



IBM Software Group – IBM Tivoli Training

IBM Tivoli® Storage Manager 5.3

Restoring Data



@business on demand.

© 2006 IBM Corporation
Updated September 26, 2006

Objectives

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

- Explain how data is restored to a backup-archive client machine
- Perform a command-line restore of a data file
- Restore a different version of the file
- Restore a file to a different location
- Explain the difference in restoring data and recovering a client system

Agenda

- Overview of Restoring Data
- Restoring with the Command Line
- Restoring a Different Version of a File
- Restoring a File to a New Location
- Recovering a Client System
- Additional Resources


Section 1

Overview of Restoring Data

Restore Overview

Restore is the process of copying a backup version of a user's file from the Tivoli Storage Manager server to the client's system over the specified communication method.

- **Data owners must issue a request to the server to restore their files.**
- **Users can restore the files of others as authorized.**
- **The Tivoli Storage Manager server sends a copy of the file back to the client, but the backup remains on the server.**
- **The user is able to restore backed up files to a specific point in time.**
- **Users can do restores through the GUI, client command line, or through a schedule.**
- **Users can restore files to a location other than the original backup location.**



Restore is the process of copying a backup version of a data owner, or user's file from the Tivoli Storage Manager server to the client's system over the specified communication method.

To restore a file or directory, you need to know what file you want to restore (such as a file name or directory), and which version of the file that is, if you want to restore a file other than the most recent one. You do not need to know where the data actually is on the server. When you request a file, Tivoli Storage Manager finds the location of that particular file from its database.

A user may only restore files that they have backed up unless they have been granted authority to restore another person's backed up files.

When a user restores a backup version of a file, Tivoli Storage Manager sends a copy of the file to the client node. The backup version remains in the Tivoli Storage Manager server.

The user can select to restore the most current copy of a file, or specify a file backed up at a particular point in time.

If more than one backup version exists, a user can restore the active backup version of the file or any inactive backup versions. By default, only active file version will be available for selection. However, inactive files can be selected and restored.

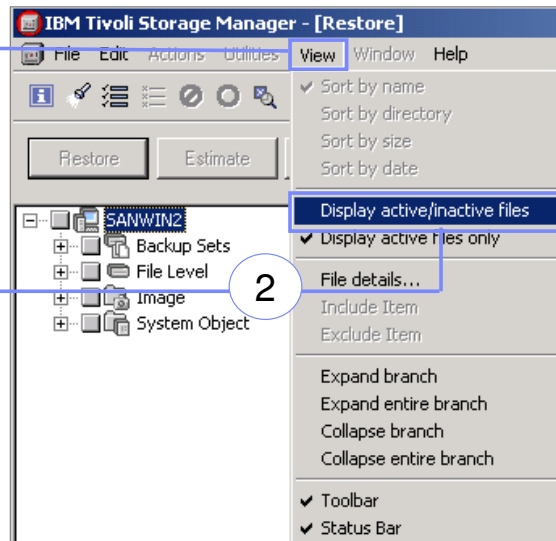
When restoring a directory its modification date and time is set to the date and time of the restore, not to the date and time the directory had when it was backed up. This is

Active and Inactive Versions of a File to Restore

Before you select the file you want to restore:

