From the page, follow the link to the requirements technote for your specific release or update level.

Use compatible hardware that is supported by the Windows operating system and the Oracle server that are being used.

Hardware requirements for the x64 platform

Use compatible hardware that is supported by the Windows operating system and the Oracle server that are being used.

Minimum software and operating system requirements

Your system must meet the minimum software and operating requirements for using Data Protection for Oracle in a Windows environment.

Virtualization support

The minimum hardware and software requirements for the Data Protection for Oracle release are available in the hardware and software requirements technote for each particular release. For current requirements, review the *Hardware and Software Requirements* technote for your version of Data Protection for Oracle. This technote is available in the *TSM for Databases - All Requirements Documents* website at <http://www.ibm.com/support/docview.wss?uid=swg21218747>. From the page, follow the link to the requirements technote for your specific release or update level.

Information about the virtualization environments that can be used with Data Protection for Oracle is available in the *IBM Tivoli Storage Manager guest support for virtual machines and virtualization* website at: <http://www.ibm.com/support/docview.wss?uid=swg21239546>.

Installing Data Protection for Oracle for Windows

Use these instructions to install Data Protection for Oracle.

Before you begin

Ensure that the system meets the prerequisites. See “Installation prerequisites” on page 5 before you install Data Protection for Oracle.

Procedure

With the Data Protection for Oracle DVD in the DVD drive, Windows automatically starts the InstallShield wizard. If the wizard is not enabled, complete the following steps.

1. Run one of the following commands, where *x* is the DVD drive letter:

- (x86): *x*:\\oracle\\windows\\x32\\client\\setup
- (x86_64): *x*:\\oracle\\windows\\x64\\client\\setup

Click **OK** to start the installation program.

2. When prompted, choose a setup language. Choose one of the following languages:
 - Chinese (Simplified).
 - Chinese (Traditional).
 - English (United States), default value.

- French (France).
 - German (Germany).
 - Italian (Italy).
 - Japanese.
 - Korean.
 - Portuguese (Brazil).
 - Spanish (Traditional Sort).
3. Accept the License Agreement and click **Finish** to complete the installation.
 4. Optional: To view the Data Protection for Oracle command output in a language other than English, install that Language Pack executable file from the product DVD. Once installed, you can view error log text, and messages in that language. Run one of the following commands, where *x* is your DVD drive and *xxx* represents the three-letter country code that is associated with that language:
 - (x86): x:\oracle\windows\x32\languages\xxx\setup
 - (x86_64): x:\oracle\windows\x64\languages\xxx\setup
 Click **OK** and follow the instructions.

Installing Data Protection for Oracle in silent mode

You can install Data Protection for Oracle in silent mode. A silent installation runs independently without any intervention so that you are not required to monitor, or provide input.

About this task

This method is useful when you must install Data Protection for Oracle on a number of different computers with identical hardware. For example, a company might have 25 Oracle servers that are installed across 25 different sites. You can create an unattended installation package and make it available to the 25 sites. This method ensures a consistent configuration and avoids different people all entering Data Protection for Oracle parameters. The installation package can be placed on a DVD and sent to each site, or it can be placed on a file server for distribution.

You can install in silent mode with one of the following methods:

Setup Program

Use the **setup** command and specify silent installation options.

Microsoft Installer (MSI)

Use **msiexec.exe** to install the package.

The following options can be used with either silent installation method.

Table 1. Silent installation options

Option	Description
/i	Specifies the program to install the product.
/l*v	Specifies verbose logging.
/qn	Runs the installation quietly, without displaying windows.
/s	Specifies silent mode.

Table 1. Silent installation options (continued)

Option	Description
/v	The setup passes the parameter string to the MSI executable program (msiexec.exe). The following syntax is required when you use the /v option: <ul style="list-style-type: none"> • Use a backslash (\) before quotation marks (") that are contained within existing quotation marks. • Do not include a space between the /v command-line option and its arguments. • Separate parameters that are entered with the /v command-line option with a space. • Create a log file by specifying the directory and file name after the command. The directory must exist before you start the installation process.
/x	Specifies that the program is to uninstall the product.
addlocal	Specifies features to install.
allusers	Specifies which users can use the installation package.
installdir	Specifies the directory where Data Protection for Oracle is to be installed.
reboot	Specifies whether to prompt the user to restart the system after silent installation. <p><i>Force</i> Always prompts user to restart after an installation.</p> <p><i>Suppress</i> Suppresses prompt to restart after an installation.</p> <p><i>ReallySuppress</i> Suppresses all restart operations, and prompts the user to restart after an installation.</p>
rebootyesno	Specifies whether to restart the system after silent installation.
transforms	Specifies the language to install.

The following components are used in this procedure and are case-sensitive.

Table 2. Silent installation components (base client only)

Component	Description
Client	Data Protection for Oracle code
License_Paid	License file is used when paid versions of Data Protection for Oracle are installed

Table 3. Silent installation components (Language Packages only)

Component	Description
LanguageFiles	Language-specific files

The following files are used in this procedure.

Table 4. Silent installation files used

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

Installing Data Protection for Oracle in silent mode with the setup executable file

Use the setup program, by running the setup.exe file, to install Data Protection for Oracle in silent mode. When you install in silent mode, you do not need to monitor the process or provide input.

Before you begin

Install Data Protection for Oracle from an account that is a member of the local Administrators group for the computer where the Oracle server is running.

About this task

The following commands can be used for installing Data Protection for Oracle in silent mode. You must substitute the appropriate feature when you install a language other than English.

Procedure

Run the following command to install Data Protection for Oracle in silent mode to the default installation directory:

```
setup.exe /s/v/qn
```

- Install Data Protection for Oracle in silent mode to the default installation directory with custom features by running the following command:

```
setup.exe /s /v"INSTALLDIR="c:\program files\tivoli\tsm\  
ADDLOCAL="Client,License_Paid\  
TRANSFORM=1033.mst /qn /!v "c:\temp\log.txt"
```

For more information about language packs, see Table 3 on page 8. In a custom installation, list all features after the **addlocal** option.

Installing Data Protection for Oracle in silent mode with MSI

You can install Data Protection for Oracle in silent mode by using the Microsoft Installer program (MSI), and the `msiexec.exe` file.

Before you begin

Prepare to install Data Protection for Oracle from an account that is a member of the local Administrator group for the computer where the Oracle server is running.

About this task

Important: Data Protection for Oracle installs the Microsoft Visual C++ 2012 Redistributable Package as a setup prerequisite. If you are installing the product with the `msiexec.exe` file, you must separately install the Microsoft Visual C++ 2012 Redistributable Package. The files are included in the installable packages:

- For Windows 32 bit, install the following file:
In the path to the Data Protection for Oracle files: `\ISSetupPrerequisites\Microsoft Visual C++ 2012 Service Update 1 Runtime Libraries (x86)\vcredist_x86.exe`
- For Windows 64 bit, install the following file:
In the path to the Data Protection for Oracle files: `\ISSetupPrerequisites\Microsoft Visual C++ 2012 Service Update 1 Runtime Libraries (x64)\vcredist_x64.exe`

When you install a language other than English, you must substitute the appropriate `.msi` package file name and Language Package feature. Table 3 on page 8 contains a list of language packs that you can install.

Procedure

The following example silently installs Data Protection for Oracle to a directory other than the default installation directory and includes custom features:

```
msiexec /i<path to msi file>\IBM Tivoli Storage Manager for Databases - Oracle.msi"
RebootYesNo="No" Reboot="Suppress" ALLUSERS=1
INSTALLDIR="c:\program files\tivoli\tsm"
ADDLOCAL="Client,License_Paid"
TRANSFORM=1033.mst /norestart /qn /l*v "c:\temp\log.txt"
```

Capturing a log of the installation

Capture a log of the installation of Data Protection for Oracle using the following procedure.

About this task

The following environmental information can be helpful:

Gather a detailed log of the installation procedure in a file called `setup.log`. To gather the information, run the setup program as follows:

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

Creating an installation package

You have a choice of making the Data Protection for Oracle installation package available for distribution on a DVD or on a server.

About this task

You can burn a DVD or place the package in a shared directory on a server. The package contains the Data Protection for Oracle code distribution files for installing the product.

Creating a package for installing in silent mode:

Use these instructions to create a silent installation package.

Before you begin

Before you create the package, decide where you are going to store the package. If you are burning a DVD, 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.

Procedure

1. Run the following command to create the tdpdpkg staging directory:
`mkdir c:\tdpdpkg`
2. From the command line, navigate to the staging directory `cd /d c:\tdpdpkg`
3. Copy the Data Protection for Oracle DVD distribution files to the staging directory: `xcopy g:*.* . /s`
4. Replace the existing `setup.bat` with the one created in the previous step: `copy c:\setup.bat`

What to do next

The installation package can be placed on a DVD or made available from a shared directory.

Playing back the installation package in silent mode

You can run an installation package from a DVD or a shared directory. This process is called *playing back* the installation.

About this task

No visual cues exist to inform you when the installation finishes, although you can add visual cues to the batch file. Therefore, ensure that you allow enough time for the installation to complete.

Procedure

- If autostart is enabled, the silent installation begins as soon as the DVD is inserted into the drive.
- If the installation package is in a shared directory, a user can run the command as follows: From a staging directory that is called `tdpdpkg` at `\\machine1\d$`, the command `net use x: \\machine1\d$` can be run from a different computer to share the drive as drive `x`.

Creating batch files for installing in silent mode

You can create a batch file to install Data Protection for Oracle in silent mode, with the parameters that you want.

About this task

The following example shows a sample script that includes specific installation options.

```
@echo off
rem =====
rem sample silent install script
rem
rem /s /v"INSTALLDIR="X:\Install Path\" /qn"
rem =====
rem Code can be added after the
rem installation completes to
rem customize the dsm.opt files
rem if needed.
rem =====
```

Chapter 3. Configuring Data Protection for Oracle

Use these instructions to configure Data Protection for Oracle for backup and restore operations.

Before you begin

Data Protection for Oracle must be installed on your system and a Tivoli Storage Manager server must be available to communicate with Data Protection for Oracle.

About this task

Review all configuration information before you run any configuration tasks.

Configuration with default settings

Use the Data Protection for Oracle quick configuration option to quickly configure with default settings and minimal configuration tasks. Setup time is minimized and you proceed quickly to a state where you can begin backing up your Oracle databases.

Before you begin

Install Data Protection for Oracle. For detailed installation instructions, see Chapter 2, “Data Protection for Oracle installation,” on page 5.

About this task

Use the instructions to configure Data Protection for Oracle on a 32-bit Windows Server system. Installation on a 64-bit Windows Server system is the same except for the installation path, which is C:\Program Files\Tivoli\TSM\Agent0BA64.

See “Configuring Data Protection for Oracle” on page 14 for detailed instructions on how to customize Data Protection for Oracle for your environment and processing needs.

Procedure

1. Change to the C:\Program Files\Tivoli\TSM\Agent0BA directory and copy the tdpo.opt.smp file to tdpo.opt. Edit the tdpo.opt file to include these options:

```
dsmi_orc_config C:\Program Files\Tivoli\TSM\Agent0BA\dsm.opt
dsmi_log C:\Program Files\Tivoli\TSM\Agent0BA
```

For more information about these options, see “Available Data Protection for Oracle options” on page 15.

2. In this directory, copy the dsm.smp file to dsm.opt. Edit the dsm.opt file to include these options:

```
COMMMethod TCPip
TCPServeraddress x.x.x.x
PASSWORDAccess generate
NODename hostname_oracle
```

Replace *x.x.x.x* with the IP address of the Tivoli Storage Manager server to which Data Protection for Oracle backs up data.

For more information about these options and the `dsm.opt` file, see “Define Tivoli Storage Manager options in the client options file” on page 18.

3. Register the node to the Tivoli Storage Manager server with the following command:

```
REG NODE hostname_oracle password maxnummp=n
```

Where *hostname* is the name of the system that Data Protection for Oracle is installed, *password* is the password for this node, and *n* is equal to the number of channels that you are planning to use.

4. Make sure that the <oracle user> has the following permissions:
 - Read and Write permission to the `C:\Program Files\Tivoli\TSM\Agent0BA` directory.
 - Read and Write permissions to `tdpoerror.log` and the directory where it is stored.
5. Change to the `C:\Program Files\Tivoli\TSM\Agent0BA` directory and run the **tdpoconf password** command as the <oracle user> to generate the password file. Enter the password from Step 3 three times.

For more information about this command, see “**password** command” on page 40.

6. Run the **tdpoconf showenvironment** command to view and confirm your configuration.

For more information about this command, see “**showenvironment** command” on page 40.

7. As Oracle user, run the RMAN backup script with the **ENV=(TDPO_OPTFILE=C:\Program Files\Tivoli\TSM\Agent0BA\tdpo.opt)** parameter specified. For example:

```
run
{
    allocate channel t1 type 'sbt_tape' parms
        'ENV=(TDPO_OPTFILE=C:\Program Files\Tivoli\TSM\Agent0BA\tdpo.opt)';

    backup
        filesperset 5
        format 'df_%t_%s_%p'
        (database);
}
```

Note, the `allocate channel` entry is divided on two lines after the `parms` option is specified to accommodate page formatting.

For more information about RMAN backup scripts, see “RMAN and Data Protection for Oracle” on page 25.

Configuring Data Protection for Oracle

After Data Protection for Oracle is successfully installed, you must complete the configuration tasks.

Procedure

1. Define Data Protection for Oracle options in the `tdpo.opt` file.
2. Register the Data Protection for Oracle node to a Tivoli Storage Manager server.
3. Define Tivoli Storage Manager options in the `dsm.opt` file.
4. Define Tivoli Storage Manager policy requirements.

5. Initialize the password with a Tivoli Storage Manager server.

Results

If you would like to configure Data Protection for Oracle using default settings, see “Configuration with default settings” on page 13 for instructions.

Define Data Protection for Oracle options in the `tdpo.opt` file

You must define options to control the way Data Protection for Oracle backs up and restores data.

About this task

The Data Protection for Oracle options file, `tdpo.opt`, contains options that determine the behavior and performance of Data Protection for Oracle. The only environment variable Data Protection for Oracle recognizes within an RMAN script is the fully qualified path name to the `tdpo.opt` file. Therefore, some RMAN scripts must be edited to use **TDPO_OPTFILE**=fully qualified path and file name of options file variable in place of other environment variables. For example:

```
allocate channel t1 type 'sbt_tape' parms
      'ENV=(TDPO_OPTFILE=C:\RMAN\scripts\tdpo.opt)'
```

For further information about RMAN scripts, see “Editing RMAN scripts” on page 26 for further information. Note, the `allocate channel` entry is divided on two lines after the `parms` option to accommodate page formatting.

If the **TDPO_OPTFILE** variable is not provided, Data Protection for Oracle uses the `tdpo.opt` file in the Data Protection for Oracle default installation directory. If this file does not exist, Data Protection for Oracle fails.

Though it is possible to change values in the options file, you cannot change the `dsmi` options until you restart Data Protection for Oracle. Data Protection for Oracle loads these values from the default `tdpo.opt` file into the default installation directory when the Oracle service is started.

Note:

- For best results, use the `tdpo.opt` file exclusively instead of default parameters.
- RMAN and the `tdpoconf` and `tdposync` utilities use the options that are defined in the `tdpo.opt` file.
- By default, the `tdpo.opt` file is in the directory where Data Protection for Oracle is installed.
- You can specify options in the `tdpo.opt` file in both uppercase or lowercase type.

Available Data Protection for Oracle options

The options that can be set in the `tdpo.opt` file for Data Protection for Oracle are described.

The following options can be set in the `tdpo.opt` file:

dsmi_log

Specify the directory that contains the Data Protection for Oracle error log file `tdpoerror.log`.

dsmi_orc_config

Specify the complete path to the Tivoli Storage Manager client user options file `dsm.opt` used during the Data Protection for Oracle session. If you do

not specify this option, Data Protection for Oracle looks for the options file in the Data Protection for Oracle installation directory. You must specify this option if your Tivoli Storage Manager client user options file is in a directory other than the Data Protection for Oracle installation directory.

tdpo_fs

Specify a file space name on the Tivoli Storage Manager server for Data Protection for Oracle backup, delete, and restore operations. The file space name can contain a string of 1-1024 characters.

- The default file space name is adsmorc.
- When you have more than one Oracle database, use this option to back up each Oracle target database to its own file space on the Tivoli Storage Manager server.
- The file space name in the include/exclude statement must match the file space name that is specified in the `tdpo_fs` option for include/exclude processing to function correctly.

tdpo_date_fmt

This option specifies the format that you want to use to display dates.

You can specify a number, 1 - 5. The default value is 1.

- 1 MM/DD/YYYY (Default value)
- 2 DD-MM-YYYY
- 3 YYYY-MM-DD
- 4 DD.MM.YYYY
- 5 YYYY.MM.DD

tdpo_num_fmt

This option specifies the format that you want to use to display numbers.

You can specify a number, 1 - 6. The default value is 1.

