IBM Software Group – IBM Tivoli Training

IBM Tivoli Storage Manager 5.3

Backup Sets

e business on demand.

Updated August 11, 2006

Objectives

Upon completion of this module, you will be able to:

- Explain the benefit of backup sets.
- Generate backup sets.
- Query backup sets.
- Restore backup sets.

Agenda

- Overview
- Generate Backup Sets
- Query Backup Sets
- Restore Backup Sets
- Additional Resources

Section 1

Overview

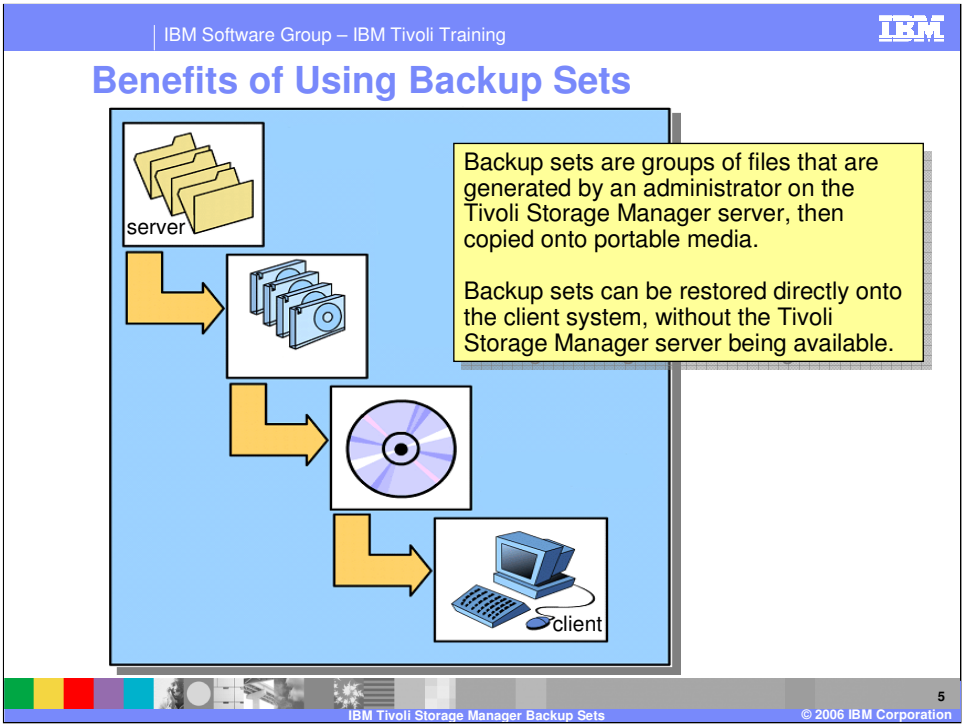

Using the most recent backup versions stored on the server, Storage Manager can

generate a complete set of client files (backup set).

These backup sets can be used to retain a copy client files for a longer period

of time (Instant Archive) or for recovery of a client system, without server or

network, when the backup set is copied onto portable media and restored locally

(Rapid Recovery).

The GENERATE BACKUPSET command runs as a background process on

the server.

The RESTORE BACKUPSET command restores a backup set from the server,

a local file, or a local tape device.

Backup sets can provide you with instant archive and rapid recovery capability .

Instant archive This capability allows an administrator to create an archive

collection from backup versions already stored on the server.

