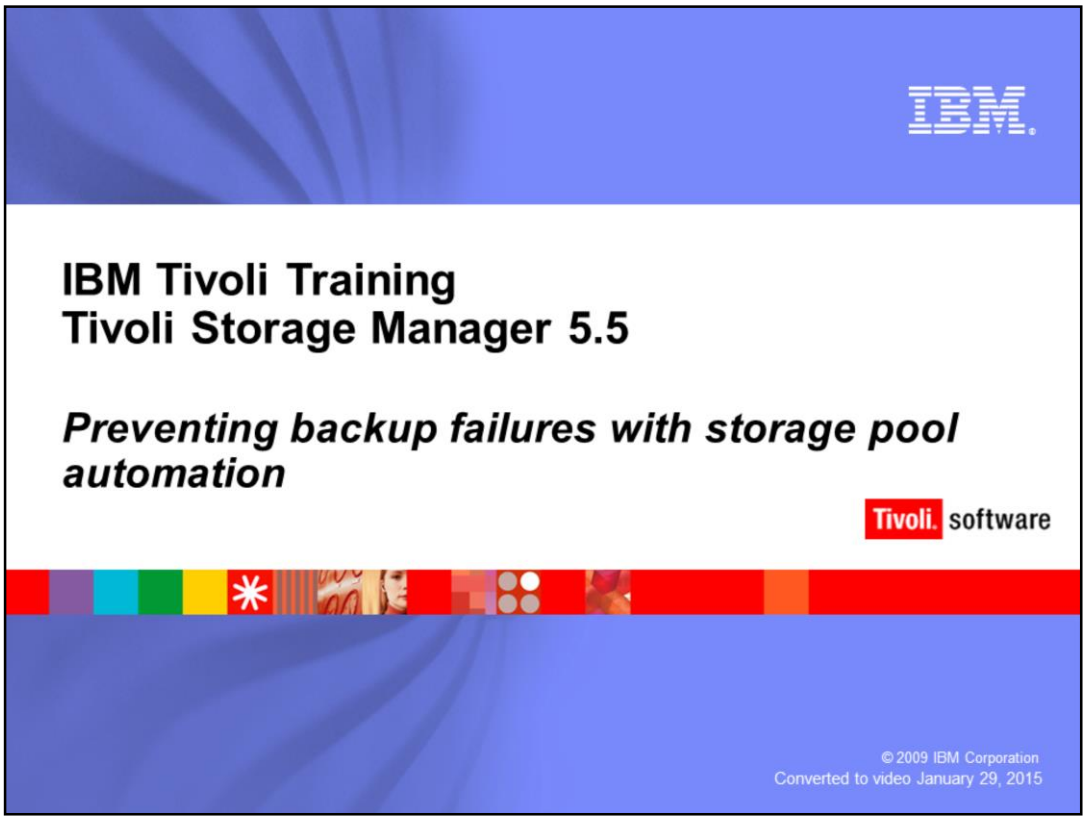

Slide 1

Welcome to the IBM Education Assistant training for IBM Tivoli Storage Manager version 5.5. This module covers the prevention of backup failures by ensuring there is available space with the automation of storage pool maintenance tasks**.** **I**n this training, Tivoli Storage Manager is also referred to as TSM.

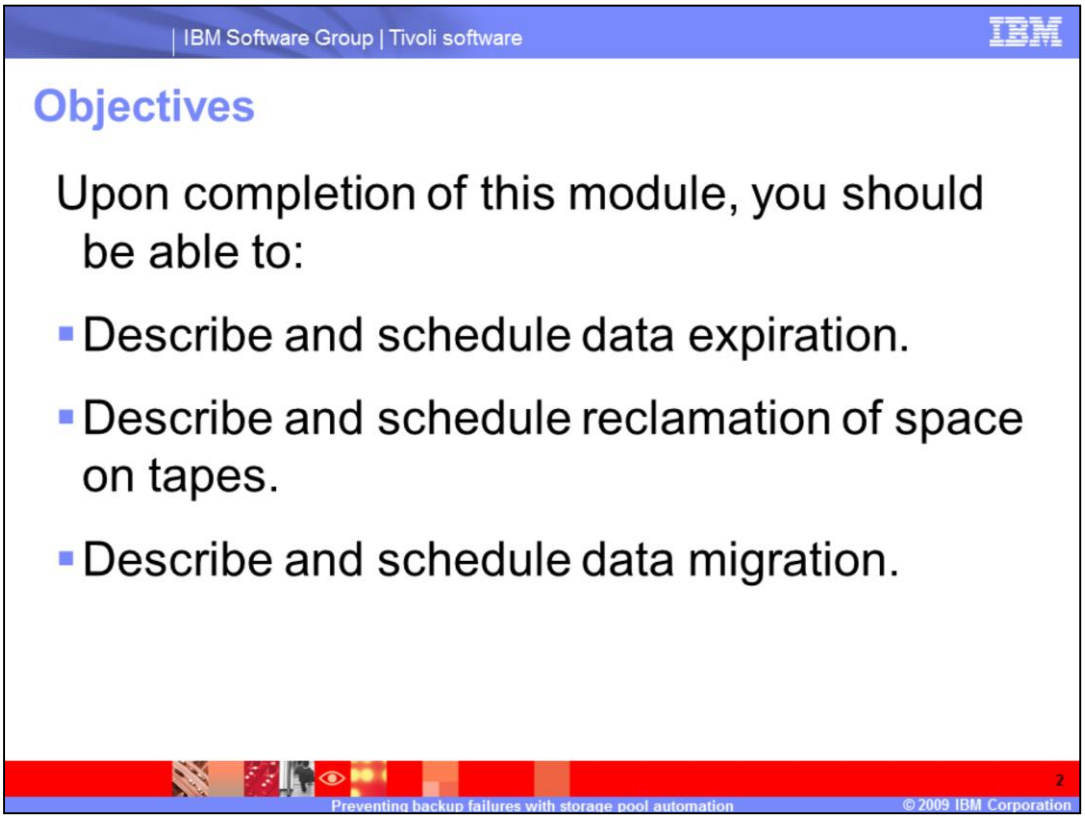

Slide 2 **Objectives**

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

Describe and schedule data expiration.

Describe and schedule reclamation of space on tapes.

Describe and schedule data migration.

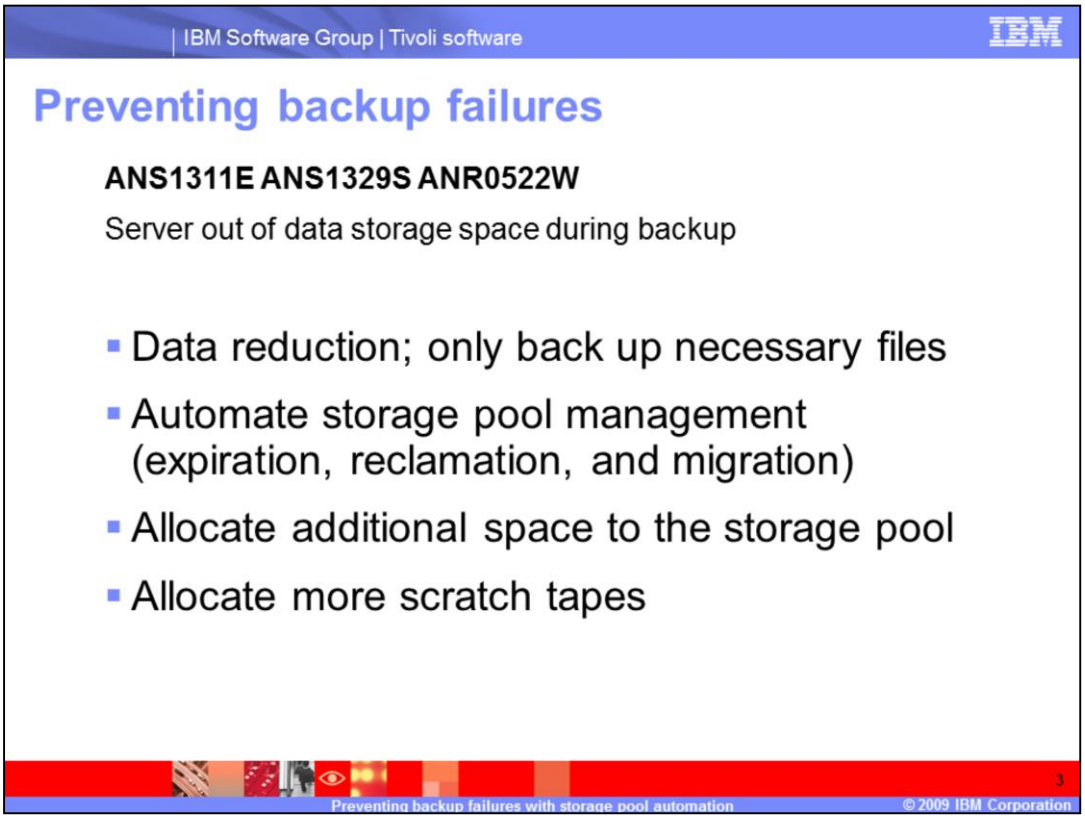

Slide 3 **Preventing Backup Failures**

Common error messages seen for backup failures, are ANS1311E, ANS1329S, and ANR0522W. These are for Server out of data storage space during backup.

Various issues such as lack of storage pool space, lack of tapes, hardware problems or configuration problems can cause these messages .