1. Click **View**.
2. Select **Display active/inactive files**.

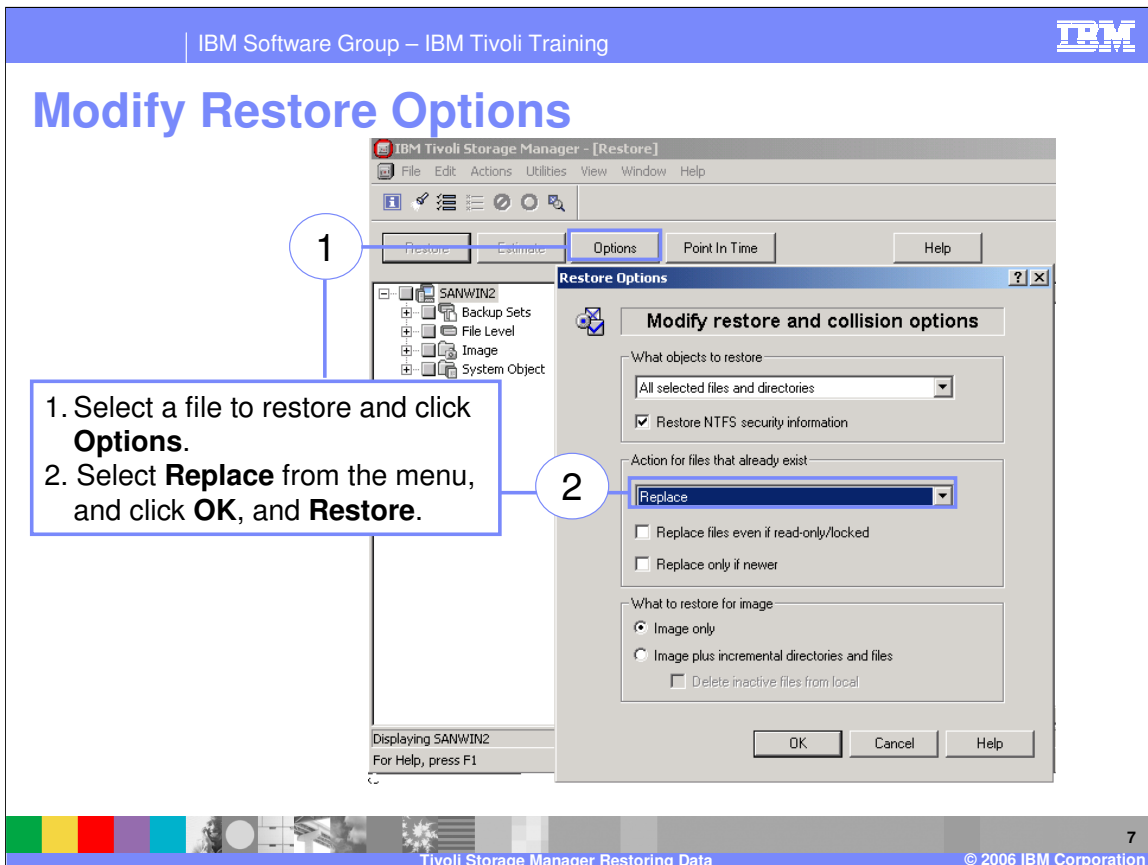
When you select the files you want to restore, you have a list of active and inactive files.



The restore GUI queries the TSM server for a list of files that have been backed up and presents them in the same format as the backup GUI. And, like the backup GUI, you can select the files you want to restore.

You can also use the **Find** function to select files. The **Find** function gives you the same options as those for doing a backup but will look on the TSM server for files that have already been backed up.

The default for restoring a file is the most recent version which is called the **ACTIVE** version. If you want to pick an object from an earlier backup, then you will have to choose an **INACTIVE** version. In order to restore an **INACTIVE** version of a file in the GUI, you must select **View** as it's shown in the slide then click **Display active/inactive files**. When you do that, all the **ACTIVE** and **INACTIVE** versions of a file are shown from which you can select the version you want.



1. Select a file to restore and click **Options**.
2. Select **Replace** from the menu, and click **OK**, and **Restore**.

The **Modify restore and collisions options** panel, shown in the slide, offers several choices for your restore process.

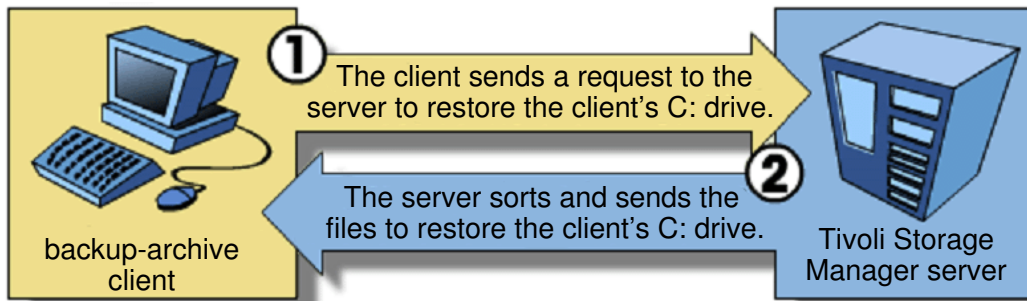
- first of all you can choose **All selected files and directories** to recover both files and the directories where the files reside.
- You can opt to chose **Files only** which means, like it says, that only files are recovered without the directories.
- and finally you can chose to recover directories only by choosing...**Directories only**.

