



IBM® Tivoli® Storage Manager 5.4

Backup set enhancements

Tivoli. software



© 2007 IBM Corporation
Updated July 19, 2007

Objectives

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

- ▶ Describe the enhancements to backup sets.
- ▶ Generate backup sets.
- ▶ Restore files for a backup set.

Contents of this presentation

- Overview of new backup set features
- Point-in-time files selection
- New data types
- Backup set stacking
- Backup set table of contents
- New and modified commands
- Troubleshooting
- Support for upgrade from IBM Tivoli Storage Manager Express

Backup sets overview

A *backup set* is a collection of data backed up from one client, which is treated as a single object in server storage

The **GENERATE BACKUPSET** command issued on the server creates copies of active versions of a client's backed up objects that are within the one or more file spaces specified. The files are then consolidated onto sequential media

The backup object types supported for backup sets include directories, files, and image data

You can generate backup sets on the server for individual client nodes or for groups of nodes. A node group is a group of client nodes that are acted upon as a single entity

To create a node group:

- ▶ Use the **DEFINE NODEGROUP** command
- ▶ Then use the **DEFINE NODEGROUPMEMBER** to add nodes to the group
- ▶ The client node for which a backup set is generated must be registered to the server

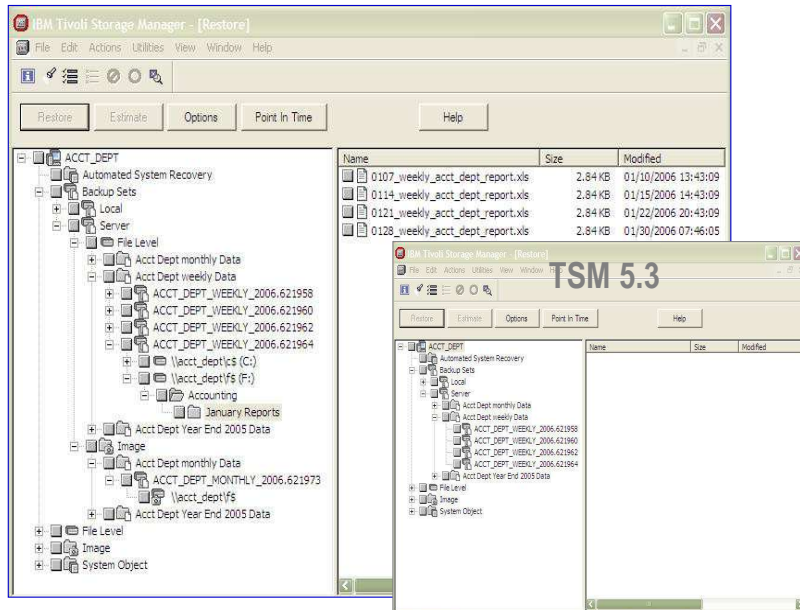
If you specify one or more node groups, the server generates a backup set for each node and places all of the backup sets together on a single set of output volumes.

Backup set enhancements overview

The following features available in IBM Tivoli Storage Manager Express are now available in the enterprise version of IBM Tivoli Storage Manager 5.4:

- ▶ Point-in-time file selection
- ▶ New data types: FILE, IMAGE
- ▶ Backup set stacking
- ▶ Backup set tables of contents (TOC) → individual file restore
- ▶ Packing lists (enhanced queries)
- ▶ Block size increased to 256K
- ▶ New and enhanced client-server protocols
 - TOC load for backup sets
 - Enhanced filespace query for backup sets
 - Classic restore-like interface from backup sets
 - No query restore using TOCSET token
 - Transparent access for data protection clients
 - Support upgrade from TSM Express to enterprise

Backup set GUI for restore



Backup set GUI for restore (cont.)

All backup set enhancements appear in the current restore dialog under the existing backup set node.

▶ Image support

- The backup sets node still expands to local and server. The server node now expands to show file level and image.
- Backup sets with file level and image data will show up under both file level and image.
- Image is not supported for the local backup set restore.

▶ Table of contents (TOC) support displays contents of backup set.

- Backup set node can now be expanded to display individual directories and files within the backup set
- A *table of contents is loading* message will be displayed with a cancel button
- TOC is not supported for local backup set restore

▶ Stacked backup sets

- Only the current node's data is displayed

Backup set enhancements requirements

Backup set enhancements (available on all supported platforms).

Limitations:

- ▶ No administration center support – command line only.
- ▶ On zOS, 256K block size requires 3590 or ECART.
- ▶ On zOS, FILE device type not compatible with open systems.
- ▶ No APPL support for enterprise; TSM Express only .
- ▶ No transparent access for data protection clients on enterprise, however, server remembers if upgraded from TSM Express.
 - Will allow DATATYPE=APPL and transparent access for data protection clients.
 - Only allows data protection for MS SQL, MS Exchange, and Continuous Data Protection.

