



FalconStor® Virtual Tape Library (VTL) Migration Scenarios



Integrating Data De-duplication into Existing Data Protection
Environments

By Jason Giza of Hitachi Data Systems

December 2010

Table of Contents

Executive Summary	3
Introduction	4
Scenario 1: Transparent Migration Using Tape Caching with Import	4
Configuration Steps	5
Scenario 2: Transparent Direct Link Access	24
Configuration Steps	24
Conclusion	25
Appendix A: About FalconStor	26

Executive Summary

Data de-duplication is an increasingly important feature that organizations are looking for when purchasing storage technology in order to minimize their storage and replication footprints and reduce operational costs. Several years ago, FalconStor Software introduced an integrated de-duplication engine into its FalconStor® Virtual Tape Library (VTL) solution. However, for companies that rely upon older VTL solutions based on FalconStor technology as part of their data protection plan, adding de-duplication can be a difficult task. It sometimes necessitates a wholesale overhaul and replacement of equipment. Hitachi Data Systems has recognized the need for a more cost-efficient solution for these organizations. It has partnered with FalconStor to create an easily adaptable process to integrate data de-duplication into these environments, leveraging existing investments without adding complexity to the management layer or changing operational processes.

Partnering with FalconStor, Hitachi Data Systems provides a unique solution that can augment an existing infrastructure, updating it with enhanced functionality, including de-duplication with minimal disruption. Working alongside FalconStor engineers (who developed this seamless functionality), Hitachi Data Systems has cultivated two options for migrating data from older VTL installations to FalconStor VTL with integrated de-duplication technology. In comparison, competitive solutions largely center on a straight "rip and replace" strategy, employing forklift migration that is disruptive, changing the operational processes and possibly requiring substantial changes to the infrastructure. With the flexibility provided by the processes outlined in this document, storage or backup administrators are able to implement custom migration plans. They can introduce the benefits of this new technology immediately, while still maintaining the value of existing investments in technology, processes, and knowledge and training of personnel.

Each migration scenario detailed in this document leverages FalconStor tape caching functionality to enable communication between VTL appliances. This allows for transparent management of the old appliance (and data residing on it) using the interface of the newer FalconStor VTL solution. The result is a simple deployment, seamless data migration without downtime and the ability to decommission the old hardware at any time, without interrupting backup operations.

Hitachi Data Systems worked hand-in-hand with FalconStor in revising its software to provide the functionality to make these data migration plans as customizable as possible. Once the software was released with the requested changes, Hitachi Data Systems personnel tested these seamless migration scenarios in the FalconStor labs (Melville, NY) with FalconStor VTL appliances and enterprise backup software.

Introduction

This application brief is designed to illustrate the methods of introducing a FalconStor VTL with integrated de-duplication into an environment with an existing VTL. It outlines two migration plans and the step-by-step configuration for the new FalconStor VTL solution in each case. Both of these migration scenarios allow for a seamless integration of the new FalconStor VTL with de-duplication, but the difference between the two lies in how data from the older VTL is retained. The first scenario nondisruptively imports data in the background from the older VTL while maintaining transparent access to that device for retrieval of data that is yet to be migrated. As soon as all of the data that is desired to be retained has been migrated from the older VTL, that appliance can be removed or retained for long-term retention. The second scenario also provides transparent access to the older VTL, but it keeps that device in place for the lifecycle of the data stored on it. When that data passes beyond the retention period based on backup policies, the older VTL appliance can be removed or retained for long-term retention.

(Note: These migration scenarios assume the new FalconStor VTL is equipped with FalconStor VTL v5.21, which includes "sync library" for migration in tape caching environments. Without the proper code update, some of the functionality needed to achieve these migration scenarios will not be available.)

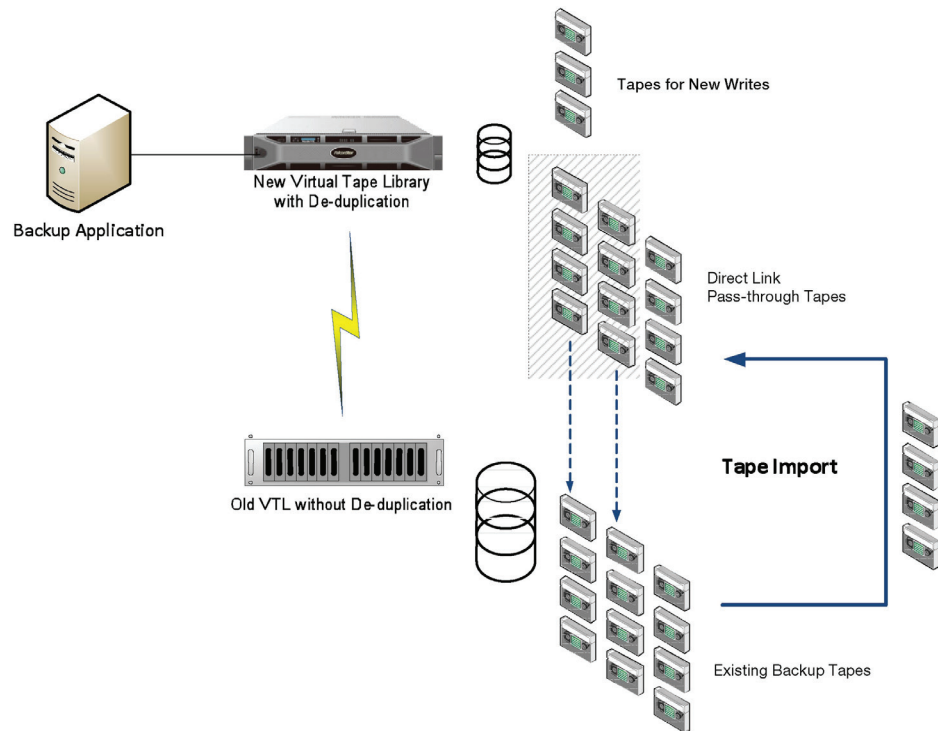
Scenario 1: Transparent Migration Using Tape Caching with Import

In this migration scenario, the old VTL remains in place (until data migration is complete) with its data intact, but its connection to the backup application is eliminated. The newer FalconStor VTL assumes the role as the only interface for the backup software, but a mechanism is introduced to create a "direct link pass-through" to the virtual tapes on the old VTL. This is done by creating virtual tapes on the new unit that are merely pointers to the location of the corresponding data on the older unit. These tapes only take up a nominal amount of space (since they are simply pointers) on the FalconStor VTL, with the true data residing on the older server. All new write operations are directed to tapes on the FalconStor VTL unit so they can take advantage of the de-duplication technology. Direct link pass-through allows users to continue to access data from the old VTL directly for restores, until the time at which this data is imported to the FalconStor VTL in the background.

This implementation allows for continued utilization of the old hardware with the option of keeping it in place until the data is aged out of the retention window or until all data is migrated (see Figure 1) using a background import process. The import process is invoked as the administrator sees fit, eventually moving all data off of the old VTL unit so it can be repurposed or retired.

The requisite downtime for this implementation is minimal, since direct link pass-through tape provides direct access to older data while migration is occurring transparently in the background. The only production downtime during the entire migration process is reconnection to the backup application. When the newer FalconStor VTL is connected (zoned) to the older VTL, the connection between the older VTL and the backup server is broken. A brief, minimal amount of downtime is required in order to present the FalconStor VTL storage resources to the backup server and to reconfigure the storage resources at the backup application level.

