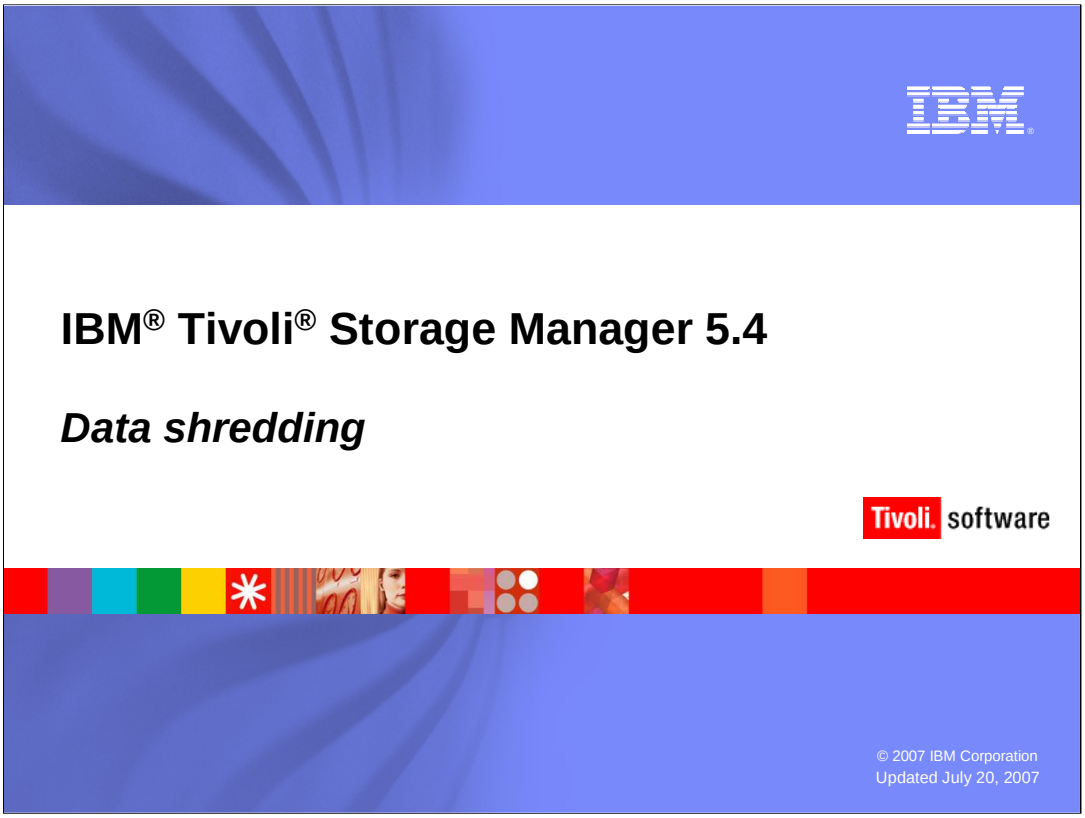
IBM
IBM® Tivoli® Storage Manager 5.4
Data shredding
Tivoli. software
© 2007 IBM Corporation
Updated July 20, 2007

Objectives

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

- Describe the data shredding process
- Shred data using the command line
- Shred data using the administration center

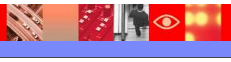

Overview of data shredding

- Shredding is the destruction of deleted data
- Normally, when an object on the server is deleted, the database entries are deleted but the data in the storage pool still exists
- Shredding causes the data in the storage pool to be overwritten one or more times after it is deleted
- This is the same concept as erase-on-scratch
- This feature addresses random-access DISK storage pools only

Requirements for data shredding

- Requires Tivoli Storage Manager 5.4 server or above; independent of storage agent or client
- No new hardware or other software is required
- Migration: Administrators can enable and disable shredding for existing pools at any time, but it is not retro-active
- The support is platform independent

Storage agent can not do random-access pools, so no interaction.

