

SYBASE®

Using Backup Server with
IBM® Tivoli® Storage Manager

Adaptive Server® Enterprise

15.5

DOCUMENT ID: DC01176-01-1550-01

LAST REVISED: November 2009

Copyright © 2009 by Sybase, Inc. All rights reserved.

This publication pertains to Sybase software and to any subsequent release until otherwise indicated in new editions or technical notes. Information in this document is subject to change without notice. The software described herein is furnished under a license agreement, and it may be used or copied only in accordance with the terms of that agreement.

To order additional documents, U.S. and Canadian customers should call Customer Fulfillment at (800) 685-8225, fax (617) 229-9845.

Customers in other countries with a U.S. license agreement may contact Customer Fulfillment via the above fax number. All other international customers should contact their Sybase subsidiary or local distributor. Upgrades are provided only at regularly scheduled software release dates. No part of this publication may be reproduced, transmitted, or translated in any form or by any means, electronic, mechanical, manual, optical, or otherwise, without the prior written permission of Sybase, Inc.

Sybase trademarks can be viewed at the Sybase trademarks page at <http://www.sybase.com/detail?id=1011207>. Sybase and the marks listed are trademarks of Sybase, Inc. ® indicates registration in the United States of America.

Java and all Java-based marks are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

Unicode and the Unicode Logo are registered trademarks of Unicode, Inc.

IBM and Tivoli are registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

All other company and product names mentioned may be trademarks of the respective companies with which they are associated.

Use, duplication, or disclosure by the government is subject to the restrictions set forth in subparagraph (c)(1)(ii) of DFARS 52.227-7013 for the DOD and as set forth in FAR 52.227-19(a)-(d) for civilian agencies.

Sybase, Inc., One Sybase Drive, Dublin, CA 94568.

Contents

About This Book	v
CHAPTER 1 Creating Backups Using the IBM Tivoli Storage Manager	1
Installing and setting up the backup system	2
Installing the backup system	3
TSM concepts and Backup Server.....	6
Logical structures on TSM.....	6
Object naming and data organization.....	6
Backing up databases and transactions	7
Using dump and load with the same and different databases...	8
Using dump and load when the source and target Adaptive Servers are different	9
Using dump and load with multiple stripes	10
Listing a server's backup objects	13
Deleting backup objects from TSM	13
Index	15



About This Book

This manual describes how to set up and administer Backup Server for integration with IBM® Tivoli® Storage Manager.

Audience

This manual is for Sybase® system administrators and database owners.

How to use this book

Use this manual in conjunction with the *System Administration Guide* for information about how to set up and administer Backup Server, and with the relevant IBM Tivoli Storage Manager documentation.

Related documents

The Adaptive Server® Enterprise documentation set consists of:

- The release bulletin for your platform – contains last-minute information that was too late to be included in the books.

A more recent version of the release bulletin may be available. To check for critical product or document information that was added after the release of the product CD, use the Sybase® Product Manuals Web site.

- The installation guide for your platform – describes installation, upgrading, and some configuration procedures for all Adaptive Server and related Sybase products.
- *New Feature Summary* – describes the new features in Adaptive Server, the system changes added to support those features, and changes that may affect your existing applications.
- *Active Messaging Users Guide* – describes how to use the Active Messaging feature to capture transactions (data changes) in an Adaptive Server Enterprise database, and deliver them as events to external applications in real time.
- *Component Integration Services Users Guide* – explains how to use Component Integration Services to connect remote Sybase and non-Sybase databases.
- The *Configuration Guide* for your platform – provides instructions for performing specific configuration tasks.
- *Glossary* – defines technical terms used in the Adaptive Server documentation.

-
- *Historical Server Users Guide* – describes how to use Historical Server to obtain performance information from Adaptive Server.
 - *Java in Adaptive Server Enterprise* – describes how to install and use Java classes as datatypes, functions, and stored procedures in the Adaptive Server database.
 - *Job Scheduler Users Guide* – provides instructions on how to install and configure, and create and schedule jobs on a local or remote Adaptive Server using the command line or a graphical user interface (GUI).
 - *Migration Technology Guide* – describes strategies and tools for migrating to a different version of Adaptive Server.
 - *Monitor Client Library Programmers Guide* – describes how to write Monitor Client Library applications that access Adaptive Server performance data.
 - *Monitor Server Users Guide* – describes how to use Monitor Server to obtain performance statistics from Adaptive Server.
 - *Monitoring Tables Diagram* – illustrates monitor tables and their entity relationships in a poster format. Full-size available only in print version; a compact version is available in PDF format.
 - *Performance and Tuning Series* – is a series of books that explain how to tune Adaptive Server for maximum performance:
 - *Basics* – contains the basics for understanding and investigating performance questions in Adaptive Server.
 - *Improving Performance with Statistical Analysis* – describes how Adaptive Server stores and displays statistics, and how to use the `set statistics` command to analyze server statistics.
 - *Locking and Concurrency Control* – describes how to use locking schemes to improve performance, and how to select indexes to minimize concurrency.
 - *Monitoring Adaptive Server with `sp_sysmon`* – discusses how to use `sp_sysmon` to monitor performance.
 - *Monitoring Tables* – describes how to query Adaptive Server monitoring tables for statistical and diagnostic information.
 - *Physical Database Tuning* – describes how to manage physical data placement, space allocated for data, and the temporary databases.