Figure 1. Tape caching with transparent import requires minimal downtime.



FalconStor VTL allows the import process to take place behind the scenes while the old unit is still available to access data via the backup application. Therefore, data restores can still occur during this process.

Tape caching with import affords the administrator a great level of flexibility for the transition, as the bulk of the impact happens when the tape caching is set up. The only downtime that occurs is when the FalconStor VTL is zoned to the backup server and introduced, and the backup software inventories the virtual media. The backup software does not see a difference between the pass-through tape (FalconStor Direct Link Tapes) and the older virtual tape containing data. Thus, older data can continue to be accessed during the transition and any subsequent imports (to replace direct link tapes) will not disrupt operations.

Configuration Steps

The first steps in the migration process are to insert the newer FalconStor VTL between the backup server and the original VTL.

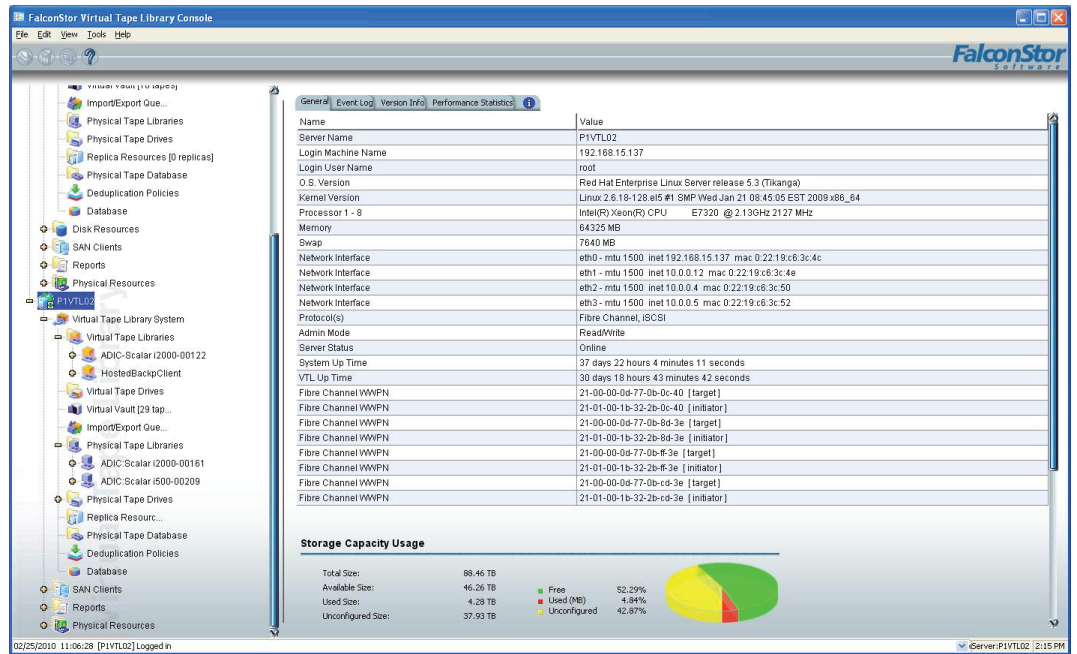
1. Prepare the Environment.

This guide assumes that both the newer FalconStor VTL and the older VTL unit are both on the same fabric. Before beginning the configuration, be sure that the World Wide Names (WWNs) of both VTLs are zoned together so that communication between the two can occur over the SAN.

2. Add Client.

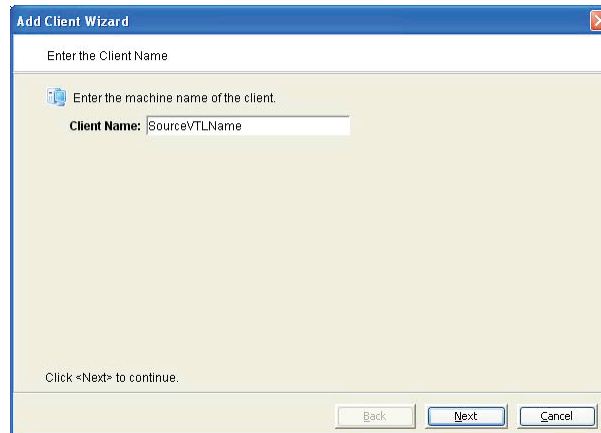
From the main FalconStor VTL console, right-click the "SAN Clients" folder underneath the new FalconStor VTL server in the menu on the left of the screen and select the option to add a new client (see Figure 2).

Figure 2. Add client on FalconStor VTL console.



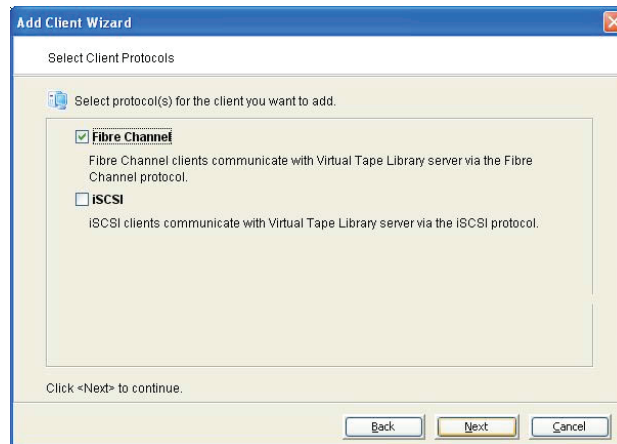
Assign the client a name (typically the VTL hostname) and click "Next" to continue (see Figure 3).

Figure 3. Enter the client name.



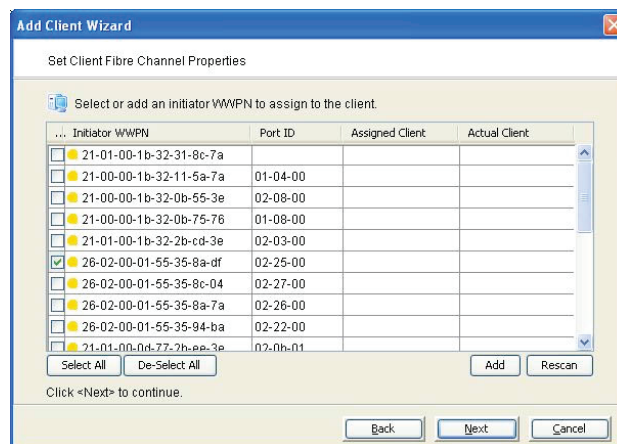
Select whether the client is attached via Fibre Channel or iSCSI. In this example, Fibre Channel is the means of attachment (see Figure 4).

Figure 4. Select the client protocol(s).



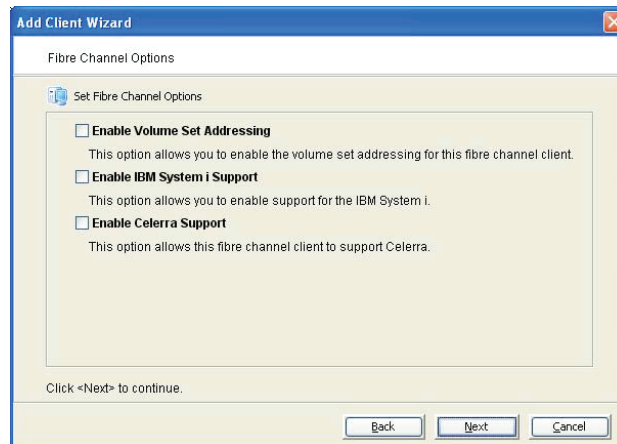
The next window will display a list of all available initiator WWNs that have been zoned to see this FalconStor VTL (see Figure 5). From this list, select the initiator address(es) that will be dedicated to handling the communication between FalconStor VTL and the older VTL. After selecting the appropriate WWNs, click "Next" to move to the final step.