Not retro-active means stuff that was already deleted is either shredded or not shredded according to the setting at the time of deletion.

Data shredding process

- Shredding is accomplished by interrupting the delete. When an object is deleted, the inventory entries are removed but the storage subsystem entries are not
- Deleted bitfiles are placed in a new *pending shred* state by creating entries for them in the *shred list*. The object is considered logically deleted at this point
- The shredder (an asynchronous thread) will scan the shred list, shred the data represented by each entry, and delete the remaining entries for the bitfile
- The object is considered physically deleted at this point. The shred is done by overwriting the original data with the user-specified number of randomly generated patterns
- The intent is to make the data more difficult to be discovered and reconstructed later

Data shredding process (cont.)

How to shred data:

- SHRED is an attribute on DEFINE STGPOOL and UPDATE STGPOOL commands
- SHREDDING server option
- SHRED DATA command
- QUERY SHREDSTATUS command

Example:

```
update stgp backuppool shred=2
```

ANR2017I Administrator SERVER_CONSOLE issued command: UPDATE STGPOOL backuppool shred=2

ANR1309W Shred value zero for storage pool DISKPOOL may render deleted data non shreddable.

ANR2202I Storage pool BACKUPPOOL updated.

in this example, 2 is the overwrite count

SHRED DATA command syntax

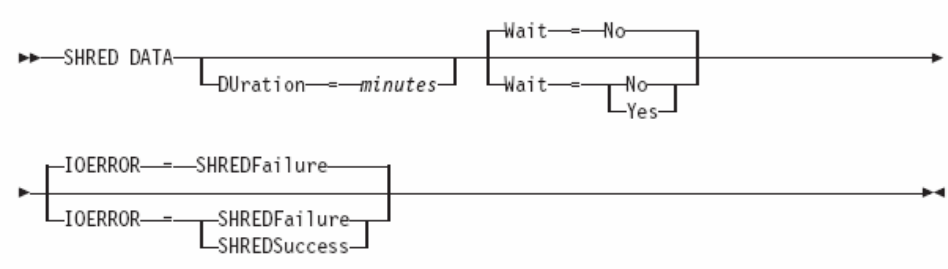

DURATION Specifies the maximum number of minutes the shredding process runs before being automatically cancelled. When the specified number of minutes elapses, the server cancels the shredding process. As soon as the process recognizes the cancellation, it ends. Because of this, the process may run longer than the value you specified for this parameter. You can specify a number from 1 to 9999. This parameter is optional. If not specified, the server will stop only after all deleted sensitive data has been shredded. **Wait** Specifies whether to wait for the server to complete processing this command in the foreground. This parameter is optional. The default is No. Possible values are: **No** Specifies that the server processes this command in the background. You can continue with other tasks while the command is being processed. Messages created from the background process are displayed either in the activity log or the server console, or both, depending on where messages are logged. To cancel a background process, use the CANCEL PROCESS command. If you cancel this process, some files might already have been shredded before the cancellation. This is the default. **Yes** Specifies that the server processes this command in the foreground. You must wait for the operation to complete before continuing with other tasks. The server displays the output messages to the administrative client when the operation completes. Messages are also displayed either in the activity log or the server console, or both, depending on where messages are logged. **IOERROR** Specifies whether an I/O error encountered while shredding the data is to be considered a successful shred. This parameter is optional. The default is SHREDFAILURE. Possible values are: **SHREDFailure** Specifies that if the server encounters an I/O error while shredding, the data will not be considered successfully shredded and the owning file will be marked as damaged. The server will attempt to shred the data again the next time the shredding process runs, giving you a chance to correct the error and ensure the data can be properly shredded. **SHREDSuccess** Specifies that if the server encounters an I/O error while shredding and the owning file had been previously marked as damaged, the data will be considered successfully shredded. You should use this option only after the server has reported I/O errors while shredding and you are unable to correct the error.