The slide also shows that you can specify what actions to take on files that already exist:

- You can choose to Replace the file without prompting,
- You can choose to be prompted before replacing the file,
- You can choose NOT to replace any files (without being prompted).
- You can also select to replace the file only if the file is newer than the one being replaced.
- You can even specify how to restore an image by choosing **image only**, or **image plus incremental directories and files**.

The options you specify are effective during the current client operation, but they are not saved for the next operation.

No Query Restore



The client uses two different methods to determine what files should be restored: *standard restore* (also called classic restore), and *no query restore*.

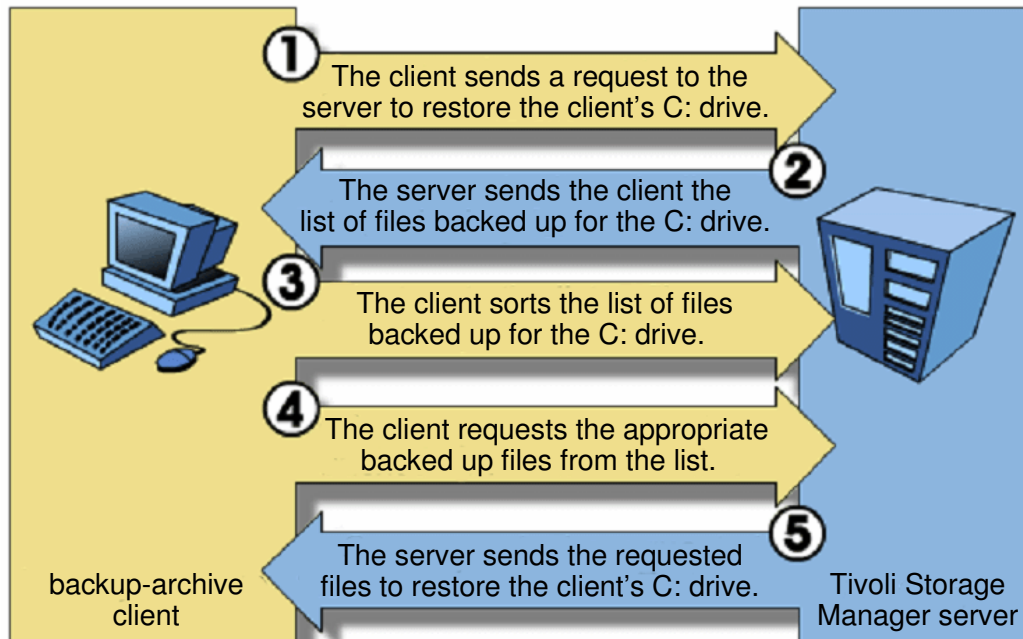
When you enter an unrestricted wildcard source file specification on the restore command and do not specify any of the options (inactive, latest, pick, fromdate, or todater), the client uses the *no query restore* method for restoring files and directories from the server. This method is called *no query restore* because instead of querying the server for each object to be restored, a single restore request is sent to the server. In this case, the server returns the files and directories to the client without further action by the client. The client merely accepts the data coming from the server and restores it to the destination named on the restore command.

So to walk through the no query process in the slide

Step 1. The client tells the server that a no query restore is going to be performed and provides the server with details about file spaces, directories, and files.

Step 2. The server sorts the data using an internal sort table which minimizes tape mounts and sends The data to be restored to the client.

Standard Restore



With the Standard restore process:

Step 1. The client queries the server for a list of files backed up for that part of the client that you're going to restore.

Step 2. The server sends a list of backed up files that match the restore criteria.

Step 3. The list of files returned from the server is sorted in client memory to determine the file restore order and to minimize tape mounts [required to perform the restore.]

Step 4. The client tells the server to restore file data and directory objects.

Step 5. The directories and files you want to restore are sent from the server to the client.

That's the standard restore method which can put some demands on the client.

Standard restore requires more interaction between the client and the server, and multiple processes cannot be used for the restore.

The *no query restore* requires less interaction between the client and the server, and the client can use multiple sessions for the restore process.

