

MANAGE DATA SMARTER™

StorFirst Apollo Administration Guide

Note to StorFirst Apollo users:

This manual provides instructions for installing, configuring, and using StorFirst Apollo software. For supported storage subsystems and the latest updates, see the *StorFirst Apollo Hardware Compatibility Guide* and *StorFirst Apollo Release Notes*.

This administration guide assumes that you are familiar with your server hardware, the Windows operating system, and your basic system administration.

Unless otherwise specified, references to Windows refer to all supported Windows operating systems.

© 2011 Seven10 Storage Software LLC. All rights reserved. This document may not be further copied, distributed, or modified without the written permission of Seven10. Unauthorized use, disclosure, and distribution are strictly prohibited. This software (including documentation) is subject to the terms and conditions set forth in the end-user license agreement or other applicable agreement, and you may use this software only if you accept all the terms and conditions of the license agreement. This software comprises proprietary and confidential information of Seven10.

StorFirst Apollo Summary

With Hitachi Content Platform (HCP), Hitachi Data Systems (HDS) has introduced a platform specifically designed for storing and managing unstructured content. HCP carries the hallmark availability and disaster recovery capabilities that HDS has marketed with their block-level storage products.

StorFirst Apollo (“Apollo”) is designed to complement HCP independently of how data is stored and retrieved on HCP. Apollo is auditable and preserves the authentication and retention functionality of HCP. Apollo is invisible to applications accessing HCP through the various supported namespace access protocols and works directly with the native REST API.

The key functions in Apollo are object replication to and object restoration from tape and virtual tape libraries (VTLs). Apollo preserves not only the data but also the complete metadata used in the implementation of HCP.

A single Apollo instance can serve multiple HCP systems. A single HCP system can simultaneously be served by several Apollo instances. Connections can be LAN or WAN, providing freedom of location. Apollo can simultaneously use a set of tape or VTL devices, ranging from a single-drive autoloader to a set of multidrive tape or VTL libraries. Apollo is a complete tape-library and VTL management solution and manages tape and VTL data in all possible locations: in drive, in library slot, and offline (on shelf or in vault).


Apollo enables these key benefits:

1. **Cost:** Cost-effective replication by using tape or VTL for disaster recovery purposes
Full site and/or device disaster recovery at much reduced cost
2. **Compliance:** Support for regulations calling for replication to removable media, including WORM tape
3. **Comfort:** Leveraging of customary procedures for tape/VTL usage and off-site storage using true, efficient HCP replication, not typical backup jobs
4. **Consolidation:** Support for many-to-one replication
5. **Migration:** Migration of older objects to tape for the transfer of data to new HCP systems

CHAPTER 1 PLANNING APOLLO

Product Concept

This simplified picture of a single Apollo instance serving HCP will make it easier to read this book. Apollo replicates the information stored either in a single HDS HCP system or in multiple HCP systems (including across multiple tenants within a system) to tape\VTL while the systems are simultaneously being used by applications on the network (this is configurable). ***Apollo replication is customized, optimized, and purpose-built for HCP in sharp contrast to any standard backup.*** Apollo is installed on a Windows server and communicates with HCP over TCP/IP using the native REST API provided by HDS.



HCP Basics (disclaimer)

The purpose of the following section is to present the most basic HCP concepts necessary for managing Apollo. You are encouraged to take advantage of the official and complete HCP documentation from HDS that in all cases takes precedence over the information presented here.

HCP Basics

Hitachi Content Platform (HCP) is the distributed, fixed-content, data storage system from Hitachi Data Systems. It provides a cost-effective, scalable, easy-to-use repository that can accommodate all types of data, from simple text files to medical images to multi-gigabyte database images. An HCP system consists of both hardware and software.

A **fixed-content storage system** is one in which the data cannot be modified. HCP uses write-once, read-many (WORM) storage technology and a variety of policies and services to ensure the integrity of the stored data and the efficient use of storage capacity. It also provides easy access to the repository for adding, retrieving, and deleting or shredding the stored data.

Objects

HCP stores **objects** in the repository. Each object is a single unit that encapsulates data HCP receives (the fixed content) with information about that data, called **metadata**. Metadata consists of system-managed properties, such as creation date, size, and retention, and, optionally, user-supplied information, called **custom metadata**, that further describes the object.

Objects are identified by name and directory path, like files in a standard file system such as Windows or UNIX. This means that access to objects is not restricted to the applications that created them.

HCP can store multiple versions of an object, thus providing a history of how the data has changed over time. Each version is an object in its own right, with system metadata and, optionally, custom metadata.

Namespaces and Tenants

An HCP repository is partitioned into namespaces. A **namespace** is a logical grouping of objects such that the objects in one namespace are not visible in any other namespace.

Namespaces provide a mechanism for separating the data stored for different applications, business units, or customers. For example, you could have one namespace for accounts receivable and another for accounts payable.

Namespaces also enable operations to work against selected subsets of objects. Thus, for example, using StorFirst Apollo, you could choose to replicate the accounts receivable and accounts payable namespaces but not the employees namespace.

Namespaces are owned and managed by administrative entities called **tenants**. A tenant typically corresponds to an organization, such as a company or a division or department within a company. A tenant can also correspond to an individual person.

An HCP system can have two types of namespaces:

- **HCP namespaces**, which are owned and managed by **HCP tenants**, support authenticated access, storage usage quotas, and a REST interface implemented by HTTP. An HCP system can have multiple HCP tenants, each of which can own multiple HCP namespaces.
- The **default namespace**, which is owned and managed by the **default tenant**, supports multiple industry-standard protocols for access to objects (including HTTP, CIFS, and NFS). An HCP system can have only one default tenant and namespace.

Access

To administer the default tenant as well as the HCP system as a whole, clients need a system-level **user account**. To administer HCP tenants, clients need a tenant-level user account. In each case, the user account specifies a username and password and assigns roles to the account user.

A **role** is a named collection of permissions that can be granted to an HCP user. Each permission in a role lets the user perform some specific interaction or set of interactions with the HCP system.

To access the content of an HCP namespace, clients need a **data access account**. The data access account specifies a username and password and assigns permissions for operations that the account user can perform in the namespace. A single data access account can assign permissions for multiple namespaces, as long as those namespaces are owned by a single tenant.

An HCP tenant can grant system-level users administrative access to itself. This enables system-level users to administer the tenant at the tenant level. It also gives those users the ability to search and read the content of namespaces owned by the tenant.

Hardware

An HCP system includes multiple servers (or blades in a blade server), called **nodes**, that are networked together. For data storage, each node can have multiple internal physical drives (HCP 300 systems) or connect to external fibre channel storage (HCP 500 systems).

Nodes can be **storage nodes** or **search nodes**. Storage nodes manage the objects that reside in the HCP system storage. Each storage node runs the complete HCP software. Search nodes support the HCP search facility.

An HCP system can be identified either by a single DNS name or by the IP address of an individual storage node. In either case, HCP distributes the processing among all the storage nodes, thereby enabling access to all objects, regardless of how the system is addressed.

Apollo vs. Backup

Apollo is tightly integrated with HCP as the optimal solution to preserve all authentication and retention management even through a complete site disaster. Conventional backup may be available in some configurations. However, using existing backup solutions may invalidate the authentication capabilities that differentiate HCP as a storage device. Further, applying recurring backup schedules to fixed-content data is wasteful.

Apollo replication is fundamentally different from traditional backup. There is no distinction in Apollo between incremental and full. Apollo replication builds a single, complete, and always-synchronized copy of all the information in HCP by incrementally capturing new objects. Objects cannot change, so there is no need to replace old with new. Simple accumulation is automatically synchronized. Initially, Apollo may need to catch up (in the case of use with a preexisting HCP), but eventually each replication session will efficiently replicate only the new objects since the last session.

Tape and VTL Capacity

Apollo works with most popular brands of tape libraries, virtual tape, drives, and physical tape formats. Please consult the release notes and *StorFirst Apollo Hardware Compatibility Guide* for up-to-date details. These choices should be made independently of Apollo with your particular vendor of choice. The total online capacity should meet or exceed the HCP storage. This calculation will be an estimate and will be affected by unpredictable factors such as content, on-line versus off-line tape volume operation, storage habits, and HCP redundancy. In configurations where retention on tape or VTL exceeds retention in HCP, additional capacity will be required over time. Fortunately, Apollo allows for easy expansion.

🔗 ***A simplified rule of thumb: use a library that has a raw capacity exceeding the raw capacity of the HCP storage.***

Tape or VTL Bandwidth

The minimum number of drives should be equal to the number of HCP namespaces, plus one (optional) drive for redundancy. A single-drive library or autoloader is supported, but at least two drives are recommended. Many current tape drives have sufficient bandwidth to keep up with the average bandwidth usage of a single HCP system. Net HCP bandwidth usage is lower for small objects and increases with object size. It is conceivable that with consistent use of large objects and fast networking, the bandwidth usage of a single HCP system could exceed the bandwidth of a single tape drive. Apollo also supports the use of multiple drives in parallel when working with a single HCP system, driving them at higher speeds (provided the network and HCP are able to keep up).

🔗 ***A simplified rule of thumb: make sure that aggregate tape or VTL drive bandwidth exceeds your measured HCP bandwidth usage, so tape or VTL speed will not be a bottleneck.***

Networking Selection

The connections between the Apollo server and the HCP units are TCP/IP based. Either LAN or WAN can be used and should provide capacity of several tens of MBs per second per HCP system. Anything less than a gigabit network backbone may be flooded to the point of becoming unusable to everything else attempting to share the network. A robust and high-performance network is therefore strongly recommended.

Server Selection

Apollo requires a dedicated Windows system. Please consult the release notes for the precise and up-to-date minimum specification.

The Apollo Database

Apollo keeps a database on hard disk of all objects stored on tape. The database entries are created during replication from HCP to tape, one entry in the Apollo database for each object. The database controls both replication and restoration. The database must be on hard disk, such as DAS or SAN.

If the database is lost, for example due to a disk crash on the Apollo server, it can be recreated by scanning the tapes from end to end. It can also be recreated from HCP by reformatting the tapes or VTL and repeating the complete replication, but that could potentially take much longer, especially if the data set is large.

CHAPTER 2 INSTALLING APOLLO

Getting Ready

Physically install your storage devices and test that Windows can see them using the instructions provided by their manufacturer. Make sure to use unique and otherwise unused SCSI addresses. A SCSI address is made up of a SCSI port number, SCSI bus number, SCSI ID, and a SCSI LUN. If your device allows, you can physically open the device and place media directly into the slots before turning on the device.

On the StorFirst Apollo server, create the software distribution folder and copy the files from the release media or the Seven10 FTP site. There are both software files and a set of documentation in PDF format. Be sure to have a PDF reader installed for access to digital documentation and online help. Read the enclosed license. Installing Apollo constitutes acceptance of its terms. Read the release notes for up-to-date product requirements.

Welcome to StorFirst

Log in as an administrator, open the distribution folder, and begin installing by double-clicking on Apollo.exe. This will be the first window to appear. Click Next.