Rapid recovery When you are away from your office without a network

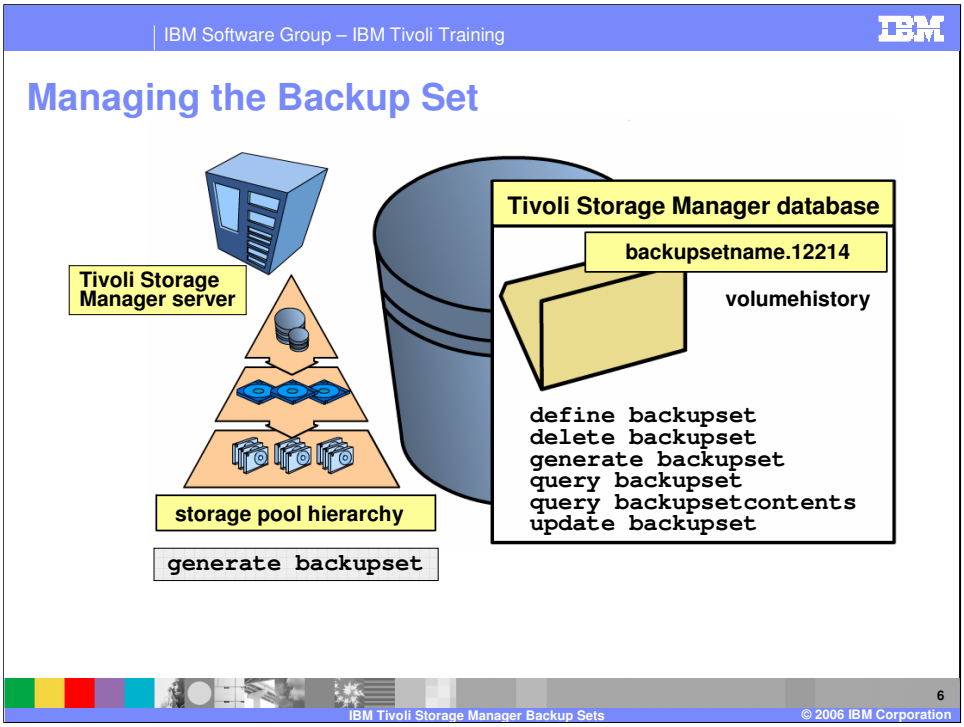

Administrators can generate multiple copies of backup sets that correspond to some

Point in time. The backup sets can be retained for various time periods. This is an

efficient way to create long-term storage of periodic backups, without requiring the

data to be sent over the network again.

While an administrator can generate a backup set from any client's backed up files,

Backup sets can only be used by a backup-archive client. You cannot generate a

backup set for a network attached storage (NAS) node.

The backup set is managed by the administrator in the Tivoli Storage Manager database.

All backup sets have an entry in the volume history file. As long as the entry is in the

volume history file, the volume cannot be reused by the Tivoli Storage Manager server.

On the command to generate the backup set, there is a parameter which defines how

Long the backup set volume will be registered in the volume history file.

The default value

Backup Set Commands

Command	Description
DEFINE BACKUPSET	Defines a previously generated backup set to a server.
DELETE BACKUPSET	Deletes a backup set.
GENERATE BACKUPSET	Generates a backup set of a client's data.
QUERY BACKUPSET	Displays information about a backup set.
QUERY BACKUPSETCONTENTS	Displays contents contained in backup sets.
UPDATE BACKUPSET	Updates a retention value associated with a backup set.

DEFINE BACKUPSET (Define a Backup Set) Use this command to define a client backup set that was previously generated on one server and make it available to the server running this command. The client node has the option of restoring the backup set from the server running this command rather than the one on which the backup set was generated. Any backup set generated on one server can be defined to another server as long as the servers share a common device type. The level of the server to which the backup set is being defined must be equal to or greater than the level of the server that generated the backup set.

You can also use the DEFINE BACKUPSET command to redefine a backup set that was deleted on a server.

GENERATE BACKUPSET (Generate a Backup Set of a Client's Data) Use this command to generate a backup set for a backup-archive client node. A *backup set* is a collection of a client's active backed up data, stored and managed as a single object, on specific media, in server storage. While the server allows you to create a backup set for any client node, a backup set can only be *used* by a backup-archive client.

Note: You cannot generate a backup set for a NAS node. The server creates copies of active versions of a client's backed up objects that are within the one or more file spaces specified with this command, and consolidates them onto sequential media. Currently, the backup object types supported for backup sets include directories and files only. The backup-archive client node can restore its backup set from the

Command Reference

Help

To get a list of commands with command parameters, issue the **`help`** command from the Tivoli Storage Manager command line, followed by the name of the command.