The *no query restore* process is useful when restoring large file systems on a client with limited memory because it avoids some processing that can affect the performance of other client applications. In addition, it can achieve a high degree of parallelism by restoring with multiple sessions from the server and storage agent simultaneously. The *no query restore* operation, however, can take much longer to complete than the *standard restore* operation.

Point in Time Restore

A Point in Time (PIT) Restore restores files to a past state that existed at a specific date and time, as opposed to the current state of the last backup.

Options for GUI and command-line point-in-time restores:

- Restore on file space, directory, or file level
- GUI window
- PITTIME hh:mm:ss
- PITDATE mm/dd/yyyy
- NLS enabled

Incremental backups notify of deleted files:


- Frequency determines point-in-time resolution for deleted files

No query restore process used:

- Fault tolerant restore
- Restartable restore

Policy definition considerations:

- Backup retention period
- Number of backup versions kept



10

Normally, when you request a restore, you get the most recent version. A point in time restore allows you to request older versions of a file that were backed before the most recent version. You can use Point in time restore to restore a file space, directory, or file to the version *equal to or before* the specified point in time.

Support for PIT restore is essential to be able to recover a file space or directory to a time when it was known to be in a good or consistent state. For example, if you discover that you've backed up a corrupted version of a file, you can restore a previously uncorrupted version using the point in time restore. Or if you've realized that you deleted the wrong version of a file after the backup has run. You can request the previous backup and restore the deleted file. When a PIT restore is performed, new files that have been created on the client after the PIT date are not deleted.

Progressive Incremental backups are necessary to capture the fact that files have been deleted.

Incremental backups should be run frequently enough to provide the necessary point-in-time resolution. If, for instance, you want to be able to restore a file for any day of the month, then you must run a backup every day. AND you must keep enough versions of the file. Ideally, the RETEXTRA and RETONLY (number of days for an inactive and deleted version) should be set to 31 and the VEREXISTS and VERDELETE (number of backed up versions) parameters should be set to NOLIMIT.

Restartable Restores


- Resumes at point of interruption

New restore states

- Active
- Restartable
- Restarts on transaction boundary

- Transparent to client

No Query Restore process used



If a restore gets interrupted, the user can restart the restore specifying the same source and destination. If the restore is started within the restart period allowed, the restore will start from where it left off.

You can restart the restore at the point of interruption by entering the restart restore command. When you restart an interrupted restore, it will restart with the first transaction not completely restored when the interruption occurred as it's shown in the slide. Since some of the files at the beginning of the transaction might have been restored before the interruption, you may receive some replace prompts for those files. When a client restore session is in a restartable state, the file space is locked in server storage and no files can be moved from sequential volumes. This prevents the data from being migrated, moved, reclaimed, or backed up by another operation. So, for instance, if you attempt to repeat the restore that was interrupted instead of restarting it, or try to back up the destination file space, the attempt will fail because you did not complete the original restore.

Restartable Restore Commands

query restore	<ul style="list-style-type: none"> ▪ Lists restartable restores ▪ Displays file specification, destination, node name, and owner
restart restore	<ul style="list-style-type: none"> ▪ Select from list of restartable restores ▪ Restore continues from point of interruption
cancel restore (administrative command)	<ul style="list-style-type: none"> ▪ Restart not performed when canceled by administrator ▪ Client cannot restart the restore
cancel sessions (administrative command)	<ul style="list-style-type: none"> ▪ Restart not performed when canceled by administrator ▪ Client can restart the restore



Here's a list of the commands that you can use with a restartable restore.

Use the **query restore** command to find out if your client has any restartable restore sessions in the server database.

The **restart restore** command will display a list of restartable restores from which you can choose.

To cancel a restore use the cancel restore command. When you cancel a restore it cannot be restarted. It's gone forever. But that also means the interrupted restore won't interfere with backups

To interrupt a restore in progress use the cancel session which is the last command in the list. When you cancel the session, the restore can be restarted.

Restoring System Objects

Certain Windows® 2000 and Windows XP system objects, such as system and boot files, and Windows Registries must be restored together in order to create a consistent system state.

Other system objects, such as event logs and replicated file systems (FRS) can be restored individually.