Troubleshooting

The SHOW SHREDLIST command reports detailed information about objects waiting to be shredded

Section

Data shredding using the administration center

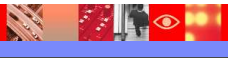

Overview of administration center changes for data shredding

- Data shredding affected a number of panels, which will display only for Tivoli Storage Manage server 5.4 or above. The shredding changes for a Tivoli Storage Manager server 5.3 or below will not be displayed
- This includes new panels and changes of existing panels to display and set shred status for storage pools and servers, and handle shred data for export, backup set, storage pool backup, and move data operations
- The new move node data and copy active data panels have shredding parameters as well
- For command storage pool updates, warning messages will be displayed in the administration center confirm panels before the command is issued

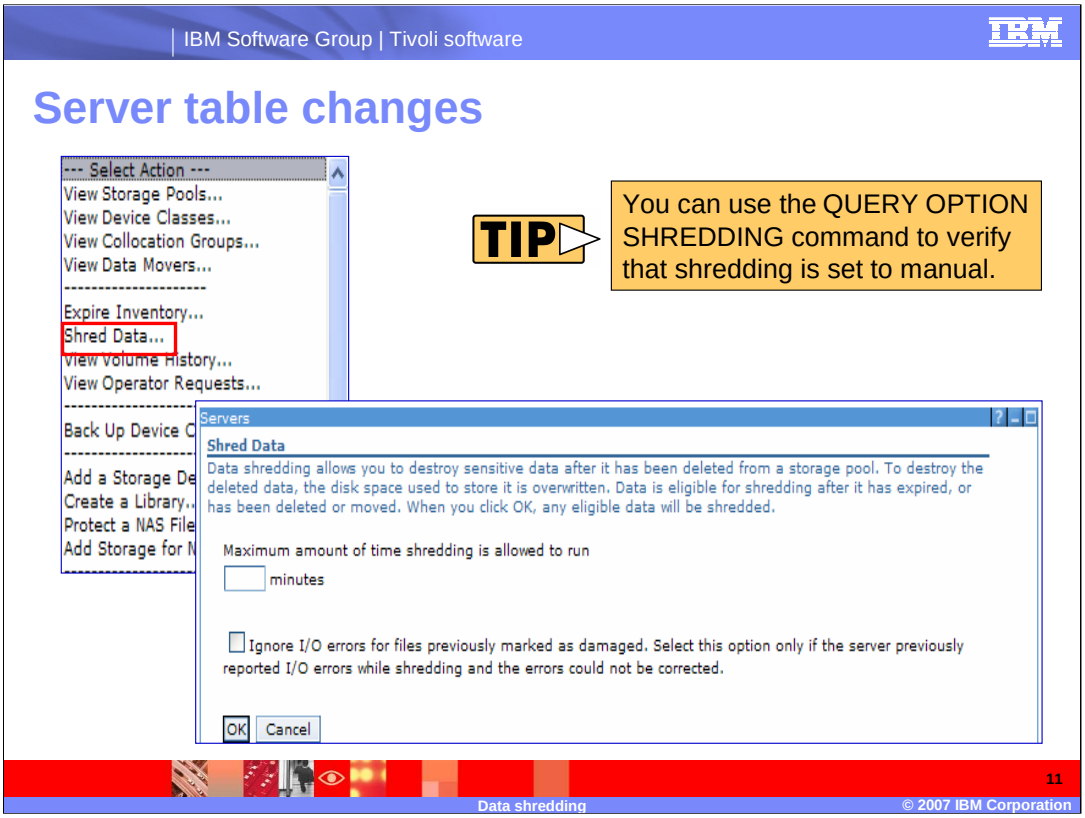

When **Shred Data**… is selected, the Shred Data form is displayed if the server is set to manual shredding.

If the server is set to automatic shredding, the following message will be issued:

The selected server is set for automatic shredding, shred data command cannot be issued. To enable manual shredding, use the server properties notebook Security tab to change the server to manually shred data.