For example:

```
help generate backupset
```

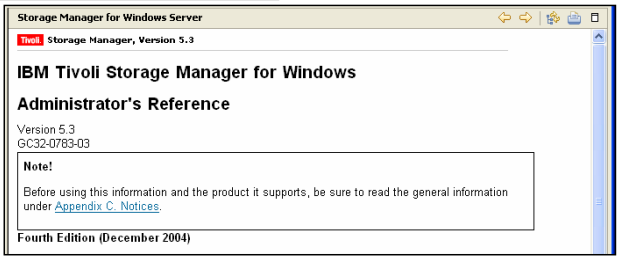

There are several methods to get command syntax assistance. Using the HELP function while working in the application is a fast way to find options and syntax.

The Administrator's Reference Guide, either the .pdf of html version, provides in depth explanations on Tivoli Storage Manager commands.

Section 2

Generate Backup Sets

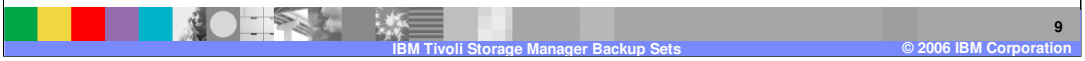

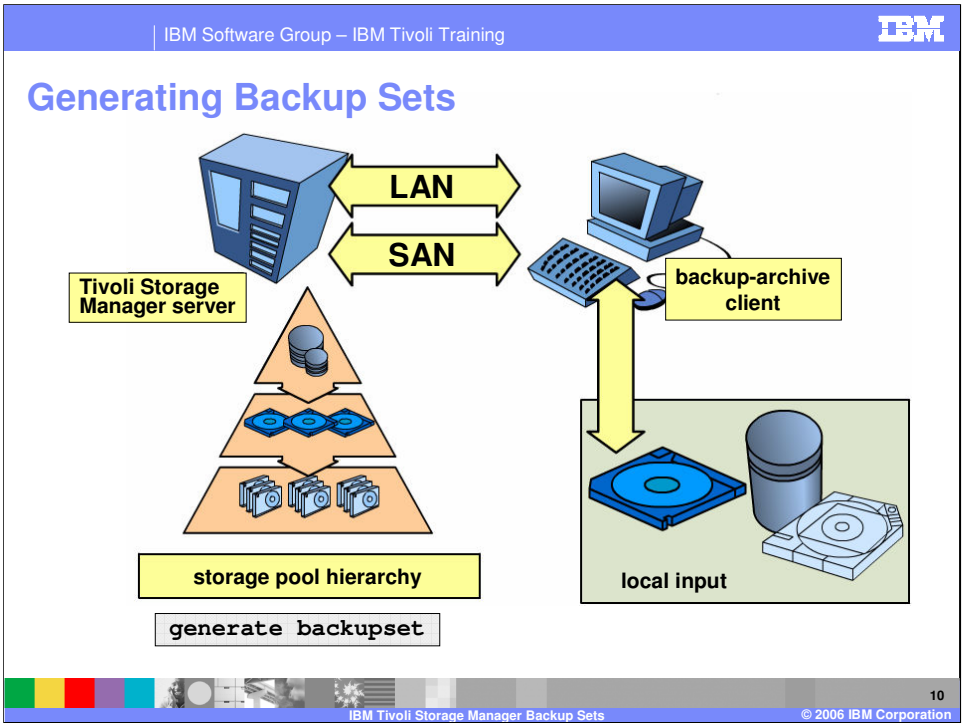

The administrator has to have system or restricted policy over the domain to which

the node is assigned. You can generate backup sets on the server for client nodes.

The client node for which a backup set is generated must be registered to the server.

An incremental backup must be completed for a client node before the server can

generate a backup set for the client node.