RESTORED TOGETHER in order to create a consistent system state:

- Active Directory (domain controller only)
- Certificate server database
- Cluster database (cluster node only)
- COM+ database
- Windows Registry
- System and boot files
- System volume

RESTORED INDIVIDUALLY. These include the following:

- Event logs (system, security and application)
- Removable Storage Management Database (RSM)
- Replicated file systems (FRS)
- Windows Management Instrumentation (WMI) repository

Section 2

Restoring with the Command Line

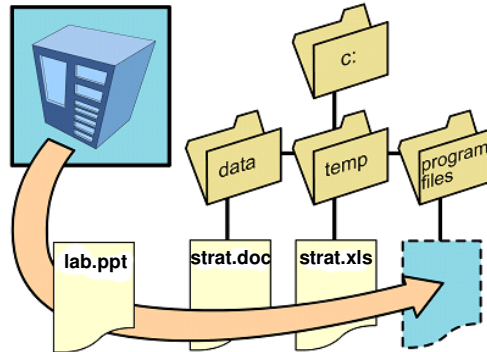
Refer to the Backup Archive docs

Restore Backed up Files Using the `dsmc restore` Command

Restore the `c:\Program Files\lab.ppt` file to its original directory.

```
dsmc restore c:\Program Files\lab.ppt
```

If you do not specify a destination, the files are restored to their original location.




Here's an example of a batch command. In this case we are restoring a specific file. The default destination is the original location of the file.

Use the `dsmc restore` Command with the PICK Option

The **pick** option creates a list of backup versions, images, or archive copies that match the file specification you enter. From the list, you can select the versions to process. Include the **inactive** option to view both active and inactive objects.

```
dsmc restore -pick -inactive c:\data\*
```



This again is an example of a command in batch mode

Here we are invoking the restore command with the pick option.

The pick option allows you to pick from a list of active and inactive objects on the server.

If you don't use the pick option with the inactive option, then it is indeterminate which version of the file will actually be restored.

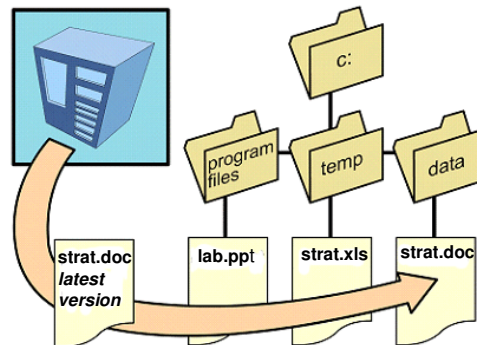
Use the LATEST Option with the `dsmc restore` Command

Use the **latest** option with the following commands to restore the most recent backup version of a file, even if the backup is inactive.

```
restore
restore group
restore was
```

Only active versions are considered for a restore unless you use either the **inactive** or the **latest** option.

```
dsmc restore -latest c:\data\*
```



Use the **latest** option with the **restore**, and **restore group** commands to restore the most recent backup version of a file, even if the backup is inactive.

IF you don't use the **inactive** or the **latest** option, then only active versions will be considered for a restore.

The **restore was** command specifies whether to restore the WebSphere® Application Server Network Deployment Manager (which contains setup, application files, and configuration information) or the Application Server from the Tivoli Storage Manager server.

Command Options

The following command options are available:

- TOTIME
- TODATE
- FROMTIME
- FROMDATE
- FROMNODE
- FROMOWNER
- PICK
- VIRTUALNODENAME
- SUBDIR



- Use the **TOTIME** option to specify a beginning time or a window of time. Use this option with the **FROMTIME** option to list files backed up between the from time and the to time.
- Use the **TODATE** option to specify the latest date of files to be listed. Tivoli Storage Manager does not list files backed up after the specified date. Use this option with the **FROMDATE** option to list files backed up between the from date and the to date. For example, you could request files that were backed up between 6:00 a.m. on June 1, 2004, and 11:59 p.m. on June 30, 2004.
- Use the **FROMNODE** option to see a list of file space names on another node on whose behalf you can restore or retrieve files.
- Use the **FROMOWNER** option to specify an alternative owner from which to restore files. The owner must grant access to use the files.
- Use the **PICK** option to display a list of backed up files that match the file specification you enter. From the list, you can select the versions to restore.
- The **VIRTUALNODENAME** option specifies the node name of your workstation when you want to restore files to a different workstation.
- The **SUBDIR** option specifies whether you want to include subdirectories of the specified directories.