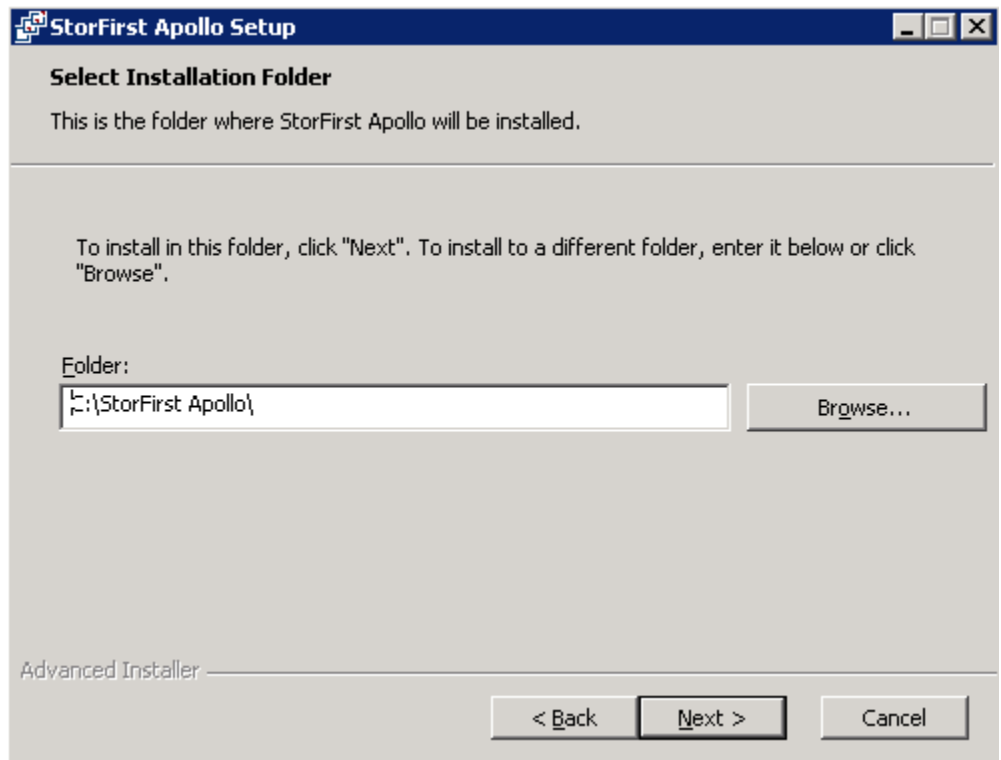


Destination Folder

This will be the location of the installed software, the logs, the help files, and the database.

The default is a StorFirst Apollo folder at the root of the drive letter with the most free disk space. You can browse to select another drive letter. It is a good idea to keep folder name unchanged.

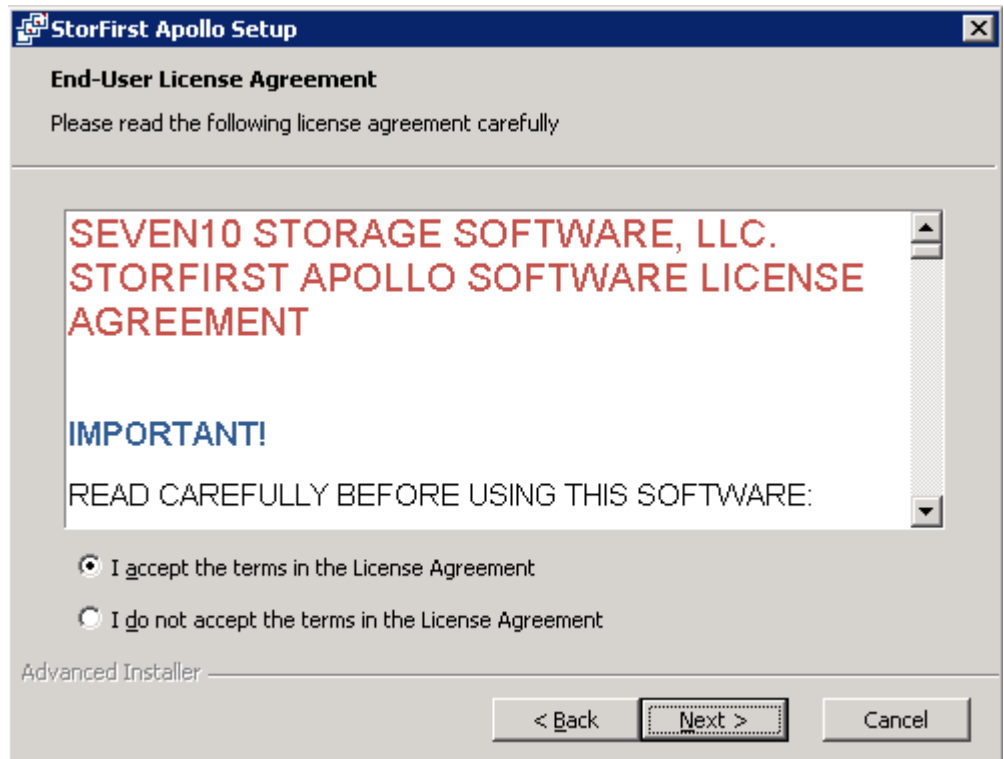
The install process automatically creates the destination folder and copies everything needed for proper operation to it. The source folder is unnecessary after completing the installation.



License Agreement

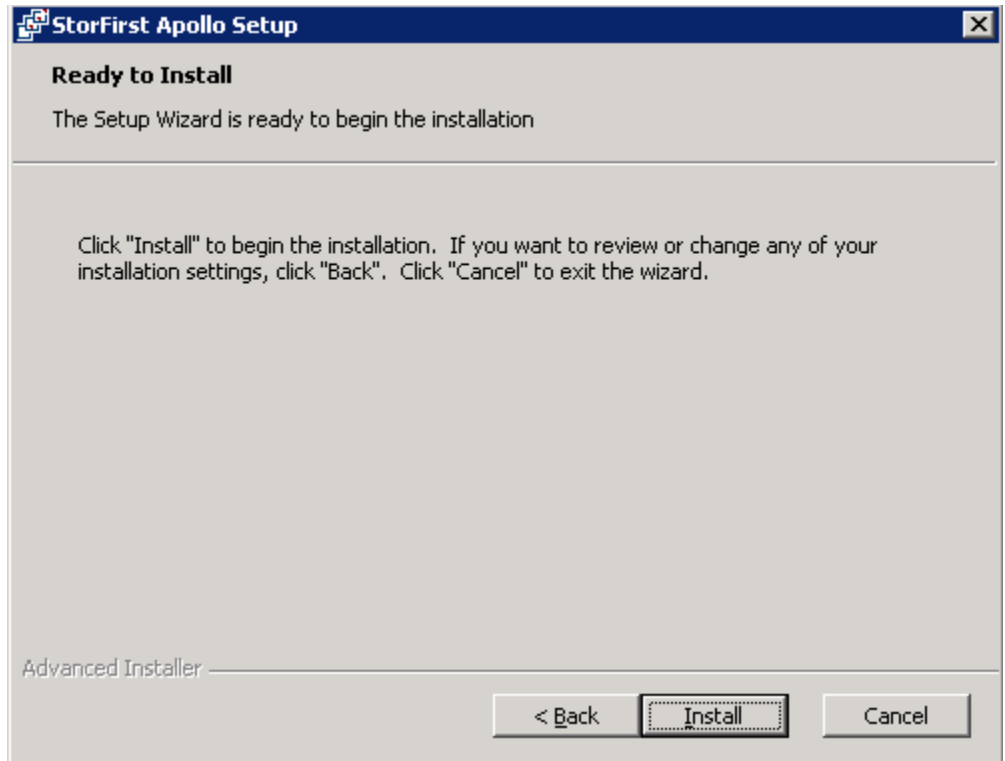
Next, you will see the StorFirst Apollo License Agreement. Please read through the agreement and select "I accept the terms in the License Agreement" to continue the installation. Click Next.

Note: If you select "I do not accept the terms in the License Agreement," you cannot continue the installation.



Ready to Install

You are now ready to install the application. Click Next to begin the installation.



StorFirst Apollo Setup Complete

The setup process is complete. Click Finish.



Reboot

You must now reboot the server in order to complete the installation process. This will load the StorFirst driver and start the background service.

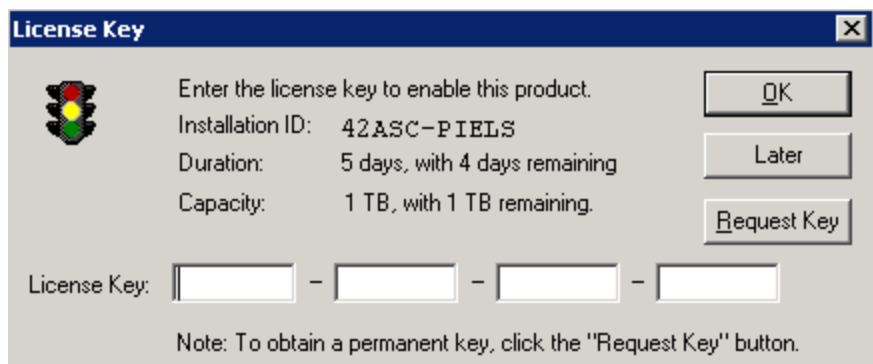
Done

When the system returns after rebooting, you will find:

1. A StorFirst Apollo shortcut icon on your desktop
2. The documentation and shortcut in Windows Start ⇒ Programs ⇒ StorFirst Apollo
3. The StorFirst Apollo service in Microsoft Services
4. In the Task Manager under the Processes tab see:
 - a. The service ApolloSvc.exe
5. The Destination folder named StorFirst Apollo in the location of your choice containing:
 - a. The StorFirst Apollo database
 - b. The software modules
 - c. The Job Log folder containing the logs of jobs that you have run and not yet deleted
 - d. The user interface log: StorFirst Apollo Administrator.log (StorFirst Apollo\Service Logs\StorFirst Apollo Adminstrator.log)
 - e. The service log: Apollosvc.log (StorFirst Apollo\Service Logs\Apollosvc.log)
 - f. The device log: DeviceList.log (StorFirst Apollo\Service Logs\DeviceList.log)
 - g. The help files (StorFirst Apollo\Docs)
 - h. The text file "errmsgeng" containing error numbers and their messages (StorFirst Apollo\Bin)

StorFirst Administrator

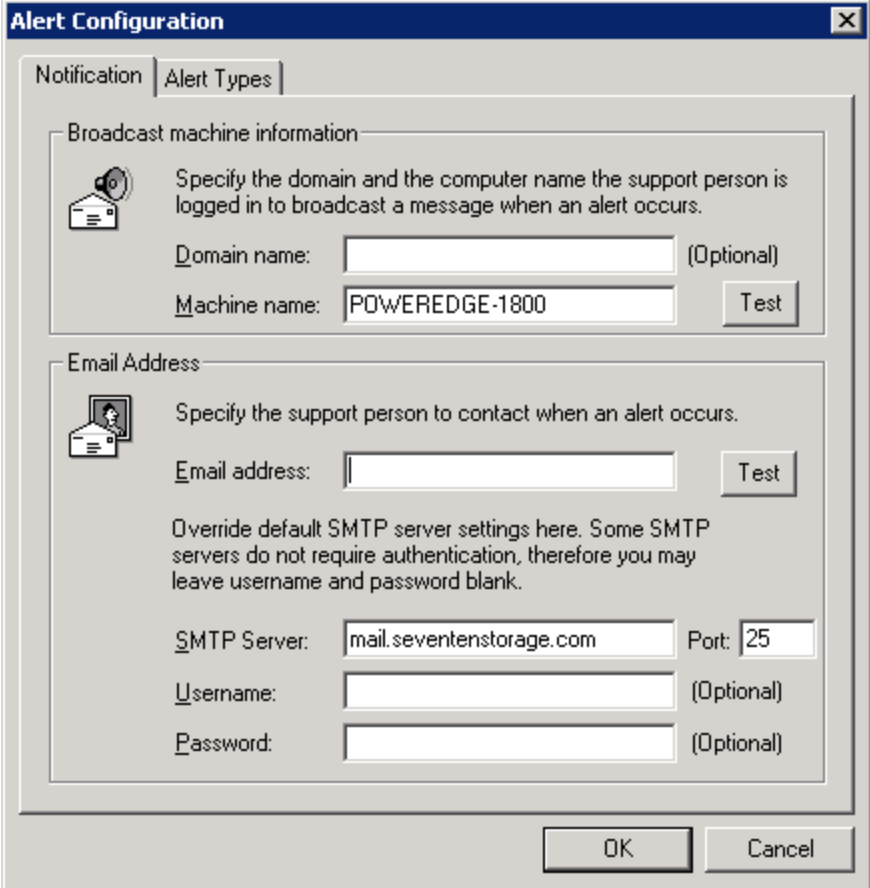
Now start the StorFirst Apollo Administrator using the StorFirst Apollo icon on your desktop. The first time you run the Administrator, the system will prompt you for a license key. You will need to enter this key before proceeding to use StorFirst Apollo. You can continue in evaluation mode for five days with capacity up to 1,000 GB by clicking Later.