Express backup sets are portable.

- ▶ Can be defined (reintroduced) on any 5.4 enterprise server.
 - This includes APPL backup sets.

All TSM Client platforms supported.

No specific hardware or software requirements.

Data Protection clients refer to IBM Tivoli Storage Manager for Microsoft SQL and IBM Tivoli Storage Manager for Microsoft Exchange

Point-in-time file selection

Historically:

- Backup sets contained all of a node's active files
- Only those files active at time of generation were eligible for inclusion

Tivoli Storage Manager 5.4:

- Backup sets allow point-in-time to be specified (default is TODAY, NOW)
- Files that were active at a specified date and time are included in backup sets
- In-progress backups do not affect selection of files



Note:

Backup sets cannot include deleted or expired files
Backup set date and time is **PITDATE** and **PITTIME**
If **PITDATE** is specified, the default **PITTIME** is midnight

New data type: File

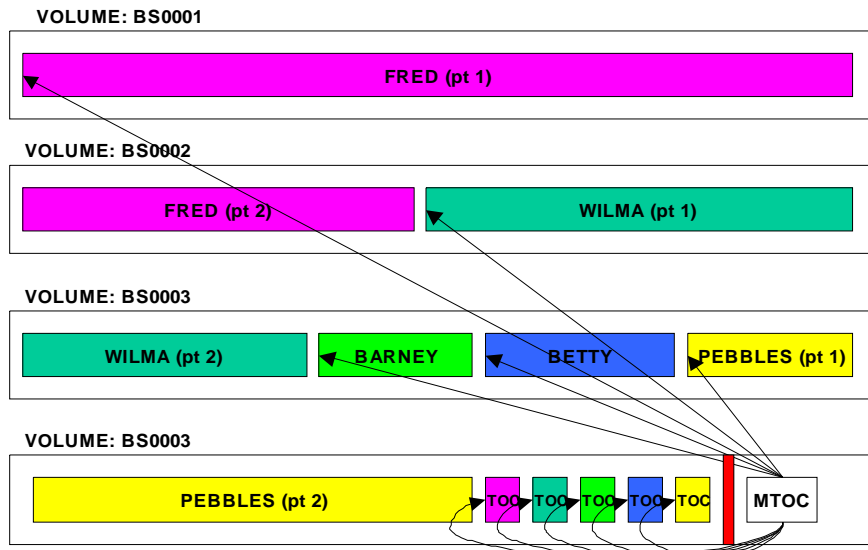
When generating backup sets, the server will search for the active version of a file in an active-data pool associated with a FILE device class, if such a pool is available. This search order is used:

1. A random-access disk (DISK) storage pool
 2. A primary or copy storage pool associated with a FILE device class
 3. An active-data pool associated with onsite removable media (tape or optical)
 4. A mounted sequential-access volume from a primary, copy, or active-data pool
 5. A sequential-access volume available in an automated library
 6. A sequential-access volume available in a manual library
- All backup set commands add DATATYPE or WHERE DATATYPE parameter
 - Data types cannot coexist; a given backup set has a specific data type
 - Node name, backup set name, and data type are needed to uniquely identify a backup set

Backup set stacking

- Multiple backup sets can reside on a single set of volumes.
- Multiple backup sets can be created at one time, but cannot append to the end of existing backup set.
- Multiple backup sets can include data from multiple nodes.
- Multiple backup sets can include different data types.
- All backup sets are given the same name.
 - ▶ Node name, backup set name, and data type
 - ▶ Used to uniquely identify a backup set
- Each backup set has its own inventory entries and table of contents (TOC).
- Each backup set can be independently queried, updated, or deleted
- Volumes do not return to scratch until all backup sets expire or are deleted.
- Local client restores are supported only for FILE data type

Stacked backup set layout



Backup set table of contents

- Tables of contents (TOC) for backup sets are created as part of backup set generation
- By default, when generating a backup set containing files and directories, the server will attempt to create a table of contents for a new backup set, but will not fail the backup set generation process if it cannot do so. You can, however, choose to require a table of contents by setting the TOC parameter to YES
- When generating a backup set that contains API data or images, the server will require that a table of contents is generated. You cannot override this default. In either case, if a table of contents is required and the server cannot create it, the backup set generation process will fail

Backup set table of contents (cont.)