- 1 1,000.00 (Default value)
- 2 1,000,00
- 3 1 000,00
- 4 1 000.00
- 5 1.000,00
- 6 1'000,00

tdpo_time_fmt

This option specifies the format that you want to use to display time.

You can specify a number, 1 - 4. The default value is 1.

- 1 23:00:00 (Default value)
- 2 23,00,00
- 3 23.00.00
- 4 12:00:00 A/P

tdpo_language

This option specifies the national language in which to display Data Protection for Oracle messages in the error log file. You can specify one of the following languages:

- chs Chinese, Simplified
- cht Chinese, Traditional
- enu American English (Default value)
- deu Standard German

- esp Standard Spanish
- fra Standard French
- ita Standard Italian
- jpn Standard Japanese
- kor Standard Korean
- ptb Brazilian Portuguese

tdpo_mgmt_class_2

This option specifies the second management class that is used for copy 2 in the RMAN duplex copy command.

tdpo_mgmt_class_3

This option specifies the third management class that is used for copy 3 in the RMAN duplex copy command.

tdpo_mgmt_class_4

This option specifies the fourth management class that is used for copy 4 in the RMAN duplex copy command. Four copies is the maximum that is allowed by RMAN.

Note: See “The Duplex Copy function” on page 28 for specific details on using management class options.

Register the Data Protection for Oracle node to a Tivoli Storage Manager server

The Data Protection for Oracle node name and password when required must be registered to the Tivoli Storage Manager server before you can begin requesting backup and restore services. The process of setting up a node name and password with the Tivoli Storage Manager server is called registration.

About this task

The following information is needed to register Data Protection for Oracle with the Tivoli Storage Manager server:

- Data Protection for Oracle node name:
The node name identifies the instance on which Data Protection for Oracle is installed. Use a separate and unique node name for Data Protection for Oracle. This prevents any confusion with an existing Tivoli Storage Manager backup-archive client on the same workstation.
- Initial password:
Specify the password that you want to use, if a password is required.

The following information is defined by the Tivoli Storage Manager administrator:

- The policy domain to which your client node belongs.
A policy domain contains policy sets and management classes that control how Tivoli Storage Manager manages the objects you back up. Rather than binding Data Protection for Oracle backups to a different management class, define a unique policy domain for Data Protection for Oracle node names. These backups can be bound to the default management class within this unique policy domain. Rather than binding a different management class for Oracle backups, specify a different domain for the backups with a separate management class.
- The authority to enable compression.

The Tivoli Storage Manager administrator can specify the server to compress files. If the Tivoli Storage Manager administrator specifies that the compression decision belongs to the client **compressionclient**, you must specify **compressionyes** in the client user options file `dsm.opt` so that the Data Protection for Oracle node is able to compress objects before it sends them to the Tivoli Storage Manager server.

- The authority to delete backup data from Tivoli Storage Manager storage. The Data Protection for Oracle node can only delete backed up data from Tivoli Storage Manager storage if the Tivoli Storage Manager administrator registers the node with `backdelete` authority. Specify the following option to allow `backdelete` authority:

`backdelete yes`

Note, when `backdelete no` is specified and a deletion request is made, the request fails and an error message displays. Therefore, specify `backdelete yes` for the object to be immediately removed from the Tivoli Storage Manager server when the next inventory expiration occurs. This expiration also makes the previously used storage space available for new use.

Define Tivoli Storage Manager options in the client options file

You must define some Tivoli Storage Manager options after the Data Protection for Oracle node is registered to the Tivoli Storage Manager server:

About this task

- These options are defined in the Tivoli Storage Manager client user options file `dsm.opt` by default.
- Note, the Tivoli Storage Manager client user options file `dsm.opt` by default, that you must edit for Data Protection for Oracle is in the directory that is specified by the `dsmi_orc_config` option. If this option is not specified, Data Protection for Oracle looks for this options file in the Data Protection for Oracle installation directory.
- Data Protection for Oracle provides sample Tivoli Storage Manager options files that you can modify for this purpose. These sample files are in the Data Protection for Oracle installation directory.
- The Tivoli Storage Manager administrator can provide you with the TCP server address **tcpserveraddress** and communication method **commmethod** for connecting Data Protection for Oracle to the Tivoli Storage Manager server.

Required options

You must set required Tivoli Storage Manager client options to operate Data Protection for Oracle.

Specify the required options in the Tivoli Storage Manager client user options file `dsm.opt` by default, in the directory that is specified by the `dsmi_orc_config` option.

nodename

Specify the Data Protection for Oracle node name that is used during operations with the Tivoli Storage Manager server. The node name can contain a string of 1 - 64 characters.

To restore data from one Oracle server to another with Data Protection for Oracle, make sure that the value of the `nodename` option in the `dsm.opt` file on the target Oracle server equals the value of the `nodename` option in the `dsm.opt` file on the source Oracle server. Data Protection for Oracle, the

Tivoli Storage Manager backup-archive client, and the Tivoli Storage Manager API must be at the same levels on both the source Oracle server and the target Oracle server. If there is a password that is associated with the `nodename` option in the `dsm.opt` file on the target Oracle server, use the `tdpoconf` utility to generate the local password file when the value of the `nodename` option changes.

passwordaccess

Specify `passwordaccess generate`. This option allows the Tivoli Storage Manager API to manage all password actions after the password is created with the **`tdpoconf password`** command. The Tivoli Storage Manager API saves the current encrypted password in the Windows registry and automatically generates a new password when the current password expires. This method of password management is useful when you are running unattended scheduled backups because it ensures that the backup never fails with an expired password.

tcpserveraddress

Specify the TCP/IP address for the Tivoli Storage Manager server to be used for Oracle backups.

commmethod

Specify the communication method for Data Protection for Oracle to communicate with the Tivoli Storage Manager server. Note, this option requires other Tivoli Storage Manager options, depending on the communication method you specify.

Other options to consider

There are other Tivoli Storage Manager client options that you can use when you are configuring Data Protection for Oracle.

You can specify other options in the Tivoli Storage Manager client user options file `dsm.opt`.

compression

Specify whether the Tivoli Storage Manager API compresses data before it sends it to the Tivoli Storage Manager server. You can specify `yes` or `no`. The default value is `No`. The value of the compression option for Data Protection for Oracle is allowed only if the Tivoli Storage Manager administrator leaves the compression decision to the node. Enabling compression affects performance in three ways:

- Processor usage is higher on the system on which Data Protection for Oracle is running.
- Network bandwidth usage is reduced because fewer bytes are transmitted.
- Storage usage on the Tivoli Storage Manager server is reduced.

When any of the following conditions exist, you should specify `yes`:

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

When any of the following conditions exist, you should specify `no`:

- The system that is running Data Protection for Oracle has a processor overload. The added processor usage as a result of enabling compression can impact other applications, including the Oracle server.

You can monitor processor and network resource usage with the Performance Monitor program that is shipped with Windows.

- You are not constrained by network bandwidth. In this case, you can achieve the best performance by specifying **compression no** and enabling hardware compaction on the tape drive, which also reduces storage requirements.
- Hardware compression is in use for the media where Data Protection for Oracle data is stored.

After a completed backup operation, view the throughput rate and the compression status for a backup object in the Tivoli Storage Manager server activity log file. Run the Tivoli Storage Manager server **QUERY ACTLOG** command in the Tivoli Storage Manager server administrative client window. The throughput rate and the compression status are not written to the activity log when activity logging is disabled on the Tivoli Storage Manager server. See the **SET ACTLOGRETENTION** command in the *Tivoli Storage Manager Administrator's Reference* for complete activity logging information.

You can also determine whether objects were compressed by running the **tdposync query** command.

deduplication

Specify whether the Tivoli Storage Manager API deduplicates data before it sends 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 Oracle applies only if the Tivoli Storage Manager administrator allows client-side data deduplication.

You can determine if objects are deduplicated by running the **tdposync query** command or by examining the Tivoli Storage Manager server activity log file.

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

The deduplication and **enableclientencryptkey** options are also mutually exclusive. Therefore, you must use either one option or the other, but not both options together.

enablelanfree

Specify whether you run backup or restore operations in a LAN-free environment if you are equipped to do so. You can specify yes or no. The default value is no. You can avoid network constraints by shifting the movement of data to a storage area network (SAN). After a completed backup operation, view the LAN-free status for a backup object in the Tivoli Storage Manager server activity log file. For more information, see the appropriate Storage Agent User's Guide.

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

include

When a management class other than the default management class is

defined within an existing policy domain, add an include statement to the client options file that is used by the Oracle node.

You must add an include statement to the dsm.opt file.

This include statement binds the Oracle backup objects to the management class that is defined for managing these objects. The include statement uses the following naming convention:

```
\FilespaceName\orcnt\ObjectName
```

The FORMAT parameter in the RMAN script can also be used to assist with object naming. For example, if the FORMAT parameters (in the RMAN script) specified the following values for databases and logs:

```
format 'DB_%u_%p_%c'
format 'LOG_%u_%p_%c'
```

The include statement in the dsm.opt file, which is used by the Oracle node, would be as follows:

```
INCLUDE \FilespaceName\orcnt\DB* mgmtclassname
INCLUDE \FilespaceName\orcnt\LOG* mgmtclassname
```

Make sure that the **FORMAT** parameter specifies a unique name for the backup. If the object name exists on the Tivoli Storage Manager server, the backup might fail with an RC=8 error that is recorded in the sbtio.log file.

enableclientencryptkey

When enableclientencryptkey is set to yes, Data Protection for Oracle provides 128-bit transparent encryption of Oracle databases during backup and restore processing. One random encryption key is generated per session and is stored on the Tivoli Storage Manager server with the object in the server database. Although Tivoli Storage Manager manages the key, a valid database must be available to restore an encrypted object.

Important: The enableclientencryptkey and deduplication options are mutually exclusive because encrypted files cannot be deduplicated. Therefore, you can use only one or the other option, but not both options together.

You can specify the databases that you want encrypted by adding an include statement with the include.encrypt option in the client user options file dsm.opt.

For example, to enable transparent encryption, do the following steps:

1. Edit the client user options file, dsm.opt.
2. Specify enableclientencryptkey yes.
3. Specify encryptiontype AES128, or DES56.
4. Specify the objects to encrypt. This example encrypts all data:

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

Thus, the encryption options would be as follows in this client user options file, dsm.opt:

```
enableclientencryptkey yes
encryptiontype aes128
include.encrypt      \adsmorc\...\*
```

See *IBM Tivoli Storage Manager Using the Application Programming Interface* for more details about the enableclientencryptkey option.

You can determine whether objects were encrypted by running the **tdposync query** command.

Define Tivoli Storage Manager policy requirements

Data Protection for Oracle requires special Tivoli Storage Manager policy domain settings.

About this task

RMAN uses the **format** parameter in the RMAN script to generate unique backup file names. Because all backup objects inserted into the Tivoli Storage Manager backup storage pool have unique file names, they never expire on the Tivoli Storage Manager server. As a result, Data Protection for Oracle requires the following Tivoli Storage Manager policy domain settings:

Backup copy group values

Data Protection for Oracle provides the **tdposync** utility to remove unwanted backup objects from the Tivoli Storage Manager server. Set the following Tivoli Storage Manager backup copy group options:

- **verdeleted** 0
- **retonly** 0

When Data Protection for Oracle marks a backup object inactive, that object is deleted from the Tivoli Storage Manager server the next time expiration processing occurs. A backup object is marked for immediate expiration when you delete it through RMAN with the Data Protection for Oracle interface or with the **tdposync** utility. Note, an inactive backup object cannot be restored through RMAN with the Data Protection for Oracle interface.

Note:

1. The Tivoli Storage Manager administrator must also register your node by specifying **backdelete yes** in order for backup objects to be deleted. However, be aware that a backup object is marked for immediate expiration when **backdelete yes** and you delete it through RMAN with the Data Protection for Oracle interface or with the **tdposync** utility. Note, when **backdelete no** is specified and a deletion request is made, the request fails and an error message displays.
2. The following backup copy group options are not applicable to Data Protection for Oracle:
 - **frequency**
 - **verexists**
 - **retextra**
 - **mode**
 - **serialization**

Data Protection for Oracle accepts default values for these options.

3. An archive copy group is not required (although it can exist) because Data Protection for Oracle stores all objects as backup objects on Tivoli Storage Manager.

Management class

Tivoli Storage Manager uses management classes to manage backups on the Tivoli Storage Manager server. When you back up a database, the default management class for your node is used. Because the policy

requirements for Data Protection for Oracle might be different from the wanted settings for the regular Tivoli Storage Manager backup-archive clients, you must have a different management class that is defined for Data Protection for Oracle. You must define a separate policy domain where the default management class has the required settings. Then, register all Data Protection for Oracle nodes to that domain.

If you choose to define a new management class within an existing policy domain, not the default management class for that domain, then you must add an `include` statement to the Data Protection for Oracle options file to bind all objects to that management class.

The following steps assign a management class name `orcbbackup` to all Oracle backups with a default file space name `adsmorc`:

1. Add this `incl excl` entry under the server stanza you use in the `dsm.opt` file:

```
incl excl C:\myfile\include.def
```

- 2.

Add the following `include` entry to the `C:\myfile\include.def` file:

```
include \adsmorc\...\* orcbbackup
```

Note: The file space name in the `include/exclude` statement must match the file space name that is defined with the `tdpo_fs` option. If a file space name other than the default value `adsmorc` is used:

- a. You must specify the file space name with the `tdpo_fs` option.
- b. You must specify the file space name that is defined in the `tdpo_fs` option in the `include/exclude` statement.

All the files that are backed up with a default file space name of `adsmorc` are assigned to management class *orcbbackup*.

Note: Data Protection for Oracle stores all objects as backup objects on Tivoli Storage Manager storage, so an archive copy group is not required, although it can exist.

See your Tivoli Storage Manager administrator or see the *Tivoli Storage Manager Administrator's Guide* for more information about defining or updating Tivoli Storage Manager policy domains and copy groups.

Initialize the password with a Tivoli Storage Manager server

The administrator must run the `tdpoconf` utility program to set the password before you use Data Protection for Oracle.

Related reference:

"`tdpoconf` utility" on page 40

Chapter 4. Protecting Oracle Server data

Use Data Protection for Oracle to back up and restore Oracle Server data.

Before you begin

Data Protection for Oracle must be installed and configured on your system and an Oracle Server must be available.

RMAN and Data Protection for Oracle

You can run full or partial, offline, or online backups with Oracle. When you identify which database to back up, Oracle locates all necessary files and sends them to the Tivoli Storage Manager server through Data Protection for Oracle.

About this task

Data Protection for Oracle provides an interface between Oracle Media Management API calls and Tivoli Storage Manager API routines.

Starting RMAN

Use RMAN to back up and restore an Oracle database.

About this task

In this example, the catalog database contains a registered target database. Start an RMAN session with this command:

```
$> rman target xxx/yyy@target rcvcat aaa/bbb@catalog  
cmdfile bkdb.scr msglog bkdb.log
```

RMAN starts in the sequence shown.

```
target xxx/yyy@target: connect to target database  
using user xxx and password yyy with connect string target  
rcvcat aaa/bbb@catalog: connect to catalog database  
using user aaa and password bbb with connect string catalog  
cmdfile bkdb.scr: run bkdb.scr script  
msglog bkdb.log: log the output messages in bkdb.log
```

Tip: In the example, RMAN creates a log file, `bkdb.log`, in the current working directory. If an error occurs, the error stack is logged to the log file.

After a completed backup or restore operation, view the throughput rate and encryption status for a backup object in the Tivoli Storage Manager server activity log file. Run the Tivoli Storage Manager server **QUERY ACTLOG** command in the Tivoli Storage Manager server administrative client window. A message similar to the following is displayed:

```

08/03/11
16:23:08
ANE4991I (Session: 67, Node: MACHINE_ORC) DP Oracle Win64 ANU0599 TDP for Oracle:
(5508): =>()
ANU2526I Backup details for backup piece \adsmorc\orcnt\df_727444762_116_1
(database "orcl").
Total bytes processed: 9961472. Deduplicated: Yes. Bytes after deduplication: 2272805.
Deduplication reduction: 77.18%. Compressed: Yes. Bytes after compression: 52253.
Compressed by: 97.70%. Encryption: None. LAN-Free: No. Total bytes sent: 52253.
Total data reduction: 99.48%. Total processing time: 00:00:01.
Throughput rate: 9728.00Kb/Sec. (SESSION: 67)

```

Editing RMAN scripts

You must edit existing RMAN scripts to use **TDPO_OPTFILE**=*fully qualified path and file name of options file* variable in place of other environment variables.

About this task

Sending options with the send command

Use the Oracle RMAN **send** command in an RMAN script to pass Tivoli Storage Manager options to the Tivoli Storage Manager API.

Before you begin

To send options from the Tivoli Storage Manager to the Tivoli Storage Manager API, you must specify the **send** command in an RMAN script.