Preparing to Generate a Backup Set on CD

1. Define a library whose type is **MANUAL**. Name the library **MANUALLIB**.

```
define library manuallib libtype=manual
```

2. Define a device class whose device type is **FILE**. Name the device class CDFILE.

```
define devclass cdfile devtype=file library=manuallib
```

3. Define a drive to associate with the library. Name the drive **CDDRIVE**.

```
define drive manuallib cddrive
```

4. Define a path to the drive with the device **x:**.

```
define path server1 cddrive srctype=server desttype=drive
   libr=manuallib device=x:
```

5. Define a device class whose device type is **FILE**. Name the device class **CDFILE**:

```
define devclass cdfile devtype=file maxcapacity=640M
   dir=/backupset
```

Considerations when Generating Backup Sets

- The Tivoli Storage Manager server adds a unique suffix to the name of your backup set, so backup sets with the same name are not overwritten.
- Backup sets are tracked as single objects.
- Backup sets have retention values to determine how long the Tivoli Storage Manager server retains the backup set in its volume history.
  - When the backup set expires, the volumes containing the backup sets return to scratch status.
  - Unless a volume is reused, the backup set will still exist, such as copies on local files or CDs.
  - If a retention value is not specified, backup sets are retained for 365 days.
- Backup sets can be redefined to the same server or a different server.
  - Redefined backup sets will be assigned a new retention period.

Backup Set Media Support

Considerations when selecting a device class for writing the backup set:

- Generate the backup set on any sequential access devices whose device types are supported on **both** the client and server machines.
- Ensure that the media type and recording format used for generating the backup set is supported by the device that will be reading the backup set.
- You can write backup sets to sequential media, such as sequential tape and device class **FILE**.
- The tape volumes containing the backup sets are not associated with storage pools and, therefore, are not migrated through the storage pool hierarchy.
- For device class **FILE**, the server creates each backup set with a file extension of *.OST*.
- You can copy **FILE** device class volumes to removable media that is associated with CD-ROM, JAZ, or ZIP devices, by using the **REMOVABLEFILE** device type.

Using Scratch Volumes

Overview of scratch volumes

- Scratch volumes are empty or contain no valid data.
- Their status is changed to PRIVATE when data is written to them.
- They can be used to satisfy any request to mount a scratch volume.

Using scratch volumes for backup sets:

- You can determine whether to use scratch volumes when you generate a backup set.
- If you do not specify volumes, the server uses scratch volumes for the backup set.
- You can use specific volumes for the backup set. If there is not enough space to store the backup set on the volumes, the server uses scratch volumes to store the remainder of the backup set.

PRIVATE VOLUMES:

May contain valid data.

Are used or owned by an application.

Can only be used to satisfy a request to mount the specified volume.

Generate Backupset Command

GENERATE BACKUPSET (Generate a Backup Set of a Client's Data)

Use this command to generate a backup set for a backup-archive client node. A *backup set* is a collection of a client's active backed up data, stored and managed as a single object, on specific media, in server storage. While the server allows you to create a backup set for any client node, a backup set can only be *used* by a backup-archive client.

Note: You cannot generate a backup set for a NAS node.

The server creates copies of active versions of a client's backed up objects that are within the one or more file spaces specified with this command, and consolidates them onto sequential media. Currently, the backup object types supported for backup sets include directories and files only.

The backup-archive client node can restore its backup set from the server as well as from the media to which the backup set was written.

This command generates a background process that can be cancelled with the CANCEL PROCESS command. If the background process created by this command is cancelled, the media may not contain a complete backup set. You can use the QUERY PROCESS command to display information about the background process that is created by this command.

Privilege Class

To issue this command, you must have system privilege or policy privilege for the domain to which the client node is assigned.

Excerpt from *IBM Tivoli Storage Manager Version 5.3 for Windows Administrator's Reference*

Generate Backupset Command Syntax

Syntax

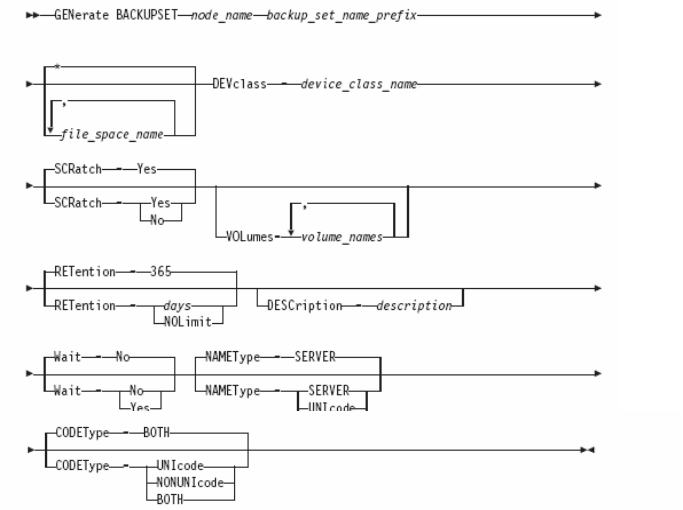

Excerpt from *IBM Tivoli Storage Manager Version 5.3 for Windows Administrator's Reference*