Figure 5. Set client Fibre Channel properties.



For this example, none of the Fibre Channel options apply, so simply click "Next" to finalize the "Add Client" process (see Figure 6).

Figure 6. Select Fibre Channel options.

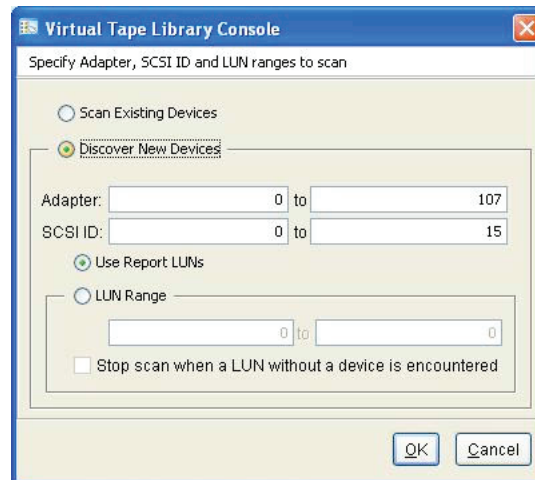


3. Scan Physical Devices.

Right click on the "Physical Resources" item underneath the new FalconStor VTL in the menu on the left of the console and choose to the option "Scan Physical Devices."

Click on the radio button to "Discover New Devices" and leave all of the ranges set at their default values (see Figure 7). Click "OK" to run the device scan.

Figure 7. Choose "Scan Existing Devices" or "Discover New Devices." Run device scan.



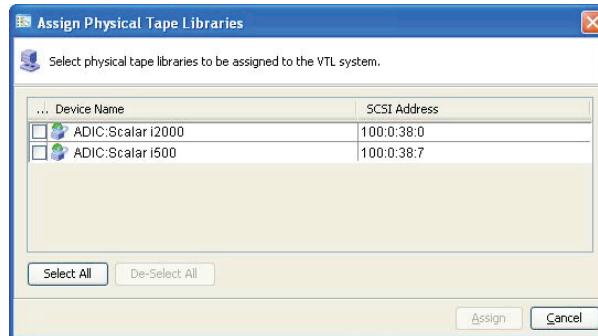
Once completed, the original VTL and all of its associated virtual tape drives should appear in the "Physical Resources" menu.

4. Assign Physical Resource(s) as a Tape Library.

From the menu on the left of the screen, right click on "Physical Tape Libraries" and select the

"Assign" command. The old tape library should appear here. Verify the SCSI address, select the correct unit and click "Assign" (see Figure 8).

Figure 8. Select physical tape libraries to be assigned to the VTL.

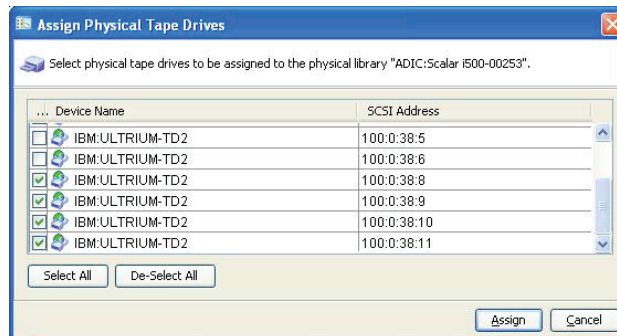


5. Assign Physical Resource(s) as Tape Drives.

In the menu on the left of the screen, locate and expand the group for Physical Tape Libraries. This should now include the physical library assigned as part of Step 4. Right-click on that corresponding physical library entry and select "Assign" once more to add the corresponding virtual tape drives to the library.

Select the entries for the virtual drives that are associated with the recently added library (see Figure 9), and click "Assign" to complete the association.

Figure 9. Select physical tape drives to be assigned to the physical library.



6. Document Details from Old Library.

Make it a point to gather as many details from the old library as possible: library type and/or model, the number and/or type of the drives, the size of the tapes, etc.

7. Create Duplicate Virtual Library to Match Old Library.

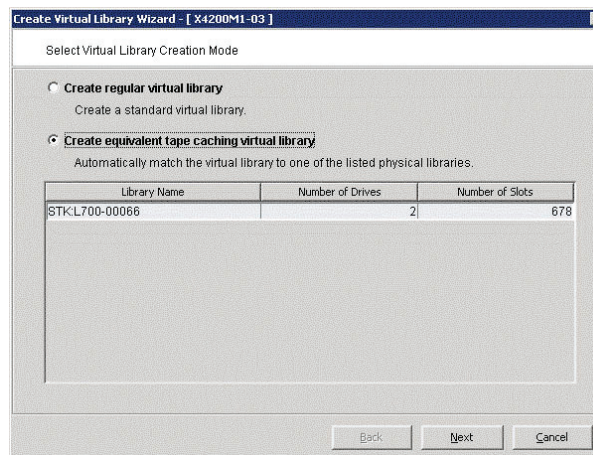
In the left hand menu, under "Virtual Tape Library System," right-click "Virtual Tape Libraries" and select the "Create New Library" option. The new library created will be a physical library with the exact characteristics as the older VTL.

By selecting the equivalent tape caching option, a virtual library identical to the one selected on the older VTL will be created automatically. It will include:

- Library type
- Type and number of tape drives inside the library
- Serial numbers for both library and tape drives

After the exact equivalent of the older VTL is automatically created (see Figure 10). Click "Next" to continue.

Figure 10. Select tape caching virtual library.



Enter the name of the new virtual library. This name is internal to the FalconStor VTL (see Figure 11). Click "Next" to continue.

Figure 11. Specify the virtual library name.

Specify Virtual Library Name and Type

Please specify a virtual library name or use the default name.

Virtual Library Name: L700_Moscow_Migration
Invalid characters for the Resource Name: < > * & \$ / \ ' "

Vendor ID	Product ID	Revision	Maximum Drives	Maximum Slots
STK	L700	3.05	20	678

Click <Next> to continue.

Back Next Cancel

You may edit the drive naming convention to match your own or leave it as the default, but do not modify the number of drives. This number is an exact reflection of the number of physical drives in the physical library of the older VTL (see Figure 12). Click "Next" to continue.

Figure 12. Verify the virtual drive information.

Enter Virtual Drive Information.

Please specify a virtual drive name prefix or use the default name prefix.