You must now contact Seven10 support for a license key. You can contact the Seven10 Customer Support Center by calling 978-725-5525, by faxing 978-725-8808, or by email at support@seven10storage.com. Support is available Monday through Friday from 9:00 AM to 5:00 PM EST.

Alert Configuration

Next, you are asked how you would like to be notified when StorFirst reports an alert. You can choose either or both of two methods to receive alerts: broadcast message or email. These are configurable on individual types of alerts.



The image shows a screenshot of the 'Alert Configuration' dialog box. It has two tabs: 'Notification' and 'Alert Types'. The 'Notification' tab is selected. The dialog is divided into two main sections: 'Broadcast machine information' and 'Email Address'. In the 'Broadcast machine information' section, there is a speaker icon and a text box for 'Domain name' (optional) and a text box for 'Machine name' containing 'POWEREDGE-1800'. A 'Test' button is next to the machine name field. In the 'Email Address' section, there is a person icon and a text box for 'Email address' with a 'Test' button. Below this, there is a note: 'Override default SMTP server settings here. Some SMTP servers do not require authentication, therefore you may leave username and password blank.' There are three fields: 'SMTP Server' containing 'mail.seventenstorage.com', 'Port' containing '25', 'Username' (optional), and 'Password' (optional). At the bottom right, there are 'OK' and 'Cancel' buttons.

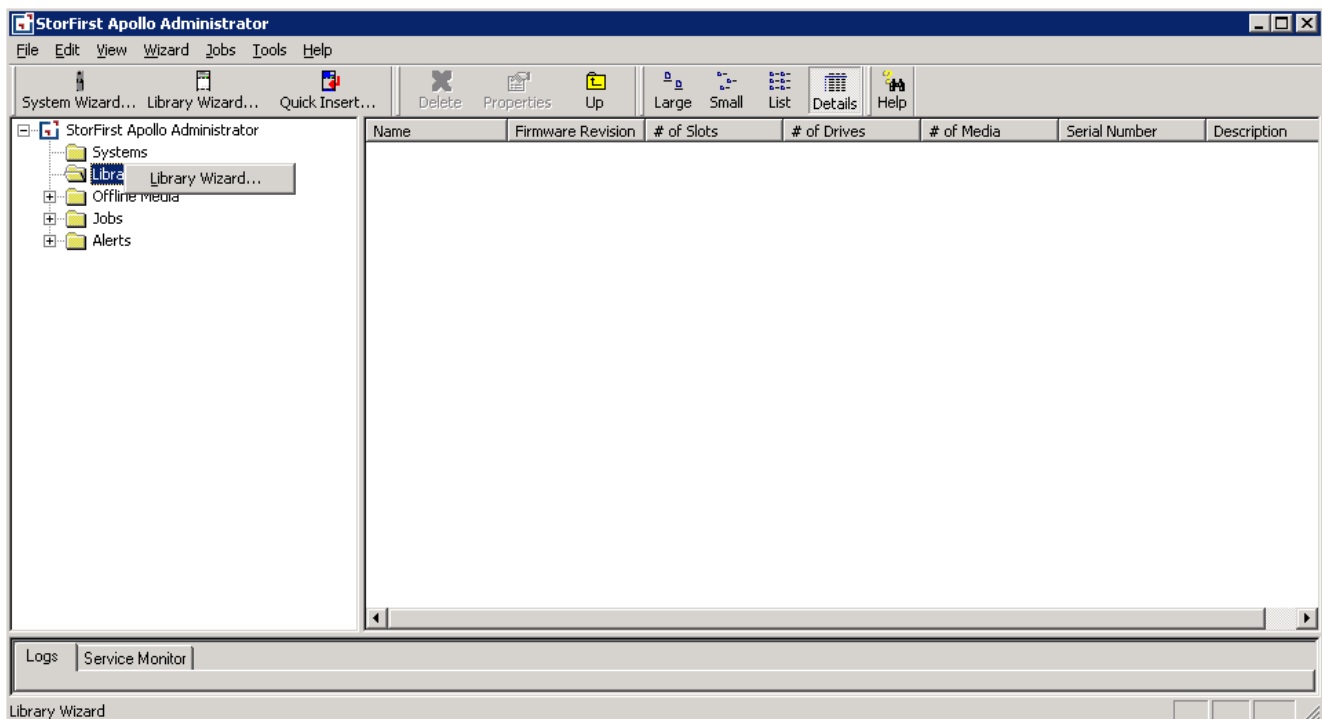
CHAPTER 3

CONFIGURING STORAGE DEVICES

Library and HCP Configuration

StorFirst Apollo uses wizards to install the software itself as well as configure devices within the software. The wizards can be activated from the Wizard menu or toolbar at any time or by right-clicking on individual folders in the tree pane. For example:

- To add a tape library → Click the Library Wizard icon or right-click Libraries in the tree pane and click Library Wizard.
- To add an HCP system → Click on the System Wizard icon or open the Wizard menu and click System Wizard.



The Library Wizard

The Library Wizard is used to configure new tape libraries. It is automatically activated after starting the Administrator. It can also be activated from the top menu at any time. When activated, the Library Wizard rescans the buses so any hot-pluggable device can be added while the system is running. A list of supported tape libraries is available in the StorFirst Apollo *Hardware Compatibility Guide*.

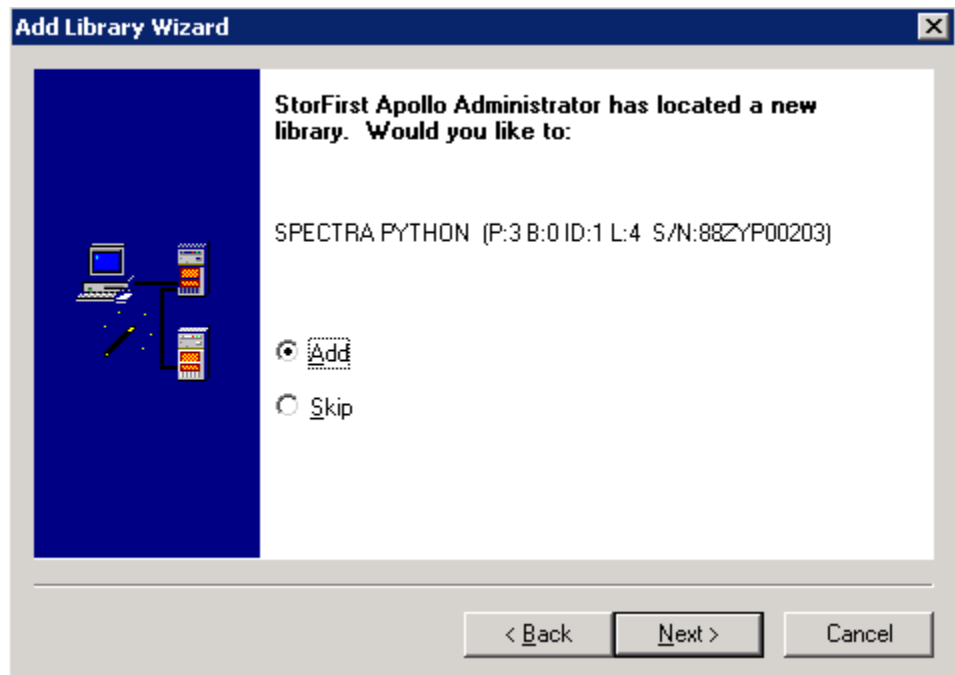


Configuring Storage Libraries

You can now configure the storage devices in the Apollo system using the Library Wizard.

You can choose to add or skip each device. If added, the device will be completely controlled by Apollo and cannot be used in any other way. If skipped, the device will be ignored by Apollo but can be added later.

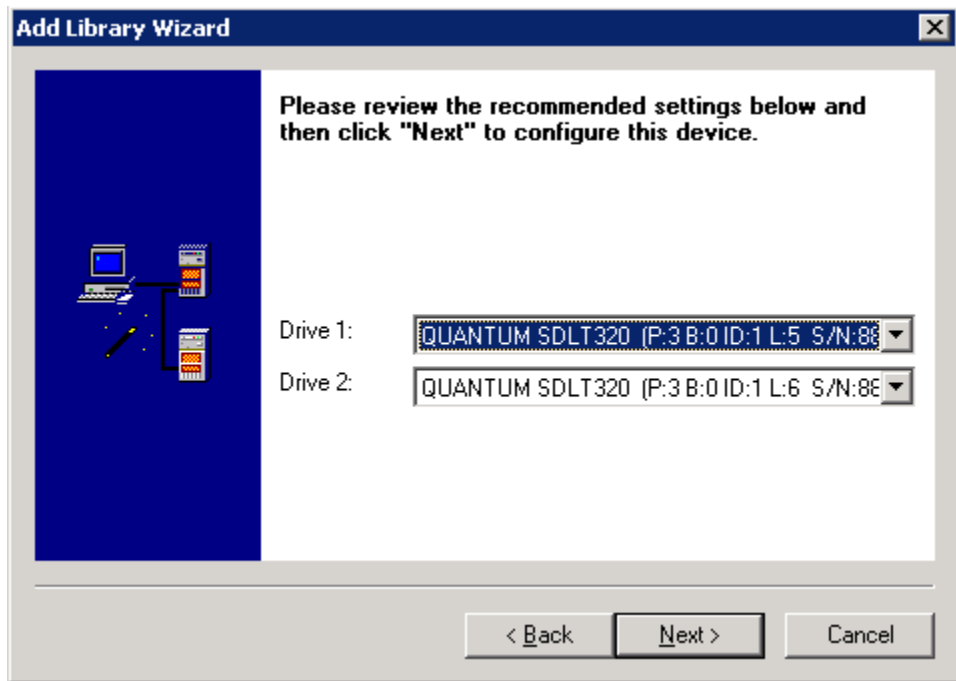
For ACSLS, see Appendix C of this guide.




Note: Do not add any Microsoft drivers for the library or drives.

Drives in Libraries

Drives are connected to the Apollo server using standard SCSI over direct cables or Fibre Channel. The library changer identifies the number of drives and the drive type and location by slot number. Apollo will normally automatically recognize the relationship between the SCSI addresses of the drives and their slot numbers. This screen is for confirmation and selection.



If Apollo cannot map drives automatically, the drive fields will contain "unknown." In this case, you should contact your service representative, who will help you through manual association of drive numbers with the drives found by scanning the buses.



A drive can be excluded from use by Apollo by selecting "unknown" in the dropdown menu at this time or by disabling it later. This enables sharing of a library between multiple

applications. All drives used by Apollo must be interchangeable and compatible with all media in the assigned slots. Two separate libraries can use different drives and different types of tape.

