



IBM® Tivoli® Storage Manager 5.4

Overview of active-data pools

Tivoli. software



© 2007 IBM Corporation
Updated July 19, 2007

Objectives

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

- Describe active-data pools (ADP)
- Explain how data is handled in an active-data pool
- Restore data to a client from an active-data pool

Contents of this presentation

- Overview of active-data pools (ADP)
- Handling data in an active-data pool
- Client restore from active-data pools

Overview

Overview of active-data pools

- **New to Tivoli Storage Manager 5.4.0**
- **Can be used with sequential-access disk (FILE)**
- **Can be used with sequential-access tape or optical**
- **Can collocate active versions of backup data**
- **Can be used for backup data only; no archive or HSM data**

Benefits of active-data pools on sequential-access disk (FILE)

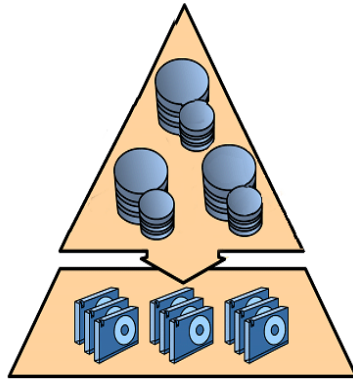
- **Optimized access to active versions for fast restore**
- **Reduced size of disk pool if only active versions are stored**
- **Reduced data movement in preparation for restore of active data**

Benefits of active-data pools on tape


- **Reduced storage requirement (onsite or offsite) while protecting data against media failure or disaster**
- **Simplified tape management because of fewer tapes**



Active-data pools and the storage hierarchy




storage hierarchy with active and inactive data



Copy storage pool:
Copies of active or inactive data from primary storage pools for offsite storage.



Active-data pools on disk:
Copies of active data for fast restore of client data.



Active-data pools on tape:
Copies of active data to reduce the number of tapes stored offsite.



Before storing data in an active-data pool

To use an existing domain:

- Use the UPDATE DOMAIN command and specify the list of active-data pools in the ACTIVEDESTINATION parameter

To restrict which nodes can write to an active-data pool:

- Define a new domain and specify the list of active-data pools in the ACTIVEDESTINATION parameter
- Update the node definitions to assign the nodes to the new domain
- Complete the policy structure

Domains are used to filter which nodes can store data into an active-data pool. Before TSM writes data to an active-data pool it verifies that the node owning the data is assigned to a domain that has the active-data pool defined in the ACTIVEDESTINATION parameter.

New parameter, ACTIVEDESTINATION, on the domain specifies the names of active-data pools where active-versions of backup data is stored for nodes assigned to the DOMAIN.

Administrators must decide whether to use an existing domain or establish a new policy in which the domain will authorize nodes to store data into an active-data pool

Storing data in an active-data pool

Using simultaneous write

The new parameter for the `DEFINE STGPOOL` and `UPDATE STGPOOL` commands is `ACTIVEDATAPOOLS`. This specifies the names of active-data pools where the server simultaneously writes data during a client backup operation.

Using the `COPY ACTIVE` command

Use this command to copy active versions of backup data from

A primary storage pool to an active-data pool. This command is similar to `BACKUP STGPOOL`.



Removing inactive files from an active-data pool

When files are deactivated, the logical occupancy of the volume it is stored on is reduced by the size of the file. This causes reclamation to run sooner for the volume.

Reclamation (reconstruction) is used to remove inactive files. This will build the new aggregate by copying just the active files and not any deleted or inactive files.

If the removed inactive files are part of an aggregate, they are represented in the database by having a zero length.

Nonaggregated inactive files are deleted during deactivation. There are no nonaggregated inactive files in an active-data pool.

The **COPY ACTIVE DATA**, **MOVE DATA**, and **MOVE NODE DATA** commands will reconstruct aggregates as they are copied or moved within an active-data pool.

The **RECONSTRUCT** option on the **MOVE DATA** and **MOVE NODE DATA** commands do not apply when dealing with active-data pools.



Using active-data pools to restore storage pools or volumes

Two new parameters on the RESTORE STGPOOL and RESTORE VOLUME commands:

- **ACTIVEDATAONLY**

- ▶ Specifies that active versions of backup files are to be restored from active-data pools only.

- **ACTIVEDATAPOOL**

- ▶ Specifies the name of the active-data pool from which the active versions of backup files are to be restored. If this parameter is not specified, files are restored from any active-data pool in which active versions of backup files can be located

Restoring from an active-data pool will cause some inactive files to be deleted from the server inventory. If the server determines that the inactive file needs to be replaced but can not find it in the active-data pool, the inactive file will be deleted.

If the inactive file has not been removed it will be restored to the primary storage pool or volume.

Client restore from active-data pools

When restoring client data, the server will be modified to select the active version of a file from an active-data pool (ADP) when appropriate. The restore order has been modified to include ADPs:

- ▶ ADP FILE
- ▶ DISK
- ▶ FILE primary or copy
- ▶ ADP sequential onsite
- ▶ Sequential onsite volume (primary, copy, or ADP), priority considers:

mounted(idle) > automated > manual

Performing a point-in-time restore from active-data pools is not allowed.

QUERY OCCUPANCY and QUERY CONTENT

QUERY OCCUPANCY shows:

- Number of files
 - ▶ Reduced when nonaggregated files are deleted during deactivation
 - ▶ Reduced by the number of logical files only when aggregate is deleted
- Physical size reflects the actual size of the aggregates
- Logical Space Occupied
 - ▶ Reduced by the size of logical file being deleted
 - ▶ Reduced by the size of logical file that is removed during reclamation (reconstruction)

QUERY CONTENT shows:

- All logical files in an aggregate, regardless if they have been removed.

It is important to maintain the referential integrity of the list of logical files in the aggregate.

Troubleshooting

Unable to store or copy data into an active-data pool.

- ▶ Check to ensure the domain the node is assigned to has the active-data pool listed in the ACTIVEDESTINATION parameter of the domain.

QUERY OCCUPANCY still shows files that have been removed.

- ▶ Files are in an aggregate that still has active files. Until the aggregate has been deleted, all logical files will be shown.

COPY ACTIVATEDATA runs slower than BACKUP STGPOOL.

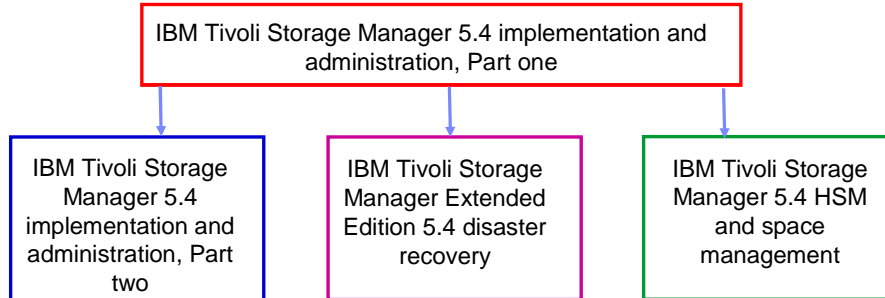
- ▶ Difference is due to additional processing needed to verify the active or inactive state of logical files in aggregates and act accordingly.

Summary

You should now be able to:

- Describe active-data pools (ADP)
- Explain how data is handled in an active-data pool
- Restore data to a client from an active-data pool

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 Active Data.ppt](mailto:iea@us.ibm.com?subject=Feedback%20about%20TSM54%20Active%20Data.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

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.