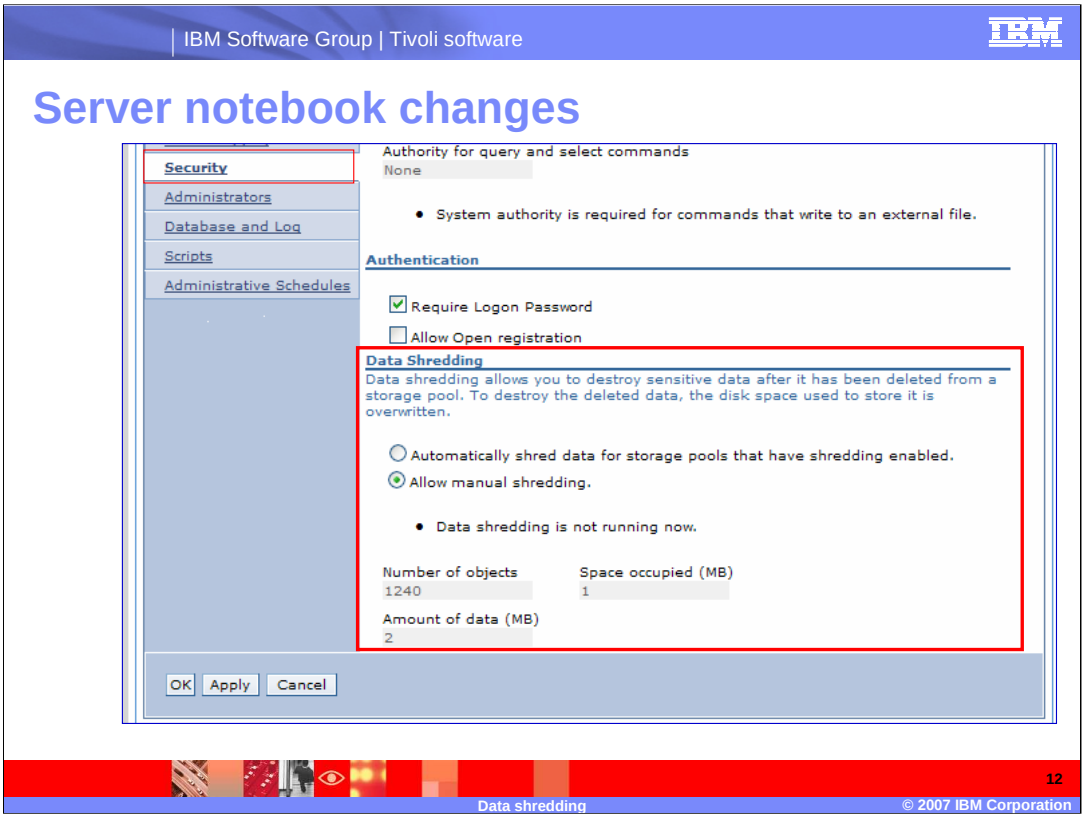

Server Properties Notebook – Security Tab: Added Data Shredding

This section sets automatic or manual shredding and displays shredding information

Storage pool table changes

The storage pool table has a new field, *Shred*.

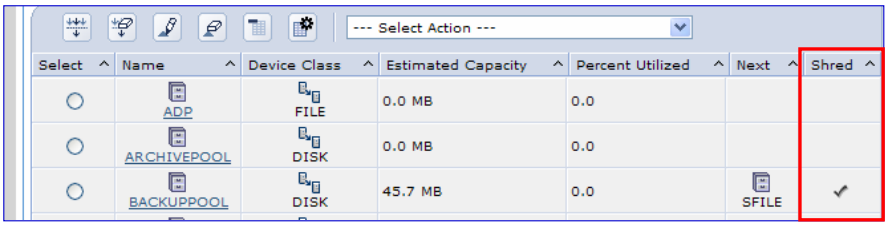

Storage pool notebook changes

Storage pool notebook has added text and a selection box in the General tab. This will only appear for DISK storage pools.