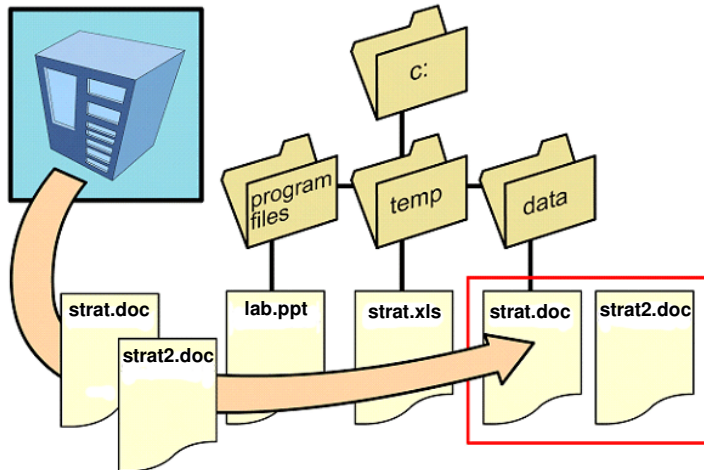
Section 3

Restoring a Different Version of a File

Use the `dsmc restore` Command with Overwrite

Restore all the files matching `strat*.doc` to their original locations, overwriting existing files.

```
dsmc restore -replace=yes c:\Data\strat*.doc
```

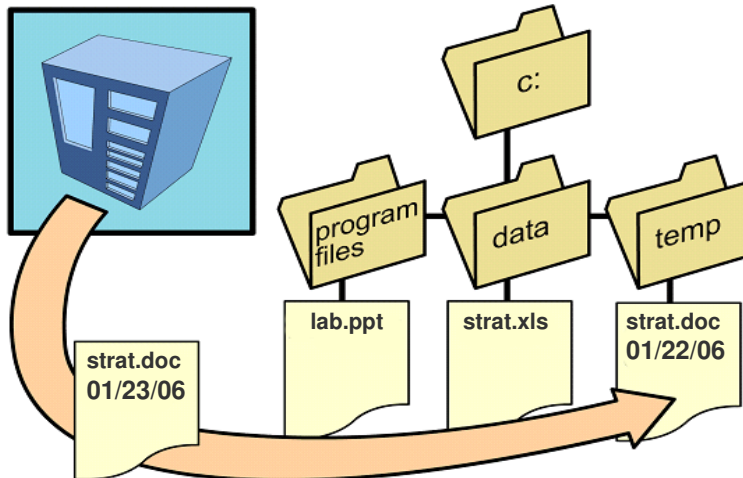


Use the `dsmc restore` command with the `replace=yes` option to restore all the files to their original locations and overwrite existing files. This is the same as the Replace option in the Backup-archive GUI.

Use the IFNEWER Option with the dsmc restore Command

The **ifnewer** option replaces an existing file with the latest backup version only if the backup version is newer than the existing file.

```
dsmc restore -ifnewer c:\temp\strat.doc
```



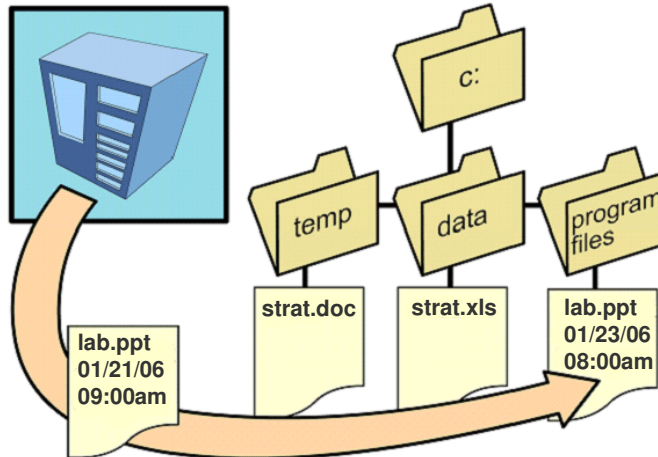
Use the **IFNEWER** option to instruct Tivoli Storage Manager to replace an existing file with the backup if the backup is newer than the existing file.