Parameters

***node_name* (Required)** Specifies the name of the client node whose data is contained in the backup set. You cannot use wildcard characters to specify a name nor can you specify a list of client node names.

***backup_set_name_prefix* (Required)** Specifies the name of the backup set for the client node. The maximum length of the name is 30 characters. When you select a name, Tivoli Storage Manager adds a suffix to construct your backup set name. For example, if you name your backup set *mybackupset*, Tivoli Storage Manager adds a unique number such as 3099 to the name. The backup set name is then identified to Tivoli Storage Manager as *mybackupset.3099*. To later display information about this backup set, you can include a wildcard with the name, such as *mybackupset.** or you can specify the fully qualified name, such as *mybackupset.3099*.

file_space_name Specifies the names of one or more file spaces that contain the data to be included in the backup set. This parameter is optional. The file space name you specify can contain wildcard characters.

You can specify more than one file space by separating the names with commas and no intervening spaces. If you do not specify a file space, data from all the client node's backed up and active file spaces is included in the backup set.

DEVclass (Required) Specifies the name of the device class for the volumes to which the backup set is written. The maximum length of the name is 30 characters.

Example of `generate backupset` Command

Generate a backup set of a file space called /test that belongs to client node XYZ. Name the backup set LAB_DATA and retain it for 90 days. Specify that volumes LABVOL4 and LABVOL5 contain the data for the backup set. The volumes are to be read by a device that is assigned to the FILE device class. Include a description.

```
generate backupset xyz lab_data /test devclass=file
  retention=90 volumes=labvol4,labvol5
  description="test lab results"
```

Using the Administration Center to Generate Backup Sets

1. Click **Policy Domains and Client Nodes** in the Tivoli Storage Manager tree.
2. In the **Servers** table, click the name of a server.
3. In the server's **Policy Domains** table, click the name of a domain.
4. In the domain's properties portlet, click **Client Nodes**.
5. In the **Client Nodes** table, click the name of a client node.
6. In the node's properties notebook, click the **Backup Sets** tab.
7. In the **Backup Sets** table, click **Select Action**, select **Generate Backup Set**.

ISC

Section 3

Query Backup Sets

Query Backupset Command Syntax

QUERY BACKUPSET (Query a Backup Set)

Use this command to display information about one or more backup sets.

Privilege Class

Any administrator can issue this command.

Syntax

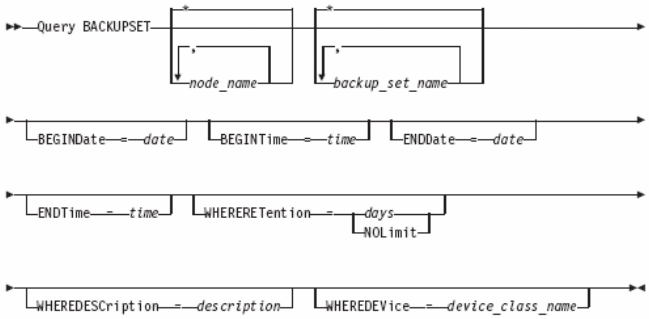

Excerpt from *IBM Tivoli Storage Manager Version 5.3 for Windows Administrator's Reference*

Parameters

node_name Specifies the name of the client node whose data is contained in the backup set to be displayed. The node name you specify can contain wildcard characters. You can specify more than one node name by separating the names with commas and no intervening spaces.

backup_set_name Specifies the name of the backup set whose information is to be displayed. The backup set name you specify can contain wildcard characters. You can specify more than one backup set name by separating the names with commas and no intervening spaces.