To prevent the Server out of data storage space errors, consider the following actions:

•Data reduction; only back up necessary files

•Automate storage pool management such as expiration, migration, and reclamation

•Allocate additional space to the storage pool

•Allocate more scratch tapes

This IBM Education Assistant module covers the prevention of backup failures by ensuring there is available space with the automation of storage pool maintenance tasks**.**

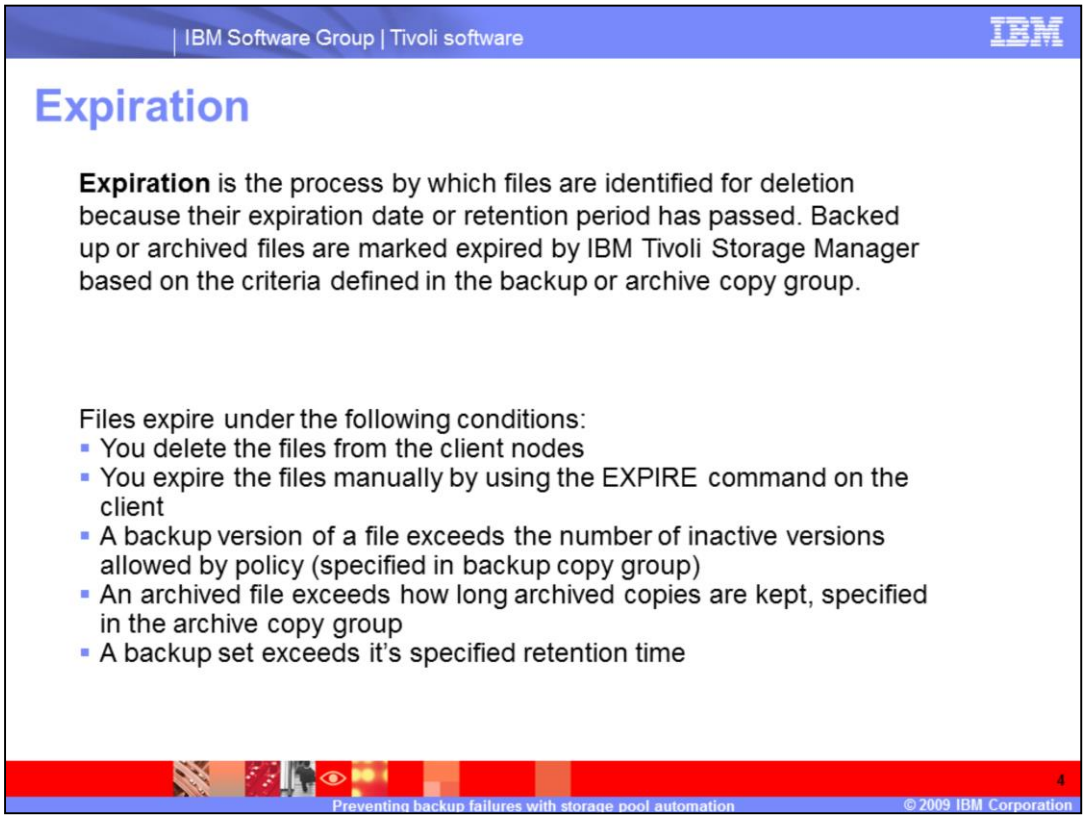

Slide 4 **Expiration**

Expiration is the process by which files are identified for deletion because their expiration date or retention period has passed. Backed up or archived files are marked expired by IBM Tivoli Storage Manager based on the criteria defined in the backup or archive copy group.

Files expire under the following conditions:

You delete the files from the client nodes

You expire the files manually by using the EXPIRE command on the client

A backup version of a file exceeds the number of inactive versions allowed by policy (specified in backup copy group)

An archived file exceeds how long archived copies are kept, specified in the archive copy group

A backup set exceeds its specified retention time

Note: A base file is not eligible for expiration until all of its dependent subfiles have been expired.

Also, an archive file is not eligible for expiration if there is a deletion hold on it. If a file is not held, it will be handled according to existing expiration processing.

The server deletes expired files from the server database only during expiration processing.

After expired files are deleted from the database, the server can reuse the space in the storage pools that was occupied by expired files. You should ensure that expiration processing runs periodically to allow the server to reuse space.

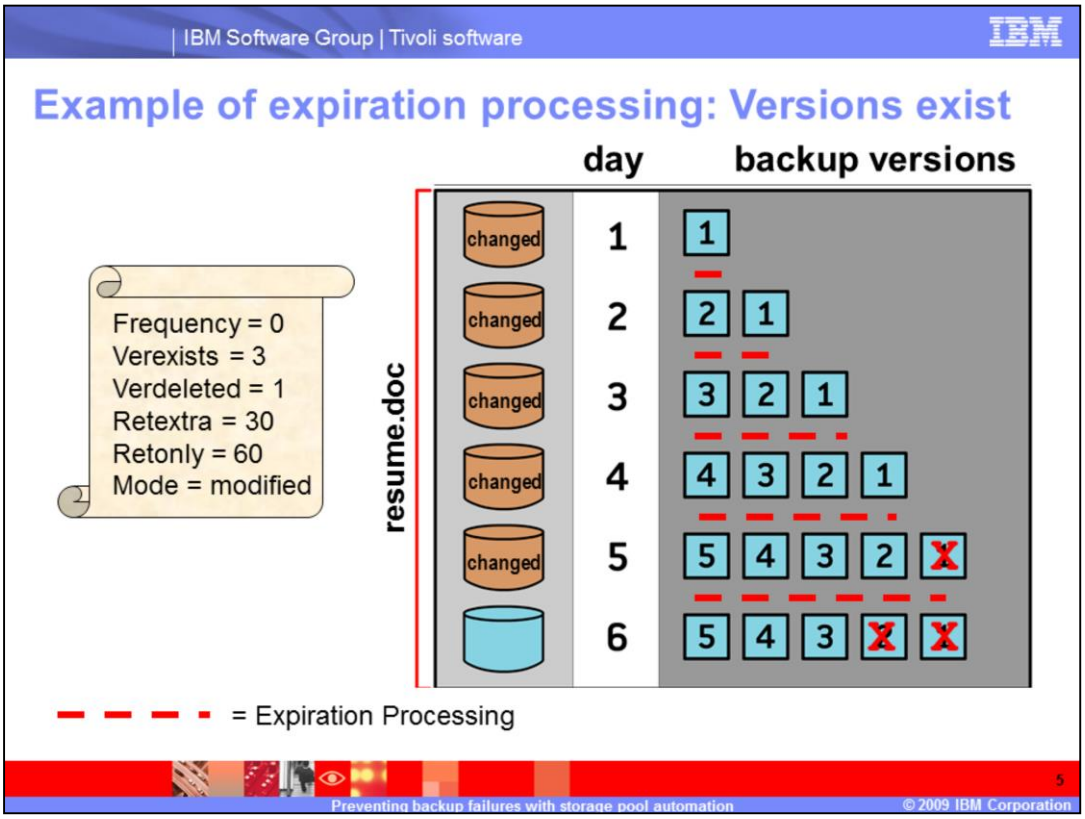

Slide 5 **Example of Expiration Processing: Versions Exist**

The slide shows the life cycle of a backed up client file as it is stored in a TSM storage pool.

A scheduled incremental backup is run every evening, followed by expire inventory.

The copy group values for this file are:

– Frequency = 0

– Verexists = 3 (Maximum)

– Verdeleted = 1

– Retextra = 30 (File exists on workstation)

– Retonly = 60 (File deleted from workstation)

– Mode = modified

On day 1 the file is created. That evening it is backed up and expire inventory is run. The file would be marked as active.

On day 2, the file is changed, backed up, and expiration is run. The backup for day 2 becomes active. The backup from day 1 becomes inactive.

On day 3, the file is changed, backed up, and expiration is run. The backup for day 3 becomes active. The backups from days 1 and 2 become inactive. Since Versions Exist is set to 3, no files have expired and nothing has been removed during expiration processing.

On day 4, the file is changed, backed up, and expiration is run. The backup for day 4 becomes active. The backups from days 1, 2, and 3 become inactive. Because we are only keeping three versions of the file, the backup from day 1 (the 4th backup) expires and is removed during expiration processing.

On day 5, the file is changed, backed up, and expiration is run becomes the active file. The backups from days 1, 2, and 3 become inactive. Because we are only keeping three versions of the file, the backup from day 2 also expires and is removed during expiration processing.

On day 6, the file is not changed or backed up. Because we have three versions of the file after expiration on day 5, there is no change. The file from day 5 is active and the files from days 4 and 3 are inactive. We have 3 versions of the file.

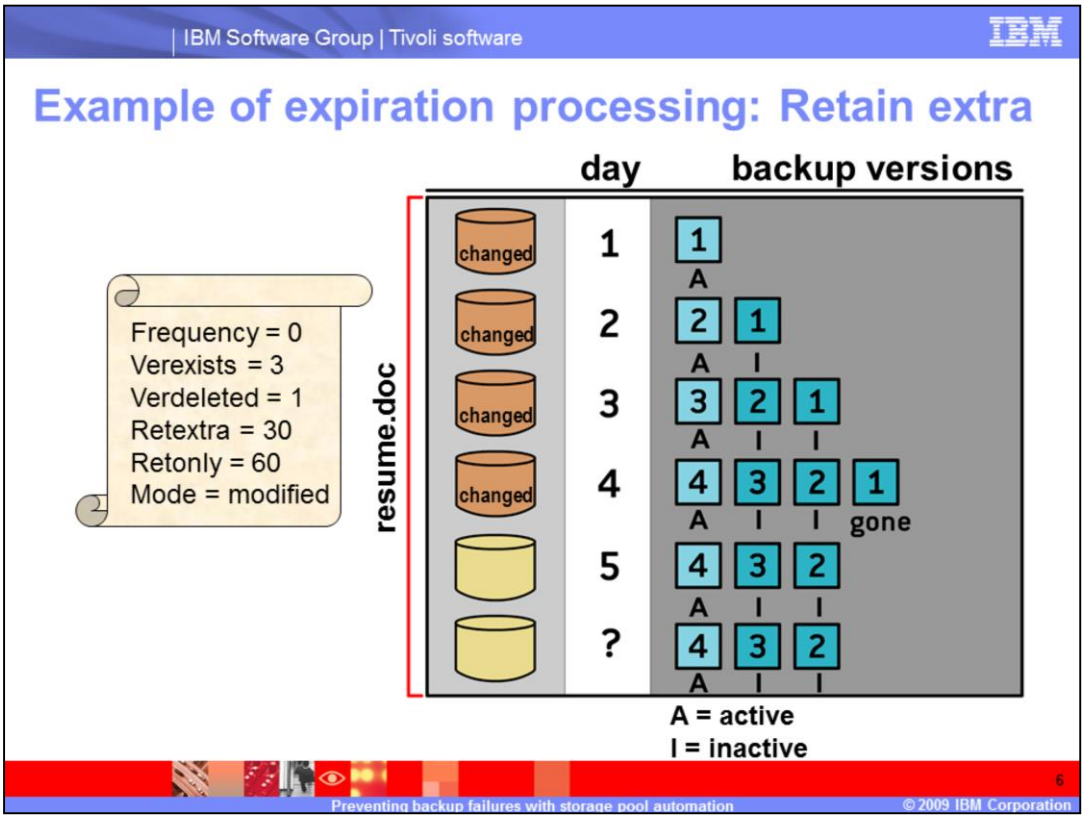

Slide 6 **Example of Expiration Processing: Retain Extra**

We are using the same copy group parameters as previously.