Remember, only active backups are considered unless you include the inactive or latest option.

Use the PITDATE and PITTIME Options

Use the **pitdate** option with the **pittime** option to establish a point in time for which you want to display or restore the latest version of your backups. If **pittime** is not specified, it defaults to **23:59:59**.

```
dsmc restore -pitt=09:00:00 -pitd=01/21/06  
c:\Program Files\lab.ppt
```



The **TODATE** and **TOTIME** or **FROMDATE** and **FROMTIME** options specify a date and time range for restore processing.

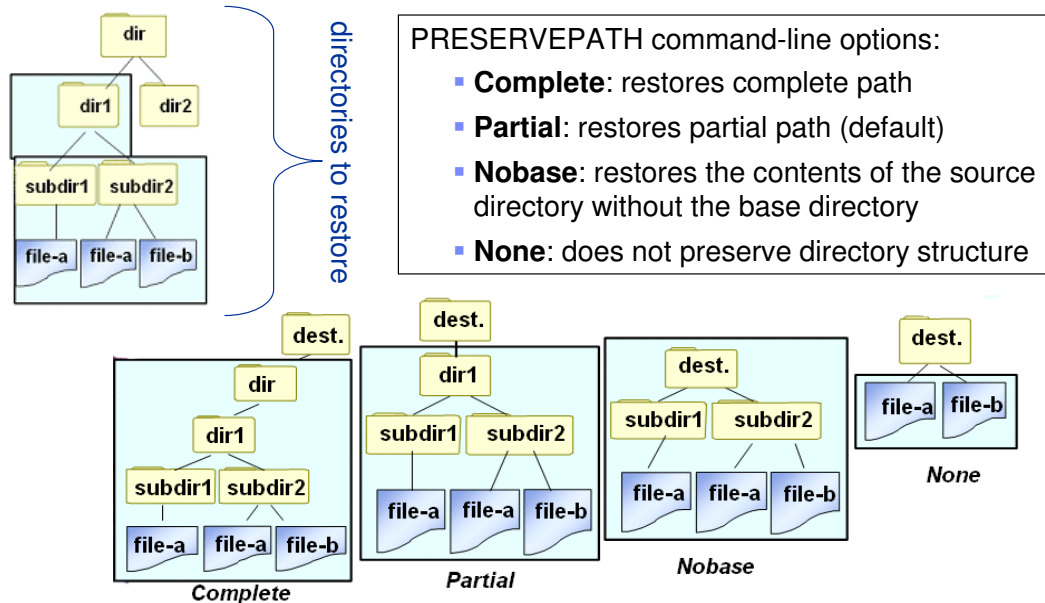
You cannot use the **no query** method with the **todate** option.

Do not select a **pittime** during a backup window.

Section 4

Restoring a File to a Different Location

Restore to New Location Using PRESERVEPATH Restore Option



The PRESERVEPATH restore option is available with the backup-archive command line. Use this option to specify how directory structures are handled when you restore objects to a new location. It creates the lowest-level source directory as a subdirectory of the target directory. Files from the source directory are then stored in the new subdirectory. The PRESERVEPATH option can be specified with one of four options.


- The **complete** option restores the entire path, starting from the root, into the specified directory. The entire path includes all the directories except the file space name.
- The **subtree or partial** option creates the lowest-level source directory as a subdirectory of the target directory. Files from the source directory are stored in the new subdirectory. This is the default. In the slide notice that folder “dir” is missing from the path.
- The **nobase** option restores the contents of the source directory without the base directory, into the specified destination directory.
- The **none** option restores all selected source files to the target directory. No part of the source path at or above the source directory is reproduced at the target. If you specify `subdir=yes` with the None parameter, then TSM restores all files in the source directories to the single target directory.

Restoring Files from Other Clients

The **FROMOWNER** and **FROMNODE** options permit one node to perform commands for another node.

Use the **FROMOWNER** option to specify an alternative owner from which to restore files. The owner must grant access to use the files.

Use the **FROMNODE** option to see a list of file space names on another node on whose behalf you can restore files.



Command example from node ABC:

```
dsmc restore -fromnode=XYZ \\XYZ\d$\invoice\* d:\invoice\
```



After a data owner grants you access to their files on the server, you can restore those files to your local system. You can display another user's file spaces on the server, or restore the other user's backup versions to your local file system.