Virtual Drive Name Prefix: HP-Ultrium 1-SCSI-
Invalid characters for the Resource Name: < > * & \$ / \ ' "

Total Virtual Drives: 20

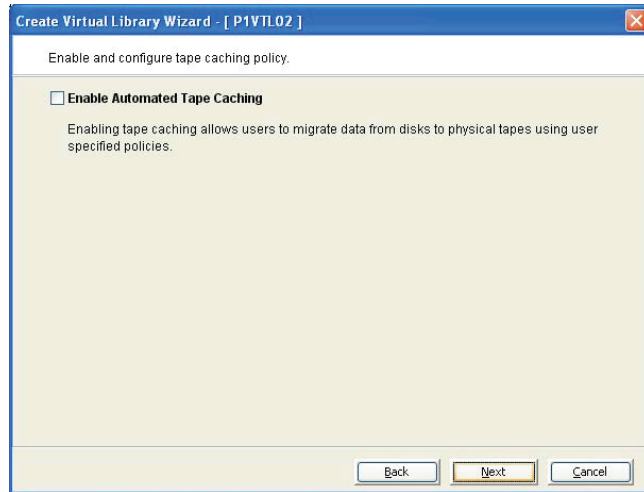
Vendor ID	Product ID	Media Type
HP	Ultrium 1-SCSI	ULTRIUM1
HP	Ultrium 2-SCSI	ULTRIUM2
HP	Ultrium 3-SCSI	ULTRIUM3
HP	Ultrium 4-SCSI	ULTRIUM4
IBM	ULTRIUM-TD1	ULTRIUM1
IBM	ULTRIUM-TD2	ULTRIUM2
IBM	ULTRIUM-TD3	ULTRIUM3
IBM	ULTRIUM-TD4	ULTRIUM4
QANTUM	DLT7000	DLTIV

Click <Next> to continue.

Back Next Cancel

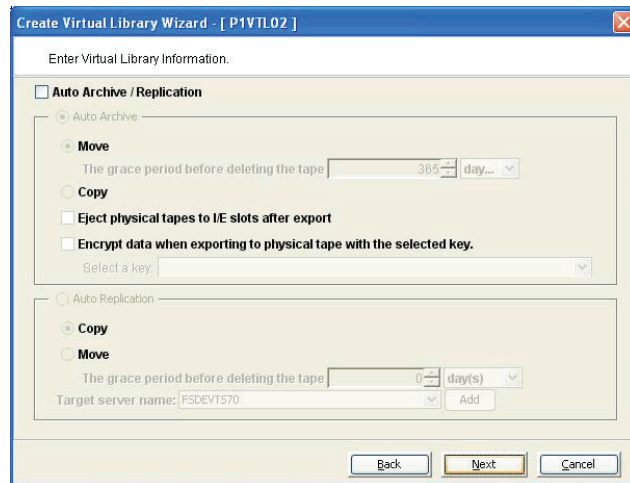
For our example, leave the "Tape Caching" box unchecked and click "Next" (see Figure 13).

Figure 13. Enable and configure the tape caching policy.



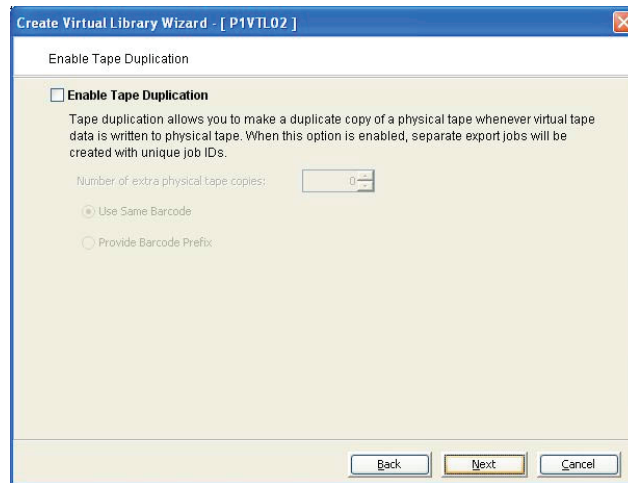
For this example, leave the "Auto Archive/Replication" options unconfigured (as default) and click "Next" (see Figure 14).

Figure 14. Enter further virtual library auto archive or replication information as necessary.



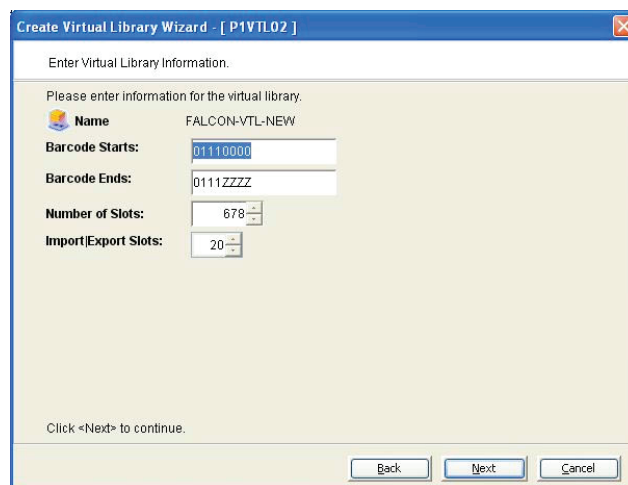
For this example, leave the box for "Enable Tape Duplication" unchecked and click "Next" to continue (see Figure 15).

Figure 15. Enable tape duplication.



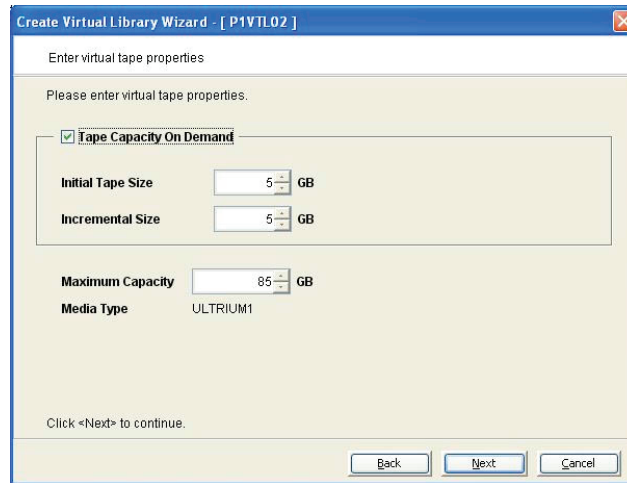
Configure the barcodes for the new library (see Figure 16). Again, match the details for the old library: virtual tapes that will be accessed by "direct tape links" or imported. Click "Next" to continue.

Figure 16. Enter barcode information.



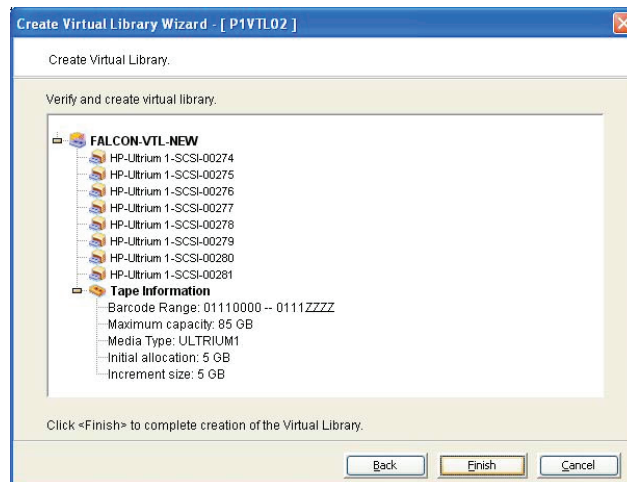
Enable the "Tape Capacity On Demand" option, allowing virtual tape to grow on demand, as needed (see Figure 17). Refer to the FalconStor installation guide for best practices in configuring any of the above settings. Click "Next" to finalize the VTL creation.

Figure 17. Enter virtual tape properties.