A scheduled incremental backup is run every evening, followed by expire inventory.

Note that backup 1 on day 4 is gone because we are keeping only 3 versions based on VERExist=3. Notice, too, that the file goes unchanged after day 4. In this case, can you predict the results for day 33?

On day 33, only one version remains. This is because RETAINEXTRA is set to 30. If you always want three versions RETAINEXTRA needs to be NO LIMIT.

Can you predict when backup 2 would expire? Backup 2 would expire on day 33 because it became inactive on day 3 and you would add 30 days based on RETEXTRA=30.

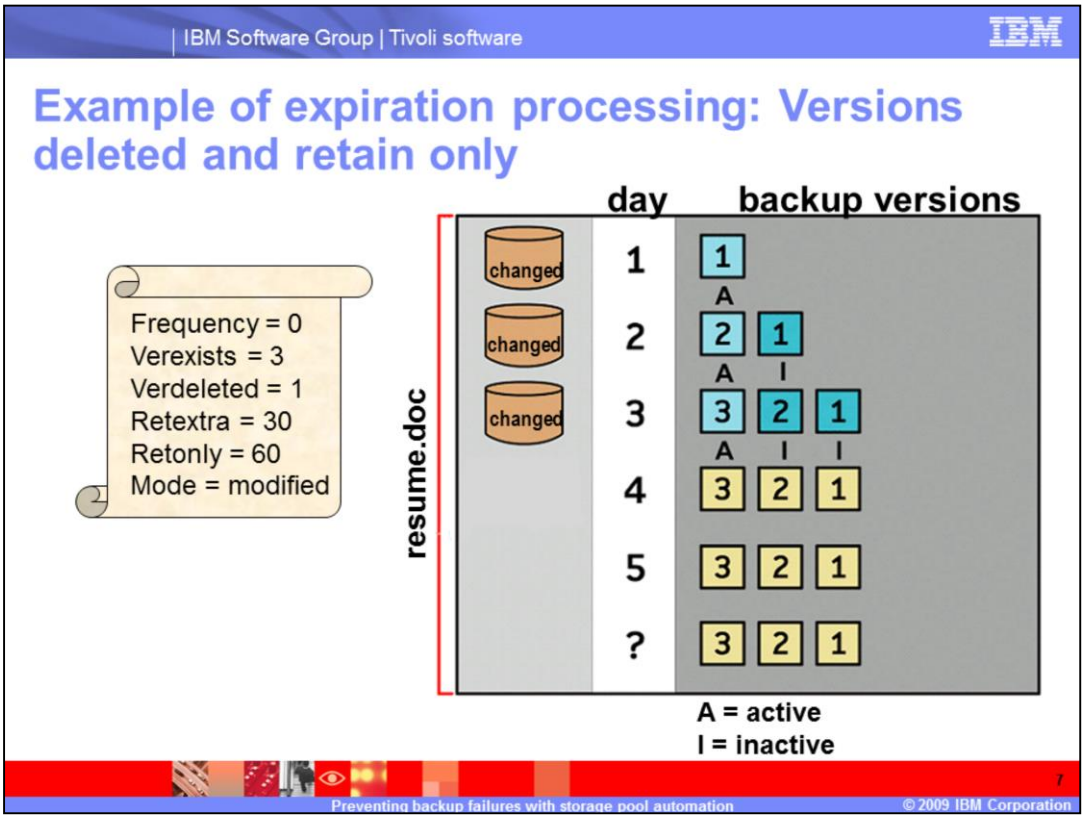

Slide 7 **Example of Expiration Processing: Versions Deleted and Retain Only**

We are using the same copy group parameters as previously.

A scheduled incremental backup is run every evening, followed by expire inventory.

What happens if the file is removed from the client workstation on the morning of day 4?

Backup 3 will become inactive if the file is deleted.

How many extra copies will be there after the incremental backup on day 4?

To determine this, notice that VERDELETED is set to 1. This means that

backups 1 and 2 will be removed.

Can you predict when the only backup will expire and be removed?

The only backup will expire and be removed on day 64 because RETONLY is set to 60.

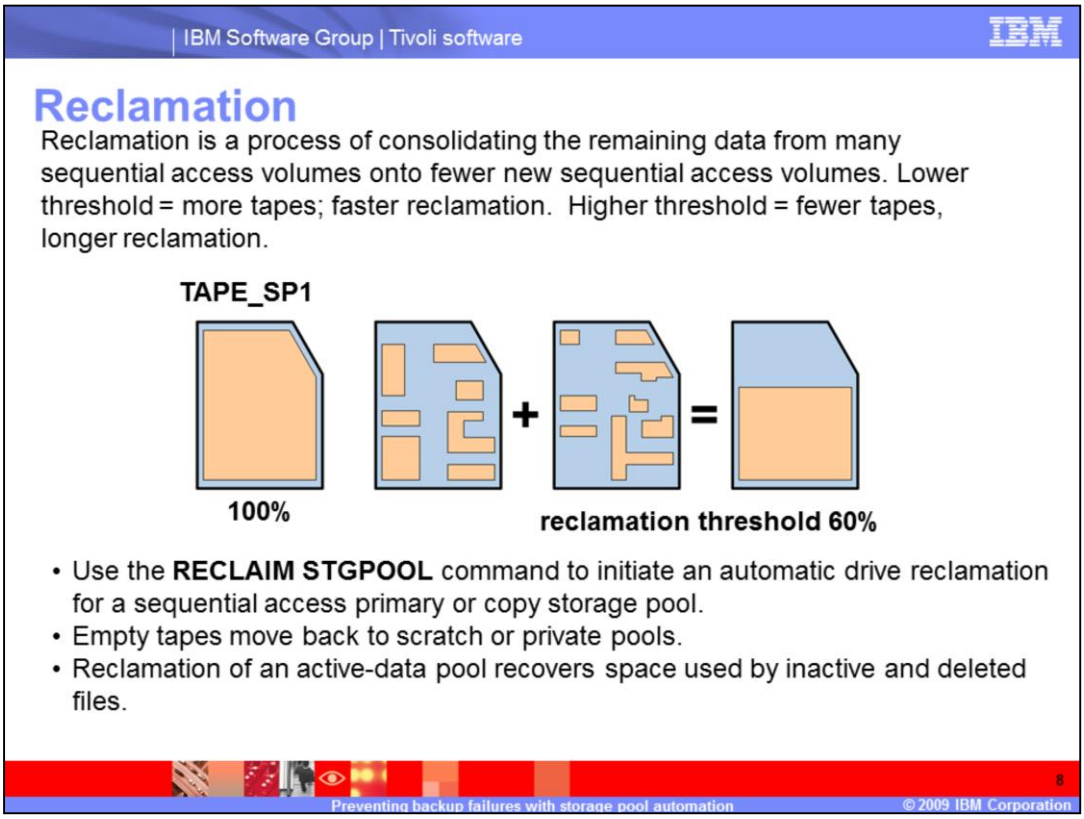

Slide 8 **Reclamation**

Reclamation is a process of consolidating the remaining data from many sequential-access volumes onto fewer, new sequential-access volumes. A lower threshold results in more tapes and faster reclamation. A higher threshold results in fewer tapes and longer reclamation.

When the percentage of reclaimable space on a volume exceeds the reclamation threshold set for the storage pool, the volume is eligible for reclamation. The server checks whether reclamation is needed at least once per hour and begins space reclamation for eligible volumes. You can set a reclamation threshold for each sequential-access storage pool when you define or update the pool. During space reclamation, the server copies files that remain on eligible volumes to other volumes.

The reclaim stgpool command enables you to manually drive reclamation for a sequential-access primary or copy storage pool.

When multiple volumes are eligible for reclamation, Tivoli Storage Manager reclaims the eligible volumes in random order.

Space within aggregate files is also reclaimed during the reclamation process. An aggregate is a physical file that contains multiple logical files backed up or archived from a client in a single transaction. Unused space from expired or deleted logical files is removed as the aggregate file is copied to another volume during reclamation.

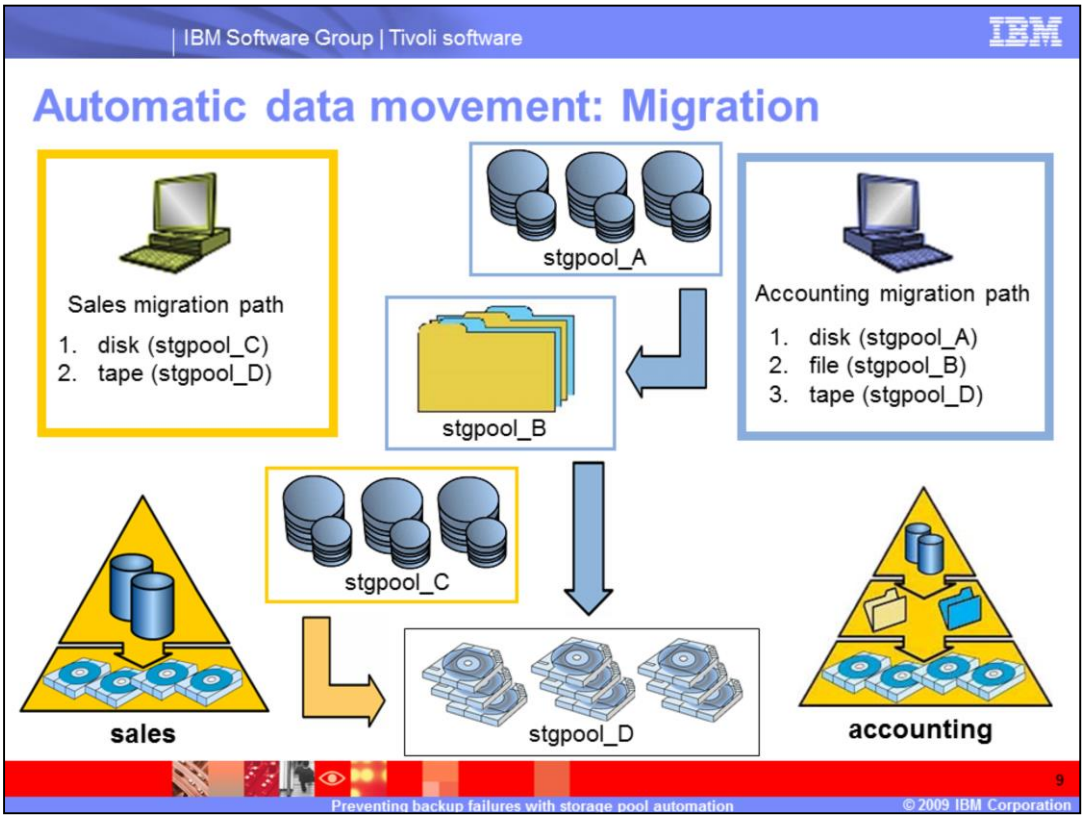

Slide 9 **Automatic Data Movement: Migration**

Automatic data movement between storage pools is used to balance the performance and cost of different storage devices while ensuring adequate free space to satisfy new space allocations. This process is known as migration.