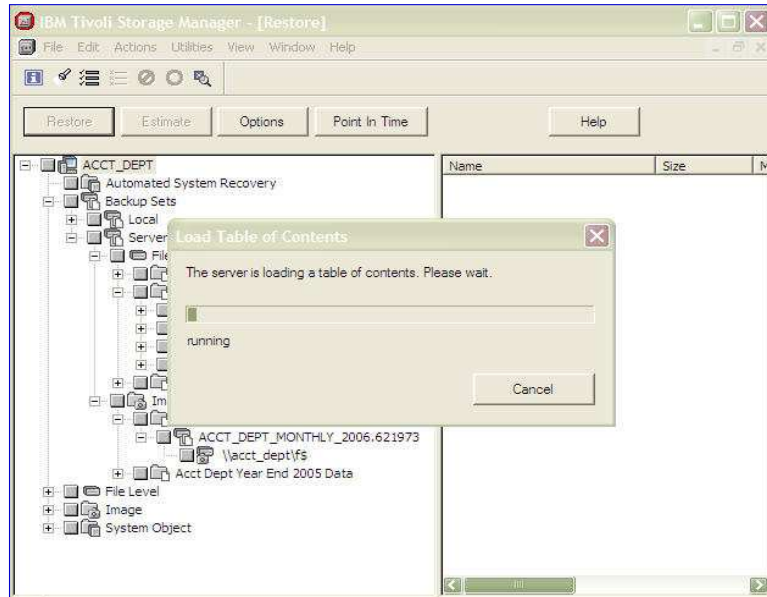
Tables of contents:

- Reside on the server even if the backup set's media has been moved offsite.
- Can be generated for existing backup sets that do not contain a table of contents.
- Can be re-generated when a backup set is defined on a new server, or if using a user-generated copy on a different medium.

The **GENERATE BACKUPSETTOC** command allows a table of contents to be created for backup sets that do not have one

GENERATE BACKUPSETTOC may be used after a backup set is added to the server with the **DEFINE BACKUPSET** command or for backup sets that were generated by an earlier release of the Tivoli Storage Manager server

Loading the table of contents



New and modified commands

■ GENERATE BACKUPSET

- ▶ **PITDATE**, **PITTIME** parameters added
- ▶ **TOC={PREFERRED,YES,NO}**, **TOCMGmtclass** parameters added
- ▶ **DATATYPE={ALL,FILE,IMAGE,APPL}** parameter added
- ▶ **NODENAME** parameter extended to allow node list, wild cards, and node groups

■ DEFINE BACKUPSET

- ▶ **TOC={PREFERRED,YES,NO}**, **TOCMGmtclass** parameters added
- ▶ **WHERE DATATYPE={ALL,FILE,IMAGE,APPL}** parameter added
- ▶ **NODENAME** parameter extended to allow node list, wild cards, and node groups
- ▶ **VOLUMELIST** can be in external file

APPL for TSM Express

TOC Specifies whether a table of contents (TOC) is saved for each file system backup.

No Specifies that table of contents information is not saved for file system backups.

Preferred Specifies that table of contents information should be saved for file system backups. However, a backup does not fail just because an error occurs during creation of the table of contents. This is the default value. **Yes** Specifies that table of contents information must be saved for each file system backup. A backup fails if an error occurs during creation of the table of contents.

WHERE DATATYPE Specifies the data type of backup set to be displayed. This parameter is optional. The default is to display all types of backup sets. To specify multiple data types, separate data types with commas and no intervening spaces.

WHERE TOCEXISTS Specifies whether a backup set must have a table of contents in order to be displayed. This parameter is optional. The default is to display all backup sets whether or not they have a table of contents.

New and modified commands (cont.)

- **DELETE BACKUPSET**
 - ▶ **WHERE DATATYPE** parameter added
- **QUERY BACKUPSET**
 - ▶ **WHERE DATATYPE, WHERE TO EXISTS** parameters added
- **QUERY BACKUPSET CONTENTS**
 - ▶ **DATATYPE={FILE,IMAGE,APPL}** parameter added
- **GENERATE BACKUPSET TOC** command added
- **DEFINE NODEGROUP** command added
- **NODEGROUPMEMBER** command added

17

Backup set enhancements

© 2007 IBM Corporation

GENERATE BACKUPSET TOC Use this command to generate a table of contents for a backup set that does not already have one. A table of contents allows the backup-archive client to display the contents of the backup set and allows the selection of individual files to be restored from the backup set.

DEFINE NODEGROUP Use this command to define a node group. A *node group* is a group of client nodes that are acted upon as if they were a single entity. A node can be a member of one or more node groups.

NODEGROUPMEMBER Use this command to add a client node to a node group. A *node group* is a group of client nodes that are acted upon as if they were a single entity.

Command line syntax for backup set

A new syntax was added to allow you to use the same commands that you used for non-backup set data, by introducing the **BACKUPSETNAME** parameter.