Review all of the details for the new library and click "Finish" to complete the process (see Figure 18). The new virtual library is identical to the physical one (old VTL); it is a total representation of the older VTL.

Figure 18. Verify and create the virtual library.



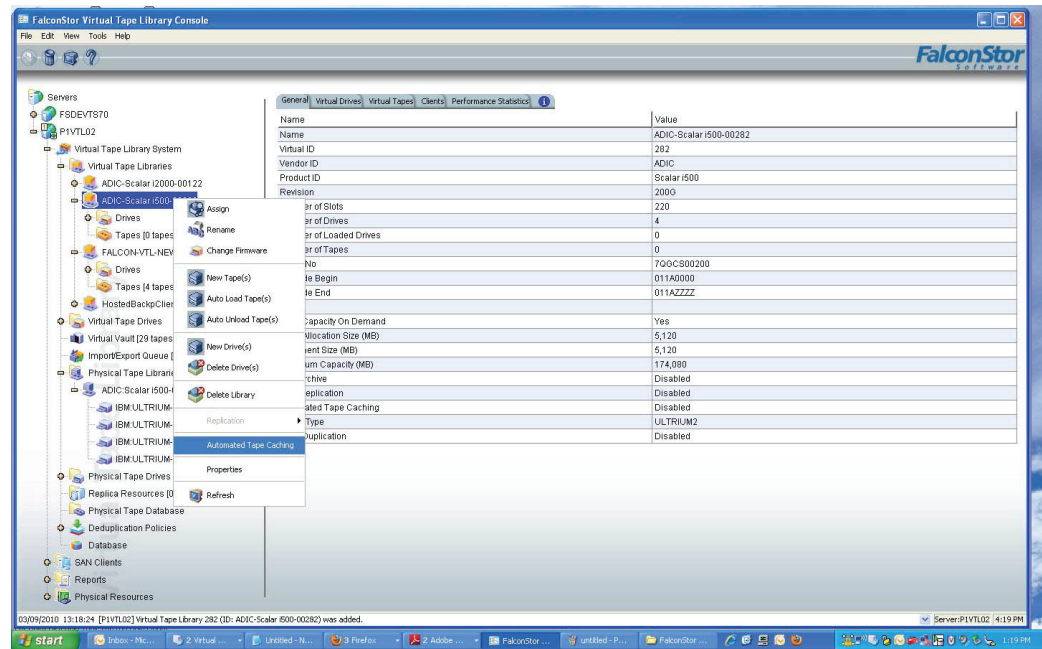
Where there are a large number of virtual libraries, drives or tapes inside the older VTL, a script can automatically perform all of the operations described above.

8. Enable Tape Caching on the Library.

Navigate the menu on the left pane of the console to expand the "Virtual Tape Library System" menu and the "Virtual Tape Libraries" submenu.

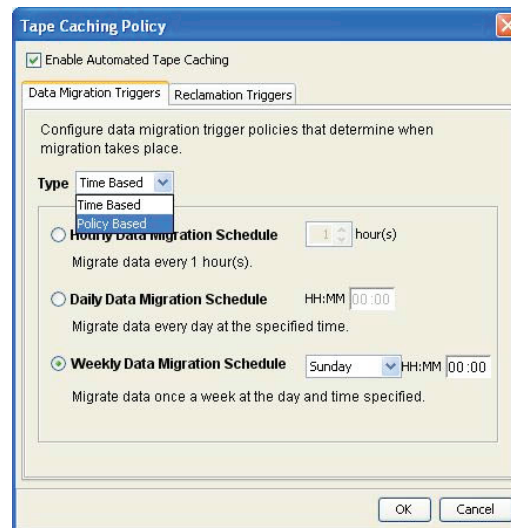
Right-click on the newly created virtual tape library and select the "Automated Tape Caching" menu (see Figure 19).

Figure 19. Expand the virtual tape libraries submenu to select "Automated Tape Caching."



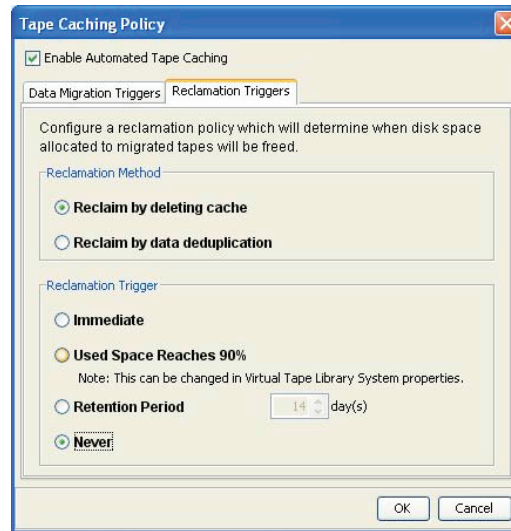
Click the checkbox labeled "Enable Automated Tape Caching" (see Figure 20). The migration trigger options can be left to the default values for the direct link caching setup.

Figure 20. Select "Enable Automated Tape Caching."



Click on the "Reclamation Triggers" tab to enable that submenu. Under the "Reclamation Trigger" options (see Figure 21), click the radio button marked "Never" and click "OK" to close the Tape Caching Policy options.

Figure 21. Choose and select reclamation options.




9. Create Direct Link Tapes.

This section describes how to create virtual tapes in the FalconStor VTL that correspond to the ones used by the older VTL. The operation can easily be performed leveraging the Sync Library option, which automatically creates counterparts for physical tapes inside the older VTL.

Navigate the menu on the left pane of the console to expand the "Virtual Tape Library System" menu, the "Virtual Tape Libraries" submenu, and then the menu for the new VTL that was just created (where tape caching was previously enabled).

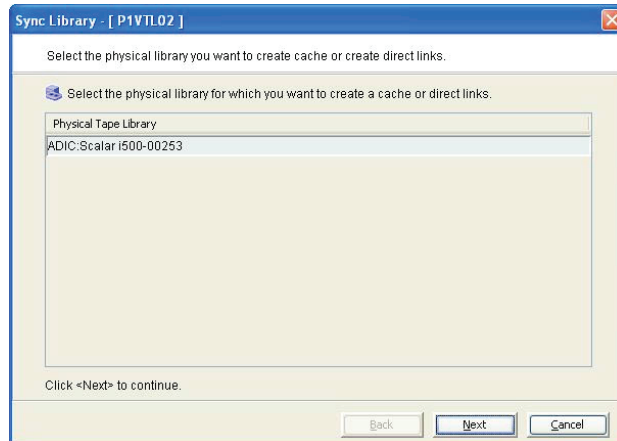
Right-click on the "Tapes" option and select the "Sync Library" function (see Figure 22).

Figure 22. Select the "Sync Library" function.



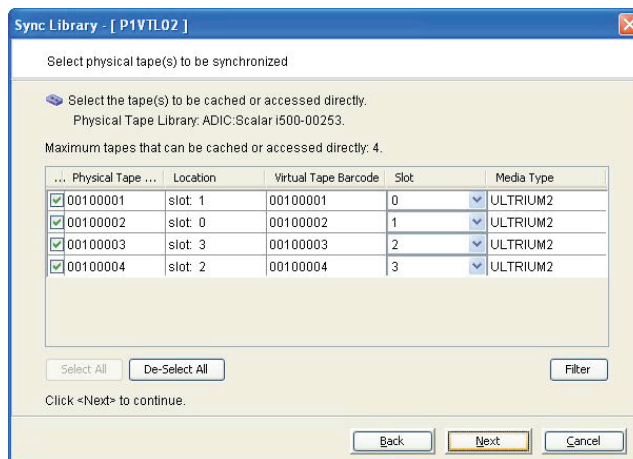
From the "Sync Library" dialog box, select the library where the virtual tapes reside in the older VTL (see Figure 23). Click "Next" to continue.