For each storage pool, you define low and high migration thresholds. Migration thresholds are based on a percentage of the storage pool's total data capacity.

The low threshold identifies the amount of free space needed to satisfy the daily processing requirements of your business.

The high threshold is used to trigger migration and ensure that enough free space is available while migration is performed. The difference between the high and low thresholds indicates the approximate amount of data that will be migrated.

To reduce tape mounts and to use the space on tape volumes most effectively (when collocation is not used), ensure that the amount of data that is migrated from a disk storage pool is a multiple of the capacity of a tape volume in the next storage pool.

This slide shows two scenarios. On the left, the Sales department's data is written first to disk (storage pool C, the primary storage pool), then migrates to tape (storage pool D). On the right, the Accounting department's data is backed up to disk (storage pool A, the primary storage pool), then migrates to file (storage pool B), then migrates to tape (storage pool D).

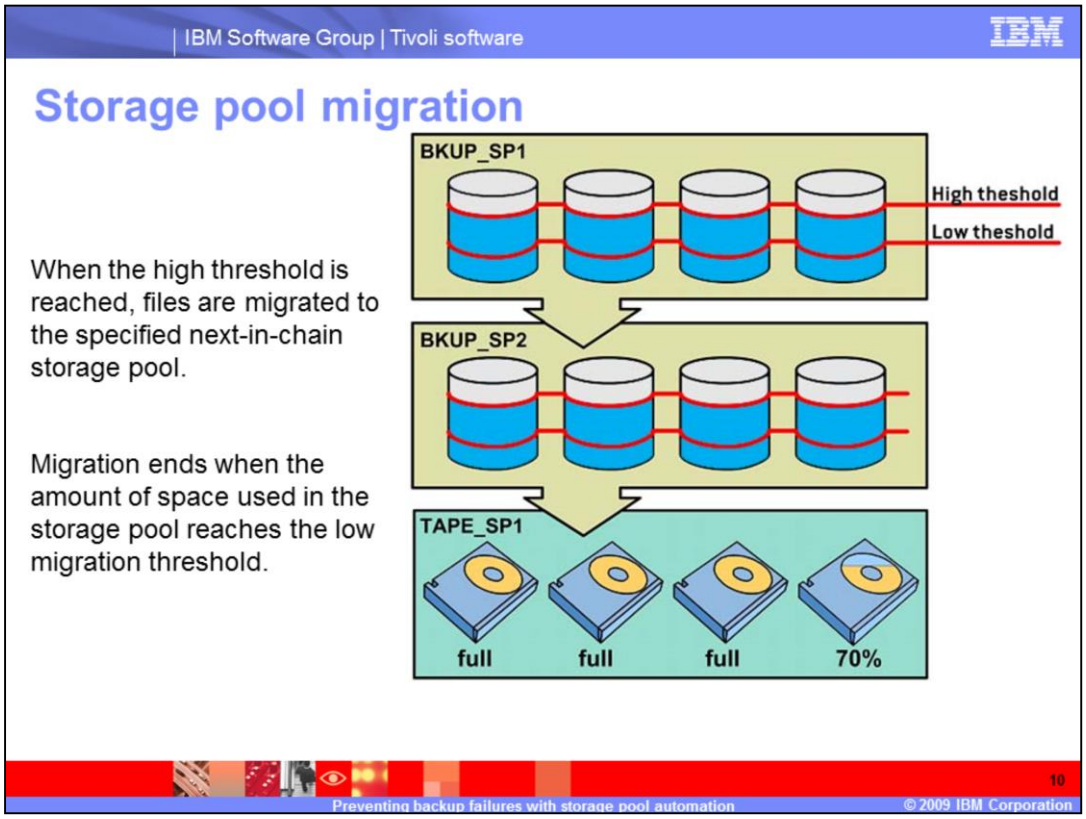

Slide 10 **Storage Pool Migration**

No migration occurs if there is no next storage pool. Tivoli Storage Manager first identifies which client node has backed up or migrated the largest single file space, or has archived files that occupy the most space. When the server identifies that client node, the server migrates all files from every file space belonging to that client.

When the high threshold is reached, files are migrated to the specified next-in-chain storage pool. Migration ends when the amount of space used in the storage pool reaches the low migration threshold.

The migration applies to those files whose number of days in the storage pool exceeds the value specified by the migration delay, or MIGDELAY parameter, specified during the define or update storage pool commands.

After the files for the first client node are migrated to the next storage pool, the server checks the low migration threshold for the storage pool to determine if the migration process should be stopped. If the amount of space used in the storage pool is now below the low migration threshold, migration ends. If not, using the same criteria as previously described, Tivoli Storage Manager chooses another client node, and the migration process continues.

If the value for migration continue, or MIGCONTINUE, has been set to YES, then Tivoli Storage Manager continues the migration process based on how long the files have been in the storage pool. This is also specified during the define or update storage pool commands.

The oldest files are migrated first until the low migration threshold is reached. If the value for MIGCONTINUE has been set to NO, then the migration process ends, and a warning message will be issued to the administrator.

If multiple migration processes are running (controlled by the MIGPROCESS parameter of the define stgpool command), the files for more than one node may be chosen for migration at the same time.

If the cache option is enabled during the define or update storage pool commands, files that are migrated remain on disk storage (that is, the files are cached) until space is needed for new files. You can enable caching by specifying CACHE=YES when you define or update a disk storage pool. When caching is enabled, the migration process leaves behind duplicate copies of files on disk after the server migrates these files to subordinate storage pools in the storage hierarchy.

The copies remain in the disk storage pool, but in a cached state, so that subsequent retrieval requests can be satisfied quickly. However, if space is needed to store new data in the disk storage pool, cached files are erased and the space they occupied is used for the new data.

The advantage of using a cache for a disk storage pool is that caching can improve how quickly the server retrieves some files. When you use a cache, a copy of the file remains on fast disk storage after the server migrates the primary file to another storage pool. You may want to consider using a disk storage pool with caching enabled for storing space-managed files that are frequently accessed by clients.

However, using a cache has some important disadvantages, such as:

• It can increase the time for client backup operations to complete.

• It can require more space for the Tivoli Storage Manager database.

The migrate stgpool command allows you to manually run migration for a random-access or sequential-access primary storage pool.

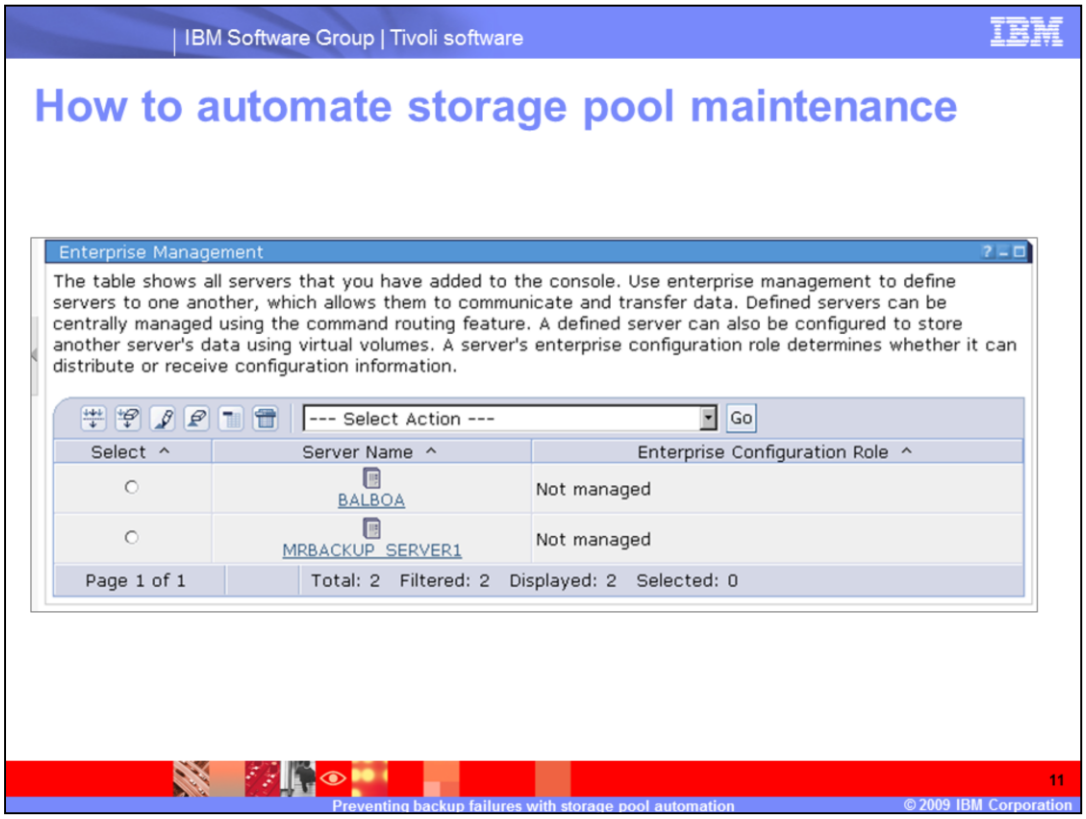

Slide 11 How to Automate Storage Pool Maintenance

Select **Server Maintenance** in the Tivoli Storage Manager tree, and the Server Maintenance page will display as shown.