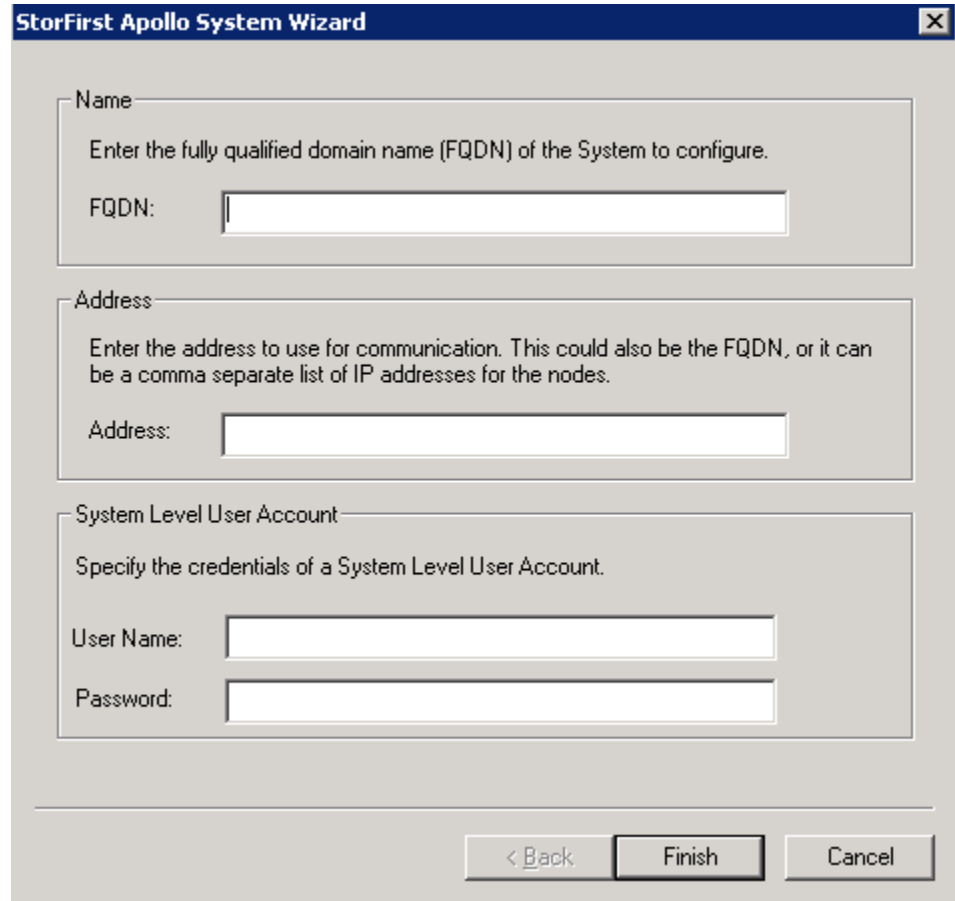
Library Configuration Complete. Click Finish.

Configuring HCP

To configure an HCP system, select the System Wizard button at the top of the screen. You could also select System Wizard under the Wizard option in the Tool Bar. Use the HCP wizard once per HCP system. For each system, enter the fully qualified domain name (FQDN), the address (which can be the FQDN or a list of IP addresses for the HCP nodes), the user name, and the password.

Click **Finish** when complete.

Click **Yes** when prompted to create new tenants.



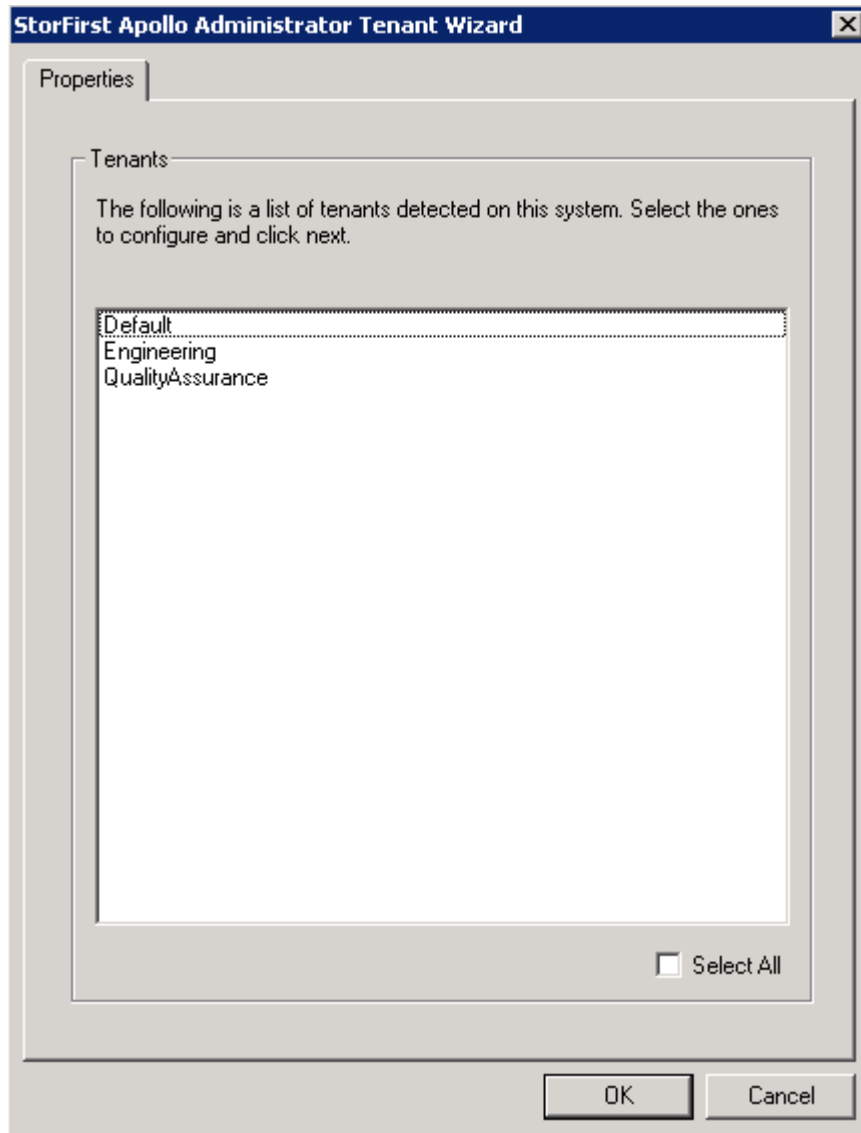
The screenshot shows the 'StorFirst Apollo System Wizard' dialog box. It has a title bar with the text 'StorFirst Apollo System Wizard' and a close button (X). The dialog is divided into three sections:

- Name:** A section with the instruction 'Enter the fully qualified domain name (FQDN) of the System to configure.' and a text input field labeled 'FQDN:'.
- Address:** A section with the instruction 'Enter the address to use for communication. This could also be the FQDN, or it can be a comma separate list of IP addresses for the nodes.' and a text input field labeled 'Address:'.
- System Level User Account:** A section with the instruction 'Specify the credentials of a System Level User Account.' and two text input fields labeled 'User Name:' and 'Password:'.

At the bottom of the dialog, there are three buttons: '< Back', 'Finish', and 'Cancel'.

Configuring a Tenant

To configure a tenant or multiple tenants for replication, select the tenant(s) you want from the list of tenants that StorFirst Apollo has detected. Then click OK.



The StorFirst Apollo Tenant Wizard will appear. The name of the tenant you selected will already be present. Enter the credentials for the tenant in the Tenant Level User Account and the Data Access Account sections. Then click Finish.

StorFirst Apollo Tenant Wizard

Tenant Name
Specify the name of the tenant you want to configure.
Engineering

Tenant Level User Account
Specify the credentials of a tenant level user account.
User Name: apollosystemadmin
Password: *****

Data Access Account
Specify the credentials of a data access account.
Note: This account should have read, write, and search permissions in all namespaces to be backed up.
User Name:
Password:

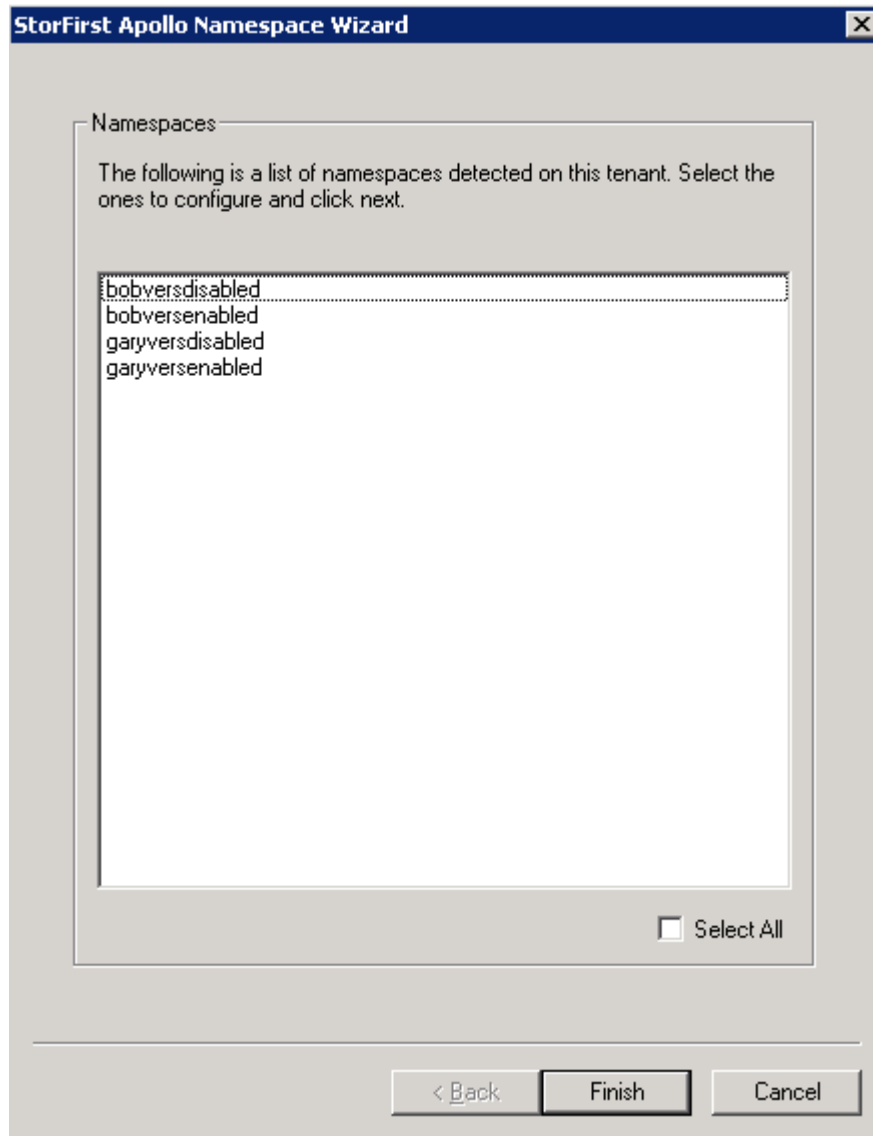
< Back Finish Cancel

Click Yes when asked if you want to configure a namespace.

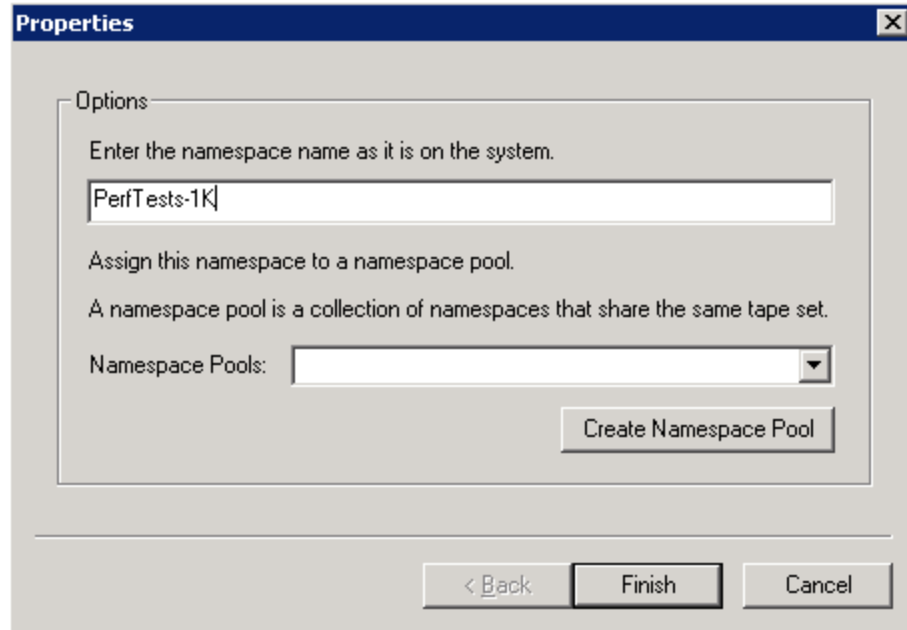
If you had chosen to configure multiple tenants, the software will ask you to configure one tenant at a time. All associated namespaces will be configured with each tenant before moving on to the next tenant.