- *Query Processing and Abstract Plans* – explains how the optimizer processes queries, and how to use abstract plans to change some of the optimizer plans.
- *Quick Reference Guide* – provides a comprehensive listing of the names and syntax for commands, functions, system procedures, extended system procedures, datatypes, and utilities in a pocket-sized book (regular size when viewed in PDF format).
- *Reference Manual* – is a series of books that contains detailed Transact-SQL[®] information:
 - *Building Blocks* – discusses datatypes, functions, global variables, expressions, identifiers and wildcards, and reserved words.
 - *Commands* – documents commands.
 - *Procedures* – describes system procedures, catalog stored procedures, system extended stored procedures, and dbcc stored procedures.
 - *Tables* – discusses system tables, monitor tables, and dbcc tables.
- *System Administration Guide* –
 - *Volume 1* – provides an introduction to the basics of system administration, including a description of configuration parameters, resource issues, character sets, sort orders, and instructions for diagnosing system problems. The second part of *Volume 1* is an in-depth discussion about security administration.
 - *Volume 2* – includes instructions and guidelines for managing physical resources, mirroring devices, configuring memory and data caches, managing multiprocessor servers and user databases, mounting and unmounting databases, creating and using segments, using the reorg command, and checking database consistency. The second half of *Volume 2* describes how to back up and restore system and user databases.
- *System Tables Diagram* – illustrates system tables and their entity relationships in a poster format. Full-size available only in print version; a compact version is available in PDF format.
- *Transact-SQL Users Guide* – documents Transact-SQL, the Sybase-enhanced version of the relational database language. This guide serves as a textbook for beginning users of the database management system, and also contains detailed descriptions of the pubs2 and pubs3 sample databases.

-
- *Troubleshooting Series* –
 - *Troubleshooting: Error Messages Advanced Resolutions* – contains troubleshooting procedures for problems you may encounter. The problems discussed here are the ones the Sybase Technical Support staff hear about most often.
 - *Troubleshooting and Error Messages Guide* – contains detailed instructions on how to resolve the most frequently occurring Adaptive Server error messages.
 - *Encrypted Columns Users Guide* – describes how to configure and use encrypted columns with Adaptive Server.
 - *In-Memory Database Users Guide* – describes how to configure and use in-memory databases.
 - *Using Adaptive Server Distributed Transaction Management Features* – explains how to configure, use, and troubleshoot Adaptive Server DTM features in distributed transaction processing environments.
 - *Using Sybase Failover in a High Availability System* – provides instructions for using Sybase Failover to configure an Adaptive Server as a companion server in a high availability system.
 - *Unified Agent and Agent Management Console* – describes the Unified Agent, which provides runtime services to manage, monitor, and control distributed Sybase resources.
 - *Utility Guide* – documents the Adaptive Server utility programs, such as `isql` and `bcp`, which are executed at the operating system level.
 - *Web Services Users Guide* – explains how to configure, use, and troubleshoot Web services for Adaptive Server.
 - *XA Interface Integration Guide for CICS, Encina, and TUXEDO* – provides instructions for using the Sybase DTM XA interface with X/Open XA transaction managers.
 - *XML Services in Adaptive Server Enterprise* – describes the Sybase native XML processor and the Sybase Java-based XML support, introduces XML in the database, and documents the query and mapping functions that are available in XML services.

Other sources of information

Use the Sybase Getting Started CD, the SyBooks™ CD, and the Sybase Product Manuals Web site to learn more about your product:

- The Getting Started CD contains release bulletins and installation guides in PDF format, and may also contain other documents or updated information not included on the SyBooks CD. It is included with your software. To read or print documents on the Getting Started CD, you need Adobe Acrobat Reader, which you can download at no charge from the Adobe Web site using a link provided on the CD.
- The SyBooks CD contains product manuals and is included with your software. The Eclipse-based SyBooks browser allows you to access the manuals in an easy-to-use, HTML-based format.

Some documentation may be provided in PDF format, which you can access through the PDF directory on the SyBooks CD. To read or print the PDF files, you need Adobe Acrobat Reader.

Refer to the *SyBooks Installation Guide* on the Getting Started CD, or the *README.txt* file on the SyBooks CD for instructions on installing and starting SyBooks.

- The Sybase Product Manuals Web site is an online version of the SyBooks CD that you can access using a standard Web browser. In addition to product manuals, you will find links to EBFs/Maintenance, Technical Documents, Case Management, Solved Cases, newsgroups, and the Sybase Developer Network.

To access the Sybase Product Manuals Web site, go to Product Manuals at <http://www.sybase.com/support/manuals/>.

Sybase certifications on the Web

Technical documentation at the Sybase Web site is updated frequently.

❖ Finding the latest information on product certifications

- 1 Point your Web browser to Technical Documents at <http://www.sybase.com/support/techdocs/>.
- 2 Click Partner Certification Report.
- 3 In the Partner Certification Report filter select a product, platform, and timeframe and then click Go.
- 4 Click a Partner Certification Report title to display the report.