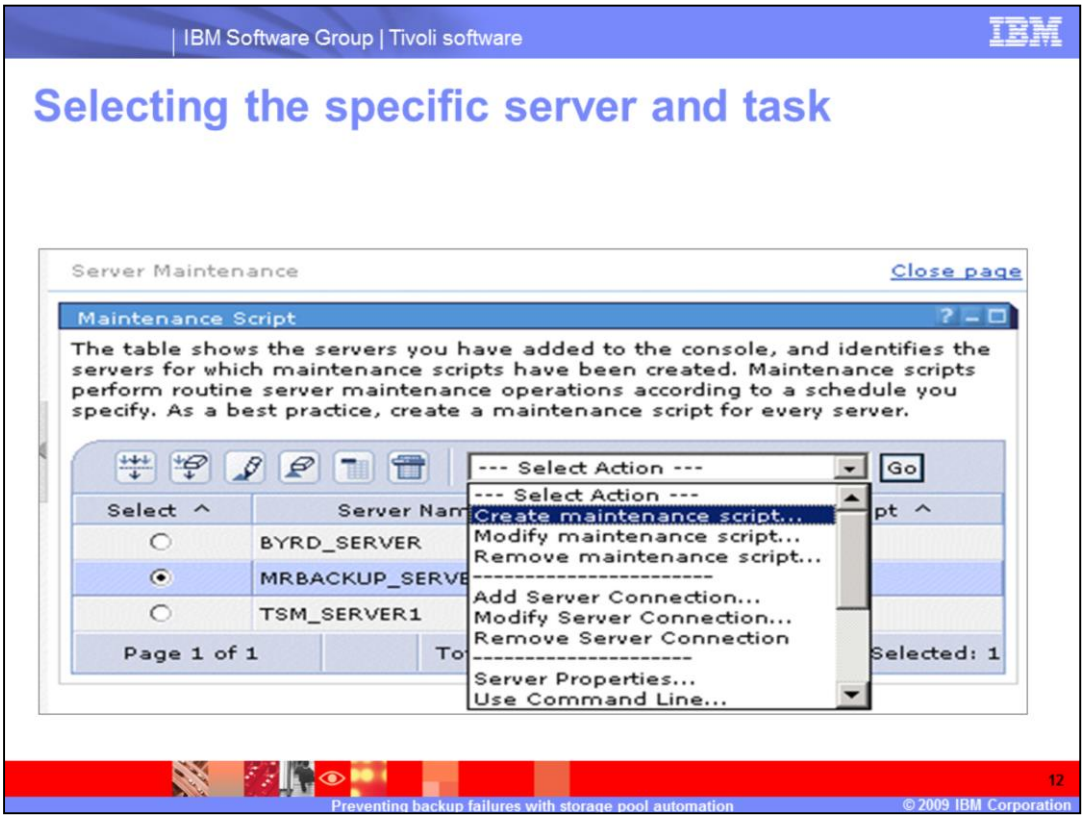

Slide 12 **Selecting the Specific Server and Task**

Select the server you want to make the script for, and from the Select Action menu select **Create maintenance script**. The options to modify and to remove a script are also available from the same menu.

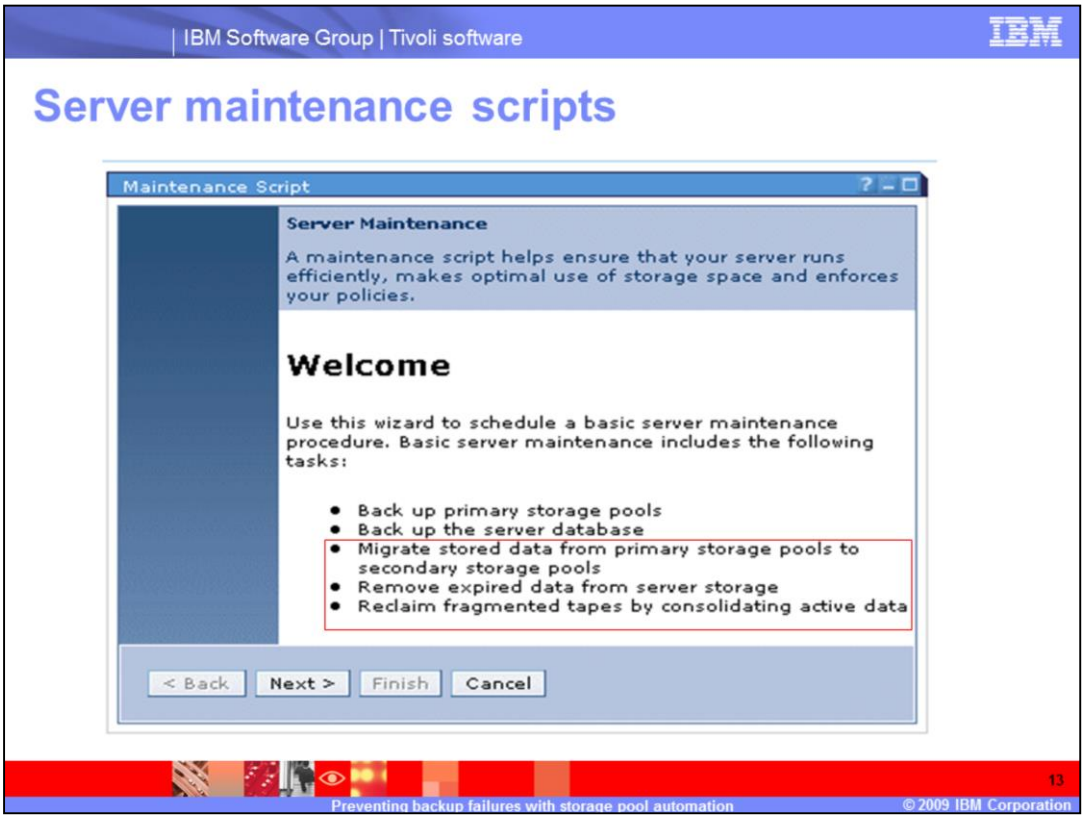

Slide 13 **Server Maintenance Scripts**

While backing up the server database and storage pools are vital server maintenance processes, this training module is focused on stored data migration, expiration, and reclamation.

Here, the Server Maintenance wizard appears. Click Next.

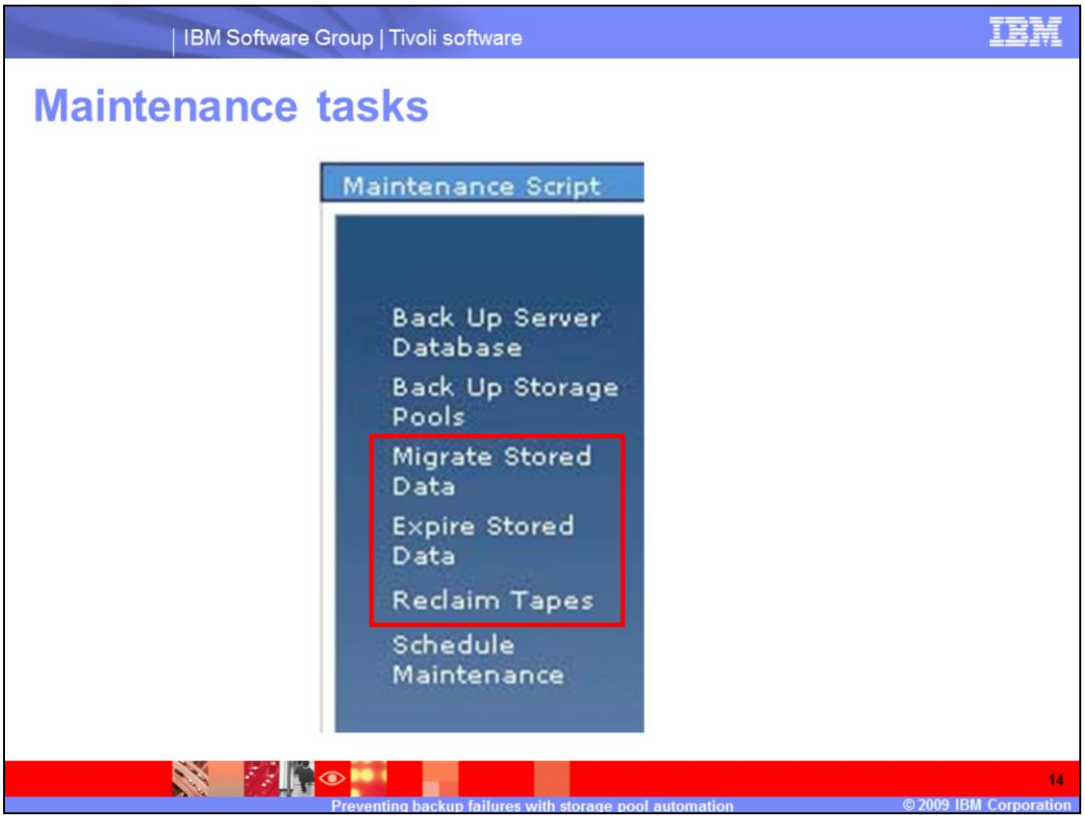

Slide 14 Maintenance Tasks

This is a close up of the left pane of the wizard. It shows the sequence of processes that will be scheduled.

Click Next to progress through the screens.

Slide 15 Migrate Stored Data

Select the storage pool to migrate, set the migration threshold, and set the amount of time the process can run before it is canceled. Click **Next** to continue.