Using the backup-archive GUI:

1. Click **Utilities** from the main window.
2. Click **Access Another Node**. The Access Another Node window appears.
3. Type the node name of the user's host machine in the Node name field.
4. Click the **Set** button.

If you are using commands, use the **fromnode** option to indicate the node. You must also use the file space name, rather than the drive letter, to select the restore drive you want to access. Include the file space name in braces and use it like a drive letter. For example, to restore xyz's files from the \invoice directory on the d-disk file space to your own \invoice directory on abc, enter: `dsmc restore -fromnode=xyz \\xyz\d$\invoice* d:\invoice\`


Section 5

Recovering a Client System

The Difference between Data Restore and System Recovery


A data restore is used when:

- Files or directories need to be replaced.
- Database data needs to be replaced (requires an API client).
- Application data, such as mail, needs to be replaced (requires an API client).
- Windows system objects need to be replaced.



System recovery, or Bare Machine Recovery is used if:

- Hardware failure makes it impossible to boot.
- The operating system is damaged.
- A system is being cloned.



Files or directories need to be replaced. Either due to damage to the original or a different version or location of the file or directory is needed.

Database data needs to be replaced. You will need IBM Tivoli Storage Manager for Databases installed.


Application data, such as mail, needs to be replaced. You will need IBM Tivoli Storage Manager for Mail, WebSphere Application Server... installed.

Windows system objects need to be replaced.

Bare Machine Recovery Overview

Bare Machine Recovery recreates a computer's:

- Hardware configurations
- Operating system
- Operating system patches
- Operating system customizations
- Disk partitions
- Device drivers
- Network settings



IBM Tivoli Storage Manager's Bare Machine Recovery Solutions

Operating Systems and Versions Supported	Cristie Bare Machine Recovery	Windows Automated System Recovery	IBM Tivoli Storage Manager for System Backup and Recovery
AIX 5.1, 5.2 and 5.3	-	-	Yes
Windows NT	Yes	-	-
Windows 2000	Yes	-	-
Windows XP	Yes	Yes	-
Windows 2003	Yes	Yes	-
Sun	Yes	-	-
Linux	Yes	-	-

There are 3 Bare Machine Recovery solutions with IBM Tivoli Storage Manager:
 Cristie Bare Machine Recovery, a third party software that works with IBM Tivoli Storage Manager

Windows Automates System Recovery, a Windows application that works with IBM Tivoli Storage Manager

IBM Tivoli Storage Manager for System Backup and Recovery, an IBM which formerly known as Sysback

This slide shows which operating systems support which solution.

Section 6

Additional Resources

IBM Tivoli Storage Manager Training

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



The screenshot shows the IBM Tivoli software training website. At the top, there is a navigation bar with the IBM logo, a search box, and a 'Search' button. Below this is a secondary navigation bar with links for Home, Products, Services & solutions, Support & downloads, and My account. The main content area is titled 'Tivoli software training' and includes a breadcrumb trail: Software > Tivoli > Training >. The page is sorted by product. A sidebar on the left contains a list of navigation options: Tivoli, Products, Solutions, Software demos, Literature, Technical resources, News, Events, Success stories, Training (with sub-links for Resources and Certification), Services, Community, How to buy, and Support. The main content area features a 'Tivoli software' logo, a heading 'Course listing sorted by product', and tabs for Product, Solution, and Training type. Below this, there is a message: 'Select title to view description and link to enrollment/worldwide schedule information. If you are unable to locate a course at the time and location you desire, including training for prior product releases, please feel free to contact your delivery management team:'. Contact information is provided for Americas (tivamedu@us.ibm.com), Asia Pacific (tivtrainingap@au1.ibm.com), and EMEA (tived@uk.ibm.com). A alphabetical index (A-Z) is displayed, followed by a section titled 'Select from this list of products:' with a dropdown menu showing 'IBM Tivoli Storage Manager' and a 'Go' button. Below the dropdown, a table header is visible with columns for 'Title' and 'Type'.

Suggested Reading

- Product Documentation
- IBM Redbooks

IBM Redbooks

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



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.

© Copyright International Business Machines Corporation 2004, 2005, 2006. All rights reserved.

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.