Configuring a Namespace

Once the StorFirst Apollo Namespace Wizard appears, the system will detect namespaces on the tenant and automatically populate them. Select the namespace(s) you want to replicate and click Finish.



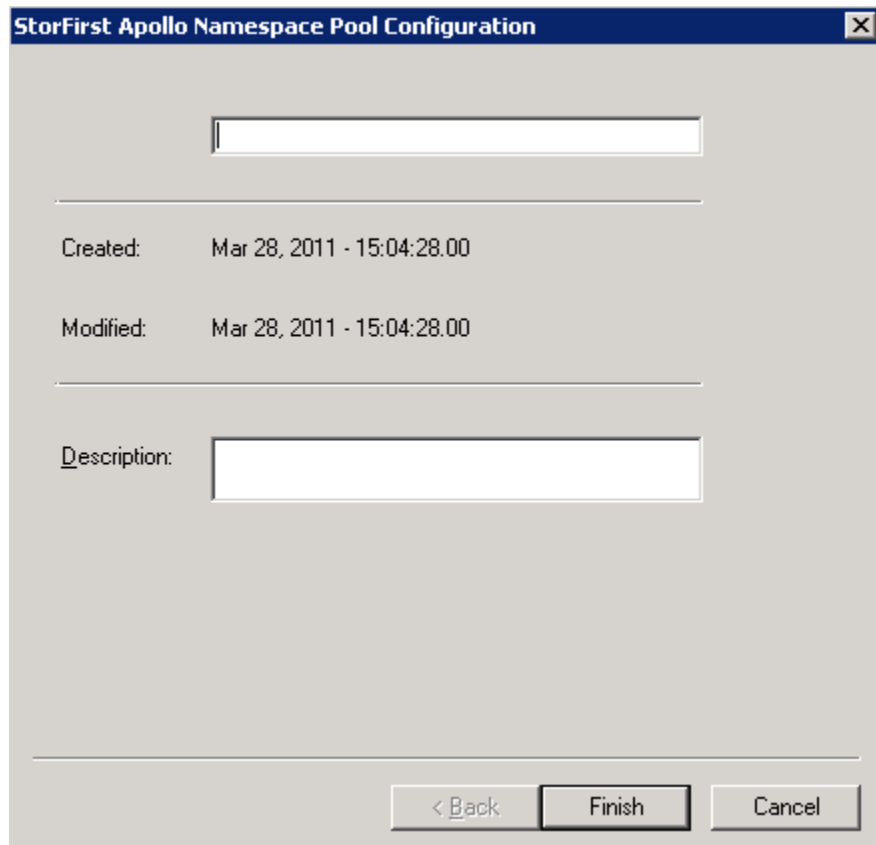
Enter the namespace name as it is in the HCP system. If this is a new installation, you must first create a namespace pool to which the configured namespace will be assigned. To do this, click Create Namespace Pool.



The screenshot shows a 'Properties' dialog box with a title bar containing a close button. The main area is titled 'Options' and contains the following text: 'Enter the namespace name as it is on the system.' Below this is a text input field containing 'PerfTests-1K'. The next line says 'Assign this namespace to a namespace pool.' followed by 'A namespace pool is a collection of namespaces that share the same tape set.' Below that is a dropdown menu labeled 'Namespace Pools:' which is currently empty. A 'Create Namespace Pool' button is located to the right of the dropdown. At the bottom of the dialog are three buttons: '< Back', 'Finish', and 'Cancel'.

Enter a name for your namespace pool and click Finish.

After creating a Namespace Pool, click Finish on the Properties window to complete your configuration of the namespace.

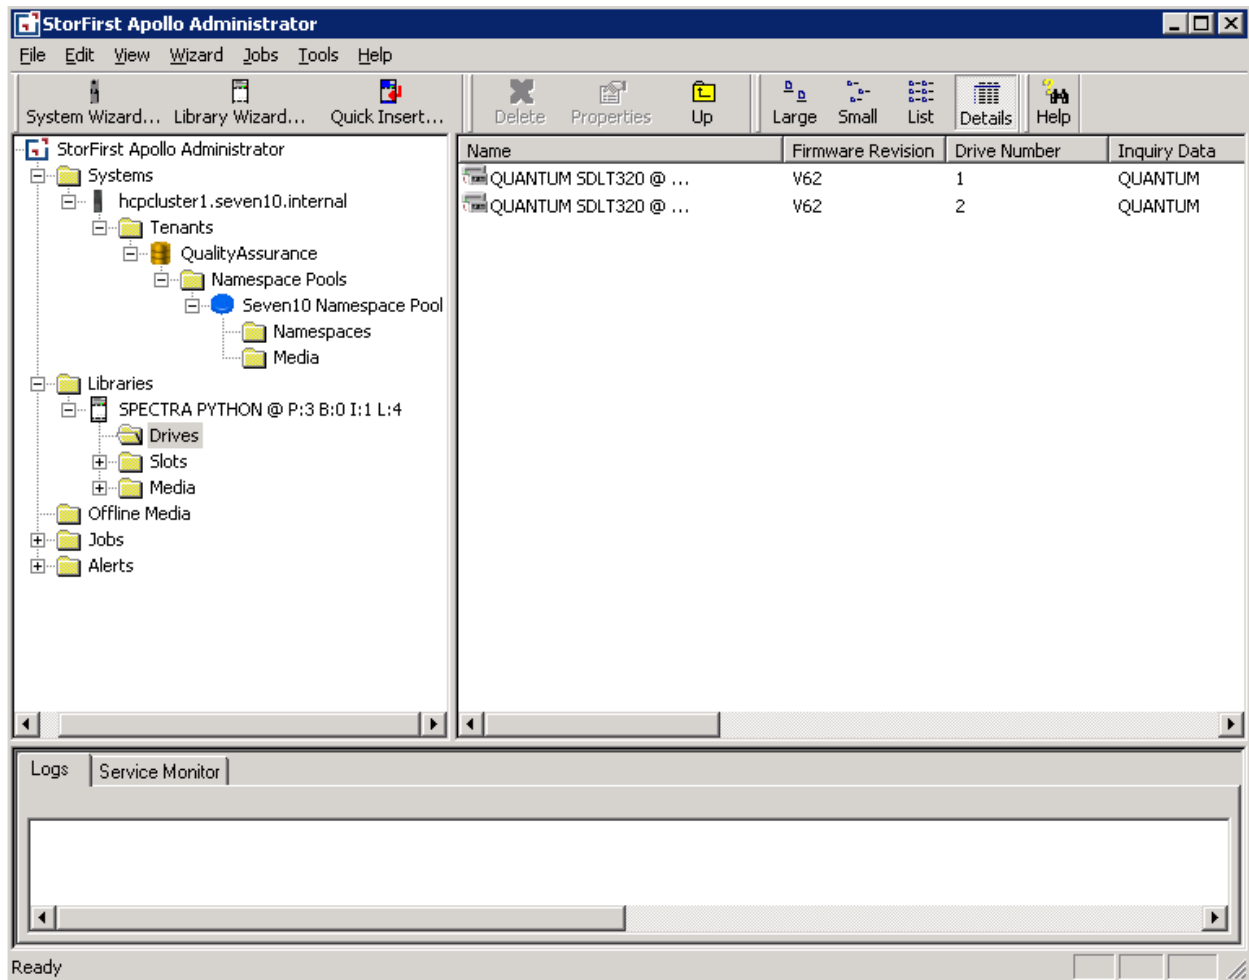


The screenshot shows a 'StorFirst Apollo Namespace Pool Configuration' dialog box with a title bar containing a close button. It features a text input field at the top for the namespace pool name. Below this is a horizontal line, followed by 'Created: Mar 28, 2011 - 15:04:28.00' and 'Modified: Mar 28, 2011 - 15:04:28.00'. Another horizontal line is below the modification date. The 'Description:' label is followed by a text input field. At the bottom are three buttons: '< Back', 'Finish', and 'Cancel'.

CHAPTER 4

THE APOLLO MANAGEMENT WINDOW

The StorFirst Apollo Administrator management window has the familiar look of Windows Explorer and works in similar ways. Expand the entries in the folder pane on the left, view content and details in the pane on the right, and right-click on objects in either pane to find properties and operations that can be used on the selected objects.



This screen capture shows an expanded view of the folders in the left pane. In the right pane, it shows some tape drives that are provisioned for use by StorFirst Apollo. The bottom pane shows the tabs for viewing detailed system information. The Logs tab details progress of the jobs you are running or have run. The Service Monitor tab contains messages from the background service.

The Folder Pane

The folder pane on the left side of the window displays various expandable folders. These folders work the same way as folders in Windows Explorer, providing a hierarchical representation of systems, libraries, media, jobs, and alerts in the StorFirst system.

Expand **Systems** and you will see HCP systems. Listed under each system are the tenants defined in the system. Under each tenant, the namespace pools are listed. Within each namespace pool, the various namespaces are presented. Right-click on any item in the folder to access commands related to that item or to view the properties of the item.

Expand **Libraries** and you will see the drives, media slots, and media (All, Primary, Duplicate, and Scratch) within each library. Right-click on any item in the folder to access commands related to that item or to view the properties of the item.

Expand **Offline Media** to see all the media that is being managed offline.

Expand **Jobs** to see Pending and Completed. The Pending folder lists all jobs that are either in progress or that are scheduled to run at a later date or time. All time-consuming operations that you initiate in the Apollo system are accomplished by creating and running jobs. The Completed subfolder displays all jobs that have been completed and are not scheduled to run again. Right-click on any job in either of these folders to run, delete, reschedule or view status logs of that job. To view the log for an individual job, right-click on a job in the folder pane and select View Log. Job logs are kept until you delete the completed job. This applies both to jobs you initiate and to jobs that are run automatically on a schedule.

Expand **Alerts** to see Unacknowledged and Acknowledged. The items listed in the Unacknowledged folder are alerts that were issued by the system and that you have not yet addressed or viewed. Those listed in the Acknowledged folder are those that have received attention and you have signed off on. Right-click on an alert to view its properties.

The Details Pane

The details pane on the right side of the window lists the details about the contents of the item selected in the folder pane. Operationally, it works just like Windows Explorer. Tasks on single items are initiated by right-clicking on the item. If you need to perform a task on several items at once, highlight them together by holding down the Shift or Ctrl key, and then right-click to select the operation.

The Status Pane

The status pane, located at the bottom of the StorFirst Apollo Administrator window contains tabs for pages that each detail different StorFirst activities.

The **Logs** page displays interspersed events from all jobs in order of occurrence. This page provides an easy way to monitor the system and can be a fast indicator of any potential issue. Individual job logs are discussed above. The status line provides detail progress information for any inventory, uninventory, replicate, or restore job in progress.

The **Service Monitor** page displays the log from all services and service errors in the system. The details listed on this page are also stored in a log file.

The Menu Bar

The top part of the StorFirst Apollo Administrator window contains the standard Windows menu bar and toolbar, combining standard Windows menus and tools with StorFirst custom menus and tools. Included in more detail below are: Jobs, Tools, and Help.

System Backup Job

To run a system backup job, click on Jobs, and then System Backup in the StorFirst Apollo tool bar. The job automatically backs up the StorFirst Apollo database so it can easily be restored in the case of a disaster.

The screenshot shows a 'System Backup' dialog box. It has a title bar with the text 'System Backup' and a close button. The main area is titled 'Options'. Under this title, there are two sections. The first section is 'Select an existing base destination directory for backup.' It contains a text input field with 'C:\Temp' and a 'Browse...' button. The second section is 'Select a number of old backups to keep. Choose zero to never automatically remove any.' It contains a text input field with the number '1'. At the bottom of the dialog, there are three buttons: '< Back', 'Next >', and 'Cancel'.

Click the browse button to select a destination directory for the backup. The directory can be on a local or on a network drive. If the directory is on a network drive, the server must be in the same domain as a network drive and use the UNC path.