Figure 23. Select the physical tape library for direct links.



Select all of the tapes that require a "pass-through" link (generally, all of them) and click "Next" to continue (see Figure 24).

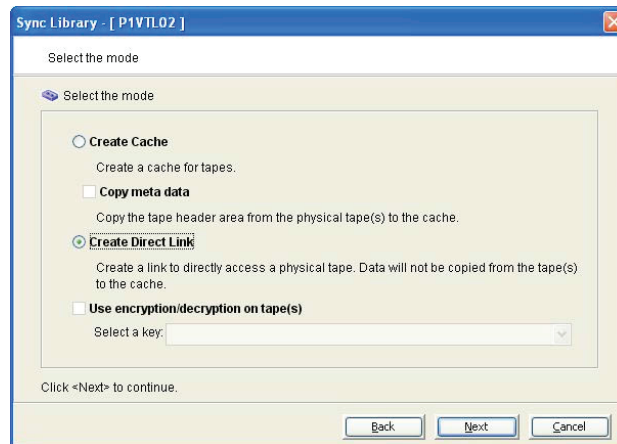
Figure 24. Select the tape(s) to be cached or accessed directly.



Select the radio button "Create Direct Link" to enable the pass-through functionality (see Figure 25). Click "Next" to move to the confirmation stage. Pass-through functionality allows the backup application to continue to access data on the older VTL, by passing through the virtual tape on

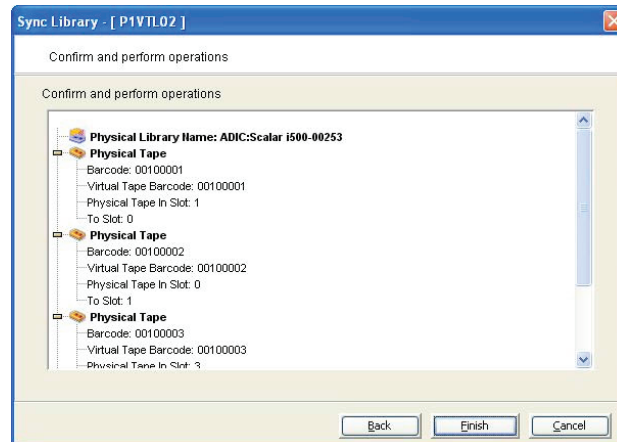
FalconStor VTL. This allows the administrator to continue to access backup data on the older VTL for restore, while the older data is being imported to FalconStor VTL in the background. This ensures complete transparency without any disruption to current backup processes.

Figure 25. Select the creation mode.



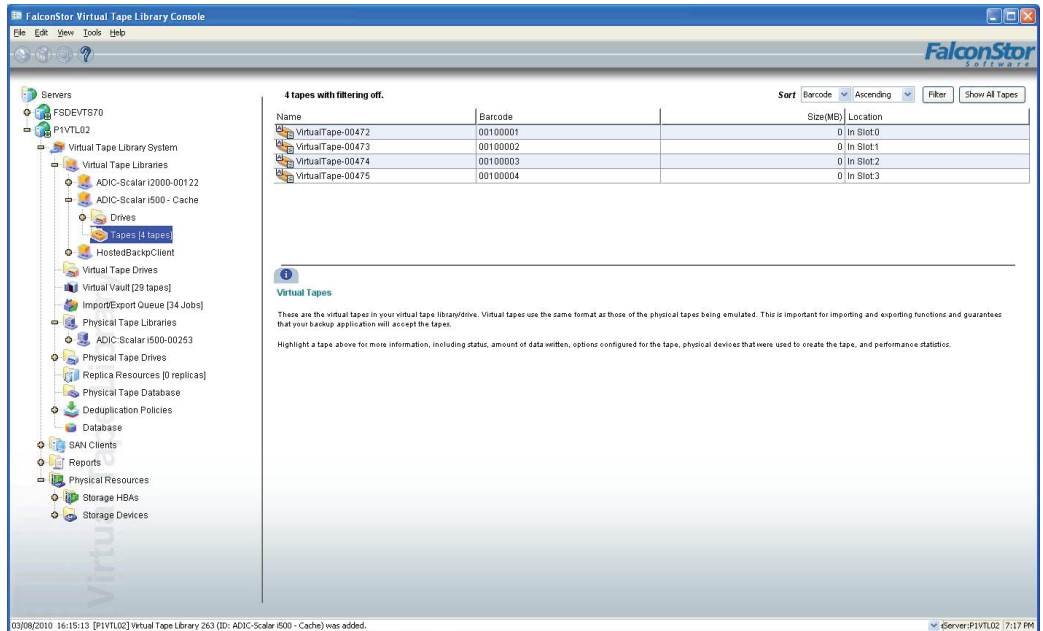
Review the information in the summary and click "Finish" to finalize the options and start the creation of the cached tapes (see Figure 26).

Figure 26. Confirm and perform operations.



After the process completes, verify that the direct links are in place by browsing to the "Tapes" menu (see Figure 27) within the new virtual library. The menu can be checked to confirm that the new "pass-through" tapes have identical serial numbers, slots and barcodes as the ones in the physical library (older VTL). The small "S" in the icon of each tape indicates that the tape is linked in direct access mode. The size of the tape is equal to 0 to indicate that the tape is in direct access mode, with no data stored in FalconStor VTL.

Figure 27. Select the tapes menu to check tape barcodes.



From this point, the virtual tapes are available for backup and restore. They will behave the same way as the original tapes on the older VTL. After the backup server is connected to the FalconStor VTL, the backup server will see no differences, since the new direct link tapes have exactly the same properties (serial number, slot and barcode) as the older VTL tapes.

10. Connect FalconStor VTL to the Backup Server.

This section describes how to remap the storage resources used by the backup server to FalconStor VTL. This step includes downtime. It should be noted that the downtime is limited to a matter of minutes. Downtime is only required to present the FalconStor VTL storage resources to the backup server (Step 10) and to re-configure the storage resources at the backup application level (Step 11 below).

Using the same process as in Step 2 earlier in this document, disconnect the backup appliance server and associated media servers from the older VTL. Then add the WWN of the backup application server and any associated media servers to the FalconStor VTL and the tapes on it. (Note: as described in Step 2, fabric must be zoned for the backup application servers to see the FalconStor VTL.) Using the direct link tapes on the FalconStor VTL, the backup server will still have access to the older VTL and can continue to access data and restore data. Once this data is transparently migrated (imported; see Step 12) into the FalconStor VTL, the direct link tape will automatically change to virtual tape and the backup server will continue to access the data with no change or differences in any backup process, policies or jobs.

Once the FalconStor VTL is zoned to the backup server, presenting the FalconStor VTL resources to the backup server is a very fast operation that is performed with the help of the FalconStor VTL console, requiring only two steps: creating the client (corresponding to the backup server) and assigning the virtual resources to the client.