❖ Finding the latest information on component certifications

- 1 Point your Web browser to Availability and Certification Reports at <http://certification.sybase.com/>.
- 2 Either select the product family and product under Search by Base Product; or select the platform and product under Search by Platform.

-
- 3 Select Search to display the availability and certification report for the selection.

❖ **Creating a personalized view of the Sybase Web site (including support pages)**

Set up a MySybase profile. MySybase is a free service that allows you to create a personalized view of Sybase Web pages.

- 1 Point your Web browser to Technical Documents at <http://www.sybase.com/support/techdocs/>.
- 2 Click MySybase and create a MySybase profile.

Sybase EBFs and software maintenance

❖ **Finding the latest information on EBFs and software maintenance**

- 1 Point your Web browser to the Sybase Support Page at <http://www.sybase.com/support>.
- 2 Select EBFs/Maintenance. If prompted, enter your MySybase user name and password.
- 3 Select a product.
- 4 Specify a time frame and click Go. A list of EBF/Maintenance releases is displayed.

Padlock icons indicate that you do not have download authorization for certain EBF/Maintenance releases because you are not registered as a Technical Support Contact. If you have not registered, but have valid information provided by your Sybase representative or through your support contract, click Edit Roles to add the “Technical Support Contact” role to your MySybase profile.

- 5 Click the Info icon to display the EBF/Maintenance report, or click the product description to download the software.

Conventions

The following sections describe conventions used in this manual.

SQL is a free-form language. There are no rules about the number of words you can put on a line or where you must break a line. However, for readability, all examples and most syntax statements in this manual are formatted so that each clause of a statement begins on a new line. Clauses that have more than one part extend to additional lines, which are indented. Complex commands are formatted using modified Backus Naur Form (BNF) notation.

Table 1 shows the conventions for syntax statements that appear in this manual:

Table 1: Font and syntax conventions for this manual

Element	Example
Command names, procedure names, utility names, and other keywords display in sans serif font.	<code>select</code> <code>sp_configure</code>
Database names and datatypes are in sans serif font.	<code>master database</code>
Book names, file names, variables, and path names are in italics.	<i>System Administration Guide</i> <i>sql.ini</i> file <i>column_name</i> \$SYBASE/ASE directory
Variables—or words that stand for values that you fill in—when they are part of a query or statement, are in italics in Courier font.	<code>select column_name</code> <code>from table_name</code> <code>where search_conditions</code>
Type parentheses as part of the command.	<code>compute row_aggregate (column_name)</code>
Double colon, equals sign indicates that the syntax is written in BNF notation. Do not type this symbol. Indicates “is defined as”.	<code>::=</code>
Curly braces mean that you must choose at least one of the enclosed options. Do not type the braces.	<code>{cash, check, credit}</code>
Brackets mean that to choose one or more of the enclosed options is optional. Do not type the brackets.	<code>[cash check credit]</code>
The comma means you may choose as many of the options shown as you want. Separate your choices with commas as part of the command.	<code>cash, check, credit</code>
The pipe or vertical bar () means you may select only one of the options shown.	<code>cash check credit</code>
An ellipsis (...) means that you can <i>repeat</i> the last unit as many times as you like.	<code>buy thing = price [cash check credit]</code> <code>[, thing = price [cash check credit]]...</code> You must buy at least one thing and give its price. You may choose a method of payment: one of the items enclosed in square brackets. You may also choose to buy additional things: as many of them as you like. For each thing you buy, give its name, its price, and (optionally) a method of payment.

- Syntax statements (displaying the syntax and all options for a command) appear as follows:

```
sp_dropdevice [device_name]
```

For a command with more options:

```
select column_name
from table_name
where search_conditions
```

In syntax statements, keywords (commands) are in normal font and identifiers are in lowercase. Italic font shows user-supplied words.

- Examples showing the use of Transact-SQL commands are printed like this:

```
select * from publishers
```

- Examples of output from the computer appear as follows:

pub_id	pub_name	city	state
0736	New Age Books	Boston	MA
0877	Binnet & Hardley	Washington	DC
1389	Algodata Infosystems	Berkeley	CA

(3 rows affected)

In this manual, most of the examples are in lowercase. However, you can disregard case when typing Transact-SQL keywords. For example, `SELECT`, `Select`, and `select` are the same.

Adaptive Server sensitivity to the case of database objects, such as table names, depends on the sort order installed on Adaptive Server. You can change case sensitivity for single-byte character sets by reconfiguring the Adaptive Server sort order. For more information, see the *System Administration Guide*.

Accessibility features

This document is available in an HTML version that is specialized for accessibility. You can navigate the HTML with an adaptive technology such as a screen reader, or view it with a screen enlarger.

Adaptive Server HTML documentation has been tested for compliance with U.S. government Section 508 Accessibility requirements. Documents that comply with Section 508 generally also meet non-U.S. accessibility guidelines, such as the World Wide Web Consortium (W3C) guidelines for Web sites.

Note You might need to configure your accessibility tool for optimal use. Some screen readers pronounce text based on its case; for example, they pronounce ALL UPPERCASE TEXT as initials, and MixedCase Text as words. You might find it helpful to configure your tool to announce syntax conventions. Consult the documentation for your tool.