About this task

Use the **send** command to set Tivoli Storage Manager options such as TCPServeraddress and TCPport to the Tivoli Storage Manager API. You can customize the actions that the script takes without updating the existing Data Protection for Oracle or Tivoli Storage Manager API options files. Any option that is sent through the **send** command overrides the option that is specified in the Data Protection for Oracle or Tivoli Storage Manager API options files.

- You can specify multiple Tivoli Storage Manager API options in the same **send** command.
- The ENABLELANFREE and DEDUPLICATION options are mutually exclusive. If both options are defined, client-side data deduplication does not occur.
- The ENABLECLIENTENCRYPTKEY and DEDUPLICATION options are also mutually exclusive. If both options are defined, client-side data deduplication does not occur.
- You can specify any Tivoli Storage Manager API option with the **send** command.

Procedure

Specify the **send** command in an RMAN script. You can specify one or more Tivoli Storage Manager options in a **send** command string. The **send** command string can contain up to 512 bytes. To back up an Oracle database to the Tivoli Storage Manager server named halley at TCP/IP port 1601, and to enable the cache for client-side data deduplication for only channel t1, specify the following statements in an RMAN script:

```

allocate channel t1 type 'SBT_TAPE';
SEND channel 't1' '-TCPSEVER=halley -TCPPOINT=1601 -ENABLEDEDUPCACHE=YES';

```

Results

Data Protection for Oracle passes the command string to the Tivoli Storage Manager API. The Tivoli Storage Manager API validates the contents of the string. If an invalid entry is detected, the API issues an ANS****E message to Data Protection for Oracle. The message returns an error condition to Oracle RMAN and stops processing.

Related tasks:

“RMAN script examples”

RMAN script examples

Sample RMAN scripts illustrate how to create parallel backup streams to Tivoli Storage Manager server storage.

Example

In these examples, to back up to Tivoli Storage Manager by using Data Protection for Oracle, you must specify type 'sbt_tape' in the RMAN script or within the global RMAN configuration settings.

Example 1:

When the Tivoli Storage Manager server and Oracle system have multiple network cards, you can back up your data with multiple network paths to improve network throughput. Your environment is set up as follows:

- The Oracle system has two network cards with two addresses, A and B.
- The Tivoli Storage Manager server also has two network cards with two addresses, C and D.
- Paths exist between A and C, B and D, but not between A and D or B and C.

Create two backup streams or Oracle channels, without using two separate options files to point to different two different addresses. Channel t1 goes to address C, channel t2 goes to address D. Be careful not to send parts of your backup to two different Tivoli Storage Manager servers because it cannot be restored.

You can maintain one Data Protection for Oracle options file and change the Tivoli Storage Manager server specification in an RMAN script in the following manner:

```
run
{
    allocate channel t1 type 'sbt_tape';
        SEND channel t1 '-TCPSEVER=<C>';
    allocate channel t2 type 'sbt_tape';
        SEND channel t2 '-TCPSEVER=<D>';

    backup
        filesperset 5
        format 'df_%t_%s_%p'
        (database);
    release channel t2;
    release channel t1;
}
```

Example 2:

This backup script allocates two parallel connections to the Tivoli Storage Manager server. The Tivoli Storage Manager server views these connections as two separate sessions:

```

run
{
  allocate channel t1 type 'sbt_tape' parms
    'ENV=(TDPO_OPTFILE=C:\oracle\scripts\tdpo.opt)';
  allocate channel t2 type 'sbt_tape' parms
    'ENV=(TDPO_OPTFILE=C:\oracle\scripts\tdpo.opt)';

  backup
    filesperset 5
    format 'df_%t_%s_%p'
    (database);

}

```

Example 3:

This restore script allocates one parallel connection to the Tivoli Storage Manager server:

```

run
{
  allocate channel t1 type 'sbt_tape' parms
    'ENV=(TDPO_OPTFILE=C:\oracle\scripts\tdpo.opt)';
  restore database;
  recover database;
  alter database open;
}

```

Note:

1. The allocate channel entry is divided on two lines after the parms option to accommodate page formatting.
2. The Oracle database must be in mount mode for the restore to succeed.

The Duplex Copy function

With Data Protection for Oracle, you can use the Oracle Server Duplex backup feature to make up to four exact duplicate copies of a backup. The backup can then be stored on different backup media.

About this task

A different management class is required for each backup copy. By default, the primary management class is the default management class on the policy domain that is defined for the Data Protection for Oracle node.

Note: It might be necessary to define the Oracle parameter value (BACKUP_TAPE_IO_SLAVES=TRUE) in the `init.ora` file of the target database for Data Protection for Oracle to use the duplex copy feature. Refer to your Oracle documentation about the use of this Oracle parameter.

For example, to create four backup copies:

Procedure

1. Specify the following option in the RMAN backup script:

```
set duplex=4
```
2. Define the following options in the `tdpo.opt` file:
 - `tdpo_mgmt_class_2`
 - `tdpo_mgmt_class_3`
 - `tdpo_mgmt_class_4`

3. Run the RMAN backup script.

Results

The following backup behavior occurs:

- The first backup copy is bound to the default management class to which the node is registered.
- The second backup copy is bound to the management class defined by the `tdpo_mgmt_class_2` option.
- The third backup copy is bound to the management class defined by the `tdpo_mgmt_class_3` option.
- The fourth backup copy is bound to the management class defined by the `tdpo_mgmt_class_4` option.

Note: Take note of the considerations provided:

- The duplex copy feature does not use *include* statements. It uses the management classes that are specified in the `tdpo.opt` file.
- You receive an error message if you specify **set duplex =4** in the RMAN backup script and do not define enough `tdpo_mgmt_class` options in the `tdpo.opt` file.
- To place duplicate copies on different media:
 - Make sure that the storage pool information for each backup copy group within the management classes is not the same.
 - Make sure that backups from these different storage pools are not moved to the same storage pool later.
- Duplicate data is sent across the network.
- If you specify **set duplex =4** and allocate one channel in the RMAN backup script, RMAN will start four sessions to the Tivoli Storage Manager server. Likewise, if you specify **set duplex =4** and allocate two channels in the RMAN backup script, RMAN will start eight sessions to the Tivoli Storage Manager server.
- The duplex copy feature sends the backup copies simultaneously. If the backup destination is tape, the number of sessions is a multiple of the duplex value. As a result, make sure that RMAN does not start more sessions than the maximum mount points allowed by the Tivoli Storage Manager server. The node definition option on the Tivoli Storage Manager server **maxnummp** determines the maximum number of mount points a client node can use on the Tivoli Storage Manager server during a backup operation. View the maximum mount points that are allowed by the Tivoli Storage Manager server for a particular node by entering the **query node** command from a Tivoli Storage Manager Administrative Client prompt:
q node f=d

See the appropriate *Tivoli Storage Manager Administrator's Reference* for more information about this option.

Review your current Oracle documentation about the duplex backup function.

Removing old backups

Data Protection for Oracle uses the Tivoli Storage Manager backup repository. Each database backup creates an object with a unique name. Since these objects have unique names, they always remain active and never expire. The database administrator (DBA) can control and coordinate copies that are removed from the Tivoli Storage Manager server with RMAN.

Before you begin

Ensure that **backdelete=yes** is specified by the Tivoli Storage Manager administrator during registration of your node. Specifying this parameter gives you permissions to delete backup objects.

About this task

Note: Make sure to use the same `tdpo.opt` file that was used for the original backup. Using this file enables the backup objects to be found on the Tivoli Storage Manager server.

Removing a backup example

A sample script for removing an old backup is provided.

About this task

To remove an old backup, issue this command from the RMAN prompt:

```
run
{
    allocate channel for delete type 'sbt_tape' parms
        'ENV=(TDPO_OPTFILE=C:\oracle\scripts\tdpo.opt)';

    change backupset backupset number delete;
}
```

Refer to the Oracle RMAN manual for more information about the **change** command and its options.

Setting up a schedule example

This example illustrates how to set up a schedule to automatically back up Oracle server databases.

About this task

For consistency, this procedure uses specific information. However, you can define a command file with any set of commands you choose. You can then use the same command file to define schedules on other Tivoli Storage Manager servers. All command information is presented as command-line interface entries.

This schedule in this procedure contains the following settings:

- The Data Protection for Oracle node name is `NodeA1`.
- The password for node name `NodeA1` is `PasswordA1`.
- The policy domain to which node name `NodeA1` is registered is `PolicyA1`.
- The schedule is a daily backup of an online Oracle database.
- The scheduled backup begins between 9:00 and 9:15 PM.

Setting up a schedule on the Tivoli Storage Manager server

Define a schedule on the Tivoli Storage Manager server to automatically run online backups of Oracle server databases.

Procedure

To set up a schedule on the Tivoli Storage Manager server:

1. Define the following schedule on the Tivoli Storage Manager server. You can enter the command on the Tivoli Storage Manager server console or on an administrative client. The administrative client does not have to be running on the same system as the Tivoli Storage Manager server.

```
define schedule PolicyA1 daily_orcbkup description="07Daily Online DB Backup"
action=command objects="C:\Program Files\Tivoli\TSM\Agent0BA\sched\orcsched.cmd"
starttime=21:00 duration=15 durunits=minutes period=1 perunits=day
dayofweek=any
```

The following message must display before you proceed to the next step:

ANR2500I Schedule daily_orcbkup defined in policy domain PolicyA1.

2. Issue the following command to associate the Data Protection for Oracle node to the backup schedule defined in step 1:

```
define association PolicyA1 daily_orcbkup NodeA1
```

Results

- A backup schedule is now defined on the Tivoli Storage Manager server.
- The backup schedule runs the command file orcsched.cmd, in the C:\Program Files\Tivoli\TSM\Agent0BA\sched\ or c:\progra~1\tivoli\tsm\agentoba\sched\ directory.
- The backup runs daily around 9:00 PM.
- The backup schedule can start on any day of the week.
- You can run the Tivoli Storage Manager **query schedule** and **query association** commands to confirm that the schedule and node association are set correctly.

Setting up a schedule on the workstation with the Oracle Server

Use this procedure to define a schedule on the workstation with the Oracle Server.

About this task

This example assumes the following setup:

- The Tivoli Storage Manager backup archive client is installed on the Oracle Server in the following directory:
C:\Program Files\Tivoli\TSM\baclient
- The Data Protection for Oracle for the Oracle server is installed in the following directory:
C:\Program Files\Tivoli\TSM\Agent0BA
- The options files in each of these directories are updated so that the communication parameters point to the Tivoli Storage Manager server.

Procedure

To set up a schedule on the workstation with the Oracle Server:

1. Log in with a Windows account that has administrative privileges.
2. Change to the backup archive client installation directory with this command:
C:\>cd \Program Files\Tivoli\TSM\baclient
3. Use the Tivoli Storage Manager Client Service Configuration Utility `dsmcutil` to install the Tivoli Storage Manager Scheduler Service. Type in the command from your current location in the backup archive client installation directory.

For example:

```
dsmcutil inst /name:"TSM Oracle Bkup Scheduler"  
/node:nodea1 /password:passworda1 /autostart:yes  
/clientdir:"C:\Program Files\Tivoli\TSM\baclient"  
/optfile:"C:\Program Files\Tivoli\TSM\AgentOBA\dsm.opt"  
/validate:yes
```

You can modify the command to contain options specific to your schedule.

This example shows the output:

```
TSM Windows NT Client Service Configuration Utility  
Command Line Interface - Version 7, Release 1, Level 0  
(C) Copyright IBM Corporation, 1990, 2013, All Rights Reserved.  
Last Updated Jun 29 2012  
TSM Api Version 7.1.0
```

```
Command: Install TSM Client Service  
Machine: ABC-ARTXNC(Local Machine)
```

Installing TSM Client Service:

```
Machine           : ABC-ARTXNC  
Service Name      : TSM Oracle Bkup Scheduler  
Client Directory  : C:\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 'TSM Oracle Bkup Scheduler' .  
Updated registry value 'ADSMClientKey' .  
Updated registry value 'OptionsFile' .  
Updated registry value 'EventLogging' .  
Updated registry value 'ClientNodeName' .
```

Generating registry password ...

Authenticating TSM password for node NODEA1 ...

```
Connecting to TSM Server via client options file  
'C:\Program Files\Tivoli\TSM\AgentOBA\dsm.opt' ...
```

Password authentication successful.

The registry password for TSM node NODEA1 has been updated.

Starting the 'TSM Oracle Bkup Scheduler' service

The service was successfully started.

The options file that is defined for Data Protection for Oracle is used by the scheduler when it validates the node and password. The options file is also used when contacting the Tivoli Storage Manager server for schedule information. This example assumes that the dsm.opt file is updated so that the communication parameters point to the Tivoli Storage Manager server to which the Oracle databases are to be backed up.

If this message displays:

```
A communications error occurred connecting to the TSM server
```

Ensure that the dsm.opt file contains entries that point to the correct Tivoli Storage Manager server. Also, ensure that the Tivoli Storage Manager server is running. If you must make a correction, enter the following command:

```
dsmcutil remove /name:"TSM Oracle Bkup Scheduler"
```

Then enter the command that is given at the beginning of this step.

4. Ensure that Data Protection for Oracle is running under the Oracle Site Services account so that it can access the Oracle backup API. The Data Protection for Oracle scheduler service must log on with that account. The scheduler service account information can be specified through the services applet in the control panel.
5. Create a batch file that is called orcsched.cmd and place it in the following location:

```
C:\Program Files\Tivoli\TSM\Agent0BA\sched\orcsched.cmd
```

Note: The scheduler service runs from the Windows system directory, so it is important that complete paths are specified for all file names and non-system commands. The scheduler service checks the Windows system directory for input and produces its output by default.

The following is an example of the orcsched.cmd file:

```
rem orcsched.cmd
rem =====

rem =====
rem setting oracle sid
rem =====
set oracle_sid=ilr

rem =====
rem set orc executable binary
rem =====
set ora_exe=c:\oracle\bin\rman.exe

%ora_exe% cmdfile 'C:\rman\bkdb.rmn'
```

The following is an example of the bkdb.rmn file:

```
connect target agnttest/agttest
run
{
allocate channel ch1 type 'SBT_TAPE' parms
'ENV=(TDPO_OPTFILE=C:\oracle\tdpo.opt)';
allocate channel ch2 type 'SBT_TAPE' parms
'ENV=(TDPO_OPTFILE=C:\oracle\tdpo.opt)';
allocate channel ch3 type 'SBT_TAPE' parms
'ENV=(TDPO_OPTFILE=C:\oracle\tdpo.opt)';
allocate channel ch4 type 'SBT_TAPE' parms
```

```
'ENV=(TDPO_OPTFILE=C:\oracle\tdpo.opt)';
backup (database);
release channel ch1;
release channel ch2;
release channel ch3;
release channel ch4;
}
```

Note: The allocate channel entry for each channel is divided on two lines after the parms option to accommodate page formatting. The bkdb.rmn file must have the allocate channel entry for each channel specified on one line of text.

6. The scheduler service is now installed, but is not yet started. To start the service, issue the following command:

```
net start "TSM Oracle Bkup Scheduler"
```

A message is displayed indicating that the Scheduler is starting up.

Note, because **autostart**=yes is specified, the Tivoli Storage Manager scheduling service is automatically started each time the Windows system is restarted.

Your system is now ready to run automatic daily online database backups.

Results

Note: If you want to use the Tivoli Storage Manager server scheduling mode:

1. Ensure that the Data Protection for Oracle options 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. Scheduling Data Protection for Oracle with the regular Windows backup archive client is an example of running more than one scheduler service.

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

2. If any changes that affect the scheduler service are made to the Data Protection for Oracle options file, the scheduler service must be restarted to pick up the changes. For example, changes to the Tivoli Storage Manager server address, schedule mode, or the TCP/IP address or port. The restart is done by entering these commands:

```
net stop "TSM Oracle Bkup Scheduler"
net start "TSM Oracle Bkup Scheduler"
```

3. The file named dsmsched.log contains status information for the Tivoli Storage Manager scheduler service. Though the default directory is \Tivoli\TSM\baclient\sched\dsmsched.log, you can specify a different directory, such as \Tivoli\TSM\Agent0BA\sched\dsmsched.log with the **schedlogname** option in the Data Protection for Oracle options file.
4. If a Tivoli Storage Manager scheduler service is already installed on your workstation for the regular backups of the Windows system, you must install another with a unique name to run the schedules that are defined for Data Protection for Oracle. The Tivoli Storage Manager scheduler service needs a different node name than the regular Tivoli Storage Manager backup archive client. If the path you want to use has a space in it, for example C:\Program Files\Tivoli\TSM\baclient, you can place quotation marks around the section of the path name that contains a space C:\"Program Files". You can also use the short form of the path name. The short form of a path name is made up of

the first six characters in the path, followed by a tilde (~) and the unique identifier. The following is an example of the short form of the path name:
C:\Progr~1\Tivoli\TSM\baclient

The orcsched.cmd file must contain these variables with the following values:

oracle_sid

Specify the Oracle system identifier for the instance to be used in the scheduled backup. A value of ilr is used in the example.

ora_exe

Specify the complete path and file name of the Oracle RMAN executable file. A value of c:\oracle\ora91\bin\rman.exe is used in the example.

cmdfile

Specify the complete path and file name of the Oracle command file. A value of C:\rman\bkdb.rmn is used in the example.