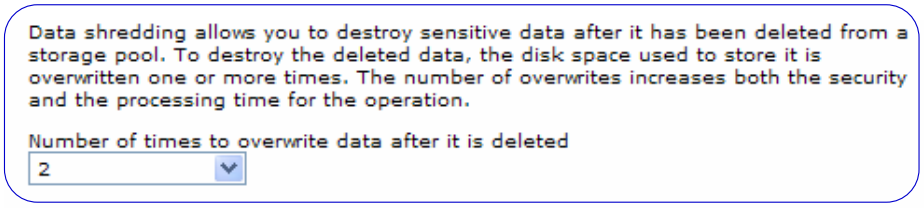

When the APPLY or OK button is pressed on any disk storage pool notebook panel, the shredding status is checked for conflicts. There are two general types of conflict that can exist, a conflict with data caching, or a data shredding conflict.

Data shredding from the general tab and caching from the migration tab are mutually exclusive selections, either shredding must be ‘Do not overwrite’, or ‘Leave a copy of the data in the storage pool after migration’ must be unchecked. If both are checked a data caching conflict exists, and the migration panel will be marked in error, and an error message will be displayed.

Storage pool notebook changes (cont.)

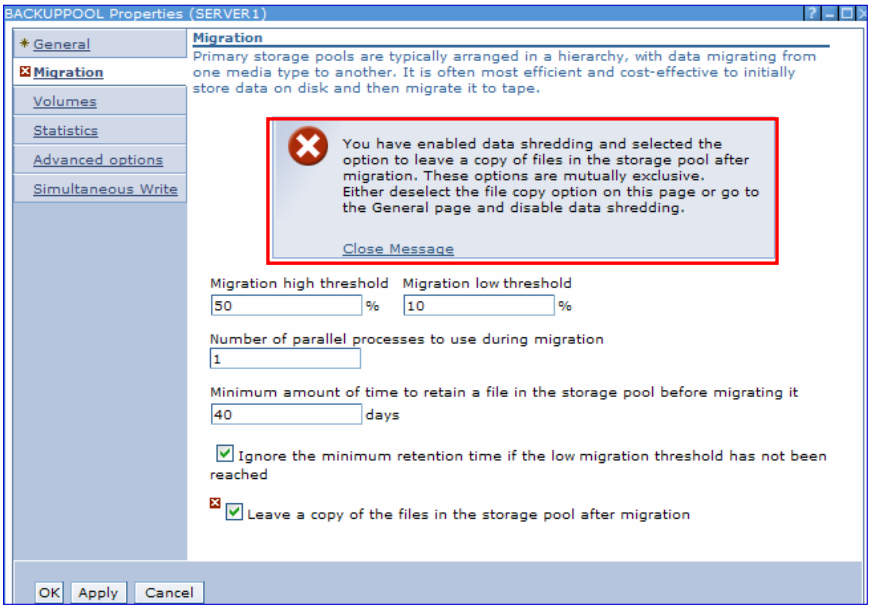

Storage pool data shredding conflict

A data shredding conflict exists when a change to storage pool properties will cause data that previously existed in a shred pool to no longer be in a location where it will be shredded.

This can occur when:

- Shredding is disabled on the storage pool.
  - Value is changed from a number to *Do not overwrite.*
- The storage pool is setup for shredding.
  - The next pool is updated, and the new next pool is not a shred pool.
  - There is a simultaneous write pool defined or updated.
- The storage pool is changed to shredding.
  - The next pool is not a shred pool.
  - There is a simultaneous write pool defined.

In these cases, a shred conflict panel will appear. Accept or Cancel buttons will display. Cancel will cause return to the notebook, Accept behavior will follow the original notebook button, OK ends notebook, Apply returns to the notebook.