Click Next. You will then choose the schedule for the system backup job and click Finish.

Note: Only one System Backup job can exist at a time.

System Configuration

In the menu bar, select Tools and then System Configuration. You'll see the Log Files Page.

The Log Files Tab

This Log Files tab in the System Configuration window controls the volume of service logs kept by the StorFirst Apollo system. It can also remove the service log from view.

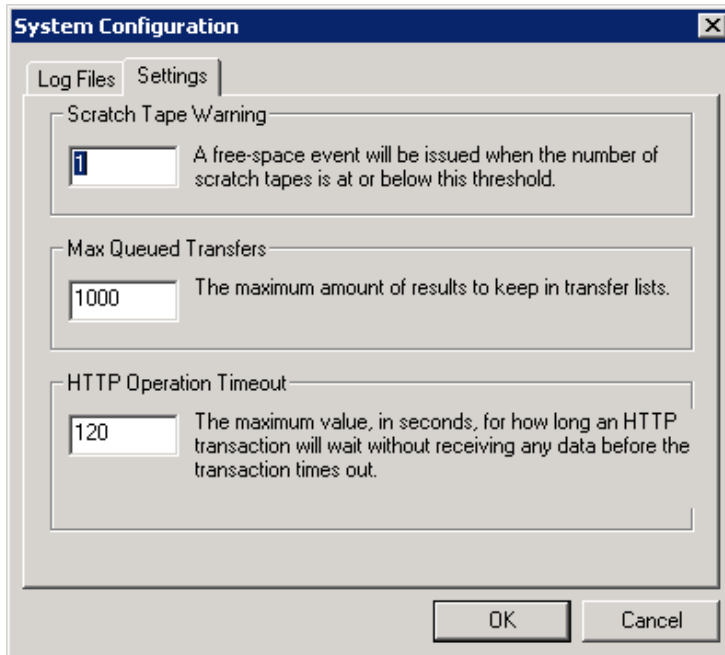
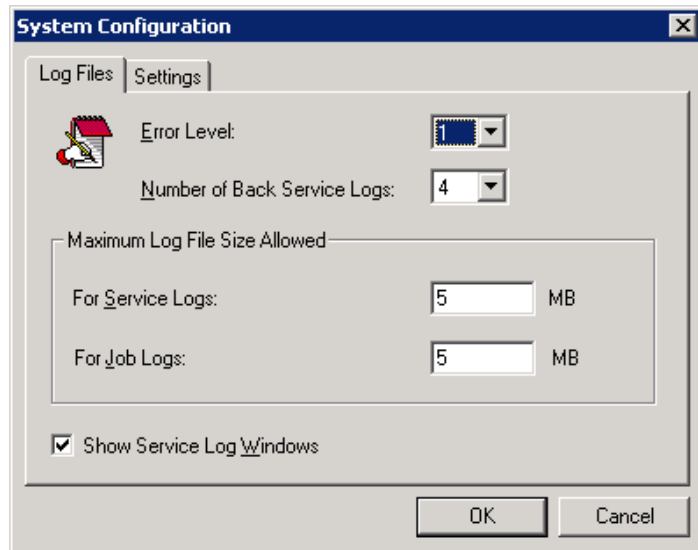
The Apollo system has three kinds of log files in addition to those shown in the Administrator window: job logs, service logs, and UI logs. Logs are saved to help resolve device errors as well as customer service call situations. They can be found in the StorFirst installation folder.

Use the defaults unless otherwise instructed.

The error level controls logging volume. The default is low, producing minimal logging. A log file will end when the specified size has been reached, and a new log-file will be started. The name of a job log file is the job name.

Disk consumption for system service logging is limited by log size multiplied by the number of such logs.

Jobs that complete and terminate when they are done will typically generate a single log file. Recurring jobs on a never-ending schedule will generate logs that use disk capacity limited by the specified size multiplied by the number of logs. Jobs are initiated manually and can be automated by scheduling them to run at specified times. The logs are kept until the job is manually deleted.



The Settings Tab

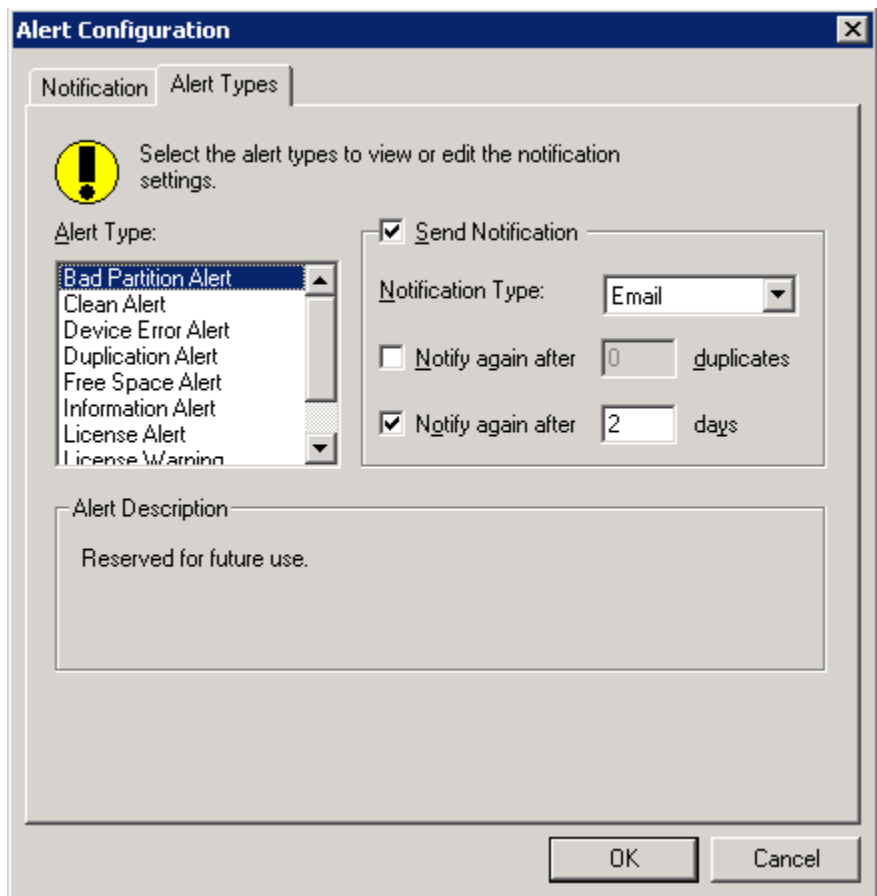
This Settings tab in the System Configuration window controls the warning issued when you are running low on scratch tapes. A free-space alert will be issued when the number of scratch tapes is at or below the number set.

Alert Configuration

In the menu bar, select Tools and then Alert Configuration. Select the Alert Types tab to view the various alerts.

Alerts represent important system events. All alerts are posted in the Unacknowledged folder and can be reviewed at the bottom of the folder pane.

Duplicate alerts are usually not delivered. Two observations of the same condition are considered duplicates. The reason duplicates are not delivered is that it is possible for the same condition to be discovered a large number of times inside a short span of time.



It is important to monitor and respond to alerts, as proper system operation may depend on timely operator assistance. Most alerts represent conditions that require physical access to the system and intervention of some sort. It's recommended that you acknowledge alerts at the time they are received. A pending and unacknowledged alert will, the next time that type of alert occurs, suppress the new alert, even if differences exist between the events that generate the alerts.

Note: *Two different conditions of the same type are considered duplicates.*

Alert Types

Configure each alert type individually, making selections for each type one at a time and repeating the process for all alert types.

However, the “Notify again after” option prevents message flooding while ensuring that conditions are not forgotten. The recommended renotification interval is one day.

There are eight alert types:

- **Bad Partition Alert** – Reserved for future implementation. Please contact StorFirst support for more details if needed.
- **Clean Alert** – Generated when a drive is reporting that it needs to be cleaned.
- **Device Error Alert** – Generated in the case of a general SCSI or SCSI timeout error for libraries, drives, or disks.
- **Duplication Alert** – Generated after a duplicate job is run. Includes instructions for any actions you may need to take.
- **Free Space Alert** – Generated if the system is running out of scratch tapes.
- **Information Alert** – Generated as an informational message. Includes instructions for any actions you may need to take.
- **License Alert** – Generated if the license for this installation has expired or exceeded total capacity.
- **License Warning** – Generated when the license for this installation is approaching the expiration date or total capacity.
- **Offline Alert** – Generated when a piece of media has been offlined and is required for a restore job. The recovery for this alert is to reinsert the media back into the library.

Alert Delivery

Select the Notification tab in the Alert Configuration window. The StorFirst Apollo Administrator software is accessible directly on the StorFirst server, which may be located far from the workplaces of the StorFirst Apollo administrators. It may be convenient to configure StorFirst to deliver alert notifications over the network to other systems that are closer to the administrators. Alerts can be delivered by email.

For email, the default email server is the Seven10 public SMTP server. In some cases your StorFirst Apollo server

Alert Configuration

Notification | Alert Types

Broadcast machine information

Specify the domain and the computer name the support person is logged in to broadcast a message when an alert occurs.

Domain name: (Optional)

Machine name:

Email Address

Specify the support person to contact when an alert occurs.

Email address:

Override default SMTP server settings here. Some SMTP servers do not require authentication, therefore you may leave username and password blank.

SMTP Server: Port:

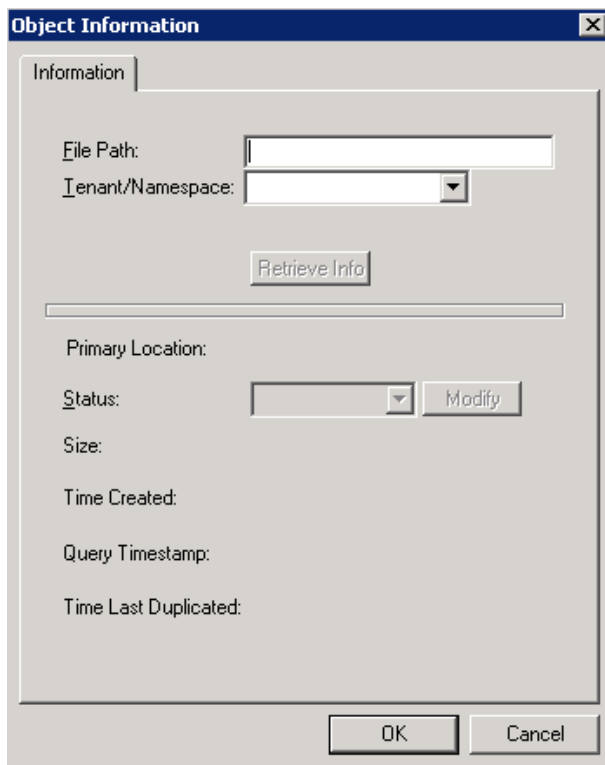
Username: (Optional)

Password: (Optional)

will be behind a firewall and will require a local SMTP server. This can be changed by entering the email address of the person to be contacted and the name of your company SMTP server. The username and password are optional.

Object Information

In the menu bar, select Tools and then Object Information. This option provides metadata information on the objects replicated by StorFirst Apollo.



Help/About

In the menu bar, select Help. There are three options:

Help Topics will direct you to the StorFirst Apollo Administration Guide.