Querying backup objects

Use the **tdposync query** command to query the Tivoli Storage Manager server for information about objects that are backed up.

About this task

When you issue the **tdposync query** command, information about a backup object is displayed. Information is listed including the size and date of the backup, and whether the object is compressed, encrypted, or deduplicated by the Tivoli Storage Manager during the backup operation.

Related tasks:

"Data deduplication with Data Protection for Oracle"

Related reference:

"Query command" on page 47

Data deduplication with Data Protection for Oracle

You can use data deduplication with Data Protection for Oracle to reduce the amount of redundant data that is backed up to the Tivoli Storage Manager server.

Overview of data deduplication

Data deduplication is a method of reducing storage needs by eliminating redundant data

Two types of data deduplication are available with Tivoli Storage Manager: client-side data deduplication and server side data deduplication.

Client-side data deduplication is a data deduplication technique that is used on the Tivoli Storage Manager API to remove redundant data during backup processing before the data is transferred to the Tivoli Storage Manager server. Using client-side data deduplication can reduce the amount of data that is sent over a local area network.

Server side data deduplication is a data deduplication technique that is done by the server. The Tivoli Storage Manager server administrator can specify the data deduplication location on either the client or server to use with the **DEDUP** parameter on the **REGISTER NODE**, or **UPDATE NODE** server command.

Setting up Data Protection for Oracle for client-side data deduplication

You must edit the client options file before Data Protection for Oracle can use client-side data deduplication through the Tivoli Storage Manager API.

About this task

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

The Tivoli Storage Manager server administrator must enable data deduplication for the Data Protection for Oracle with the appropriate server command. For example:

```
UPDATE NODE ORACLE_NODE DEDUPLICATION=CLIENTORSERVER
```

The Tivoli Storage Manager server administrator must enable data deduplication on the storage pool where the Oracle data is stored with the following server command:

```
UPDATE STGPPOOL BACKUP_POOL DEDUPLICATION=YES
```

For more information, see the Tivoli Storage Manager information center at <http://pic.dhe.ibm.com/infocenter/tsminfo/v7r1/index.jsp>, and search on "API deduplication".

Results

After you created backups with client-side data deduplication enabled, you can use the **tdposync query** command to verify that client deduplication occurred during the backup operation. For detailed statistics, you can also query the Tivoli Storage Manager server activity log for the total data reduction.

You can also use the performance monitor feature in the Tivoli Storage Manager server to verify the percentage of data that has been deduplicated. The performance monitor feature is part of the Tivoli Storage Manager Administration Center. The data deduplication statistics are displayed graphically in the Performance GUI in the Administration Center.

Considerations:

- The **deduplication** and **enablelanfree** options are mutually exclusive. Therefore, you can use either one option or the other, but not both options together.
- The **deduplication** and **enableclientencryptkey** options are also mutually exclusive. Therefore, you can use either one option or the other, but not both options together.
- A local deduplication cache is an optimization that can reduce network traffic between the Tivoli Storage Manager server and the client. Client-side data deduplication can occur with or without it. Do not use the deduplication cache with Data Protection for Oracle for the following reasons:
 - The cache cannot be used when multiple processes, such as concurrent backups or Tivoli Storage Manager API applications, transfer content concurrently. Data Protection for Oracle backup operations that use multiple channels use multiple processes.
 - It is possible that the client deduplication cache can become out of sync with the server-deduplicated disk storage pool. This state can be the result of object expiration, file space deletion, and overflow to an associated tape

storage pool. When the client cache contains entries that are no longer in the Tivoli Storage Manager server deduplicated pool, the cache is reset and the backup operations fails. The Tivoli Storage Manager API does not attempt the backup again.

- When Tivoli Storage Manager server expiration or a similar process that removes deduplicated data extents runs concurrently with a deduplicated backup, the backup might fail. Backup operations with client-side deduplication enabled fails with the following messages:
 - Return code=254
 - Error message: ANS7899E The client referenced a deduplicated extent that does not exist on the TSM server.

Related tasks:

“Determining total data reduction”

Related reference:

“Query command” on page 47

Determining total data reduction

You can determine the percentage of total data reduction by querying the Tivoli Storage Manager server activity log.

About this task

Look for message number ANU2526I, which displays the data deduplication statistics, as shown in the following example:

```
ANE4991I (Session: 67, Node: HALLEY_ORC) DP Oracle Win64 ANU0599 TDP for Oracle: (5508): =>()
ANU2526I Backup details for backup piece \adsmorc\orcnt\df_727444762_116_1 (database "orcl").
Total bytes processed: 9961472. Deduplicated: Yes. Bytes after deduplication: 2272805.
Deduplication reduction: 77.18%. Compressed: Yes. Bytes after compression: 52253. Compressed by: 97.70%.
Encryption: None. LAN-Free: No. Total bytes sent: 52253. Total data reduction: 99.48%.
Total processing time: 00:00:01. Throughput rate: 9728.00Kb/Sec. (SESSION: 67)
```

In the following example, the Oracle database backup piece size is 9,961,472 bytes. Then, it was deduplicated and the number of bytes after deduplication is 2,272,805.

The total data reduction is calculated as follows:

- The percentage of data that is deduplicated is as follows:
$$\text{Deduplication reduction} = (1 - 2272805 / 9961472) = 0.7718$$
- After data deduplication, the object was compressed. The number of bytes before compression is the number of bytes after deduplication. The data was compressed to 52,253 bytes. Therefore,
$$\text{Compressed by} = (1 - 52253 / 2272805) = 0.9770$$
- The total bytes sent to the server equals the number of bytes after compression. The formula for total data reduction is as follows:
$$\begin{aligned} \text{Total data reduction} &= (1 - \text{bytes after compression} / \text{bytes processed}) \\ &= (1 - 52253 / 9961472) = 0.9948 \end{aligned}$$

Results

If there is no deduplication, the number of bytes after deduplication equals the number of bytes processed. If there is no compression, the number of bytes after compression equals the number of bytes after deduplication.

If you want to find out data reduction across multiple backup pieces, you can add up the numbers and calculate the ratios.

You can also use the performance monitor feature in the Tivoli Storage Manager server to verify the percentage of data that has been deduplicated. The performance monitor feature is part of the Tivoli Storage Manager Administration Center. The data deduplication statistics are displayed graphically in the Performance GUI in the Administration Center.

For more information about deduplication statistics, see the topic "Querying a storage pool for statistics about data deduplication" in the Tivoli Storage Manager Information Center at <http://pic.dhe.ibm.com/infocenter/tsminfo/v7r1/index.jsp>.

Chapter 5. Commands and utilities for Data Protection for Oracle

The Data Protection for Oracle commands and utilities are used to protect Oracle Server data.

tdpoconf and tdposync utilities

Set up and maintain Data Protection for Oracle with the tdpoconf and tdposync utilities. Find the utilities in the directory where Data Protection for Oracle is installed.

Use the Data Protection for Oracle utilities to do the following tasks:

- Set up and maintain Data Protection for Oracle with the tdpoconf utility. The utility is also used for password maintenance.
- Synchronize the RMAN catalog and Oracle control file by using tdposync. The utility is used to delete Oracle backups that are stored on the Tivoli Storage Manager.
- Query objects that are backed up on the Tivoli Storage Manager by using the tdposync utility.

Command line syntax and characteristics

Guidelines for the command line syntax for the Data Protection for Oracle utilities.

The Data Protection for Oracle utilities use the following command line syntax:

tdpoconf command 0 or more optional parameters

tdposync command 0 or more optional parameters

The command-line parameters have the following characteristics:

- Minimum abbreviations for keywords are indicated in uppercase.
- Optional parameters begin with a dash (-).
- Optional parameters can display in any order.
- Some keyword parameters require a value that is separated by the equal sign (=).
- If a parameter requires more than one value, the values are separated with commas.
- A space separates the invocation from the command and the command from any optional parameters.
- Each parameter is separated from others by a space.
- If a parameter value includes spaces, the entire parameter must be enclosed in double quotation marks.

tdpoconf utility

The `tdpoconf` utility provides setup tasks for configuring Data Protection for Oracle. The utility uses the `tdpo.opt` file that is stored in the installation directory to centralize information for setup purposes.


Use the following commands with the `tdpoconf` utility:

- PASSWord
- SHOWENVironment

password command

Use the **password** command to create a password or change an existing password on the Tivoli Storage Manager server. You are prompted to enter both the old and new passwords when you use this utility to change the password.

Syntax



Optional parameters

-TDPO OPTFILE=pathname

This parameter specifies the fully qualified path name to the `tdpo.opt` file. If you choose not to specify this option, the default path is used.

Example

An output example of the **tdpoconf password** command is provided:

[illegible]

Please enter current password:

Please enter new password:

Please reenter new password for verification:

```
ANU0260I Password successfully changed.
```


showenvironment command

Use the **showenvironment** command to query the Tivoli Storage Manager server with the options that are set in `-TDPO_OPTFILE`, the `tdpo.opt` file in the default installation directory, or the default values set by Data Protection for Oracle.

The screen output displays information about the Tivoli Storage Manager API and Tivoli Storage Manager server. This command is useful if you are troubleshooting setup errors for Data Protection for Oracle. If the password file is not initialized properly, the output of the `tdpoconf showenvironment` command reports the error.

Tip: To ensure that the environment is set up correctly before you use RMAN, direct the setup output to a file with the `-outfile` option.

Syntax



Optional parameters

-TDPO_OPTFILE=*pathname*

This parameter specifies the fully qualified path name to the `tdpo.opt` file. The options file is used by the utilities and the Data Protection for Oracle library.

-OUTfile=*pathname*

This parameter specifies the fully qualified path name to the output file. The formatted text of this file is the same content that in the output on screen.

Example

The following example shows the output of the **tdpoconf showenvironment** command:

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

```
Data Protection for Oracle Information
Version:          7
Release:          1
Level:            0
Sublevel:         0
Platform:         64bit DP Oracle Win64
```

```
Tivoli Storage Manager Server Information
Server Name:      DSMSERV
Server Address:   TSMSEVER
Server Type:      Linux/x86_64
Server Port:      1500
Communication Method: TCP/IP
```

```
Session Information
Node Name:        NODE_ORC
Node Type:        DP Oracle Win64
DSMI_DIR:         C:\Program Files\Common Files\Tivoli\TSM\api64
DSMI_ORC_CONFIG:  n:\Program Files\tivoli\tsm\agentoba64\dsm.opt
TDPO_OPTFILE:     N:\Program Files\Tivoli\TSM\Agent0BA64\tdpo.opt
Compression:      TRUE
License Information: License file exists and contains valid license data.
```

tdposync utility

The **tdposync** utility checks for items on the Tivoli Storage Manager server that are not in the RMAN catalog or Oracle control file. With this utility, you can repair these discrepancies by removing unwanted objects from the Tivoli Storage Manager, and reclaim space on the server.

Attention: Deleted files and inactive files cannot be restored. When you are using this utility to delete files, ensure that you do not log in to the wrong node name. You might query a different database than intended, and delete files in error. Ensure that the node name in the PICK window is the one you need. See “Optional parameters” on page 44 and “PICK window” on page 46 for further details.

When you run an RMAN deletion script, entries are deleted in the RMAN recovery catalog or Oracle control file before confirmation from the Tivoli Storage Manager server. In cases where objects are not found on the Tivoli Storage Manager server, RMAN tries to delete backup sets from the Tivoli Storage Manager server and fails. However, the entries in the RMAN catalog or control file for these objects are still removed. When they are deleted, RMAN can no longer identify these backups through the catalog or control file even though the file exists on the Tivoli Storage Manager server. This utility therefore synchronizes the contents of the servers.

When the RMAN catalog or control file contains backups that are marked as expired, RMAN still considers these objects as existing. If you run the **tdposync** utility against these objects, it recognizes these objects in the RMAN catalog or control file and on the Tivoli Storage Manager server and considers them to be in sync. Therefore, you must delete these objects from the RMAN catalog or control file for them to be deleted from the Tivoli Storage Manager server. Use the Oracle **crosscheck** command to verify whether the backups exist. Then, use the Oracle **delete expired** command to remove their record from the RMAN catalog or control file.

When you start **tdposync**, the following processing takes place:

1. Prompts you for the RMAN catalog owner ID or the Oracle database user name, password, and connect string.
2. Gathers information for the Oracle servers.
3. Queries the Oracle backup catalog and the Tivoli Storage Manager server.
4. Displays a list of files that exist on the Tivoli Storage Manager server but not in the RMAN catalog or Oracle control file.
5. Prompts you to take one of the following actions:
 - Delete any files found causing the discrepancy.
 - Delete all files.
 - Exit the program without deleting files from the Tivoli Storage Manager server.

tdposync considerations

To run the **tdposync** utility successfully, resynchronize the Oracle catalogs with the target databases. If you are using multiple Oracle catalogs, use the **numcatalogs** parameter. Each Oracle database must be backed up to the Tivoli Storage Manager server.

The following information must be considered before you use the **tdposync** command:

- Resynchronize Oracle catalogs with the target databases before you run the **tdposync syncdb** command. First, connect to the target database and the catalog database. The following is an example:

```
$ rman target xxx/yyy@targetdb rcvcat xxx/yyy@catalogdb
```

When you are connected to both databases, type **resync catalog** at the RMAN prompt.

- By default, Data Protection for Oracle prompts you to synchronize with one Oracle catalog at a time. If you use multiple Oracle catalogs to back up multiple target databases to the same file space, the same node name, and the same owner name on the same Tivoli Storage Manager server, you must use **-numcatalogs=number**. This action is necessary so that **tdposync** has all the information to correctly query both Oracle and the Tivoli Storage Manager.

Similarly, if you use Oracle control files to back up multiple target databases to the same file space, the same node name, and the same owner name on the same , you must use **-numinstances=number**.

If, for example, you back up only one target database by using two catalogs, do not specify this option. However, if you back up two target databases by using two catalogs, one catalog for each, to the same under the same file space, node name, and owner name, you must specify **numcatalogs**. If you fail to provide information for the second target database by not specifying two catalogs, that database is displayed as eligible for deletion. For more information, see “Optional parameters” on page 44.

Restriction: Failure to provide all pertinent and correct information can result in erroneous output. To prevent the erroneous output, see the next consideration.

- If you have more than one Oracle database, back up each Oracle target database to its own file space on the Tivoli Storage Manager server. To back up each Oracle target database to its own file space, use the **tdpo_fs** option in the **tdpo.opt** file. For best results, use a separate Data Protection for Oracle options file for each database that you back up to Tivoli Storage Manager. In this way, it is only necessary to synchronize one catalog at a time, one for each target database. The possibility of showing wrong information in the PICK window is minimized.

Tip: Make sure to use the same **tdpo.opt** file that was used for the original backup.

- If the information for **sqlplus** that you provide to **tdposync** is incorrect, such as logon, password, or connect string information, **sqlplus** stops at its logon screen. You must log on again at the prompt by using the RMAN catalog owner ID, password, and connect string. For example:

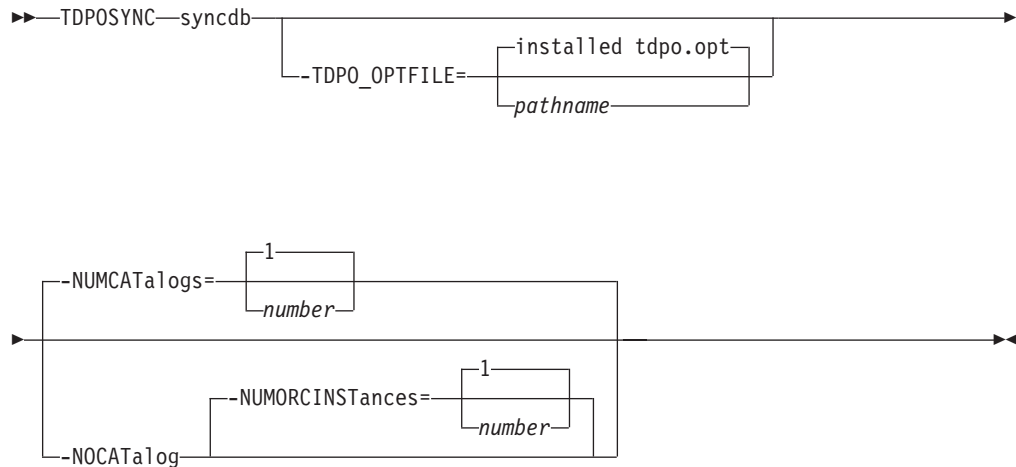
```
login/password@connectstring
```

where **connectstring** represents the Oracle database in which to connect. The **connectstring** is also sometimes referred to as the Transparent Network Substrate (TNS) alias. When the correct input is entered, **tdposync** proceeds.

syncdb command

The **syncdb** command synchronizes Oracle catalog databases or the Oracle control file with the Tivoli Storage Manager server.