Conflict forms

There are two versions of the conflict form for disabling shredding, depending on whether or not the pool is also defined as a next pool for another storage pool.

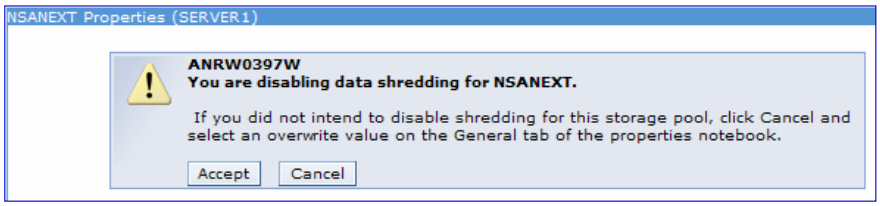

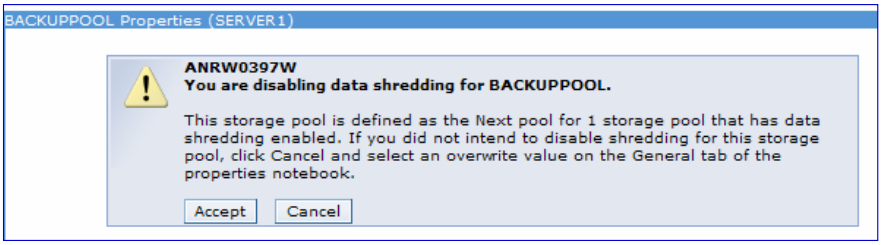

Conflict forms (cont.)

There are several versions of the conflict form for enabling shredding, depending the defined next pool or simultaneous write pools. The basic form is the *enabling* message, followed by the pools causing a conflict, so it can display the next pool, simultaneous write pools, or both.

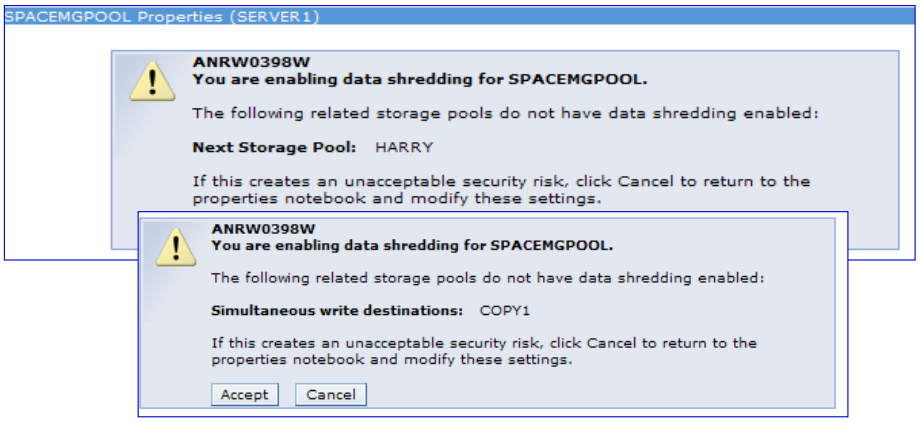

Conflict forms (cont.)

There are also several versions of the conflict form for changing shredding related pools. The basic form is the *modifying* message, followed by the pools causing a conflict, then a *fix* or accept instructions. If the only conflict is the next pool, a drop down menu to change that pool is also displayed.