- dsmc query filespace -backupsetname=weekly_acct_dept.1234
- dsmc query backup f:* -backupsetname=weekly_acct_dept.1234 -subdir=yes
- dsmc query image -backupsetname=monthly_acct_dept.3212
- dsmc query systemobject -backupsetname=weekly_acct_dept.1234
- dsmc query systemstate -backupsetname=weekly_acct_dept.1234
- dsmc query systemservices -backupsetname=weekly_acct_dept.1234
- dsmc restore f:\dir1* -backupsetname=monthly_acct_dept.3212 -subdir=yes
- dsmc restore image f: -backupsetname=monthly_acct_dept.3212
- dsmc restore systemobject -backupsetname=weekly_acct_dept.1234
- dsmc restore systemstate -backupsetname=weekly_acct_dept.1234
- dsmc restore systemservices -backupsetname=weekly_acct_dept.1234
- dsmc query backupset -backupsetname=weekly_acct_dept.1234 -location=server
- dsmc restore backupset f:* -backupsetname=weekly_acct_dept.1234 -location=server -subdir=yes

The 5.3.0 command line backup set commands are still valid

```
dsmc query backupset <backupsetname> -location=<server|file|tape>
```

```
dsmc restore backupset <backupsetname> <srcfilespec> <destfilespec> -  
location=<server|file|tape>
```

Local backup set considerations

- Local backup set support same functionality as in 5.3.
- TOC support not implemented locally.
- Tivoli Storage Manager 5.4 backup set enhancements not supported for local restore.

Local backup set command examples

Backupsetname is the backup set file name or the name of the tape device.

- `dsmc query filespace -
backupsetname=e:\mybsets\weekly_acct_dept.ost -
location=file`
- `dsmc query backup f:* -backupsetname=\\.\tape0 -
location=tape -subdir=yes`
- `dsmc restore f:\dir1* -backupsetname=\\.\tape0 -
location=tape -subdir=yes`
- `dsmc query backupset -
backupsetname=e:\mybsets\weekly_acct_dept.ost -
location=file`
- `dsmc restore backupset f:* -
backupsetname=e:\mybsets\weekly_acct_dept.ost -
location=file -subdir=yes`

Troubleshooting

- Common usability problems:
 - ▶ Too many volume mounts
 - Storage pool volumes during GENERATE BACKUPSET
 - Backup set volumes during image restore, DEFINE BACKUPSET
 - ▶ Restores single-threaded on backup set volume
 - ▶ Using PIT could cause backup sets to expire too early

Support for IBM Tivoli Storage Manager Express

Data type APPL - new for Tivoli Storage Manager Express

- Includes only data backed up by API
- Limited to data from data protection for Microsoft SQL, data protection for Microsoft Exchange, and continuous data protection
- If you are upgrading from Tivoli Storage Manager Express, backup sets can also contain data from data protection for Microsoft SQL and data protection for Microsoft Exchange servers

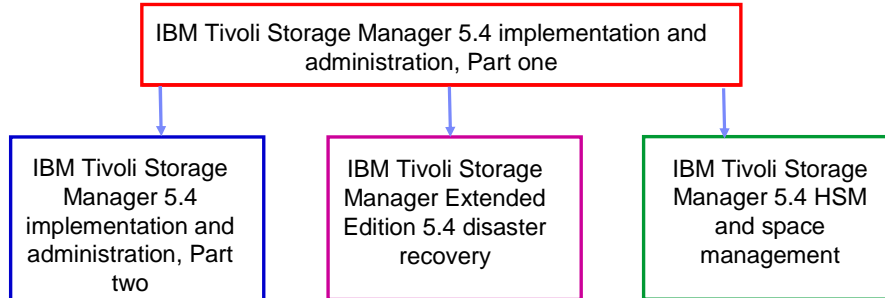
Image backups - new for Tivoli Storage Manager Express

- Includes new and modified files backed up after image backup
- Includes information about files deleted after image backup
- Excludes API data

Summary

- You should now be able to:
 - ▶ Describe the enhancements to backup sets
 - ▶ Generate backup sets
 - ▶ Restore files for a backup set

IBM Tivoli Storage Manager 5.4 curriculum roadmap for implementers and administrators



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

Feedback

Your feedback is valuable

You can help improve the quality of IBM Education Assistant content to better meet your needs by providing feedback.

- Did you find this module useful?
- Did it help you solve a problem or answer a question?
- Do you have suggestions for improvements?

Click to send e-mail feedback:

mailto:iea@us.ibm.com?subject= Feedback about TSM54_backup_set.ppt

You can help improve the quality of IBM Education Assistant content by providing feedback.

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 Tivoli

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

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 or changes in the products or programs 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 NON-INFRINGEMENT. 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 (for example, 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 2007. 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.