Syntax



Optional parameters

-TDPO_OPTFILE=pathname

This parameter specifies the fully qualified path name to the **tdpo.opt** file. This file is the options file that is used by the utilities and the Data Protection for Oracle library. This file contains the information for the Tivoli Storage Manager server name and address that **tdposync** needs for synchronizing.

Note: For **syncdb TDPO_OPTFILE**, you must specify the same options file values that were used to do the original backup operations.

-NUMCATalogs=number

This parameter specifies the number of Oracle catalog databases that you want to synchronize and prompts you for information for each catalog.

Specify this option only when you use multiple Oracle catalogs to back up multiple target databases to the same Tivoli Storage Manager server under the same file space, node name, and owner name.

According to the number you specify for **-numcatalogs**, you are prompted for the user name, password, and connect string for each. If you do not specify **-numcatalogs**, the default is 1, and you are prompted only once.

You are prompted for start and end dates for your query. Then you are prompted for the following information for each catalog:

- Catalog # User Name:
- Catalog # Password:
- Catalog # Connect String:

You are also prompted for the following date information to narrow your search:

- From Date: MM/DD/YYYY
- To Date: MM/DD/YYYY

If no dates are specified, Data Protection for Oracle displays all objects that are not in sync.

-NOCATalog

This parameter specifies that the **tdposync** utility uses the backup history information that is stored in the Oracle control file rather than a catalog database to reconcile the Tivoli Storage Manager database with the RMAN backup history.

-NUMORCINSTANCES=number

This parameter specifies the number of Oracle instances that you want to synchronize and prompts you for information for each instance.

Specify this option only when you use multiple Oracle instances to back up multiple target databases to the same Tivoli Storage Manager server under the same file space, node name, and owner name.

According to the number you specify for **-numorcinstances**, you are prompted for the user name, password, and connect string for each instance. If you do not specify a value for **-numorcinstances**, the default is 1, and you are prompted only once.

For each Oracle instance, the following information is requested:

- Oracle Database # User Name
- Oracle Database # Password
- Oracle Database # Connect String

You are also prompted for the following date information to narrow your search:

- From Date: MM/DD/YYYY
- To Date: MM/DD/YYYY

If no dates are specified, Data Protection for Oracle shows all objects that are not in sync.

Example

Synchronize the Tivoli Storage Manager database with the RMAN backup history and the Oracle control file using the **tdposync syncdb** command. The following output is displayed:

Command: **TDPOSYNC syncdb -nocatalog -numorcinstances=2**

Output:

```
IBM Tivoli Storage Manager for Databases:
Data Protection for Oracle
Version 7, Release 1, Level 0.0
(C) Copyright IBM Corporation 1997, 2013. All rights reserved.
From Date (01/01/1990): 01/01/2013
To Date (01/05/2013): 12/12/2013
```

```
Oracle Database 1 User Name: OrcUser1
Oracle database 1 Password: OrcUser1pw
Oracle database 1 Connect String: Oracle_DB_A
```

```
Oracle Database 2 User Name: OrcYser2
Oracle database 2 Password: OrcUser2pw
Oracle database 2 Connect String: Oracle_DB_B
```

When you specify the **syncdb -numcatalogs** parameter, you are prompted for input for each catalog on the node. This example assumes two separate RMAN catalogs, **rman** and **rman2**:

```
C:\Tivoli\Tsm\Agent0ba>tdposync syncdb -numcatalogs=2 -TDPO_OPTFILE=c:\RMAN\scripts\tdpo.opt
```

Output:

IBM Tivoli Storage Manager for Databases:

Data Protection for Oracle

Version 7, Release 1, Level 0.0

(C) Copyright IBM Corporation 1997, 2013. All rights reserved.

From Date (01/01/1990): 01/01/2013

To Date (01/05/2013): 12/12/2013

Catalog 1 User Name: rman

Catalog 1 Password: rman

Catalog 1 Connect String: rman

Catalog 2 User Name: rman2

Catalog 2 Password: rman2

Catalog 2 Connect String: rman2

Note: **From Date** implicitly searches from time=00:00:01 and **To Date** implicitly searches to time=23:59:59 of the same day.

PICK window:

The PICK window provides information to help you decide if the files that are displayed are out of synchronization with the Oracle catalog or control file.

The following information is provided:

- The node with which you are querying the Tivoli Storage Manager server
- The date of the file backup
- The size of the backup
- The backup name \h1\orcnt\11

Attention: Use caution when you are selecting files for deletion. If you are unsure that the files in question are out of synchronization, do further research before you delete them. Deleted files cannot be restored.

Example

The PICK window shows the node names, and names the files that are backed up. The following example shows the output that is displayed for a node called AGENT_NODE:

```

Synchronize utility PICK Window
Node Name:                AGENT_NODE

Backup Date                Size                Backup Name
-----
1. 01/09/2013 09:19:59    50.35MB    \adsmorc\orcnt\1kc2cnfv_1_1
2. 01/02/2013 11:36:20    427.08MB   \adsmorc\orcnt\4kc3cnfv_1_1
3. 01/02/2013 07:14:30    508.00MB   \adsmorc\orcnt\4qcgdhfr_1_1
4. 01/02/2013 07:21:38    763.94MB   \adsmorc\orcnt\4ocf8999_1_1
5. 01/09/2013 11:00:11     69.12MB    \adsmorc\orcnt\4ocf8999_1_2
6. 01/09/2013 11:00:12    950.09MB   \adsmorc\orcnt\4ocf8999_1_3
7. 01/09/2013 11:00:13    656.69MB   \adsmorc\orcnt\4rch25jk_1_1
8. 01/09/2013 11:00:14    135.36MB   \adsmorc\orcnt\4rch25jk_1_2
9. 01/09/2013 11:00:15    298.01MB   \adsmorc\orcnt\4rch25jk_1_30

0-----10-----20-----30-----40-----50-----60-----70
<U>=Up <D>=Down <T>=Top <B>=Bottom <R>=Right <L>=Left
<G>=Goto Line # <#>=Toggle Entry <+>=Select A1<->=Deselect A11
<#:#>=Select A Range <#:#->=Deselect A Range <O>=Ok <C>=Cancel
pick>

```

Files that are selected for deletion are marked by a plus (+). To delete selected files:

1. Enter **OK** at the PICK prompt.
A warning message is shown confirming the deletion of the selected files.
2. Enter **Yes** to delete the selected files from the Tivoli Storage Manager server.


Query command

Use this command to query the Tivoli Storage Manager server for information about objects that are backed up. You can obtain information such as whether an object is compressed, encrypted, or deduplicated by the client during a backup operation.

The **query** command uses the options that are set in the **-TDPO_OPTFILE** parameter, the **tdpo.opt** file in the default installation directory, or the default values set by Data Protection for Oracle to query the Tivoli Storage Manager server.

When you issue the **tdposync query** command, you are prompted to enter date range for the query. The screen output displays information about the objects that were backed up to the Tivoli Storage Manager server between the start and end dates that you specified.

Syntax



Optional parameters

-TDPO OPTFILE =*pathname*

This parameter specifies the fully qualified path name to the `tdpo.opt` file. This file is the options file that is used by the utilities and the Data Protection for Oracle library. The file contains the information for the Tivoli Storage Manager server and the server address that **tdposync** command must use for synchronizing.

When you specify the **query TDPO_OPTFILE** command, you must specify the same options file values that were used for the original backup operations. If you do not specify the **TDPO_OPTFILE** path, the default value in the default Oracle installation path (/Program Files/Tivoli/TSM/Agent0BA64/tdpo.opt) is used.

Description of the output fields

Name Object name on the Tivoli Storage Manager server; for instance, /fs/h1/11.

Owner

The name of the user who backed up the object.

Size The size of the object size on the Tivoli Storage Manager server.

Creation Date / Time

The date and time the object was backed up.

Compressed

Lists whether an object was compressed during the backup operation.

Encryption Type

Lists the type of encryption that was used during the backup operation.
The possible values are as follows:

None The object was not encrypted.

AES-128

The object was encrypted by using AES-128 encryption.

DES-56

The object was encrypted by using DES-56 encryption.

Client-deduplicated

Lists whether an object underwent client-side data deduplication.

Examples

Use the `tdposync query` command to find information about backed up objects, encryption type and data deduplication.

Query the Tivoli Storage Manager server for information about objects that are backed up

The command to be run is `tdposync query`.

The following output is displayed:

```

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

From Date (01/01/2013):
To Date (07/02/2013):

Backup Object Information
-----

Name ..... \adsmorc\orcnt\df_722435657_35_1
Owner.....
Size ..... 2,010 KB
Creation Date / Time ..... 07/02/2013 10:08:20
Compressed ..... Yes
Encryption Type ..... None
Client-deduplicated ..... No

Backup Object Information
-----

...

```

Finding the encryption type

When you issue the **tdposync query** command, the entire list of backup object information is printed to the command prompt window without page separators, scrolling, or canceling capability. Redirect the output of the query to a file and find out the encryption type that was used for the backups from the previous week.

Command: `echo from_date > in.txt & echo to_date >> in.txt & tdposync query < in.txt > out.txt`

Open the file `out.txt` with a text editor and search for Encryption Type to determine the type of encryption that was used.

Finding data deduplication information

Determine the data deduplication reduction for a particular node by querying the Tivoli Storage Manager server activity log for the ANU2526I message.

Related tasks:

“Data deduplication with Data Protection for Oracle” on page 35

Appendix A. Tivoli support information

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

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

Communities and other learning resources

In addition to product documentation, many forms of assistance are available to help you get started as you deploy and use the Tivoli Storage Manager family of products. These resources can also help you to solve problems that you might have.

You can use forums, wikis, and other social media tools to ask questions, talk to experts, and learn from others.

User groups

Tivoli Global Storage Virtual User Group

Access this user group at <http://www.tivoli-ug.org/storage>.

This group makes it possible for individuals from many different industries and types of organizations to share information and work directly with the IBM product experts. Local chapters also exist where members meet in person to share experiences and hear from guest speakers.

ADSM.ORG

Access this mailing list at <http://adsm.org>.

This independently managed Storage Management discussion forum started when Tivoli Storage Manager was known as ADSTAR Distributed Storage Manager (ADSM). The members of this forum have many years of experience with Tivoli Storage Manager in almost every type of IT environment.

To subscribe to the forum, send an email to listserv@vm.marist.edu. The body of the message must contain the following text: `SUBSCRIBE ADSM-L your_first_name your_family_name`.

Tivoli Storage Manager community on Service Management Connect

Access Service Management Connect at <http://www.ibm.com/developerworks/servicemanagement>. In the Storage Management community of Service Management Connect, you can connect with IBM in the following ways:

- Become involved with transparent development, an ongoing, open engagement between users and IBM developers of Tivoli products. You can access early designs, sprint demonstrations, product roadmaps, and prerelease code.
- Connect one-on-one with the experts to collaborate and network about Tivoli and the Tivoli Storage Manager community.
- Read blogs to benefit from the expertise and experience of others.

- Use wikis and forums to collaborate with the broader user community.

Tivoli Storage Manager wiki on developerWorks®

Access this wiki at <https://www.ibm.com/developerworks/servicemanagement/sm/index.html>.

Find the latest best practices, white papers, and links to videos and other resources. When you log on, you can comment on content, or contribute your own content.

Tivoli Support Technical Exchange

Find information about upcoming Tivoli Support Technical Exchange webcasts at http://www.ibm.com/software/sysmgmt/products/support/supp_tech_exch.html. Replays of previous webcasts are also available.

Learn from technical experts who share their knowledge and then answer your questions. The sessions are designed to address specific technical issues and provide in-depth but narrowly focused training.

Other social media sites

LinkedIn

You can join groups on LinkedIn, a social media site for professionals. For example:

- **Tivoli Storage Manager Professionals:** <http://www.linkedin.com/groups/Tivoli-Storage-Manager-Professionals-54572>
- **TSM:** <http://www.linkedin.com/groups?gid=64540>

Twitter

Follow @IBMStorage on Twitter to see the latest news about storage and storage software from IBM.

Tivoli education resources

Use these education resources to help you increase your Tivoli Storage Manager skills:

Tivoli Education and Certification website

View available education at <http://www.ibm.com/software/tivoli/education>.

Use the Search for Training link to find local and online offerings of instructor-led courses for Tivoli Storage Manager.

Education Assistant

Access resources at <http://publib.boulder.ibm.com/infocenter/ieduasst/tivv1r0/index.jsp>.

Scroll to view the list of available training videos. Recorded product demonstrations are also available on a YouTube channel.

Searching knowledge bases

If a problem occurs while you are using one of the Tivoli Storage Manager family of products, you can search several knowledge bases.

Begin by searching the Tivoli Storage Manager Information Center at <http://pic.dhe.ibm.com/infocenter/tsminfo/v7r1>. Within the information center, you can enter words, phrases, or message numbers in the **Search** field to find relevant topics.

Searching the Internet

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

To search multiple Internet resources, go to the IBM support website at <http://www.ibm.com/support/entry/portal/>. You can search for information without signing in.

Sign in using your IBM ID and password if you want to customize the site based on your product usage and information needs. If you do not already have an IBM ID and password, click **Sign in** at the top of the page and follow the instructions to register.

From the support website, you can search various resources:

- IBM technotes.
- IBM downloads.
- IBM Redbooks® publications.
- IBM Authorized Program Analysis Reports (APARs). Select the product and click **Downloads** to search the APAR list.

Using IBM Support Assistant

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

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

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

You can find more information and download the IBM Support Assistant web page at <http://www.ibm.com/software/support/isa>.

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

Finding product fixes

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

Procedure

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

If you previously customized the site based on your product usage:

1. Click the link for the product, or a component for which you want to find a fix.
2. Click **Downloads**, and then click **Search for recommended fixes**.

If you have not previously customized the site:

Click **Downloads** and search for the product.

Receiving notification of product fixes

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

Procedure

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

Contacting IBM Software Support

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

Procedure

1. Ensure that you have completed the following prerequisites:
 - a. Set up a subscription and support contract.
 - b. Determine the business impact of the problem.
 - c. Describe the problem and gather background information.
2. Follow the instructions in “Submitting the problem to IBM Software Support” on page 56.

Setting up and managing support contracts

You can set up and manage your Tivoli support contracts by enrolling in IBM Passport Advantage®. The type of support contract that you need depends on the type of product you have.

Procedure

Enroll in IBM Passport Advantage in one of the following ways:

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

Determining the business impact

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

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

Describing the problem and gathering background information

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

To save time, know the answers to these questions:

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

Submitting the problem to IBM Software Support

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

Online

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

By telephone

For critical, system-down, or severity 1 issues, you can call 1-800-IBMSERV (1-800-426-7378) in the United States. For the telephone number to call in your country, go to the IBM Software Support Handbook web page at <http://www14.software.ibm.com/webapp/set2/sas/f/handbook/home.html> and click **Contacts**.

Appendix B. Accessibility features for the Tivoli Storage Manager product family

Accessibility features help users who have a disability, such as restricted mobility or limited vision to use information technology products successfully.

Accessibility features

The IBM Tivoli Storage Manager family of products includes the following accessibility features:

- Keyboard-only operation using standard operating-system conventions
- Interfaces that support assistive technology such as screen readers

The command-line interfaces of all products in the product family are accessible.

Tivoli Storage Manager Operations Center provides the following additional accessibility features when you use it with a Mozilla Firefox browser on a Microsoft Windows system:

- Screen magnifiers and content zooming
- High contrast mode

The Operations Center and the Tivoli Storage Manager Server can be installed in console mode, which is accessible.

The Tivoli Storage Manager Information Center is enabled for accessibility. For information center accessibility information, see “Accessibility features in the information center” (http://pic.dhe.ibm.com/infocenter/tsminfo/v7r1/topic/com.ibm.help.ic.doc/iehs36_accessibility.html).

Vendor software

The Tivoli Storage Manager product family includes certain vendor software that is not covered under the IBM license agreement. IBM makes no representation about the accessibility features of these products. Contact the vendor for the accessibility information about its products.

IBM and accessibility

See the IBM Human Ability and Accessibility Center (<http://www.ibm.com/able>) for information about the commitment that IBM has to accessibility.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan*

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who want to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Corporation
2Z4A/101
11400 Burnet Road
Austin, TX 78758
U.S.A.*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample

programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows: © (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, and [ibm.com](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>.

Intel, Itanium, and Pentium are registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and 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 company, product, or service names may be trademarks or service marks of others.

Privacy policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at <http://www.ibm.com/privacy> and IBM's Online Privacy Statement at <http://www.ibm.com/privacy/details> the section entitled "Cookies, Web Beacons and Other Technologies" and the "IBM Software Products and Software-as-a-Service Privacy Statement" at <http://www.ibm.com/software/info/product-privacy>.

Glossary

This glossary provides terms and definitions for Tivoli Storage Manager, Tivoli Storage FlashCopy Manager, and associated products.

The following cross-references are used in this glossary:

- *See* refers you from a nonpreferred term to the preferred term or from an abbreviation to the spelled-out form.
- *See also* refers you to a related or contrasting term.