11. Re-inventory Backup Software. Reconfigure Backup Tape Resources.

Disconnect the old tape library, add the new library at the operating system level of backup application servers, and then add it to backup software. Finally, re-inventory the media within the backup application so it can access the data on the new FalconStor VTL appropriately.

This step ensures that the new FalconStor VTL and tapes are correctly detected by the operating system hosting the backup application. It also ensures the older VTL devices are no longer present to the operating system or backup application.

The operation is as simple as rescanning the devices from the Microsoft Windows® device manager using a Microsoft Windows operating system. For Linux and UNIX operating systems, follow the operating system procedure for scanning ports for new hardware devices.

The second step is to reconfigure the backup application to use the new (FalconStor VTL) hardware resources. Backup application reconfiguration is specific to each backup application, but usually relatively simple.

As an example, for Symantec NetBackup, select the "Configure Storage Device" wizard. During the configuration process, NetBackup will detect the previous device configuration (older VTL) and the new path to the FalconStor VTL. This is normal since NetBackup retains previous paths and devices. The FalconStor VTL migration process retains the exact same serial numbers for devices as used by the older VTL and NetBackup identifies storage resources by serial number. Therefore, NetBackup will transparently use the new path to the FalconStor VTL for all libraries and tape drives, requiring no change to any backup processes or jobs. All jobs will continue to function transparently, as if they were continuing to write to or restore from the older VTL.

At this point, the downtime period is over, the system is back in production and the older VTL will no longer be accessed by the backup server.

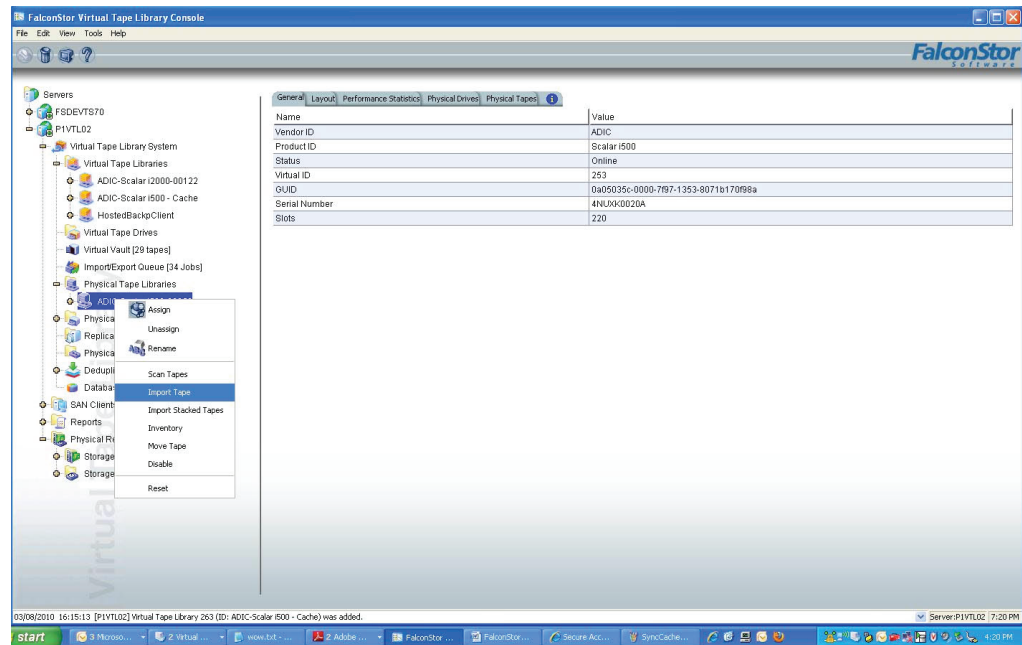
12. Import Tapes to New Library.

The data import process is entirely managed by FalconStor VTL using its own resources and does not require any action from the backup server or from the older VTL. It is a seamless and transparent process that can be scheduled, can import individual tapes or groups of tapes, and can perform a tape copy rather than a data move. In this fashion, the original data is retained on the older VTL until an administrator deletes the data or the backup retention policies automatically delete or move the data to physical tape. Using the FalconStor direct tape management function, tape libraries can be connected directly to the FalconStor VTL. Data movement to physical tape can occur directly from the FalconStor VTL to physical tape libraries without using any backup server resources.

The tape import process can also be performed simultaneously with backup and restore operations, assuming the operations are performed on different tapes. If the same tape is used, the import will be cancelled because backup and restore have the highest priority.


Navigate the menu on the left pane to expand the "Physical Tape Libraries" menu to list your newly added FalconStor VTL (see Figure 28). The "Physical Tape Library" in this case is the older VTL. Right-click on it and select the "Import" option.

Figure 28. Import tapes to the newly created library.



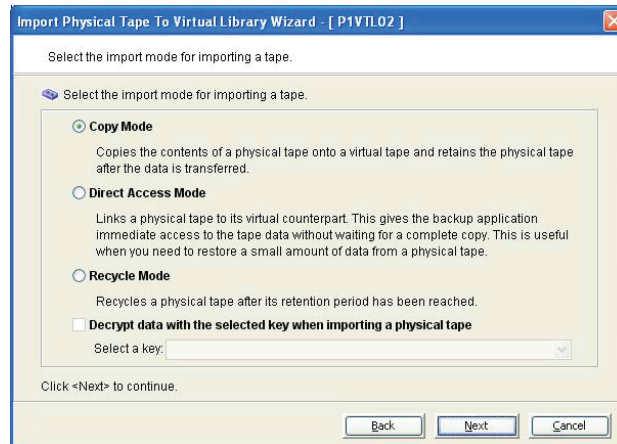
From the list of available virtual libraries to import into, select the library that was just created and click "Next" (see Figure 29).

Figure 29. Select the virtual library to which tape will be imported.



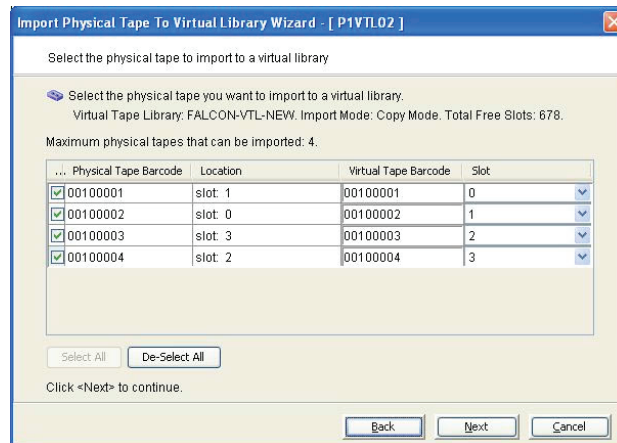
Make sure the "Copy Mode" radio button is selected, which ensures that the data is only copied to the new virtual tape, yet still retained on the older VTL. It remains this way until the time at which administrator deletes the data or the backup retention policy deletes or moves the data to physical tape (see Figure 30). Click "Next" to continue.

Figure 30. Select the import mode.