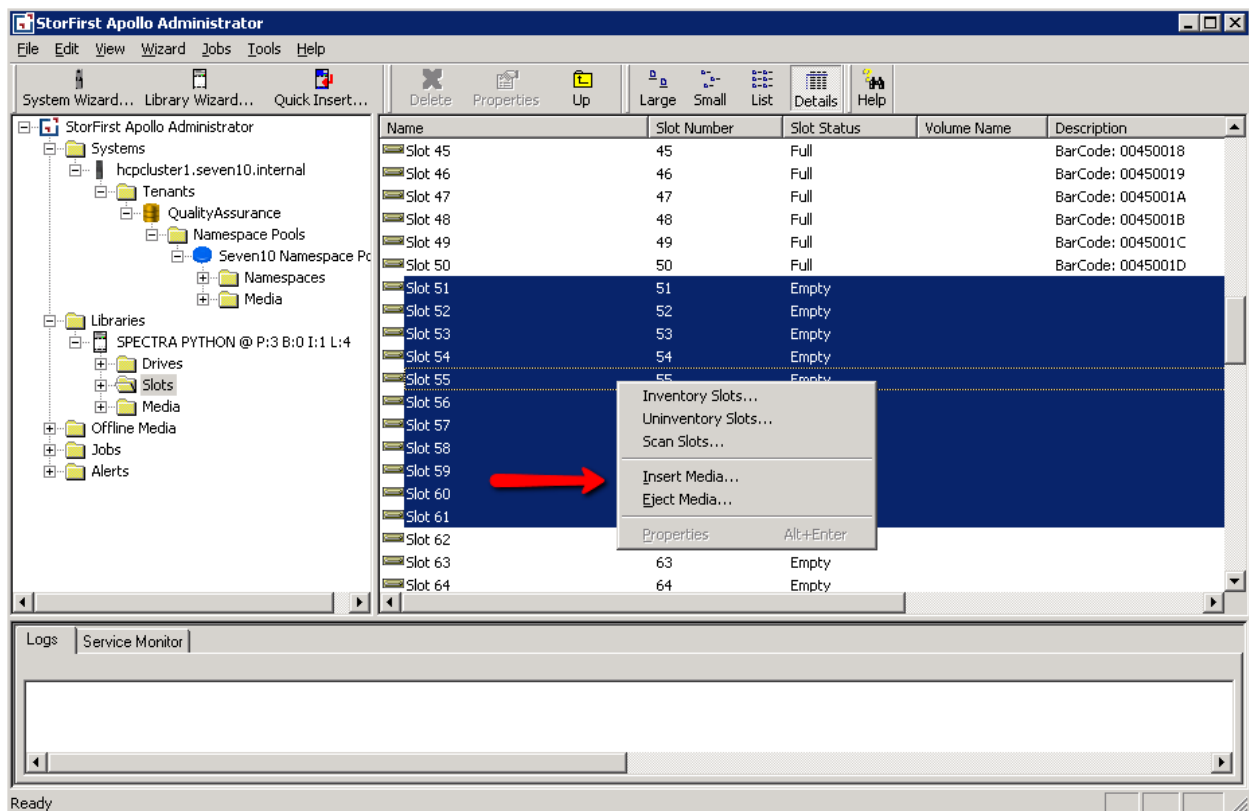
License Key will show you your installation ID, license key, and how much licensed capacity you have remaining.

About StorFirst Apollo Administrator provides the build and version of the software that you're using.


CHAPTER 5 TAPES & INVENTORY

Apollo keeps unused tapes in a namespace pool called the **Scratch**. In the folder pane in the StorFirst Apollo window, tapes that are partially or fully written are below the HCP system that is using the tape. The **Scratch** pool is stocked by formatting media using an inventory job. From there, media is provisioned automatically by replicate jobs as needed and placed in the appropriate HCP Namespace Pool media folder. Scratch media can also be assigned to namespaces.

If the library is empty, start by inserting media. Select slots, right-click, and select Insert Media.



Apollo will control the library and instruct you through the process.



Load the tapes into the library as instructed by the tape library's operational guide. Many libraries require

you to use an operator panel to open the door for media insertion.

After completion of the media insert operation, the slots will be full and the system will automatically launch an inventory job to deal with the just inserted tapes; this can be delayed by clicking Cancel. Full slots can be inventoried anytime. Select a slot or set of slots, right-click, and select Inventory Slots.

Inventory Job

Apollo protects its own data tapes from accidental formatting by not allowing existing data tapes from an Apollo instance to be inventoried. Any tape with valid data from this Apollo instance will be recognized and formatting refused. Any other tape will be overwritten, including tapes from other Apollo instances and tapes that were purged from this instance (see “Uninventory Job Offline” p. 31 for more information).

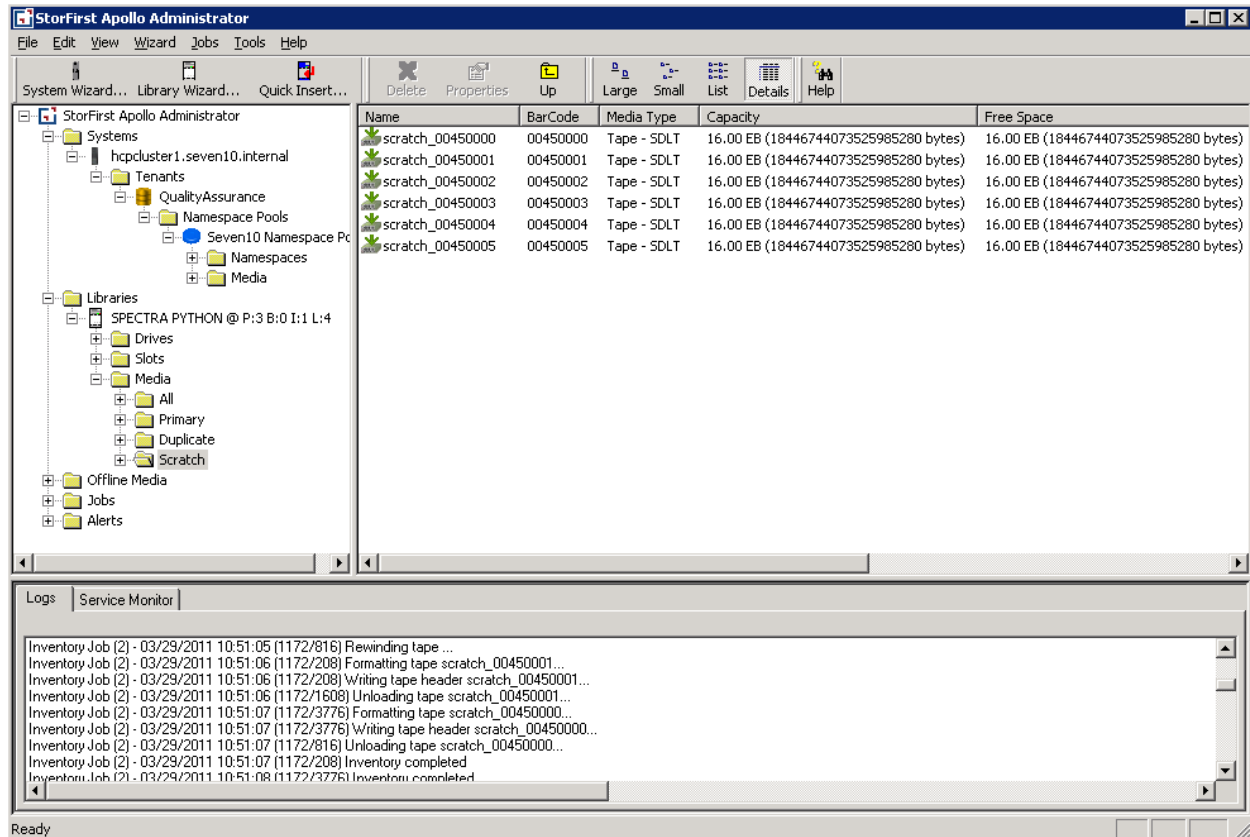
Tapes can be formatted with sizes ranging from 1KB to 2MB for better granularity, space saving, and performance, depending on the set of HCP namespaces. The general recommendation is to set the block size at half of the average object size of the HCP namespace.

Note: *Formatting will erase all prior data that may have existed on the tape.*

Select the Logs tab in the bottom pane of the StorFirst Apollo Administrator window and watch the progress of the inventory job. This job can also be found in the Pending job folder during execution. This job is likely to take some time, as every tape is loaded into a drive and a header is written.

The inventory job delivers the new media to the Scratch pool (or configured HCP namespace) where they

will remain until they are provisioned by some job. The same tapes can also be found in the All folder under the Media folder under the device.



After completion, the inventory job moves to the completed folder and remains there until manually deleted.

Tape Identification

All tapes must have readable bar codes. The initial tape name consists of “scratch” and the bar code. The final tape name is assigned when the tape is provisioned for use. Tapes are then identified by a time stamp and the bar code in the following format:

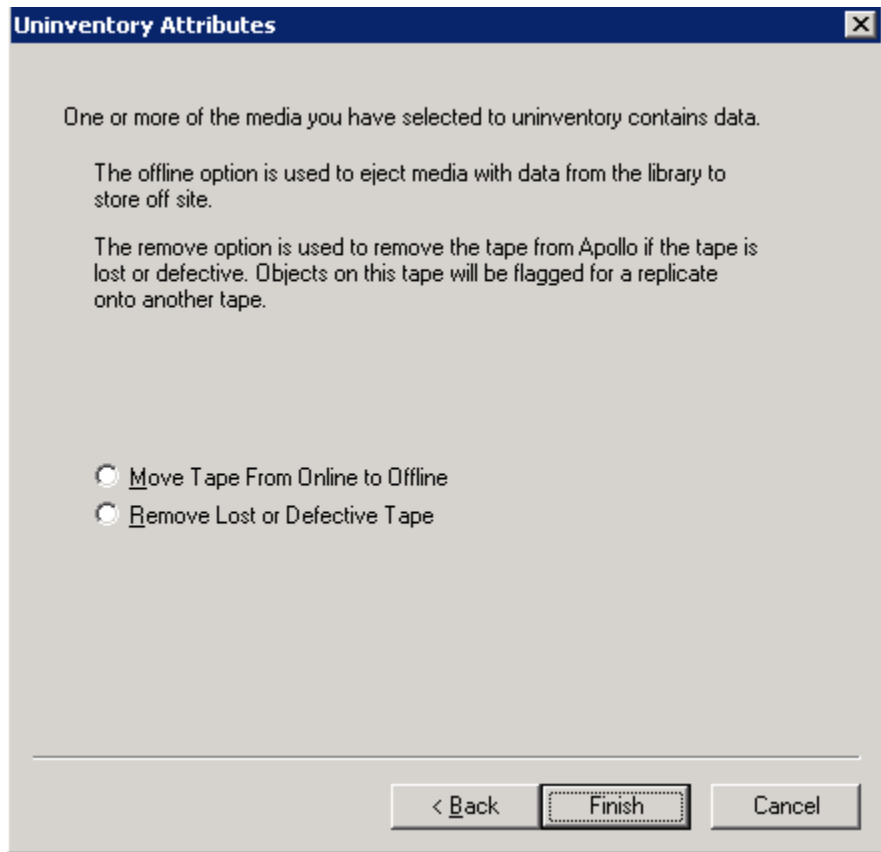
yyymmddhhmm barcode

As tapes are loaded from slot to drive, the bar code is checked using the picker robot, and the volume label is read from the tape and verified. In addition to the name, the volume label identifies the tape as a proper Apollo-formatted tape and also identifies the HCP system that provided the objects. All operator references regarding offline tapes include the full name and thereby both identifiers.

Uninventory Job Offline

Data tapes are sometimes removed after completion for offline storage to make room for new blanks, increasing the effective capacity of the system and/or providing the added protection of off-site vault storage. The uninventories job prepares tapes for ejection. Select tapes, right-click, and select Eject. The Eject command automatically initiates an uninventories job. Follow the instructions on the screen and in your library manual.

For tapes that are removed for offline storage, select "Move Tape From Online to Offline." Apollo retains all internal information about the offline tape content in its database. A replicate job would know not to replicate the data again, and a restore job will prompt the operator to reload the tape. Offline tapes are kept visible in the offline media folder so they can be recognized when reintroduced by inventory.



Remove Lost or Defective Tapes

Tapes may go bad, become lost, or be physically destroyed outside the library. In such cases, these tapes should be removed from the library. This type of uninventories job removes all associations of objects to tape(s). Removal of a tape will cause replicate jobs to re-replicate the objects that were lost by removed tapes. Remove applies to both online and offline tapes.