For other terms and definitions, see the IBM Terminology website at www.ibm.com/software/globalization/terminology.

A

absolute mode

In storage management, a backup copy-group mode that specifies that a file is considered for incremental backup even if the file has not changed since the last backup. See also mode, modified mode.

access control list (ACL)

In computer security, a list associated with an object that identifies all the subjects that can access the object and their access rights.

access mode

An attribute of a storage pool or a storage volume that specifies whether the server can write to or read from the storage pool or storage volume.

ACK See acknowledgment.

acknowledgment (ACK)

The transmission of acknowledgment characters as a positive response to a data transmission.

ACL See access control list.

activate

To validate the contents of a policy set and then make it the active policy set.

active-data pool

A named set of storage pool volumes that contain only active versions of client

backup data. See also server storage, storage pool, storage pool volume.

active file system

A file system to which space management has been added. With space management, tasks for an active file system include automatic migration, reconciliation, selective migration, and recall. See also inactive file system.

active policy set

The activated policy set that contains the policy rules currently in use by all client nodes assigned to the policy domain. See also policy domain, policy set.

active version

The most recent backup copy of a file stored. The active version of a file cannot be deleted until a backup process detects that the user has either replaced the file with a newer version or has deleted the file from the file server or workstation. See also backup version, inactive version.

activity log

A log that records normal activity messages that are generated by the server. These messages include information about server and client operations, such as the start time of sessions or device I/O errors.

adaptive subfile backup

A type of backup that sends only changed portions of a file to the server, instead of sending the entire file. Adaptive subfile backup reduces network traffic and increases the speed of the backup.

administrative client

A program that runs on a file server, workstation, or mainframe that administrators use to control and monitor the server. See also backup-archive client.

administrative command schedule

A database record that describes the planned processing of an administrative command during a specific time period. See also central scheduler, client schedule, schedule.

administrative privilege class

See privilege class.

administrative session

A period of time during which an administrator user ID communicates with a server to perform administrative tasks. See also client node session, session.

administrator

A person responsible for administrative tasks such as access authorization and content management. Administrators can also grant levels of authority to users.

agent node

A client node that has been granted proxy authority to perform operations on behalf of another client node, which is the target node.

aggregate

An object, stored in one or more storage pools, consisting of a group of logical files that are packaged together. See also logical file, physical file.

aggregate data transfer rate

A performance statistic that indicates the average number of bytes that were transferred per second while processing a given operation.

application client

A program that is installed on a system to protect an application. The server provides backup services to an application client.

archive

To copy programs, data, or files to another storage media, usually for long-term storage or security. See also retrieve.

archive copy

A file or group of files that was archived to server storage

archive copy group

A policy object containing attributes that control the generation, destination, and expiration of archived files. See also copy group.

archive-retention grace period

The number of days that the storage manager retains an archived file when the server is unable to rebind the file to an appropriate management class. See also bind.

association

The defined relationship between a client

node and a client schedule. An association identifies the name of a schedule, the name of the policy domain to which the schedule belongs, and the name of a client node that performs scheduled operations.

audit

To check for logical inconsistencies between information that the server has and the actual condition of the system. The storage manager can audit information about items such as volumes, libraries, and licenses. For example, when a storage manager audits a volume, the server checks for inconsistencies between information about backed-up or archived files that are stored in the database and the actual data that are associated with each backup version or archive copy in server storage.

authentication rule

A specification that another user can use to either restore or retrieve files from storage.

authority

The right to access objects, resources, or functions. See also privilege class.

authorization rule

A specification that permits another user to either restore or retrieve a user's files from storage.

authorized user

A user who has administrative authority for the client on a workstation. This user changes passwords, performs open registrations, and deletes file spaces.

AutoFS

See automounted file system.

automatic detection

A feature that detects, reports, and updates the serial number of a drive or library in the database when the path from the local server is defined.

automatic migration

The process that is used to automatically move files from a local file system to storage, based on options and settings that are chosen by a root user on a workstation. See also demand migration, threshold migration.

automounted file system (AutoFS)

A file system that is managed by an

automounter daemon. The automounter daemon monitors a specified directory path, and automatically mounts the file system to access data.

B

backup-archive client

A program that runs on a workstation or file server and provides a means for users to back up, archive, restore, and retrieve files. See also administrative client.

backup copy group

A policy object containing attributes that control the generation, destination, and expiration of backup versions of files. A backup copy group belongs to a management class. See also copy group.

backup retention grace period

The number of days the storage manager retains a backup version after the server is unable to rebind the file to an appropriate management class.

backup set

A portable, consolidated group of active versions of backup files that are generated for a backup-archive client.

backup set collection

A group of backup sets that are created at the same time and which have the same backup set name, volume names, description, and device classes. The server identifies each backup set in the collection by its node name, backup set name, and file type.

backup version

A file or directory that a client node backed up to storage. More than one backup version can exist in storage, but only one backup version is the active version. See also active version, copy group, inactive version.

bind To associate a file with a management class name. See also archive-retention grace period, management class, rebind.

C

cache To place a duplicate copy of a file on random access media when the server migrates a file to another storage pool in the hierarchy.

cache file

A snapshot of a logical volume created by Logical Volume Snapshot Agent. Blocks are saved immediately before they are modified during the image backup and their logical extents are saved in the cache files.

CAD See client acceptor daemon.

central scheduler

A function that permits an administrator to schedule client operations and administrative commands. The operations can be scheduled to occur periodically or on a specific date. See also administrative command schedule, client schedule.

client A software program or computer that requests services from a server. See also server.

client acceptor

A service that serves the Java applet for the web client to web browsers. On Windows systems, the client acceptor is installed and run as a service. On AIX®, UNIX, and Linux systems, the client acceptor is run as a daemon.

client acceptor daemon (CAD)

See client acceptor.

client domain

The set of drives, file systems, or volumes that the user selects to back up or archive data, using the backup-archive client.

client node

A file server or workstation on which the backup-archive client program has been installed, and which has been registered to the server.

client node session

A session in which a client node communicates with a server to perform backup, restore, archive, retrieve, migrate, or recall requests. See also administrative session.

client option set

A group of options that are defined on

the server and used on client nodes in conjunction with client options files.

client options file

An editable file that identifies the server and communication method, and provides the configuration for backup, archive, hierarchical storage management, and scheduling.

client-polling scheduling mode

A method of operation in which the client queries the server for work. See also server-prompted scheduling mode.

client schedule

A database record that describes the planned processing of a client operation during a specific time period. The client operation can be a backup, archive, restore, or retrieve operation, a client operating system command, or a macro. See also administrative command schedule, central scheduler, schedule.

client/server

Pertaining to the model of interaction in distributed data processing in which a program on one computer sends a request to a program on another computer and awaits a response. The requesting program is called a client; the answering program is called a server.

client system-options file

A file, used on AIX, UNIX, or Linux system clients, containing a set of processing options that identify the servers to be contacted for services. This file also specifies communication methods and options for backup, archive, hierarchical storage management, and scheduling. See also client user-options file, options file.

client user-options file

A file that contains the set of processing options that the clients on the system use. The set can include options that determine the server that the client contacts, and options that affect backup operations, archive operations, hierarchical storage management operations, and scheduled operations. This file is also called the dsm.opt file. For AIX, UNIX, or Linux systems, see also client system-options file. See also client system-options file, options file.

closed registration

A registration process in which only an administrator can register workstations as client nodes with the server. See also open registration.

collocation

The process of keeping all data belonging to a single-client file space, a single client node, or a group of client nodes on a minimal number of sequential-access volumes within a storage pool.

Collocation can reduce the number of volumes that must be accessed when a large amount of data must be restored.

collocation group

A user-defined group of client nodes whose data is stored on a minimal number of volumes through the process of collocation.

commit point

A point in time when data is considered to be consistent.

communication method

The method by which a client and server exchange information. See also Transmission Control Protocol/Internet Protocol.

communication protocol

A set of defined interfaces that permit computers to communicate with each other.

compression

A function that removes repetitive characters, spaces, strings of characters, or binary data from the data being processed and replaces characters with control characters. Compression reduces the amount of storage space that is required for data.

configuration manager

A server that distributes configuration information, such as policies and schedules, to managed servers according to their profiles. Configuration information can include policy and schedules. See also enterprise configuration, managed server, profile.

conversation

A connection between two programs over a session that allows them to communicate with each other while processing a transaction. See also session.

copy backup

A full backup in which the transaction log files are not deleted so that backup procedures that use incremental or differential backups are not disrupted.

copy group

A policy object containing attributes that control how backup versions or archive copies are generated, where backup versions or archive copies are initially located, and when backup versions or archive copies expire. A copy group belongs to a management class. See also archive copy group, backup copy group, backup version, management class.

copy storage pool

A named set of volumes that contain copies of files that reside in primary storage pools. Copy storage pools are used only to back up the data that is stored in primary storage pools. A copy storage pool cannot be a destination for a backup copy group, an archive copy group, or a management class (for space-managed files). See also destination, primary storage pool, server storage, storage pool, storage pool volume.

D**daemon**

A program that runs unattended to perform continuous or periodic functions, such as network control.

damaged file

A physical file in which read errors have been detected.

database backup series

One full backup of the database, plus up to 32 incremental backups made since that full backup. Each full backup that is run starts a new database backup series. A number identifies each backup series. See also database snapshot, full backup.

database snapshot

A complete backup of the entire database to media that can be taken off-site. When a database snapshot is created, the current database backup series is not interrupted. A database snapshot cannot have incremental database backups associated with it. See also database backup series, full backup.

data center

In a virtualized environment, a container that holds hosts, clusters, networks, and data stores.

data deduplication

A method of reducing storage needs by eliminating redundant data. Only one instance of the data is retained on storage media. Other instances of the same data are replaced with a pointer to the retained instance.

data manager server

A server that collects metadata information for client inventory and manages transactions for the storage agent over the local area network. The data manager server informs the storage agent with applicable library attributes and the target volume identifier.

data mover

A device that moves data on behalf of the server. A network-attached storage (NAS) file server is a data mover.

data storage-management application-programming interface (DSMAPI)

A set of functions and semantics that can monitor events on files, and manage and maintain the data in a file. In an HSM environment, a DSMAPI uses events to notify data management applications about operations on files, stores arbitrary attribute information with a file, supports managed regions in a file, and uses DSMAPI access rights to control access to a file object.

data store

In a virtualized environment, the location where virtual machine data is stored.

deduplication

The process of creating representative records from a set of records that have been identified as representing the same entities.

default management class

A management class that is assigned to a policy set. This class is used to govern backed up or archived files when a file is not explicitly associated with a specific management class through the include-exclude list.

demand migration

The process that is used to respond to an

out-of-space condition on a file system for which hierarchical storage management (HSM) is active. Files are migrated to server storage until space usage drops to the low threshold that was set for the file system. If the high threshold and low threshold are the same, one file is migrated. See also automatic migration, selective migration, threshold migration.

desktop client

The group of backup-archive clients that includes clients on Microsoft Windows, Apple, and Novell NetWare operating systems.

destination

A copy group or management class attribute that specifies the primary storage pool to which a client file will be backed up, archived, or migrated. See also copy storage pool.

device class

A named set of characteristics that are applied to a group of storage devices. Each device class has a unique name and represents a device type of disk, file, optical disk, or tape.

device configuration file

1. For a storage agent, a file that contains the name and password of the storage agent, and information about the server that is managing the SAN-attached libraries and drives that the storage agent uses.
2. For a server, a file that contains information about defined device classes, and, on some servers, defined libraries and drives. The information is a copy of the device configuration information in the database.

disaster recovery manager (DRM)

A function that assists in preparing and using a disaster recovery plan file for the server.

disaster recovery plan

A file that is created by the disaster recover manager (DRM) that contains information about how to recover computer systems if a disaster occurs and scripts that can be run to perform some recovery tasks. The file includes information about the software and

hardware that is used by the server, and the location of recovery media.

domain

A grouping of client nodes with one or more policy sets, which manage data or storage resources for the client nodes. See also policy domain.

DRM See disaster recovery manager.

DSMAPI

See data storage-management application-programming interface.

dynamic serialization

Copy serialization in which a file or folder is backed up or archived on the first attempt regardless of whether it changes during a backup or archive. See also shared dynamic serialization, shared static serialization, static serialization.

E

EA See extended attribute.

EB See exabyte.

EFS See Encrypted File System.

Encrypted File System (EFS)

A file system that uses file system-level encryption.

enterprise configuration

A method of setting up servers so that the administrator can distribute the configuration of one of the servers to the other servers, using server-to-server communication. See also configuration manager, managed server, profile, subscription.

enterprise logging

The process of sending events from a server to a designated event server. The event server routes the events to designated receivers, such as to a user exit. See also event.

error log

A data set or file that is used to record error information about a product or system.

estimated capacity

The available space, in megabytes, of a storage pool.

event An occurrence of significance to a task or system. Events can include completion or

failure of an operation, a user action, or the change in state of a process. See also enterprise logging, receiver.

event record

A database record that describes actual status and results for events.

event server

A server to which other servers can send events for logging. The event server routes the events to any receivers that are enabled for the sending server's events.

exabyte (EB)

For processor, real and virtual storage capacities and channel volume, 2 to the power of 60 or 1 152 921 504 606 846 976 bytes. For disk storage capacity and communications volume, 1 000 000 000 000 000 000 bytes.

exclude

The process of identifying files in an include-exclude list. This process prevents the files from being backed up or migrated whenever a user or schedule enters an incremental or selective backup operation. A file can be excluded from backup, from space management, or from both backup and space management.

exclude-include list

See include-exclude list.

expiration

The process by which files, data sets, or objects are identified for deletion because their expiration date or retention period has passed.

expiring file

A migrated or premigrated file that has been marked for expiration and removal from storage. If a stub file or an original copy of a premigrated file is deleted from a local file system, or if the original copy of a premigrated file is updated, the corresponding migrated or premigrated file is marked for expiration the next time reconciliation is run.

extend

To increase the portion of available space that can be used to store database or recovery log information.

extended attribute (EA)

Names or value pairs that are associated with files or directories. There are three

classes of extended attributes: user attributes, system attributes, and trusted attributes.

external library

A collection of drives that is managed by the media-management system other than the storage management server.

F

file access time

On AIX, UNIX, or Linux systems, the time when the file was last accessed.

file age

For migration prioritization purposes, the number of days since a file was last accessed.

file device type

A device type that specifies the use of sequential access files on disk storage as volumes.

file server

A dedicated computer and its peripheral storage devices that are connected to a local area network that stores programs and files that are shared by users on the network.

file space

A logical space in server storage that contains a group of files that have been backed up or archived by a client node, from a single logical partition, file system, or virtual mount point. Client nodes can restore, retrieve, or delete their file spaces from server storage. In server storage, files belonging to a single file space are not necessarily stored together.

file space ID (FSID)

A unique numeric identifier that the server assigns to a file space when it is stored in server storage.

file state

The space management mode of a file that resides in a file system to which space management has been added. A file can be in one of three states: resident, premigrated, or migrated. See also migrated file, premigrated file, resident file.

file system migrator (FSM)

A kernel extension that intercepts all file system operations and provides any space

management support that is required. If no space management support is required, the operation is passed to the operating system, which performs its normal functions. The file system migrator is mounted over a file system when space management is added to the file system.

file system state

The storage management mode of a file system that resides on a workstation on which the hierarchical storage management (HSM) client is installed. A file system can be in one of these states: native, active, inactive, or global inactive.

frequency

A copy group attribute that specifies the minimum interval, in days, between incremental backups.

FSID See file space ID.

FSM See file system migrator.

full backup

The process of backing up the entire server database. A full backup begins a new database backup series. See also database backup series, database snapshot, incremental backup.

fuzzy backup

A backup version of a file that might not accurately reflect what is currently in the file because the file was backed up at the same time as it was being modified.

fuzzy copy

A backup version or archive copy of a file that might not accurately reflect the original contents of the file because it was backed up or archived the file while the file was being modified.

G

GB See gigabyte.

General Parallel File System (GPFS™)

A high-performance shared-disk file system that can provide data access from nodes in a clustered system environment. See also information lifecycle management.

gigabyte (GB)

For processor storage, real and virtual storage, and channel volume, 10 to the

power of nine or 1,073,741,824 bytes. For disk storage capacity and communications volume, 1,000,000,000 bytes.

global inactive state

The state of all file systems to which space management has been added when space management is globally deactivated for a client node.

Globally Unique Identifier (GUID)

An algorithmically determined number that uniquely identifies an entity within a system. See also Universally Unique Identifier.

GPFS See General Parallel File System.

GPFS node set

A mounted, defined group of GPFS file systems.

group backup

The backup of a group containing a list of files from one or more file space origins.

GUID See Globally Unique Identifier.