For information about how Sybase supports accessibility, see Sybase Accessibility at <http://www.sybase.com/accessibility>. The Sybase Accessibility site includes links to information on Section 508 and W3C standards.

If you need help

Each Sybase installation that has purchased a support contract has one or more designated people who are authorized to contact Sybase Technical Support. If you cannot resolve a problem using the manuals or online help, please have the designated person contact Sybase Technical Support or the Sybase subsidiary in your area.



Creating Backups Using the IBM Tivoli Storage Manager

This manual describes how to set up and use the IBM Tivoli Storage Manager to backup Adaptive Server databases. Use this manual in conjunction with the Adaptive Server *System Administration Guide* and the relevant Tivoli Storage Manager documentation.

Adaptive Server support for the Tivoli Storage Manager is a licensed option.

Topic	Page
Installing and setting up the backup system	2
TSM concepts and Backup Server	6
Backing up databases and transactions	7
Listing a server's backup objects	13
Deleting backup objects from TSM	13

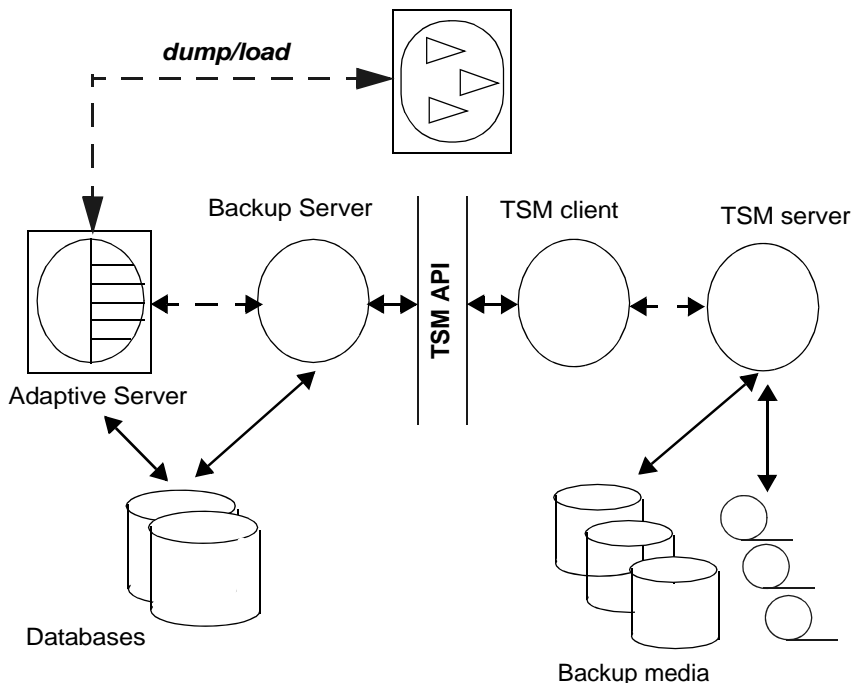
The Tivoli Storage Manager (TSM) is a third-party client/server program that provides storage management services to the licensed user. The Adaptive Server Backup Server supports tape drives and disk files as backup media. TSM works with Backup Server to enable a greater variety of backup media. You can store Adaptive Server backups on any backup media supported by TSM.

Adaptive Server handles the dump and load of databases and transactions to and from TSM; TSM handles storage and retrieval from the storage media.

TSM provides storage management services, but you can administer all your Adaptive Server backup and restore operations, including queries for backup objects and deletion of backup objects, from Backup Server.

Figure 1-1 shows how Adaptive Server works with TSM to back up databases and transactions.

Figure 1-1: Integration of Adaptive Server, Backup Server, and TSM



The Adaptive Server and TSM integrated architecture includes Adaptive Server, Backup Server, the TSM client, the TSM server, and the TSM client API. The TSM client API enables the client application to call TSM functions.

The Sybase interface with the TSM API provides I/O routines that allow Backup Server to communicate with TSM. When dump or load commands are executed, Backup Server dynamically loads this API module, which in turn loads a shared library supplied by the TSM client API component.

Installing and setting up the backup system

This section describes how to install and configure Backup Server and TSM so they can work together to back up databases and transactions.

Install the TSM client, TSM client API, and the TSM server according to instructions in the TSM documentation. Adaptive Server, Backup Server, and the TSM client API product must be installed on the same machine, and they should be of the same bit architecture as that of Backup Server; the TSM server can be installed on a different machine.

Installing the backup system

After installing Adaptive Server and Backup Server:

- 1 Install and configure the TSM client API on the same machine as Adaptive Server and Backup Server. Install and configure the TSM server on the same or different machine. See the TSM documentation for instructions.
- 2 Set these TSM-specific environment variables on the Backup Server machine:
 - `DSMI_DIR` – TSM installation directory location
 - `DSMI_LOG` – TSM error log location
 - `DSMI_CONFIG` – TSM configuration file
 - `DSMI_LIB` – TSM shared library fully qualified path name

For example, assuming that Adaptive Server, Backup Server, and the TSM client API are installed on a Linux AMD 64-bit machine

- `DSMI_DIR` – `/opt/tivoli/tsm/client/api/bin64`
 - `DSMI_LOG` – `/opt/tivoli/tsm/client/error_log`
 - `DSMI_CONFIG` – `/opt/tivoli/tsm/client/api/bin64/dsm.opt`
 - `DSMI_LIB` – `/opt/tivoli/tsm/client/api/bin64/libAPITSM64.so`
- 3 Update the TSM system file (`dsm.sys`) and the TSM configuration file (`dsm.opt`). At a minimum, set these parameters:
 - `SERVERNAME` – TSM server name.
 - `COMMETHOD` – TSM server communication type.
 - `TCPPORT` – TSM server port number.
 - `TCPSEVERADDRESS` – fully qualified name of TSM server.
 - `NODENAME` – registered name of TSM client.