BEGINDate Specifies the beginning date of the range in which the backup set to be displayed must have been created. This parameter is optional. You can use this parameter with the BEGINTIME parameter to specify a range for the date and time. If you specify a begin date without a begin time, the time will be at 12:00 a.m. (midnight) on the date you specify.

BEGINTime Specifies the beginning time of the range in which the backup set to be displayed must have been created. This parameter is optional. You can use this parameter with the BEGINDATE parameter to specify a range for the date and time. If you specify a begin time without a begin date, the date will be the current date at the time you specify.

Query Backupset Contents Command Syntax

QUERY BACKUPSETCONTENTS (Query Contents of a Backup Set)

Use this command to display information about the files and directories contained in a client node's backup set.

Be aware that processing this command can use considerable network resources and mount points.

Privilege Class

To issue this command, you must have system privilege or policy privilege for the domain to which the client node is assigned.

Syntax

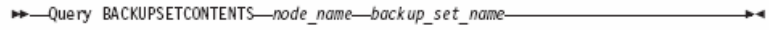

Excerpt from *IBM Tivoli Storage Manager Version 5.3 for Windows Administrator's Reference*

***node_name* (Required)** Specifies the name of the client node whose data is contained in the backup set to display. The name you specify cannot contain wildcard characters nor can it be a list of node names separated by commas.

***backup_set_name* (Required)** Specifies the name of the backup set to display. The name you specify cannot contain wildcard characters nor can it be a list of node names separated by commas.

Section 4

Restore Backup Sets

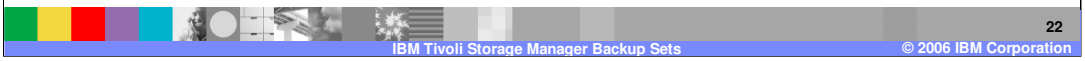

Restore Backupset Command Syntax

Syntax

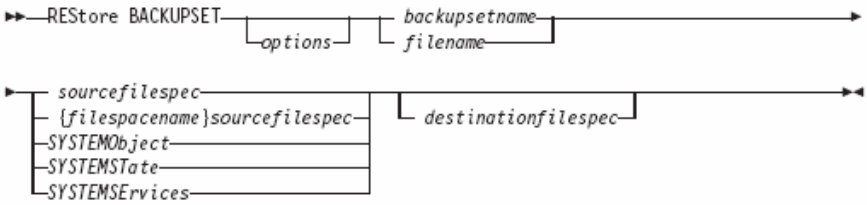

Excerpt from *IBM Tivoli Storage Manager for Windows: Backup-Archive Clients Installation and User's Guide*