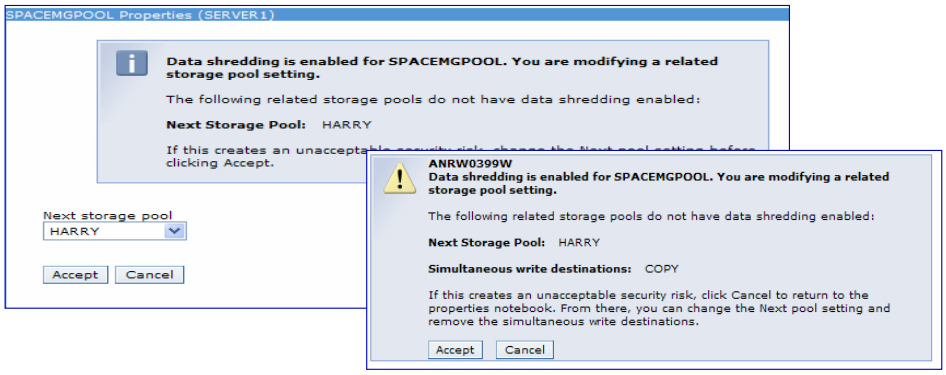

Warning message

If the original pool is a shred pool, and the location where the data will be backed up, copied, or moved is not shred, a warning message will be displayed.

Example of a warning message:

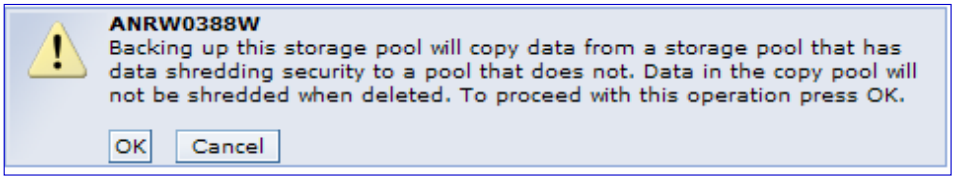

Warning message occurs on the following:

The Backup Stgpool form

Copy Active Data form

Move Data form

Maintenance script

Copy Active Data and Backup Stgpools tabs

add and modify a relationship forms

Move data panel changes

Panel is updated to show ***Shredding Not Enabled*** and ***Shredding Enabled*** pools in *Select Storage Pool* drop down menu,

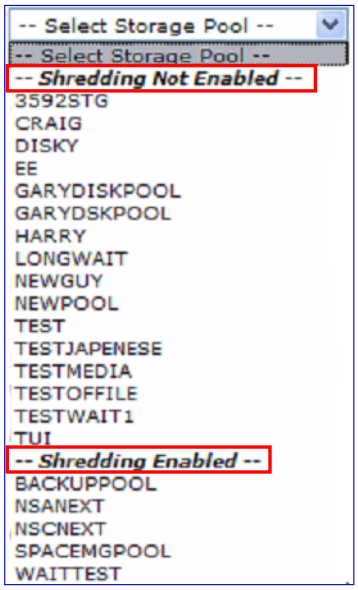

NOTE: Storage pools displayed for move data operations will have the same data format as the original pool.
The Tivoli Storage Manager 5.4 administration center will do this for both Tivoli Storage Manager 5.4 and previous servers.

If the source pool is shred enabled, the shred pools will be listed first, if it is not, as shown here, the shredding not enabled pools will display first.

Wizard changes

The export node wizard, export server wizard, and generate backupset wizard have a radio button group added if any shred pools exist on the server.

The default value is no, and for export, the radio button group is disabled until a selection is made to export client data.

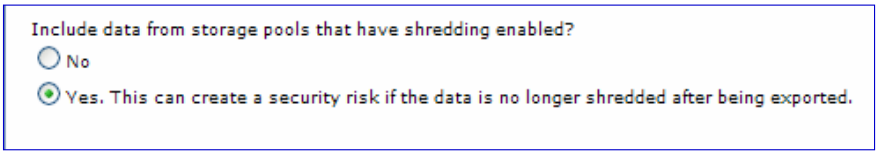

The example shown is the radio button group from the Export Node wizard.