- **PASSWORDACCESS** – access method type of TSM password. Set to “generate”.
- **PASSWORDDIR** – directory location of TSM password.

Note To improve performance, set the appropriate TSM Client configuration parameters. See the TSM documentation for instructions.

- 4 Create the TSM client node at the TSM server. See the TSM documentation for instructions.
- 5 Use the `sybtsmpasswd` utility to:
 - Register the TSM client node with TSM server.
 - Generate the TSM-encrypted password file on the TSM client node.

`sybtsmpasswd` generates the TSM password file (*TSM.PWD*) in the directory specified by the **PASSWORDDIR** parameter in the TSM configuration file. (See step 3.)

Execute `sybtsmpasswd` while logged in as the operating system user “root.” `sybtsmpasswd` prompts for the existing password and the new password for the registered TSM client node.

Note Execute `sybtsmpasswd` only when you set up TSM or change the TSM client node name, user name, or password.

- 6 Start Backup Server.
- 7 From the TSM server, grant the TSM client node name permission to delete backup copies at the TSM server. See the TSM documentation for instructions.

This enables these Sybase operations at the TSM client node:

- `sp_deletesmobj`
- dump database *database_name* to “`syb_tsm::object_name`” with `init`

After completing these steps, you can execute these commands and stored procedures:

- dump database
- dump transaction
- load database

- load transaction
- sp_querysmobj
- sp_deletesmobj

Configuring TSM to allow different source and target machines

TSM does not automatically allow cross-client or cross-owner backup and restore operations.

For example, suppose you want to dump from machine “node1” and load to machine “node2”. The source client NODENAME is “node1”; the target client NODENAME is “node2”.

To configure TSM to allow this:

- 1 On the TSM client node, set the ASNODENAME parameter to the same value for both source and target client node machines. For example:

```
ASNODENAME MyCluster
```

- 2 On the TSM server node:
 - a Add the client node name. For example: “MyCluster”.
 - b Register the common client node name “MyCluster” association with the source “node1” and target “node2” client nodes. See the TSM documentation for instructions on how to grant proxy authority to a client node.
 - c Grant permission to delete backup copies from the “MyCluster” client node. See the TSM documentation for instruction on how to update TSM client node properties.

Configuring TSM data compression

TSM provides a configuration parameter that enables backup data compression. If compression is enabled, the TSM client compresses the backup data before sending it to the TSM server. This decreases the amount of backup data sent to TSM server and the storage space it occupies at the TSM server.

To enable TSM compression for Sybase database backups, set the “COMPRESSION” parameter to “YES” in the TSM client configuration file (*dsm.sys*). The default value is “NO,” indicating no compression.

See the IBM Tivoli Storage Manager documentation.

TSM concepts and Backup Server

This section describes some key TSM concepts.

Logical structures on TSM

TSM abstracts the details of storage devices into a logical construct that requires few API calls. It allows client applications like Backup Server to use TSM API calls.

TSM uses these hierarchical constructs to structure backup data. Typically, you can use the default values supplied by TSM. Default values can be changed only by the TSM administrator at the TSM server; they cannot be changed from Backup Server.

In descending order, the constructs are:

- Policy domain – the base element of the logical structure, the policy domain associates a TSM client (a node) with a policy set.
- Policy set – each policy domain has one active policy set that determines how TSM server handles data from the TSM clients belonging to the policy domain. Each policy set contains a default management class and one or more additional management classes.
- Management class – backup copy groups that contain information about expiration or retention policies for backup objects.
- Copy group – defines the expiration or retention policies, such as event-based retention, for backup objects.

Note If Adaptive Server backups require special expiration policies, or do not use the default copy groups, the TSM administrator can create custom classes or groups on the TSM server. You cannot create these special management classes and copy groups from Backup Server.

Object naming and data organization

The TSM database schema is optimized for efficient storage and retrieval of data objects. Adaptive Server database or transaction data is backed up as named objects on the TSM server. Backup Server assigns the object names.

Data is organized hierarchically on the TSM server at these levels:

- Filespace
- High-level name
- Low-level name

Backup Server prints the dump's backup object names when you successfully execute `dump database` or `dump transaction`. See “Using dump and load with the same and different databases” on page 8.

Backing up databases and transactions

When you execute the `dump` or `load` command, Backup Server invokes the Sybase interface with TSM API, which provides communication with TSM. When you use the `dump` command, you specify an object name that is uniquely associated with the backup object. This object name is the same as the TSM object name, and should later be used to specify the same database or transaction dump when you execute the `load` command.