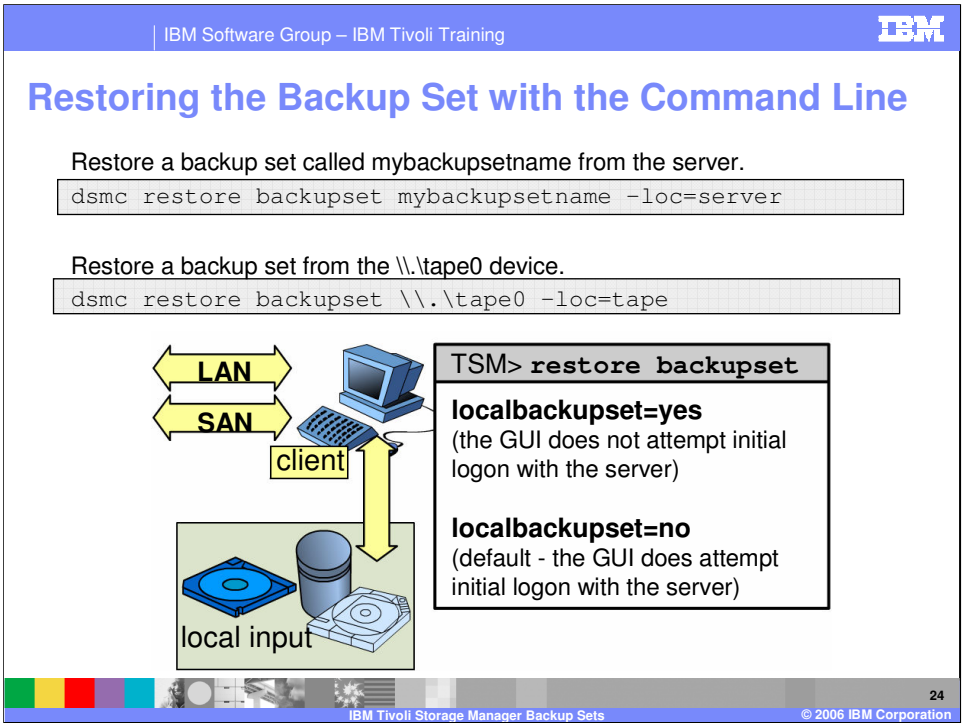

The **LOCALBACKUPSET** option specifies whether the Tivoli Storage Manager GUI

bypasses initial logon with the Tivoli Storage Manager server to restore a local backup set

on a stand-alone workstation. You can use this option on the command line or place it in

your client options file (**dsm.opt**).

If you set the **LOCALBACKUPSET** option to **YES**, the GUI does not attempt initial logon

with the server. In this case, the GUI only enables the restore functionality.

If you set the **LOCALBACKUPSET** option to **NO** (the default), the GUI attempts initial

logon with the server and enables all GUI functions.

Restoring the Backup Set with the Backup-Archive GUI

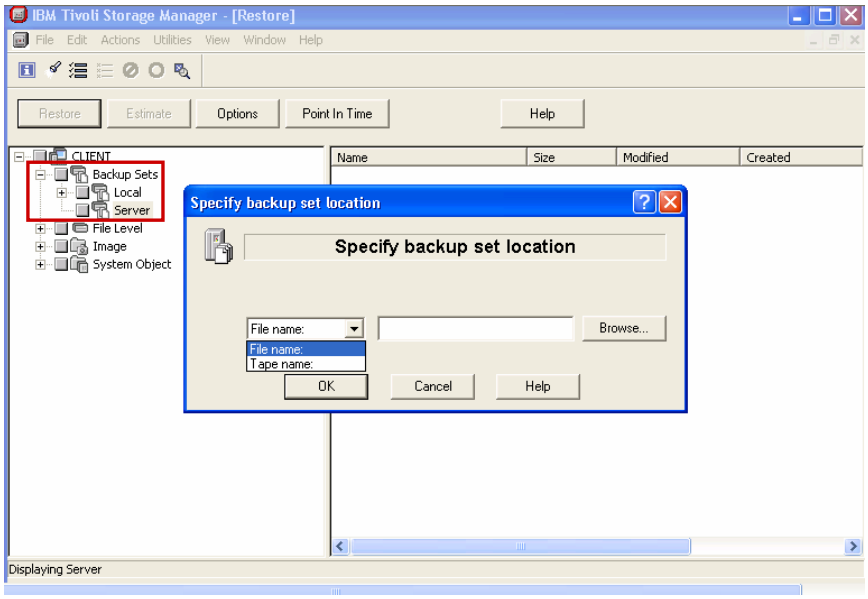

specify backup set location

navigation tree

Restoring Backup Set Considerations

- To restore files, specify the directories or selected files, or select the files from a list or graphical user interface (GUI) window.
- By default, only ACTIVE file versions will be available for selection. INACTIVE versions can be specified.
- Files can be restored to their original location or a specified directory. Collision options control whether existing files of the same name are replaced.

Remember, you can restore backup sets from the following locations:

- Portable media on a device attached to your client workstation
- A backup set file on the your client workstation
- Directly from the server