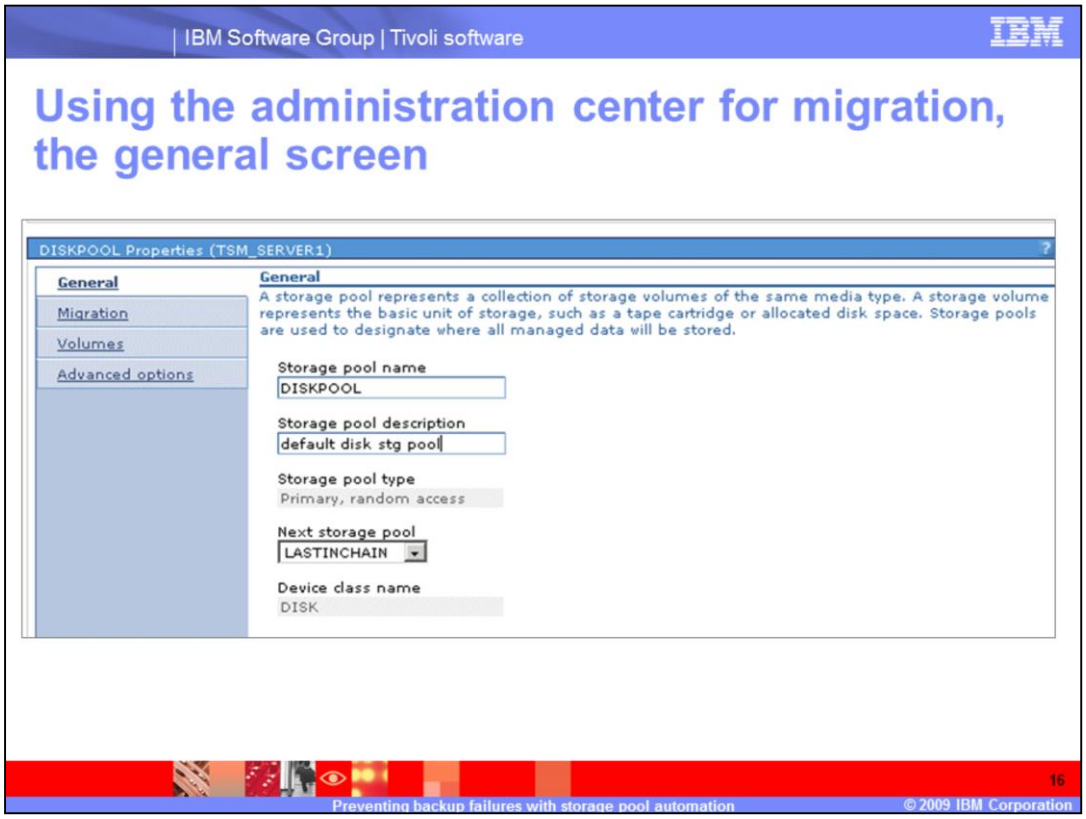

Slide 16 **Using the Administration Center for Migration, the General Screen**

Select the Next storage pool.

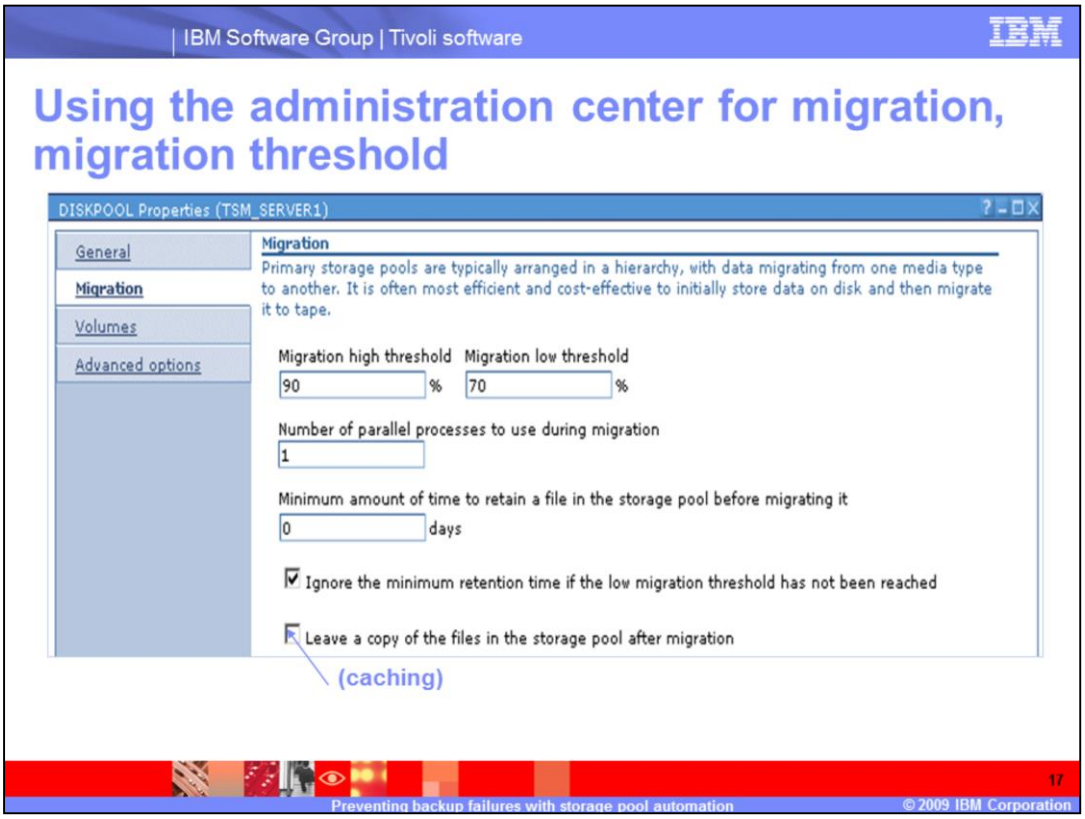

Slide 17 **Using the Administration Center for Migration, Migration Threshold**

Click Storage Devices in the Tivoli Storage Manager tree.

In the Servers table, select a server.

Click Select Action, select View Storage Pools.

In the server's storage pools portlet, click a storage pool name.

In the storage pool's properties notebook, click the Migration tab.

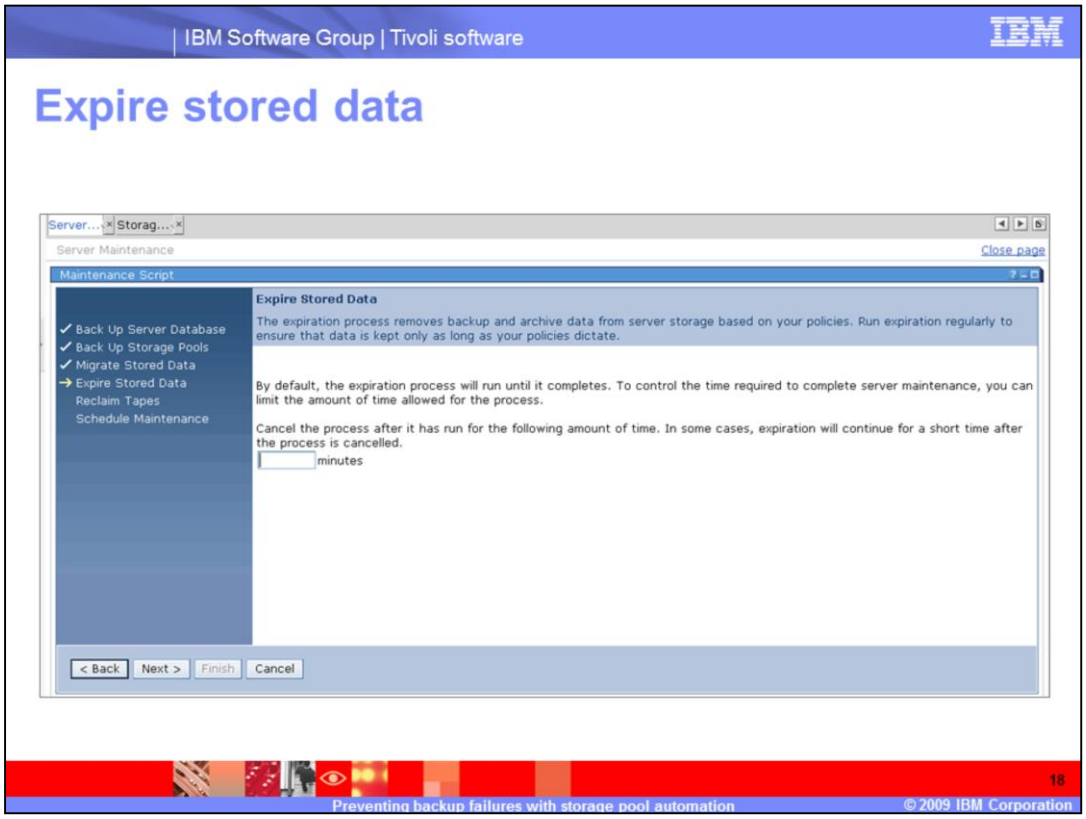

Slide 18 **Expire Stored Data**

If you want to limit the amount of time allowed for the expiration process, set that value here, and click **Next** to continue.