In general, you can use the same options with the `dump` and `load` commands with TSM as you use with Backup Server when TSM is not configured.

The `dump` and `load` syntax specific to TSM is:

```
dump {database | transaction} database_name to
"syb_tsm::object_name"

load {database | transaction} database_name from
"syb_tsm::[-S source_server_name]
[-D source_database_name::]object_name"
```

See the *Reference Manual: Commands* for complete syntax and usage information.

The keyword “`syb_tsm`” invokes the Sybase interface with the TSM API (`libsyb_tsm`). When loading a database or transaction, specify a server or database name only when the current server or database is not the same as the server or database associated with the backup object.

Using dump and load with the same and different databases

To dump and then load a database and transaction to a different database, when both databases are on the same Adaptive Server, use the `-D` option with the load command.

For example, to dump a database “testdb” to a TSM backup object named “obj1.1”, first dump “obj1.1” to the same database, then load “obj1.1” to a different database,

```
dump database testdb to "syb_tsm::obj1.1"
```

Backup Server prints the following, which provides the backup object identifiers:

```
Backup Server session id is: 5. Use this value when
executing the 'sp_volchanged' system stored procedure
after fulfilling any volume change request from the
Backup Server.
```

```
Backup Server: 4.132.1.1: Attempting to open byte stream
device:'syb_tsm::obj1.1::00'
```

```
Backup Server: 6.28.1.1: Dumpfile name
'testdb091840CA13 ' section number 1 mounted on byte
stream 'syb_tsm::obj1.1::00'
```

```
Backup Server: 4.188.1.1: Database testdb: 854
kilobytes (84%) DUMPED.
```

```
Backup Server: 3.43.1.1: Dump phase number 1 completed.
Backup Server: 3.43.1.1: Dump phase number 2 completed.
```

```
Backup Server: 3.43.1.1: Dump phase number 3 completed.
Backup Server: 4.188.1.1: Database testdb: 870
kilobytes (100%) DUMPED.
```

```
Backup Server: 3.42.1.1: DUMP is complete (database
testdb).
```

Additional message in Backup Server error log indicating corresponding TSM backup object name and ids. Please check the complete Backup Server error log attached as well.

Check the Backup Server error log to see the TSM backup object name and id:

```
Jul  3 14:30:19 2009: A00: Database 'testdb' dumped.
Tivoli Storage Manager backup object name: fs =
/demo_srv1, high = /testdb.DB, low = /obj1.1.0,
Copyids: (0, 36913).
```

Backup Server also prints the backup object identifiers to the error log. In this example, they are:

- Filespace (fs) – demo_srv1.
- High-level name (high) – testdb.DB.
- Low-level name (low) – obj1.1.0.
- TSM copyid (copyids) – 0, 36913. This is a unique identifier assigned by the TSM server. It has two parts: a high copyid and a low copyid.

To load the backup object “obj1.1.0” to the same database, “testdb,” enter:

```
load database testdb from "syb_tsm::obj1.1"
```

Alternatively, to load the backup object “obj1.1” to another database, “anotherdb,” enter:

```
load database anotherdb from  
"syb_tsm::-D testdb::obj1.1"
```

Using dump and load when the source and target Adaptive Servers are different

To dump and then load the database or transaction to a different database on a different Adaptive Server, use both the -S and -D options with the load command. Make sure that you:

- Configure the TSM server and TSM client to enable the load command for different source and target machines. See “Configuring TSM to allow different source and target machines” on page 5.
- Start Backup Server using the same login on both the source and target machines.
- Use the -S and -D options with the load database or load transaction command, where -S specifies the source server and -D specifies the source database.

```
load database database_name from "syb_tsm::  
-S server_name -D database_name::object_name"
```

For example, to load the backup object “obj1.3” to “anotherdb” where the source server is “myserver” and the database is “testdb,” enter:

```
load database anotherdb from  
"syb_tsm::-S myserver -D testdb::obj1.3"
```

Using dump and load with multiple stripes

When multiple stripes are used with dump database or dump transaction, each stripe creates a separate backup object at the TSM.

You can use the same or different backup object names for multiple stripes when TSM handles the backup media. In this example, each stripe has the same name. TSM uses a low-level identifier to distinguish between each backup object, for example, obj.2. See the error log later in this section.

Note If TSM is supported at your site, you can use the same or different names for each stripe. If TSM is not supported at your site, you must use different names for each stripe.

Using the same object name for each stripe

For example, to dump “testdb” using multiple stripes with the same stripe object name “obj”, enter:

```
dump database testdb to "syb_tsm::obj"  
stripe on "syb_tsm::obj"  
stripe on "syb_tsm::obj"  
stripe on "syb_tsm::obj"  
stripe on "syb_tsm::obj"
```

In this example, Backup Server provides uniqueness for TSM backup objects by appending the stripe identifier to the supplied backup object name “obj”. Check the Backup Server error log to verify that the TSM low-level backup object name is unique for each stripe.