H

hierarchical storage management (HSM)

A function that automatically distributes and manages data on disk, tape, or both by regarding devices of these types and potentially others as levels in a storage hierarchy that range from fast, expensive devices to slower, cheaper, and possibly removable devices. The objectives are to minimize access time to data and maximize available media capacity. See also hierarchical storage management client, recall, storage hierarchy.

hierarchical storage management client (HSM client)

A client program that works with the server to provide hierarchical storage management (HSM) for a system. See also hierarchical storage management, management class.

HSM See hierarchical storage management.

HSM client

See hierarchical storage management client.

I

ILM See information lifecycle management.

image A file system or raw logical volume that is backed up as a single object.

image backup
A backup of a full file system or raw logical volume as a single object.

inactive file system
A file system for which space management has been deactivated. See also active file system.

inactive version
A backup version of a file that is either not the most recent backup version, or that is a backup version of a file that no longer exists on the client system. Inactive backup versions are eligible for expiration processing according to the management class assigned to the file. See also active version, backup version.

include-exclude file
A file containing statements to determine the files to back up and the associated management classes to use for backup or archive. See also include-exclude list.

include-exclude list
A list of options that include or exclude selected files for backup. An exclude option identifies files that should not be backed up. An include option identifies files that are exempt from the exclusion rules or assigns a management class to a file or a group of files for backup or archive services. See also include-exclude file.

incremental backup
The process of backing up files or directories, or copying pages in the database, that are new or changed since the last full or incremental backup. See also selective backup.

individual mailbox restore
See mailbox restore.

information lifecycle management (ILM)
A policy-based file-management system for storage pools and file sets. See also General Parallel File System.

inode The internal structure that describes the individual files on AIX, UNIX, or Linux

systems. An inode contains the node, type, owner, and location of a file.

inode number
A number specifying a particular inode file in the file system.

IP address
A unique address for a device or logical unit on a network that uses the Internet Protocol standard.

J

job file
A generated file that contains configuration information for a migration job. The file is XML format and can be created and edited in the hierarchical storage management (HSM) client for Windows client graphical user interface. See also migration job.

journal-based backup
A method for backing up Windows clients and AIX clients that exploits the change notification mechanism in a file to improve incremental backup performance by reducing the need to fully scan the file system.

journal daemon
On AIX, UNIX, or Linux systems, a program that tracks change activity for files residing in file systems.

journal service
In Microsoft Windows, a program that tracks change activity for files residing in file systems.

K

KB See kilobyte.

kilobyte (KB)
For processor storage, real and virtual storage, and channel volume, 2 to the power of 10 or 1,024 bytes. For disk storage capacity and communications volume, 1,000 bytes.

L

LAN See local area network.

LAN-free data movement

The movement of client data between a client system and a storage device on a storage area network (SAN), bypassing the local area network.

LAN-free data transfer

See LAN-free data movement.

leader data

Bytes of data, from the beginning of a migrated file, that are stored in the file's corresponding stub file on the local file system. The amount of leader data that is stored in a stub file depends on the stub size that is specified.

library

1. A repository for demountable recorded media, such as magnetic disks and magnetic tapes.
2. A collection of one or more drives, and possibly robotic devices (depending on the library type), which can be used to access storage volumes.

library client

A server that uses server-to-server communication to access a library that is managed by another storage management server. See also library manager.

library manager

A server that controls device operations when multiple storage management servers share a storage device. See also library client.

local

1. Pertaining to a device, file, or system that is accessed directly from a user system, without the use of a communication line. See also remote.
2. For hierarchical storage management products, pertaining to the destination of migrated files that are being moved. See also remote.

local area network (LAN)

A network that connects several devices in a limited area (such as a single building or campus) and that can be connected to a larger network.

local shadow volume

Data that is stored on shadow volumes localized to a disk storage subsystem.

LOFS See loopback virtual file system.

logical file

A file that is stored in one or more server storage pools, either by itself or as part of an aggregate. See also aggregate, physical file, physical occupancy.

logical occupancy

The space that is used by logical files in a storage pool. This space does not include the unused space created when logical files are deleted from aggregate files, so it might be less than the physical occupancy. See also physical occupancy.

logical unit number (LUN)

In the Small Computer System Interface (SCSI) standard, a unique identifier used to differentiate devices, each of which is a logical unit (LU).

logical volume

A portion of a physical volume that contains a file system.

logical volume backup

A back up of a file system or logical volume as a single object.

Logical Volume Snapshot Agent (LVSA)

Software that can act as the snapshot provider for creating a snapshot of a logical volume during an online image backup.

loopback virtual file system (LOFS)

A file system that is created by mounting a directory over another local directory, also known as mount-over-mount. A LOFS can also be generated using an automounter.

LUN See logical unit number.

LVSA See Logical Volume Snapshot Agent.

M

macro file

A file that contains one or more storage manager administrative commands, which can be run only from an administrative client using the MACRO command. See also Tivoli Storage Manager command script.

mailbox restore

A function that restores Microsoft Exchange Server data (from IBM Data Protection for Microsoft Exchange backups) at the mailbox level or mailbox-item level.

managed object

A definition in the database of a managed server that was distributed to the managed server by a configuration manager. When a managed server subscribes to a profile, all objects that are associated with that profile become managed objects in the database of the managed server.

managed server

A server that receives configuration information from a configuration manager using a subscription to one or more profiles. Configuration information can include definitions of objects such as policy and schedules. See also configuration manager, enterprise configuration, profile, subscription.

management class

A policy object that users can bind to each file to specify how the server manages the file. The management class can contain a backup copy group, an archive copy group, and space management attributes. See also bind, copy group, hierarchical storage management client, policy set, rebind.

maximum transmission unit (MTU)

The largest possible unit of data that can be sent on a given physical medium in a single frame. For example, the maximum transmission unit for Ethernet is 1500 bytes.

MB See megabyte.

media server

In a z/OS® environment, a program that provides access to z/OS disk and tape

storage for Tivoli Storage Manager servers that run on operating systems other than z/OS.

megabyte (MB)

For processor storage, real and virtual storage, and channel volume, 2 to the 20th power or 1,048,576 bytes. For disk storage capacity and communications volume, 1,000,000 bytes.

metadata

Data that describes the characteristics of data; descriptive data.

migrate

To move data to another location, or an application to another computer system.

migrated file

A file that has been copied from a local file system to storage. For HSM clients on UNIX or Linux systems, the file is replaced with a stub file on the local file system. On Windows systems, creation of the stub file is optional. See also file state, premigrated file, resident file, stub file.

migration

The process of moving data from one computer system to another, or an application to another computer system.

migration job

A specification of files to migrate, and actions to perform on the original files after migration. See also job file, threshold migration.

migration threshold

High and low capacities for storage pools or file systems, expressed as percentages, at which migration is set to start and stop.

mirroring

The process of writing the same data to multiple disks at the same time. The mirroring of data protects it against data loss within the database or within the recovery log.

mode

A copy group attribute that specifies whether to back up a file that has not been modified since the last time the file was backed up. See also absolute mode, modified mode.

modified mode

In storage management, a backup copy-group mode that specifies that a file

is considered for incremental backup only if it has changed since the last backup. A file is considered a changed file if the date, size, owner, or permissions of the file have changed. See also absolute mode, mode.

mount limit

The maximum number of volumes that can be simultaneously accessed from the same device class. The mount limit determines the maximum number of mount points. See also mount point.

mount point

A logical drive through which volumes are accessed in a sequential access device class. For removable media device types, such as tape, a mount point is a logical drive associated with a physical drive. For the file device type, a mount point is a logical drive associated with an I/O stream. See also mount limit.

mount retention period

The maximum number of minutes that the server retains a mounted sequential-access media volume that is not being used before it dismounts the sequential-access media volume.

mount wait period

The maximum number of minutes that the server waits for a sequential-access volume mount request to be satisfied before canceling the request.

MTU See maximum transmission unit.

N

Nagle algorithm

An algorithm that reduces congestion of TCP/IP networks by combining smaller packets and sending them together.

named pipe

A type of interprocess communication that permits message data streams to pass between peer processes, such as between a client and a server.

NAS file server

See network-attached storage file server.

NAS file server node

See NAS node.

NAS node

A client node that is a network-attached

storage (NAS) file server. Data for the NAS node is transferred by a NAS file server that is controlled by the network data management protocol (NDMP). A NAS node is also called a NAS file server node.

native file system

A file system that is locally added to the file server and is not added for space management. The hierarchical storage manager (HSM) client does not provide space management services to the file system.

native format

A format of data that is written to a storage pool directly by the server. See also non-native data format.

NDMP

See Network Data Management Protocol.

NetBIOS (Network Basic Input/Output System)

A standard interface to networks and personal computers that is used on local area networks to provide message, print-server, and file-server functions. Application programs that use NetBIOS do not have to handle the details of LAN data link control (DLC) protocols.

network-attached storage file server (NAS file server)

A dedicated storage device with an operating system that is optimized for file-serving functions. A NAS file server can have the characteristics of both a node and a data mover.

Network Basic Input/Output System

See NetBIOS.

Network Data Management Protocol (NDMP)

A protocol that allows a network storage-management application to control the backup and recovery of an NDMP-compliant file server, without installing vendor-acquired software on that file server.

network data-transfer rate

A rate that is calculated by dividing the total number of bytes that are transferred by the data transfer time. For example, this rate can be the time that is spent transferring data over a network.

node A file server or workstation on which the

backup-archive client program has been installed, and which has been registered to the server.

node name

A unique name that is used to identify a workstation, file server, or PC to the server.

node privilege class

A privilege class that gives an administrator the authority to remotely access backup-archive clients for a specific client node or for all clients in a policy domain. See also privilege class.

non-native data format

A format of data that is written to a storage pool that differs from the format that the server uses for operations. See also native format.

O

offline volume backup

A backup in which the volume is locked so that no other system applications can access it during the backup operation.

online volume backup

A backup in which the volume is available to other system applications during the backup operation.

open registration

A registration process in which users can register their workstations as client nodes with the server. See also closed registration.

operator privilege class

A privilege class that gives an administrator the authority to disable or halt the server, enable the server, cancel server processes, and manage removable media. See also privilege class.

options file

A file that contains processing options. See also client system-options file, client user-options file.

originating file system

The file system from which a file was migrated. When a file is recalled, it is returned to its originating file system.

orphaned stub file

A file for which no migrated file can be found on the server that the client node is

contacting for space management services. For example, a stub file can be orphaned when the client system-options file is modified to contact a server that is different than the one to which the file was migrated.

P

packet In data communication, a sequence of binary digits, including data and control signals, that are transmitted and switched as a composite whole.

page A defined unit of space on a storage medium or within a database volume.

partial-file recall mode

A recall mode that causes the hierarchical storage management (HSM) function to read just a portion of a migrated file from storage, as requested by the application accessing the file.

password generation

A process that creates and stores a new password in an encrypted password file when the old password expires. Automatic generation of a password prevents password prompting.

path An object that defines a one-to-one relationship between a source and a destination. Using the path, the source accesses the destination. Data can flow from the source to the destination, and back. An example of a source is a data mover (such as a network-attached storage [NAS] file server), and an example of a destination is a tape drive.

pattern-matching character

See wildcard character.

physical file

A file that is stored in one or more storage pools, consisting of either a single logical file, or a group of logical files that are packaged together as an aggregate. See also aggregate, logical file, physical occupancy.

physical occupancy

The amount of space that is used by physical files in a storage pool. This space includes the unused space that is created when logical files are deleted from aggregates. See also logical file, logical occupancy, physical file.

plug-in

A separately installable software module that adds function to an existing program, application, or interface.

policy domain

A grouping of policy users with one or more policy sets, which manage data or storage resources for the users. The users are client nodes that are associated with the policy domain. See also active policy set, domain.

policy privilege class

A privilege class that gives an administrator the authority to manage policy objects, register client nodes, and schedule client operations for client nodes. Authority can be restricted to certain policy domains. See also privilege class.

policy set

A group of rules in a policy domain. The rules specify how data or storage resources are automatically managed for client nodes in the policy domain. Rules can be contained in management classes. See also active policy set, management class.

premigrated file

A file that has been copied to server storage, but has not been replaced with a stub file on the local file system. An identical copy of the file resides both on the local file system and in server storage. Premigrated files occur on UNIX and Linux file systems to which space management has been added. See also file state, migrated file, resident file.

premigrated files database

A database that contains information about each file that has been premigrated to server storage.

premigration

The process of copying files that are eligible for migration to server storage, but leaving the original file intact on the local file system.

premigration percentage

A space management setting that controls whether the next eligible candidates in a file system are premigrated following threshold or demand migration.

primary storage pool

A named set of volumes that the server uses to store backup versions of files, archive copies of files, and files migrated from client nodes. See also copy storage pool, server storage, storage pool, storage pool volume.

privilege class

A level of authority that is granted to an administrator. The privilege class determines which administrative tasks the administrator can perform. See also authority, node privilege class, operator privilege class, policy privilege class, storage privilege class, system privilege class.

profile

A named group of configuration information that can be distributed from a configuration manager when a managed server subscribes. Configuration information can include registered administrator IDs, policies, client schedules, client option sets, administrative schedules, storage manager command scripts, server definitions, and server group definitions. See also configuration manager, enterprise configuration, managed server.

profile association

On a configuration manager, the defined relationship between a profile and an object such as a policy domain. Profile associations define the configuration information that is distributed to a managed server when it subscribes to the profile.

Q**quota**

1. For HSM on AIX, UNIX, or Linux systems, the limit (in megabytes) on the amount of data that can be migrated and premigrated from a file system to server storage.
2. For HSM on Windows systems, a user-defined limit to the space that is occupied by recalled files.

R

randomization

The process of distributing schedule start times for different clients within a specified percentage of the schedule's startup window.

raw logical volume

A portion of a physical volume that is comprised of unallocated blocks and has no journaled file system (JFS) definition. A logical volume is read/write accessible only through low-level I/O functions.

rebind

To associate all backed-up versions of a file with a new management class name. For example, a file that has an active backup version is rebound when a later version of the file is backed up with a different management class association. See also bind, management class.

recall To copy a migrated file from server storage back to its originating file system using the hierarchical storage management client. See also selective recall.

receiver

A server repository that contains a log of server and client messages as events. For example, a receiver can be a file exit, a user exit, or the server console and activity log. See also event.

reclamation

The process of consolidating the remaining data from many sequential-access volumes onto fewer, new sequential-access volumes.

reclamation threshold

The percentage of space that a sequential-access media volume must have before the server can reclaim the volume. Space becomes reclaimable when files are expired or are deleted.

reconciliation

The process of ensuring consistency between the original data repository and the larger system where the data is stored for backup. Examples of larger systems where the data is stored for backup are storage servers or other storage systems.

During the reconciliation process, data that is identified as no longer needed is removed.

recovery log

A log of updates that are about to be written to the database. The log can be used to recover from system and media failures. The recovery log consists of the active log (including the log mirror) and archive logs.

register

To define a client node or administrator ID that can access the server.

registry

A repository that contains access and configuration information for users, systems, and software.

remote

For hierarchical storage management products, pertaining to the origin of migrated files that are being moved. See also local.

resident file

On a Windows system, a complete file on a local file system that might also be a migrated file because a migrated copy can exist in server storage. On a UNIX or Linux system, a complete file on a local file system that has not been migrated or premigrated, or that has been recalled from server storage and modified. See also file state.

restore

To copy information from its backup location to the active storage location for use. For example, to copy information from server storage to a client workstation.

retention

The amount of time, in days, that inactive backed-up or archived files are kept in the storage pool before they are deleted. Copy group attributes and default retention grace periods for the domain define retention.

retrieve

To copy archived information from the storage pool to the workstation for use. The retrieve operation does not affect the archive version in the storage pool. See also archive.

root user

A system user who operates without restrictions. A root user has the special rights and privileges needed to perform administrative tasks.

S

SAN See storage area network.

schedule

A database record that describes client operations or administrative commands to be processed. See also administrative command schedule, client schedule.

scheduling mode

The type of scheduling operation for the server and client node that supports two scheduling modes: client-polling and server-prompted.

scratch volume

A labeled volume that is either blank or contains no valid data, that is not defined, and that is available for use. See also volume.

script A series of commands, combined in a file, that carry out a particular function when the file is run. Scripts are interpreted as they are run. See also Tivoli Storage Manager command script.