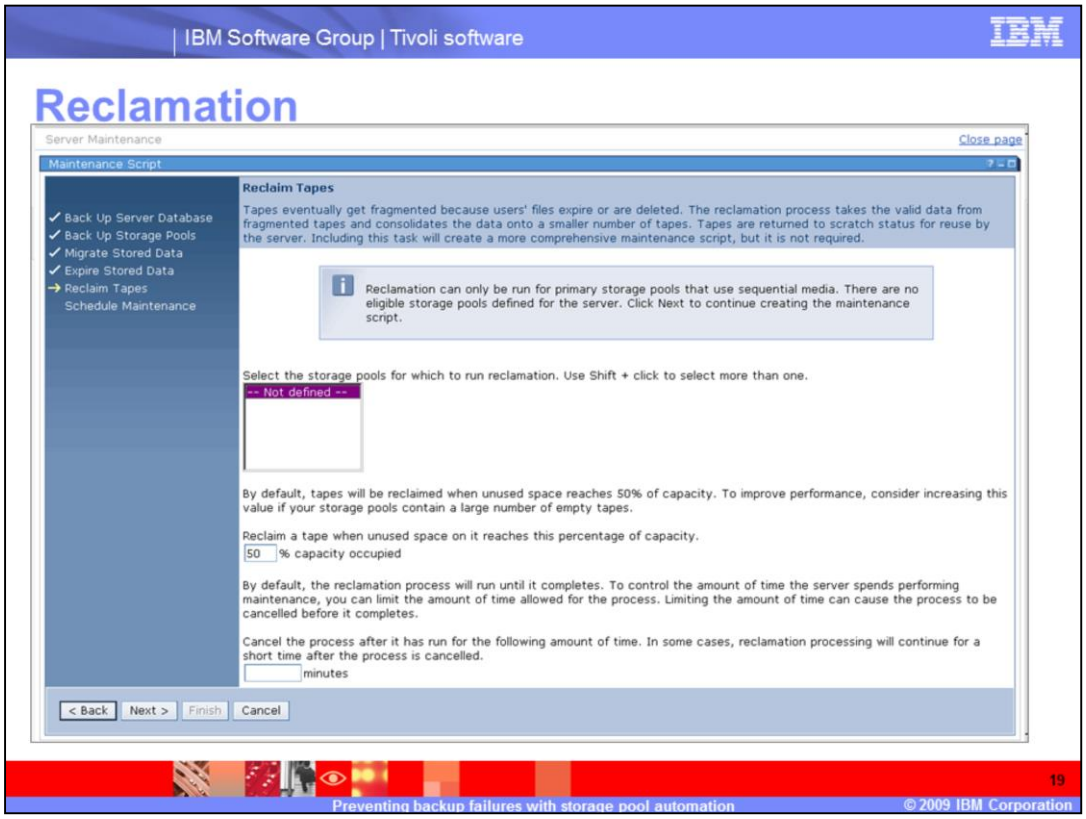

Slide 19 **Reclamation**

Select the storage pool to run reclamation for, set the percent of unused space which will be the trigger to start the reclamation process, and set the amount of time the process is allowed to run.

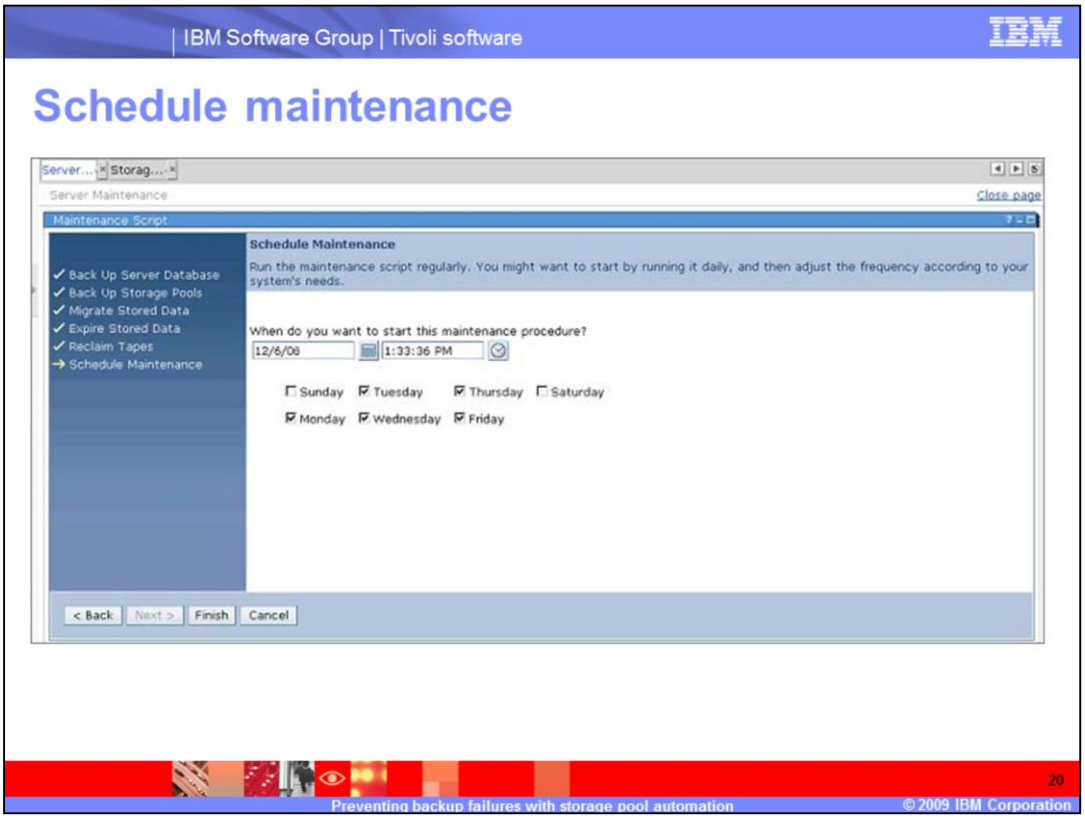

Slide 20 **Schedule Maintenance**

Set the schedule for the script, including date and time, and click **Finish**.

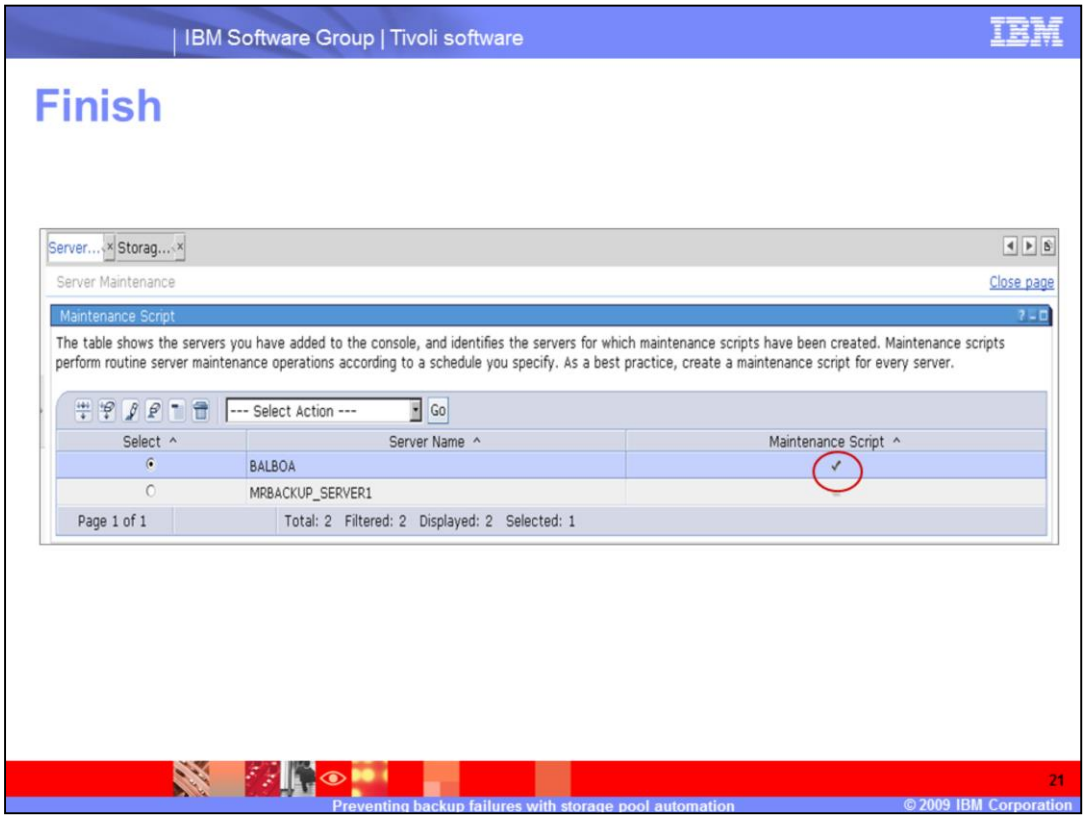

Slide 21 **Finish**

The check mark under the Maintenance Script heading indicates a maintenance script has been created for this server.

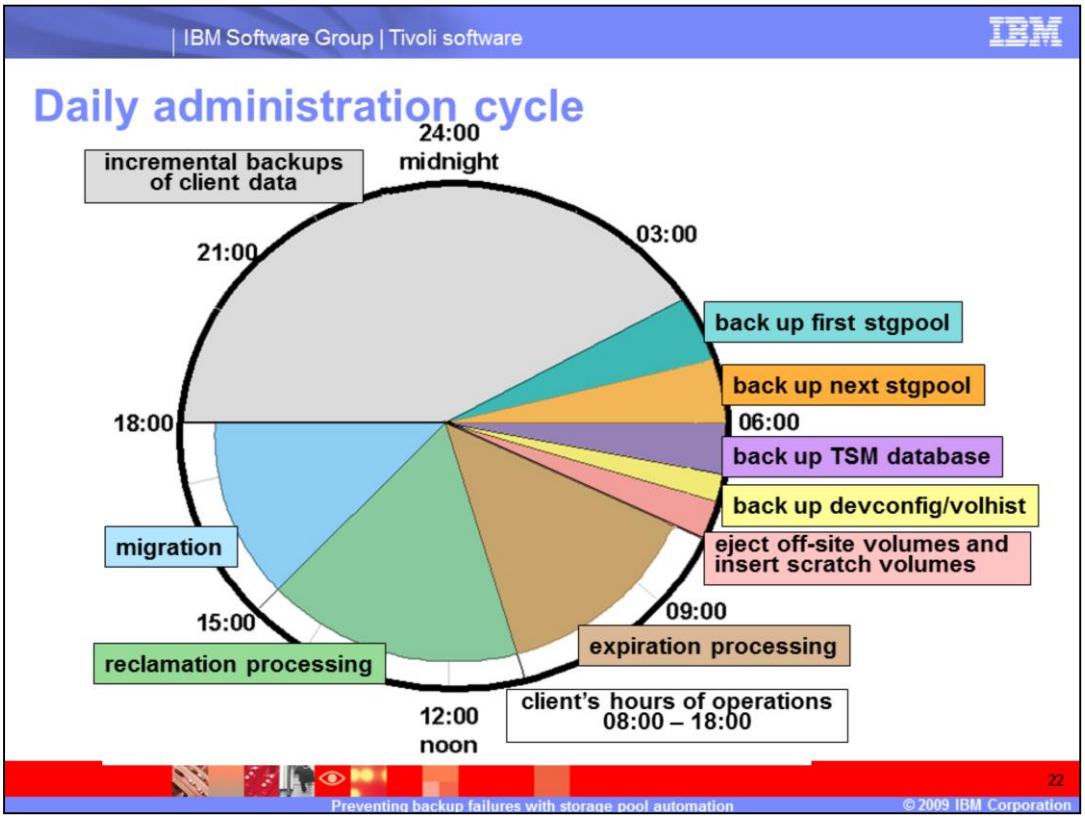

Slide 22 **Daily Administration Cycle**

This is one example of the daily operations of an IBM Tivoli Storage Manager administrator. In this example, data files are backed up during the night. Once the backups are complete, the administrative tasks begin.