Section 5

Additional Resources

IBM Tivoli Storage Manager Training

http://www-306.ibm.com/software/tivoli/education/edu_prd.html

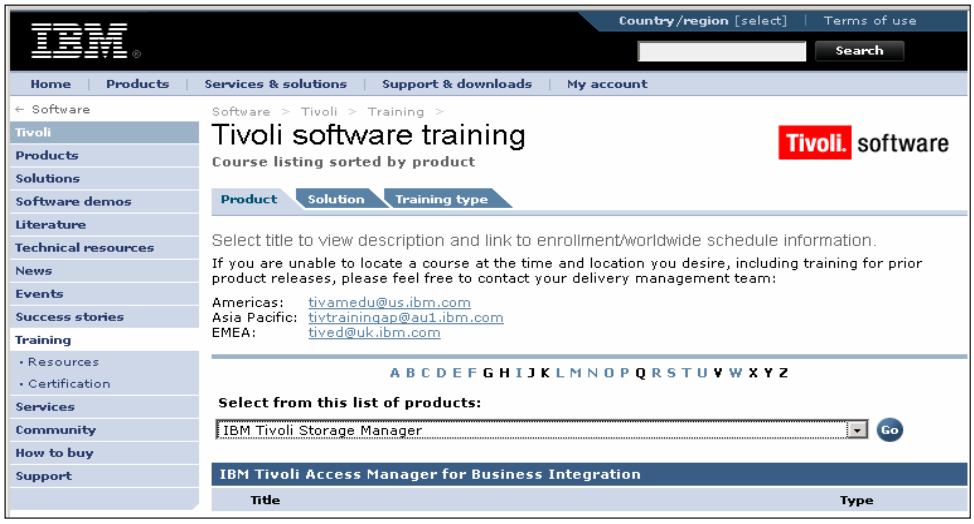

Suggested Reading

- Product Documentation
- IBM Redbooks

IBM Redbooks

http://www.redbooks.ibm.com/

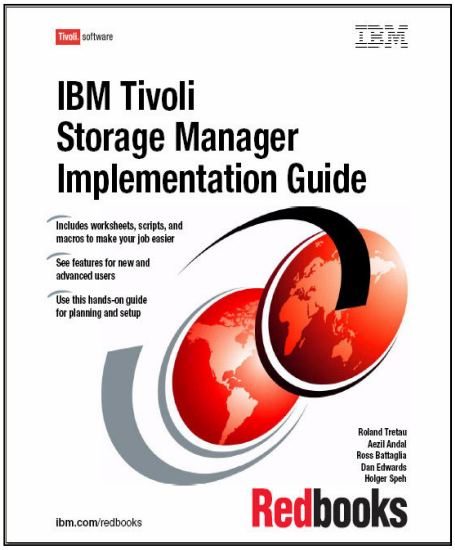

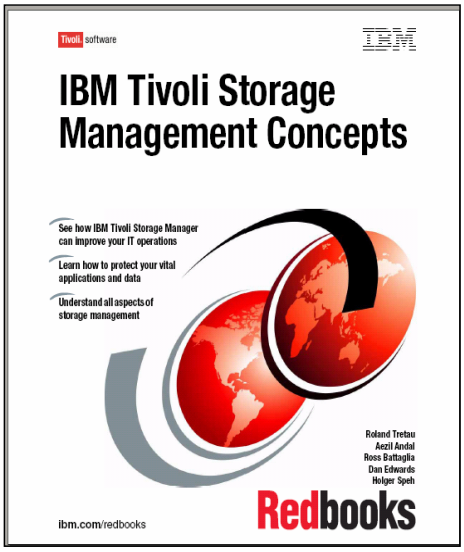

Template Revision: 7/18/2005 4:30 PM

Trademarks, Copyrights, and Disclaimers

The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

IBM	CICS	IMS	MQSeries	Tivoli
IBM(logo)	Cloudscape	Informix	OS/390	WebSphere
e(logo)business	DB2	iSeries	OS/400	xSeries
AIX	DB2 Universal Database	Lotus	pSeries	zSeries

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are registered trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel 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.

Linux is a registered trademark of Linus Torvalds.

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

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements and/or changes in the product(s) and/or program(s) described herein at any time without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectual property rights, may be used instead.

Information is provided "AS IS" without warranty of any kind. THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted, if at all, according to the terms and conditions of the agreements (e.g., IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. 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 in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. IBM makes no representations or warranties, express or implied, regarding non-IBM products and services.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

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

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Note to U.S. Government Users - Documentation related to restricted rights-Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract and IBM Corp.