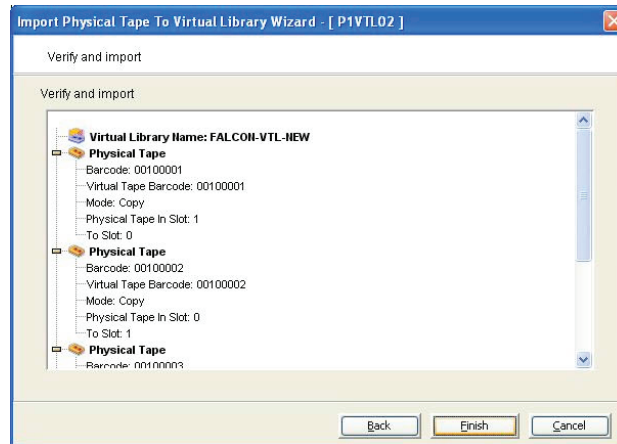
Select the specific tapes to import (see Figure 31). The tapes selected should be the same list of tapes that have direct-link access, previously created (in Step 9). Also, do not modify the virtual tape barcode so that the imported tape will match the direct access tape. Click "Next" to continue.

Figure 31. Select the physical tape that will be imported to the virtual library.



Confirm the details in the summary screen and click "Finish" to finalize the import (see Figure 32).

Figure 32. Verify the virtual tape library details and finalize the import process.



When an import of a tape with the same barcode is initiated and all the data is copied, the new tape automatically changes from a direct link tape to a virtual tape. Visually on the console, the "S" on the tape icon is no longer seen, indicating the tape is a fully operational virtual tape. Since all the data from the older VTL tape is now stored inside FalconStor VTL, the data size is set to the actual value (no longer set to 0).

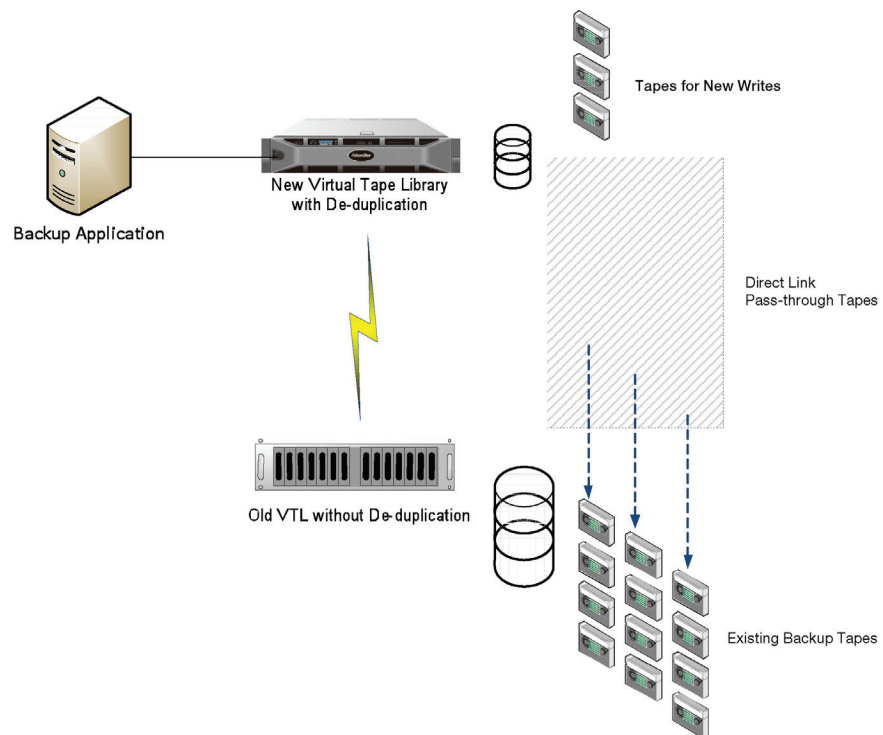
Any backup or restore is performed to the barcode, for that imported tape will now use the data on the FalconStor VTL. Administrators may choose to delete the tape on the old VTL once the import is verified complete to ensure that it is not accidentally imported again in the future.

Scenario 2: Transparent Direct Link Access

In this migration scenario, the old VTL remains in place with its data intact, but its connection to the backup application is eliminated. The FalconStor VTL with de-duplication assumes the role as the only interface for the backup software. Similar to Scenario 1, a mechanism is introduced to create a "direct link pass-through" to the tapes on the old VTL (see Figure 33). This ensures that data can still be accessed and restored from the older VTL until the time that data is deleted or migrated.

This implementation allows for continued utilization of the old hardware with the option of keeping it in place until the data is aged out of the retention window or until the administrator migrates (imports) the remainder of data off of it (see Scenario 1).

Figure 33. Tape caching with direct link pass-through tapes ensures that data can still be accessed and restored from the older VTL until the time that data is deleted or migrated.



Configuration Steps

1-11. Complete Steps 1 through 11 from Scenario 1.

The only difference in this scenario is the elimination of Step 12, Import Tapes to New Library. In the majority of cases, Scenario 1 is the preferred choice of operation in order to remove the older VTL, avoiding associated operating costs and maintenance. The one notable exception may be where data residing on the older VTL is nearing the end of its retention policy and is scheduled to be deleted or moved to physical tape. In this case it may be easier to continue to access that data on

the older VTL through the pass-through direct link access tape. The administrator would then use the FalconStor VTL for new backup data using de-duplication to retain backups in a disk format for longer periods of time.

Conclusion

Partnering with FalconStor, Hitachi Data Systems provides a unique solution and flexible approach to helping organizations leveraging first-generation VTLs. Our solution helps them to upgrade their infrastructure and add de-duplication with minimal disruption and downtime. With a focus on flexibility that gives storage administrators options to tailor a plan to fit their needs, the processes detailed in this document represent a compelling solution for such migration. Solution benefits range from straightforward implementation and single management point, to the ability to granularly transfer single virtual tapes at a time. The introduction of the new FalconStor VTL with the processes described within this document results in a seamless transition to a de-duplication environment. The transition is completed in a customizable manner that allows for the creation and execution of a plan to fit any timetable.

For more information on FalconStor solutions, visit www.falconstor.com/vtl.

Appendix A: About FalconStor

FalconStor Software, Inc. (NASDAQ: FALC), the provider of TOTALLY Open™ Data Protection solutions, delivers the most comprehensive suite of products for data protection and storage virtualization. Based on the award-winning IPStor® platform, products include the industry-leading Virtual Tape Library (VTL) with deduplication, Continuous Data Protector (CDP), File-interface Deduplication System (FDS), and Network Storage Server (NSS), each enabled with WAN-optimized replication for disaster recovery and remote office protection, and the HyperFS™ file system. FalconStor solutions are available from major OEMs and solution providers and are deployed by thousands of customers worldwide, from small businesses to Fortune 1000 enterprises.

For more information, visit www.falconstor.com or contact your local FalconStor representative.

 **Hitachi Data Systems Corporation**

Corporate Headquarters

750 Central Expressway
Santa Clara, California 95050-2627 USA
www.hds.com

Regional Contact Information

Americas: +1 408 970 1000 or info@hds.com
Europe, Middle East and Africa: +44 (0) 1753 618000 or info.emea@hds.com
Asia Pacific: +852 3189 7900 or hds.marketing.apac@hds.com

Hitachi is a registered trademark of Hitachi, Ltd., in the United States and other countries. Hitachi Data Systems is a registered trademark and service mark of Hitachi, Ltd., in the United States and other countries.

All other trademarks, service marks and company names in this document or website are properties of their respective owners.

Notice: This document is for informational purposes only, and does not set forth any warranty, expressed or implied, concerning any equipment or service offered or to be offered by Hitachi Data Systems Corporation.

© Hitachi Data Systems Corporation 2010. All Rights Reserved. WP-380-A DG December 2010