The storage pools are backed up in order, starting with the primary storagepool. Next, the database is backed up, followed by the devconfig and volhist files being backed up.

The administrator would then remove off-site volumes and insert the scratch volumes.

Finally, expiration, then reclamation, and finally, migration are run.

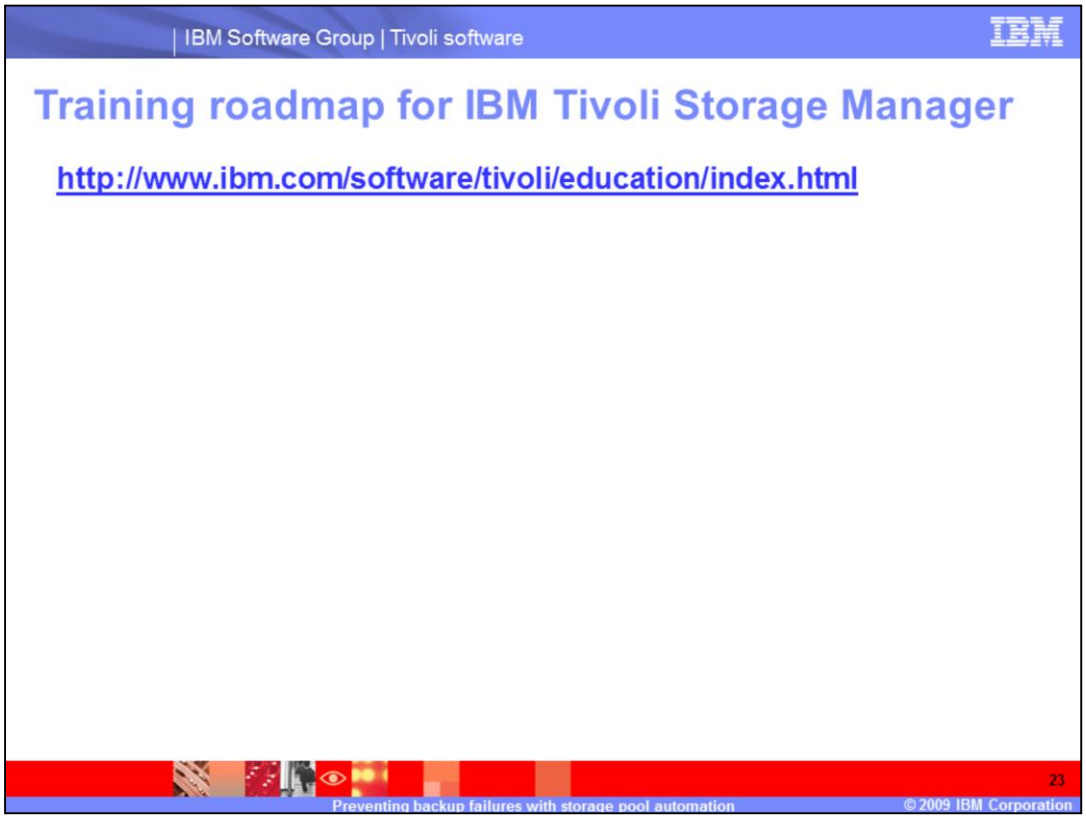

Slide 23 **Training Roadmap for *IBM Tivoli Storage Manager***

If you go to www.ibm.com/software/tivoli/education/edu_prd.html this will take you to the training page, where you can access the training list for Tivoli Storage Manager version 5.5.

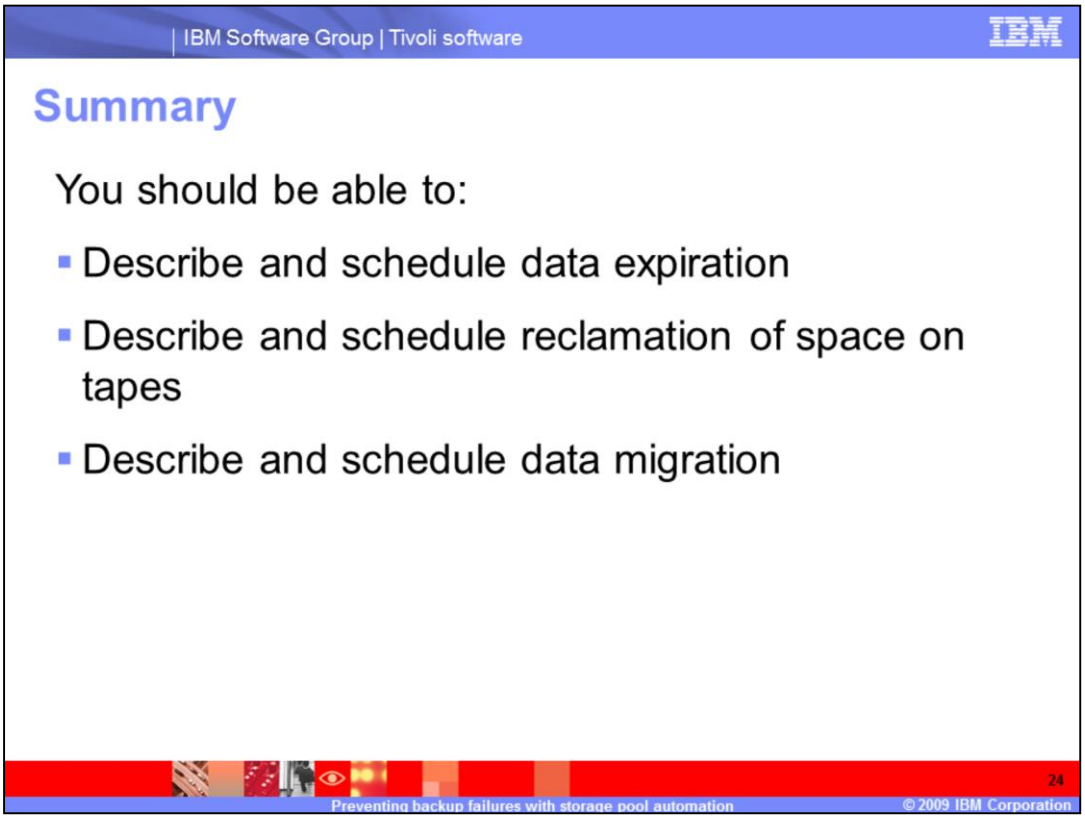

Slide 24 **Summary**

You should be able to:

Describe and schedule data expiration.

Describe and schedule reclamation of space on tapes.

Describe and schedule data migration.

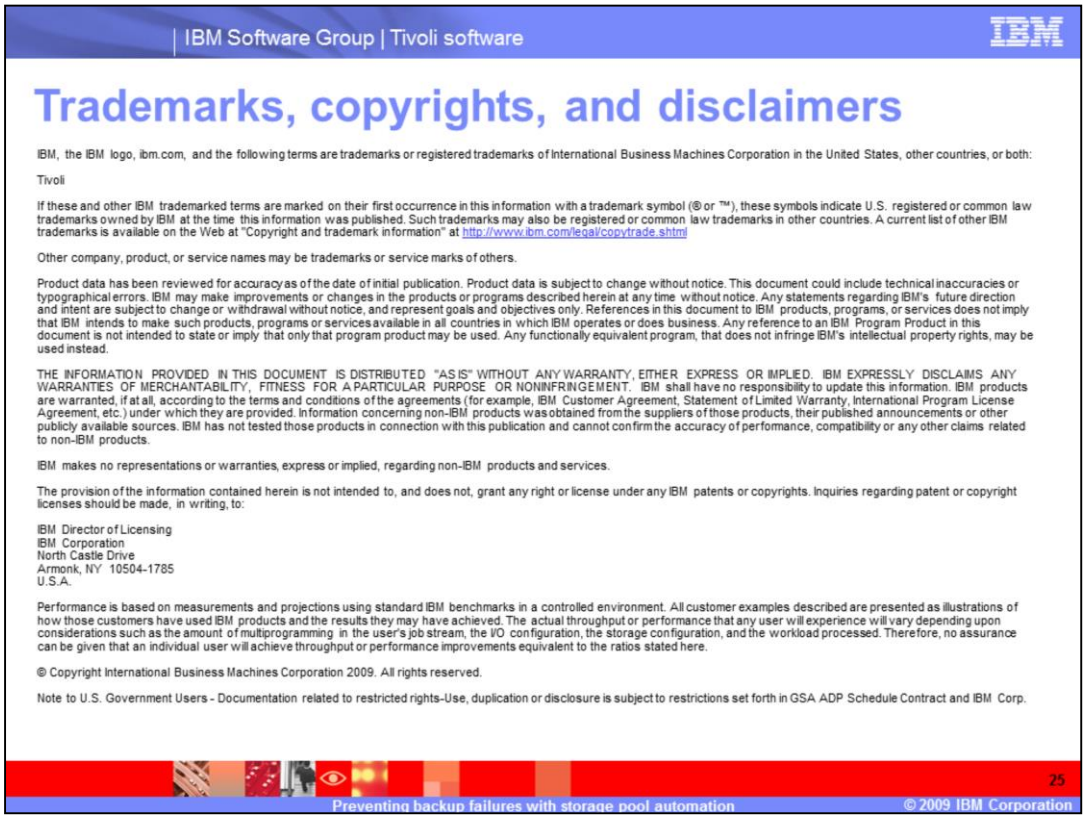

Slide 25 (Copyright)

This concludes the IBM Education Assistant training for IBM Tivoli Storage Manager version 5.5, the prevention of backup failures by ensuring there is available space with the automation of storage pool maintenance tasks**.**