Secure Sockets Layer (SSL)

A security protocol that provides communication privacy. With SSL, client/server applications can communicate in a way that is designed to prevent eavesdropping, tampering, and message forgery.

selective backup

The process of backing up certain files or directories from a client domain. The files that are backed up are those that are not excluded in the include-exclude list. The files must meet the requirement for serialization in the backup copy group of the management class that is assigned to each file. See also incremental backup.

selective migration

The process of copying user-selected files from a local file system to server storage and replacing the files with stub files on the local file system. See also demand migration, threshold migration.

selective recall

The process of copying user-selected files from server storage to a local file system. See also recall, transparent recall.

serialization

The process of handling files that are modified during backup or archive processing. See also shared dynamic serialization, shared static serialization, static serialization.

server A software program or a computer that provides services to other software programs or other computers. See also client.

server options file

A file that contains settings that control various server operations. These settings affect such things as communications, devices, and performance.

server-prompted scheduling mode

A client/server communication technique where the server contacts the client node when tasks must be done. See also client-polling scheduling mode.

server storage

The primary, copy, and active-data storage pools that are used by the server to store user files such as backup versions, archive copies, and files migrated from hierarchical storage management client nodes (space-managed files). See also active-data pool, copy storage pool, primary storage pool, storage pool volume, volume.

session

A logical or virtual connection between two stations, software programs, or devices on a network that allows the two elements to communicate and exchange data for the duration of the session. See also administrative session.

session resource usage

The amount of wait time, processor time, and space that is used or retrieved during a client session.

shadow copy

A snapshot of a volume. The snapshot can be taken while applications on the system continue to write data to the volumes.

shadow volume

The data stored from a snapshot of a volume. The snapshot can be taken while applications on the system continue to write data to the volumes.

shared dynamic serialization

A value for serialization that specifies that a file must not be backed up or archived if it is being modified during the operation. The backup-archive client retries the backup or archive operation a number of times; if the file is being modified during each attempt, the backup-archive client will back up or archive the file on its last try. See also dynamic serialization, serialization, shared static serialization, static serialization.

shared library

A library device that is used by multiple storage manager servers. See also library.

shared static serialization

A copy-group serialization value that specifies that a file must not be modified during a backup or archive operation. The client attempts to retry the operation a number of times. If the file is in use during each attempt, the file is not backed up or archived. See also dynamic serialization, serialization, shared dynamic serialization, static serialization.

snapshot

An image backup type that consists of a point-in-time view of a volume.

space-managed file

A file that is migrated from a client node by the hierarchical storage management (HSM) client. The HSM client recalls the file to the client node on demand.

space management

See hierarchical storage management.

space monitor daemon

A daemon that checks space usage on all file systems for which space management is active, and automatically starts threshold migration when space usage on a file system equals or exceeds its high threshold.

sparse file

A file that is created with a length greater than the data it contains, leaving empty spaces for the future addition of data.

special file

On AIX, UNIX, or Linux systems, a file that defines devices for the system, or temporary files that are created by processes. There are three basic types of special files: first-in, first-out (FIFO); block; and character.

SSL See Secure Sockets Layer.

stabilized file space

A file space that exists on the server but not on the client.

stanza A group of lines in a file that together have a common function or define a part of the system. Stanzas are usually separated by blank lines or colons, and each stanza has a name.

startup window

A time period during which a schedule must be initiated.

static serialization

A copy-group serialization value that specifies that a file must not be modified during a backup or archive operation. If the file is in use during the first attempt, the backup-archive client cannot back up or archive the file. See also dynamic serialization, serialization, shared dynamic serialization, shared static serialization.

storage agent

A program that enables the backup and restoration of client data directly to and from storage attached to a storage area network (SAN).

storage area network (SAN)

A dedicated storage network tailored to a specific environment, combining servers, systems, storage products, networking products, software, and services.

storage hierarchy

A logical order of primary storage pools, as defined by an administrator. The order is typically based on the speed and capacity of the devices that the storage pools use. The storage hierarchy is defined by identifying the next storage pool in a storage pool definition. See also storage pool.

storage pool

A named set of storage volumes that is the destination that is used to store client

data. See also active-data pool, copy storage pool, primary storage pool, storage hierarchy.

storage pool volume

A volume that has been assigned to a storage pool. See also active-data pool, copy storage pool, primary storage pool, server storage, volume.

storage privilege class

A privilege class that gives an administrator the authority to control how storage resources for the server are allocated and used, such as monitoring the database, the recovery log, and server storage. See also privilege class.

stub A shortcut on the Windows file system that is generated by the hierarchical storage management (HSM) client for a migrated file that allows transparent user access. A stub is the sparse file representation of a migrated file, with a reparse point attached.

stub file

A file that replaces the original file on a local file system when the file is migrated to storage. A stub file contains the information that is necessary to recall a migrated file from server storage. It also contains additional information that can be used to eliminate the need to recall a migrated file. See also migrated file, resident file.

stub file size

The size of a file that replaces the original file on a local file system when the file is migrated to server storage. The size that is specified for stub files determines how much leader data can be stored in the stub file. The default for stub file size is the block size defined for a file system minus 1 byte.

subscription

In a storage environment, the process of identifying the subscribers to which the profiles are distributed. See also enterprise configuration, managed server.

system privilege class

A privilege class that gives an administrator the authority to issue all server commands. See also privilege class.

T

tape library

A set of equipment and facilities that support an installation's tape environment. The tape library can include tape storage racks, mechanisms for automatic tape mounting, a set of tape drives, and a set of related tape volumes mounted on those drives.

tape volume prefix

The high-level-qualifier of the file name or the data set name in the standard tape label.

target node

A client node for which other client nodes (called agent nodes) have been granted proxy authority. The proxy authority allows the agent nodes to perform operations such as backup and restore on behalf of the target node, which owns the data.

TCA See trusted communications agent.

TCP/IP

See Transmission Control Protocol/Internet Protocol.

threshold migration

The process of moving files from a local file system to server storage based on the high and low thresholds that are defined for the file system. See also automatic migration, demand migration, migration job, selective migration.

throughput

In storage management, the total bytes in the workload, excluding overhead, that are backed up or restored, divided by elapsed time.

timeout

A time interval that is allotted for an event to occur or complete before operation is interrupted.

Tivoli Storage Manager command script

A sequence of Tivoli Storage Manager administrative commands that are stored in the database of the Tivoli Storage Manager server. The script can run from any interface to the server. The script can include substitution for command parameters and conditional logic. See also macro file, script.

tombstone object

A small subset of attributes of a deleted object. The tombstone object is retained for a specified period, and at the end of the specified period, the tombstone object is permanently deleted.

Transmission Control Protocol/Internet Protocol (TCP/IP)

An industry-standard, nonproprietary set of communication protocols that provides reliable end-to-end connections between applications over interconnected networks of different types. See also communication method.

transparent recall

The process that is used to automatically recall a migrated file to a workstation or file server when the file is accessed. See also selective recall.

trusted communications agent (TCA)

A program that handles the sign-on password protocol when clients use password generation.

U

UCS-2 A 2-byte (16-bit) encoding scheme based on ISO/IEC specification 10646-1. UCS-2 defines three levels of implementation: Level 1-No combining of encoded elements allowed; Level 2-Combining of encoded elements is allowed only for Thai, Indic, Hebrew, and Arabic; Level 3-Any combination of encoded elements are allowed.

UNC See Universal Naming Convention.

Unicode

A character encoding standard that supports the interchange, processing, and display of text that is written in the common languages around the world, plus many classical and historical texts.

Unicode-enabled file space

Unicode file space names provide support for multilingual workstations without regard for the current locale.

Universally Unique Identifier (UUID)

The 128-bit numeric identifier that is used to ensure that two components do not have the same identifier. See also Globally Unique Identifier.

Universal Naming Convention (UNC)

The server name and network name combined. These names together identify the resource on the domain.

UTF-8 Unicode Transformation Format, 8-bit encoding form, which is designed for ease of use with existing ASCII-based systems. The CCSID value for data in UTF-8 format is 1208. See also UCS-2.

UUID See Universally Unique Identifier.

V**validate**

To check a policy set for conditions that can cause problems if that policy set becomes the active policy set. For example, the validation process checks whether the policy set contains a default management class.

version

A backup copy of a file stored in server storage. The most recent backup copy of a file is the active version. Earlier copies of the same file are inactive versions. The number of versions retained by the server is determined by the copy group attributes in the management class.

virtual file space

A representation of a directory on a network-attached storage (NAS) file system as a path to that directory.

virtual mount point

A directory branch of a file system that is defined as a virtual file system. The virtual file system is backed up to its own file space on the server. The server processes the virtual mount point as a separate file system, but the client operating system does not.

virtual volume

An archive file on a target server that represents a sequential media volume to a source server.

volume

A discrete unit of storage on disk, tape or other data recording medium that supports some form of identifier and parameter list, such as a volume label or input/output control. See also scratch volume, server storage, storage pool, storage pool volume.

volume history file

A file that contains information about volumes that have been used by the server for database backups and for export of administrator, node, policy, or server data. The file also has information about sequential-access storage pool volumes that have been added, reused, or deleted. The information is a copy of volume information that is recorded in the server database.

Volume Shadow Copy Service (VSS)

A set of Microsoft application-programming interfaces (APIs) that are used to create shadow copy backups of volumes, exact copies of files, including all open files, and so on.

VSS See Volume Shadow Copy Service.

VSS Backup

A backup operation that uses Microsoft Volume Shadow Copy Service (VSS) technology. The backup operation produces an online snapshot (point-in-time consistent copy) of Microsoft Exchange data. This copy can be stored on local shadow volumes or on Tivoli Storage Manager server storage.

VSS Fast Restore

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

VSS Instant Restore

An operation that restores data from a local snapshot. The snapshot is the VSS backup that resides on a local shadow volume. The restore operation retrieves the data by using a hardware assisted restore method (for example, a FlashCopy operation).

VSS offloaded backup

A backup operation that uses a Microsoft Volume Shadow Copy Service (VSS) hardware provider (installed on an alternate system) to move IBM Data Protection for Microsoft Exchange data to the Tivoli Storage Manager server. This type of backup operation shifts the backup load from the production system to another system.

VSS Restore

A function that uses a Microsoft Volume Shadow Copy Service (VSS) software provider to restore VSS Backups (IBM Data Protection for Microsoft Exchange database files and log files) that reside on Tivoli Storage Manager server storage to their original location.

W**wildcard character**

A special character such as an asterisk (*) or a question mark (?) that can be used to represent one or more characters. Any character or set of characters can replace the wildcard character.

workload partition (WPAR)

A partition within a single operating system instance.

workstation

A terminal or personal computer at which a user can run applications and that is usually connected to a mainframe or a network.

worldwide name (WWN)

A 64-bit, unsigned name identifier that is unique.

WPAR See workload partition.

WWN See worldwide name.

Index

A

- About this publication v
- accessibility features 57
- archive copy group 22
- automated failover
 - overview 3
- autostart 31

B

- backdelete
 - and Data Protection for Oracle node 18
 - and Tivoli Storage Manager policy 22
- backup copy group values 22
- backups
 - removing 30
- bkdb.log 25
- bkdb.scr
 - and the Tivoli Storage Manager scheduler 31

C

- client deduplication
 - considerations 36
 - setting up 36
- command line syntax
 - characteristics 39
- commands
 - Oracle
 - change 30
 - tdpoconf password 40
 - tdpoconf showenvironment 40
 - tdposync
 - query 47
 - syncdb 44
 - Tivoli Storage Manager server
 - query association 31
 - query node 29
 - query schedule 31
- commmethod
 - description 19
- compression 19
- configure
 - Quick configuration 13
 - with default settings 13
- configuring
 - Data Protection for Oracle 13
- configuring Data Protection for Oracle 14
- control file 42, 44
- conventions
 - typeface vi
- Conventions vi
- customer support
 - contacting 54

D

- data deduplication
 - overview 35

- data deduplication (*continued*)
 - using 35
- data deduplication reduction
 - determining 37
- Data Protection for Oracle
 - and Oracle databases 25
 - configuring 13, 18
 - creating an installation package 11
 - creating an installation package on DVD or file server 11
 - installing 5, 6
 - overview 1, 2
 - protecting data 25
 - recommendations 43
 - Recovery Manager (RMAN) 3
 - reference 39
 - silent installation 7
 - silent installation with batch file 12
 - silent installation with msiexec.exe 10
 - silent installation with setup.exe 9
 - supported Oracle versions 2
 - Tivoli Storage Manager policy requirements 22
 - updates ix
 - version migration 3
- Data Protection for Oracle installing in silent mode
 - installation
 - capturing a log of the installation 10
 - installing
 - capturing a log 10
- Data Protection for Oracle silent installation
 - playing back the installation 11
- Data Protection for Oracle utilities
 - using 39
- deduplication
 - using 35
- defining a schedule
 - on the Tivoli Storage Manager server 31
 - on the workstation with the Oracle Server 31
- disability 57
- dsm.opt
 - description 18
 - recommended options 19
 - required options 18
- dsmi_log 15
- dsmi_orc_config 15
- dsmsched.log 31
- duplex copy
 - considerations 29
 - overview 28

E

- enablelanfree 20
- example
 - tdposync query command 47
- examples
 - duplex copy 28
 - include/exclude 22
 - invoking RMAN 25
 - removing backups 30
 - RMAN script
 - send command 27

examples (*continued*)

- RMAN scripts 27
 - tdpoconf password command 40
 - tdpoconf showenvironment command 40
 - tdposync syncdb command 44
 - pick window 46
 - Tivoli Storage Manager scheduler 30
- expiration of objects 22

F

- failover
 - Data Protection for Oracle 3
- fixes, obtaining 54

G

- glossary 63

I

- IBM Support Assistant 53
- includexl
 - and Tivoli Storage Manager policy 22
- include
 - and duplex copy 29
 - and Tivoli Storage Manager policy 22
 - description 20
- installation
 - instructions
 - node name registration 17
 - prerequisites 5
- installing
 - creating an installation package 11
 - creating an installation package on DVD or file server 11
 - Data Protection for Oracle 5
 - silently 7
 - silently with batch file 12
 - silently with msixec.exe 10
 - silently with setup.exe 9
- installing Data Protection for Oracle 6
- Internet, searching for problem resolution 53, 54

K

- keyboard 57
- knowledge bases, searching 53

L

- LAN-free data transfer
 - description 3
 - options 20

M

- management class
 - for automatic expiration 22
- maxnummp 29
- migration considerations 3
- Minimum software requirements 6

N

- New in this version ix
- nocatalog
 - and tdposync syncdb command 44
- node name
 - registration 17
- nodename 18
- numcatalogs
 - and tdposync syncdb command 44
- numorcinstances
 - and tdposync syncdb command 44

O

- operating system requirements 6
- options 15
- Oracle RMAN send command
 - using 26
- orcsched.cmd file 31
- outfile
 - and tdpoconf showenvironment command 41
- overview
 - data deduplication 35
 - Data Protection for Oracle 1

P

- Passport Advantage 55
- passwordaccess 19
- performance monitor on Tivoli Storage Manager server 36
- pick window 46
- policy domain 22
- prerequisites 5
- problem determination
 - describing problem for IBM Software Support 55
 - determining business impact for IBM Software Support 55
 - submitting a problem to IBM Software 56
- protecting data
 - Data Protection for Oracle 25
- publications
 - download v

Q

- querying backup objects 35

R

- reference
 - Data Protection for Oracle 39
- retonly
 - and Tivoli Storage Manager policy 22
- RMAN
 - description 3
 - invoking 25
 - scripts 26
 - send command 26

S

- schedlogname 31
- scripts 26
- send command
 - in an RMAN script 26

- send command (*continued*)
 - sample script 27
 - using 26
- set duplex 29
- setting up client deduplication 36
- software support
 - describing problem for IBM Software Support 55
 - determining business impact for IBM Software Support 55
 - submitting a problem 56
- Software Support
 - contacting 54
- support contract 55
- support information 51
- support subscription 55
- syntax diagrams vi

T

- tcpclientaddress 31
- tcpclientport 31
- tcpserveraddress 19
- tdpo_date_fmt 16
- tdpo_fs 16
- tdpo_language 16
- tdpo_mgmt_class_2 17
- tdpo_mgmt_class_3 17
- tdpo_mgmt_class_4 17
- tdpo_num_fmt 16
- TDPO_OPTFILE
 - and tdpoconf password command 40
 - and tdpoconf showenvironment command 41
 - and tdposync syncdb command 44, 47
 - example 15
- tdpo_time_fmt 16
- tdpo.opt
 - and version migration 3
 - description 15
- tdpoconf 39
 - and password initialization 23
 - and tdpo.opt 15
 - description 40
 - password command 40
 - example 40
 - syntax diagram 40
 - TDPO_OPTFILE 40
 - showenvironment command 40
 - example 40
 - outfile 41
 - syntax diagram 40
 - TDPO_OPTFILE 41
 - tdpoconf utility 40
- tdpoerror.log
 - how to specify 15
- tdposync 39
 - and tdpo.opt 15
 - considerations 43
 - description 42
 - query command 47
 - example 47
 - syntax diagram 47
 - TDPO_OPTFILE 47
 - syncdb command 44
 - example 44
 - nocatalog 44
 - numcatalogs 44
 - numorcinstances 44
 - pick window 46

- tdposync (*continued*)
 - syncdb command (*continued*)
 - syntax diagram 44
 - TDPO_OPTFILE 44
- tdposync syncdb command
 - pick window 46
- Tivoli Storage Manager
 - functions 1
 - management class 22
 - policy requirements 22
 - services 1
- typeface conventions vi

U

- using data deduplication 35

V

- verdeleted
 - and Tivoli Storage Manager policy 22
- Virtualization support 6

W

- Windows x86
 - hardware 5
- Windows x86_64
 - hardware 5



Product Number: 5608-E04

Printed in USA