Summary

You should now be able to:

- Describe the data shredding process
- Shred data using the command line
- Shred data using the Administration Center

IBM Tivoli Storage Manager 5.4 curriculum roadmap for implementers and administrators

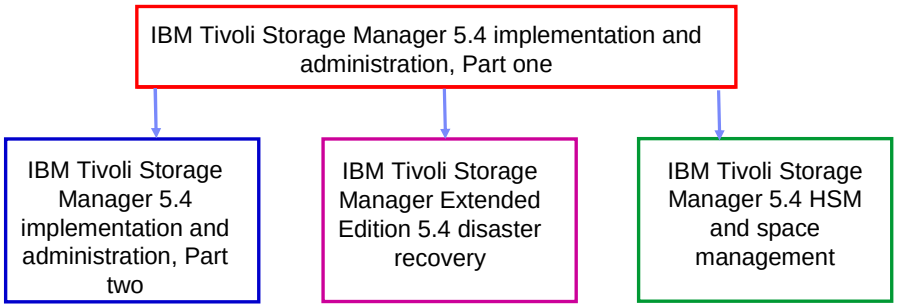

http://www.ibm.com/software/tivoli/education/edu_prd.html

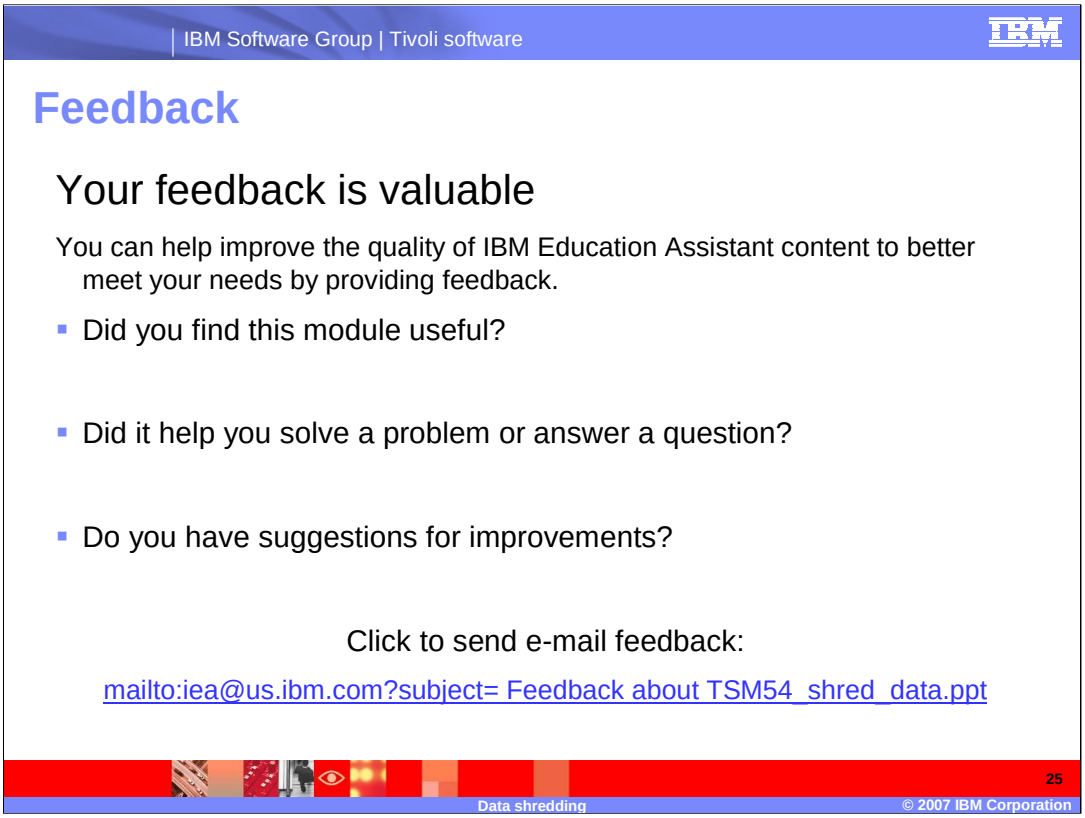

You can help improve the quality of IBM Education Assistant content by providing feedback.

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 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 (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.

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.