Backup Server prints:

```
Backup Server session id is: 5. Use this value when executing the  
'sp_volchanged' system stored procedure after fulfilling any volume change  
request from the Backup Server.  
Backup Server: 4.132.1.1: Attempting to open byte stream device:  
'syb_tsm::obj::00'  
Backup Server: 4.132.1.1: Attempting to open byte stream device:  
'syb_tsm::obj::01'  
Backup Server: 4.132.1.1: Attempting to open byte stream device:  
'syb_tsm::obj::02'  
Backup Server: 4.132.1.1: Attempting to open byte stream device:  
'syb_tsm::obj::03'  
Backup Server: 4.132.1.1: Attempting to open byte stream device:  
'syb_tsm::obj::04'  
Backup Server: 6.28.1.1: Dumpfile name 'testdb092740CDAD ' section number 1  
mounted on byte stream 'syb_tsm::obj::03'  
Backup Server: 6.28.1.1: Dumpfile name 'testdb092740CDAD ' section number 1
```



```

mounted on byte stream 'syb_tsm::obj::00'
Backup Server: 6.28.1.1: Dumpfile name 'testdb092740CDAD ' section number 1
mounted on byte stream 'syb_tsm::obj::01'
Backup Server: 6.28.1.1: Dumpfile name 'testdb092740CDAD ' section number 1
mounted on byte stream 'syb_tsm::obj::02'
Backup Server: 6.28.1.1: Dumpfile name 'testdb092740CDAD ' section number 1
mounted on byte stream 'syb_tsm::obj::04'
Backup Server: 4.188.1.1: Database testdb: 24 kilobytes (87%) DUMPED.
Backup Server: 4.188.1.1: Database testdb: 106 kilobytes (88%) DUMPED.
Backup Server: 4.188.1.1: Database testdb: 646 kilobytes (97%) DUMPED.
Backup Server: 4.188.1.1: Database testdb: 760 kilobytes (99%) DUMPED.
Backup Server: 4.188.1.1: Database testdb: 874 kilobytes (100%) DUMPED.
Backup Server: 3.43.1.1: Dump phase number 1 completed.
Backup Server: 3.43.1.1: Dump phase number 2 completed.
Backup Server: 3.43.1.1: Dump phase number 3 completed.
Backup Server: 4.188.1.1: Database testdb: 898 kilobytes (100%) DUMPED.
Backup Server: 3.42.1.1: DUMP is complete (database testdb).

```

Check the Backup Server error log to see the TSM backup object names and IDs. Check the complete Backup Server error log attached as well. The Backup Server displays this information:

```

Oct 1 14:37:43 2009: A00: Database 'testdb' dumped. Tivoli Storage Manager
backup object name: fs = /demo_srv1, high = /testdb.DB, low = /obj.0,
Copyids: (0, 80898).
Oct 1 14:37:43 2009: A03: Database 'testdb' dumped. Tivoli Storage Manager
backup object name: fs = /demo_srv1, high = /testdb.DB, low = /obj.3,
Copyids: (0, 80897).
Oct 1 14:37:43 2009: A02: Database 'testdb' dumped. Tivoli Storage Manager
backup object name: fs = /demo_srv1, high = /testdb.DB, low = /obj.2,
Copyids: (0, 80899).
Oct 1 14:37:43 2009: A01: Database 'testdb' dumped. Tivoli Storage Manager
backup object name: fs = /demo_srv1, high = /testdb.DB, low = /obj.1,
Copyids: (0, 80900).

```

To load “obj” to the same database, enter:

```

load database testdb from "syb_tsm::obj"
stripe on "syb_tsm::obj"
stripe on "syb_tsm::obj"
stripe on "syb_tsm::obj"
stripe on "syb_tsm::obj"

```

Using different object names for each stripe

To dump “testdb” using multiple stripes with different strip object names, enter:

```

dump database testdb to "syb_tsm::obj0"
stripe on "syb_tsm_obj1"

```

```
stripe on "syb_tsm_obj2"  
stripe on "syb_tsm_obj3"  
stripe on "syb_tsm_obj4"
```

In this example, each stripe has a unique backup object name; Backup Server appends a stripe identifier to each one.

Backup Server prints:

```
Backup Server session id is: 13. Use this value when executing the  
'sp_volchanged' system stored procedure after fulfilling any volume change  
request from the Backup Server.  
Backup Server: 4.132.1.1: Attempting to open byte stream device:  
'syb_tsm::obj0::00'  
Backup Server: 4.132.1.1: Attempting to open byte stream device:  
'syb_tsm::obj1::01'  
Backup Server: 4.132.1.1: Attempting to open byte stream device:  
'syb_tsm::obj2::02'  
Backup Server: 4.132.1.1: Attempting to open byte stream device:  
'syb_tsm::obj3::03'  
Backup Server: 4.132.1.1: Attempting to open byte stream device:  
'syb_tsm::obj4::04'  
Backup Server: 6.28.1.1: Dumpfile name 'testdb092740D1A6 ' section number 1  
mounted on byte stream 'syb_tsm::obj1::01'  
Backup Server: 6.28.1.1: Dumpfile name 'testdb092740D1A6 ' section number 1  
mounted on byte stream 'syb_tsm::obj0::00'  
Backup Server: 6.28.1.1: Dumpfile name 'testdb092740D1A6 ' section number 1  
mounted on byte stream 'syb_tsm::obj3::03'  
Backup Server: 6.28.1.1: Dumpfile name 'testdb092740D1A6 ' section number 1  
mounted on byte stream 'syb_tsm::obj4::04'  
Backup Server: 6.28.1.1: Dumpfile name 'testdb092740D1A6 ' section number 1  
mounted on byte stream 'syb_tsm::obj2::02'  
Backup Server: 4.188.1.1: Database testdb: 24 kilobytes (87%) DUMPED.  
Backup Server: 4.188.1.1: Database testdb: 106 kilobytes (88%) DUMPED.  
Backup Server: 4.188.1.1: Database testdb: 220 kilobytes (90%) DUMPED.  
Backup Server: 4.188.1.1: Database testdb: 334 kilobytes (92%) DUMPED.  
Backup Server: 4.188.1.1: Database testdb: 874 kilobytes (100%) DUMPED.  
Backup Server: 3.43.1.1: Dump phase number 1 completed.  
Backup Server: 3.43.1.1: Dump phase number 2 completed.  
Backup Server: 3.43.1.1: Dump phase number 3 completed.  
Backup Server: 4.188.1.1: Database testdb: 898 kilobytes (100%) DUMPED.  
Backup Server: 3.42.1.1: DUMP is complete (database testdb).
```

The Backup Server error log displays the TSM backup object name and its IDs.

```
Oct 1 14:54:33 2009: A00: Database 'testdb' dumped. Tivoli Storage Manager  
backup object name: fs = /demo_srv1, high = /testdb.DB, low = /obj0.0,  
Copyids: (0, 80904).
```

```
Oct 1 14:54:33 2009: A03: Database 'testdb' dumped. Tivoli Storage Manager
backup object name: fs = /demo_srv1, high = /testdb.DB, low = /obj3.3,
Copyids: (0, 80903).
Oct 1 14:54:33 2009: A02: Database 'testdb' dumped. Tivoli Storage Manager
backup object name: fs = /demo_srv1, high = /testdb.DB, low = /obj2.2,
Copyids: (0, 80905).
Oct 1 14:54:33 2009: A01: Database 'testdb' dumped. Tivoli Storage Manager
backup object name: fs = /demo_srv1, high = /testdb.DB, low = /obj1.1,
Copyids: (0, 80902).
Oct 1 14:54:33 2009: A04: Database 'testdb' dumped. Tivoli Storage Manager
backup object name: fs = /demo_srv1, high = /testdb.DB, low = /obj4.4,
Copyids: (0, 80906).
Oct 1 14:54:34 2009: Backup Server: 3.42.1.1: DUMP is complete (database
testdb).
```

To load this dump to the same database, enter:

```
load database testdb from "syb_tsm::obj0"
stripe on "syb_tsm::obj1"
stripe on "syb_tsm::obj2"
stripe on "syb_tsm::obj3"
stripe on "syb_tsm::obj4"
```

Listing a server's backup objects

Use the `sp_querysmobj` stored procedure to retrieve a list of a server's backup objects. See the *Reference Manual: Procedures* for complete syntax and usage information.

Deleting backup objects from TSM

Use the `sp_deletesmobj` stored procedure to delete some or all of the current server's backup objects from TSM. See the *Reference Manual: Procedures* for complete syntax and usage information.

Index

Symbols

- ::= (BNF notation)
 - in SQL statements xi
- , (comma)
 - in SQL statements xi
- { } (curly braces)
 - in SQL statements xi
- () (parentheses)
 - in SQL statements xi
- [] (square brackets)
 - in SQL statements xi

B

- backup system with load database syntax for 7
- backup system with load transaction syntax for 7
- backup system with TSM
 - configuring for different source and target machines 5
 - data organization 6
 - deleting backup objects 13
 - dump and load when databases are different 8
 - dump database syntax for 7
 - dump transaction syntax for 7
 - environment variables for 3
 - installing and setting up 2
 - listing backup objects 13
 - logical structures 6
 - object naming 6
 - syb_tsm keyword 7
 - sybtsmpasswd utility 4
 - using multiple stripes 10
 - when source and target servers are different 9
- Backus Naur Form (BNF) notation x, xi
- BNF notation in SQL statements x, xi
- brackets. *See* square brackets []

C

- case sensitivity
 - in SQL xii
- comma (.)
 - in SQL statements xi
- conventions
 - See also* syntax
 - Transact-SQL syntax x
 - used in the Reference Manual x
- curly braces ({}) in SQL statements xi

D

- data compression, configuring 5
- data hierarchy 7, 9
- deleting backup objects 13
- dump database command 7, 8, 10
- dump transaction command 7

L

- libsyb_tsm module 2
- listing backup objects 13
- load database command 7
- load transaction command 7
- logical structure, TSM 6

M

- multiple stripes, backups with 10

O

- object names 7

Index

P

parentheses ()
 in SQL statements xi

S

sp_deletesmobj stored procedure 4, 13
sp_querysmobj stored procedure 4, 13
square brackets []
 in SQL statements xi
syb_tsm keyword 7
sybtspasswd utility 4
symbols
 in SQL statements x, xi
syntax conventions, Transact-SQL x

T

Tivoli Storage Manager (TSM) 1–13
 and Adaptive Server 1
 integration with Adaptive Server 1