Purge from the Database

Purge from the database is done by first uninventories all the tapes associated with the namespace and then right-clicking on the namespace and selecting Uninventories Namespace. This type of uninventories job

removes all references to the objects from the database. The tapes associated with the namespace can now be formatted for reuse or inventoried as a database restoration option.

Note: The *Uninventory Namespace* option should be used only if you are removing the HCP namespace.

Reinstating Tapes for Database Recovery

In normal operation, this form of inventory can be used to reverse an accidental purge operation. Replication may already have picked up some objects, but no harm is done, except for wasting space on tape, because duplicates are ignored.

If this operation is incorrectly used during normal operation, it will often fail. The error message will indicate one of the following reasons:

- This is not an Apollo tape.
- This is a known offline tape.

Reinserting Offline Tapes

Reloading a tape using the “Insert Media” option will change the slot status to full. These tapes may be subsequently recognized as online by an inventory job working on the filled slots.

Apollo prevents accidental formatting of offline tapes. Apollo rejects random tapes presented as offline.

Inventory Attributes

Inventory Attributes:

Reinstate media in Database Recovery Mode.

Missing Primary Tape Override - Select this option if wish to inventory any duplicates found even if their primary tapes are missing.

Reinstate media from Offline to Online.

Format and Erase all previously recorded data.

During format, media will be added to the pool indicated below. If this job encounters WORM media, nothing will be written to the tape until it is used for replicate.

Add To Pool:

Block Size: KB

Note: This block size may be reduced if it is larger than supported by the drive or HBA. Check log file after formatting for an error code 5.

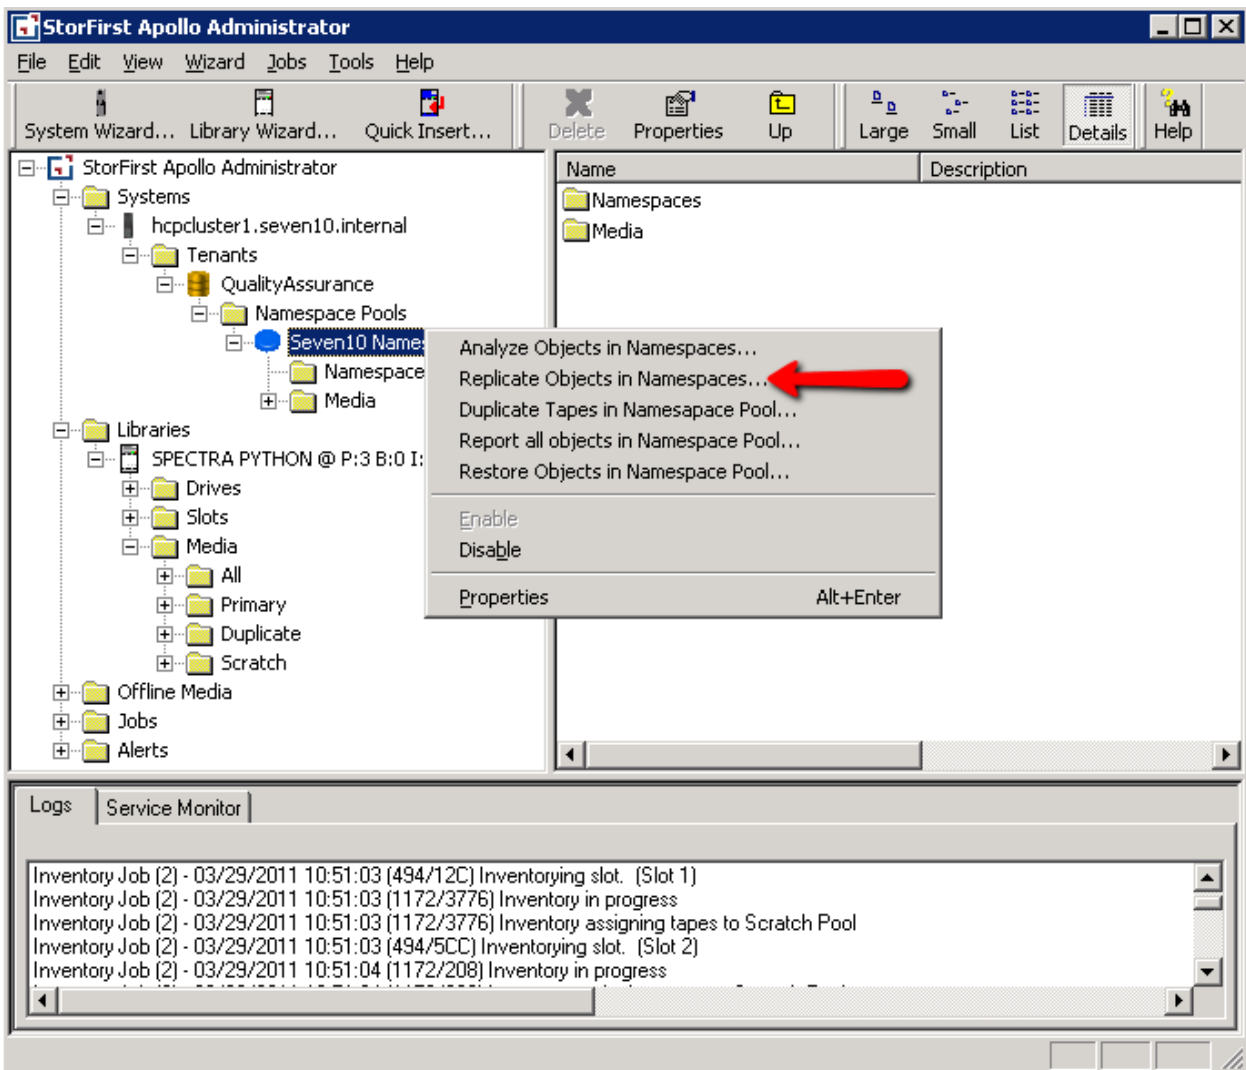
< Back Next > Cancel

In the event of a library replacement, this inventory job performs the very useful procedure of recognizing all the tapes after their relocation.

CHAPTER 6 THE REPLICATE JOB

Everything is now in place: HCP, the tape/VTL library, and a set of blank tapes. Apollo replicates in lights-out environments with all functions automated. Apollo replication is fundamentally different from traditional backup. Apollo replication builds a single, complete, and always synthesized copy of all the information in HCP by incrementally capturing new objects. Objects cannot change so there is no need to replace old with new. Simple accumulation is all that is necessary.

Start by right-clicking on the HCP namespace pool you want and selecting Replicate Objects in Namespaces.



This command will create the ongoing replicate job that will be the normal state of operation unless interrupted by some exception event. A predictable exception event could be Apollo running low on media. Equipment, site, or network failures represent other, less predictable, exceptions. HCP object storage may have begun before installation of Apollo, and HCP may or may not contain objects already. Regardless, we will create a replicate job that will replicate all the objects from HCP to tape. Initially, the job may be in catch-up mode. Eventually the job will settle into a mode of incremental replication while also monitoring deletion of objects on HCP.

This job will usually be a permanent part of the system, so you should name it for easy recognition. You may have several jobs on different schedules or serving different HCP systems depending on configuration or windows of opportunity.

Replicate Job Attributes

It is easy to measure the actual performance of a replicate job by using the information displayed in the job status line. The measured rate of replication can be used to calibrate the replication schedule if, for example, the average daily archiving volume is known. Increasing the number of drives will increase performance if HCP and the network can keep up. Vary the number of drives and compare the results. Environments with large object sizes (> 5MB) will also make good use of multiple drives.

HCP uses a multiple nodes for high availability and scalable performance. Performance is gained only to the extent that the clients are multithreaded. To that goal, the replicate job offers multiple connections, each with its own thread. While the aggregate performance is increased by adding threads, each thread will consume a portion of the total available HCP resources. Therefore, decrease the connections to limit the share of resources taken up by Apollo versus other concurrent HCP clients. If Apollo will be scheduled to run replication during off hours, use the maximum setting; otherwise, reduce the setting accordingly (for scheduling, see

Replicate Attributes

Maximum Number of Connection Threads: Maximum number of connection threads relates to the system performance sharing. Consult the documentation for recommendations.

Maximum Number of Concurrent Drives: This option will limit the number of drives used by this job, allowing sharing of available drives for other jobs. Note: You cannot have more drives than connections.

Continuous Replication: Continuous Replication Job will run until canceled or time expires, even if the task is complete.

Query Options:

Include Deletes Query the namespaces for deleted objects.

Reset Incremental Replication Query the namespaces from the earliest object available. NOTE: This option will force a rescan that will be very time consuming.

Date/Time Range Override:

Start Query for objects between the specified start and end dates. Do not use these options unless you need to query for specific date ranges only.

End

Duplicate Tape(s) Option: This option will start an automatic tape duplication job on the namespace pool immediately following the completion of this job. NOTE: For this option to have an effect, there must be a namespace pool tape duplicate job created for this namespace pool using the automatic schedule option.

Generate Event Upon Completion: This option generates an information alert upon completion of the job. The last known status of the job is emailed to the support contact.

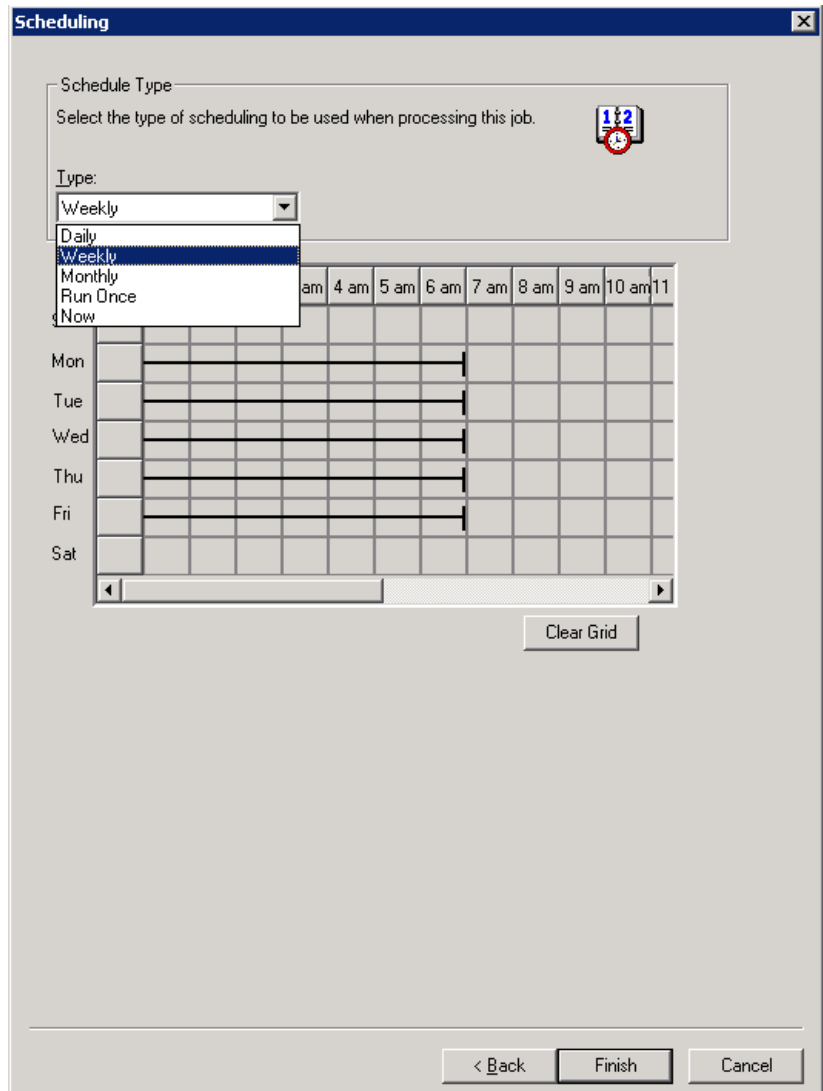
< Back Next > Cancel

“Job Scheduling” section p. 36).

Incremental replication implies that as Apollo scans an HCP system, it will not waste time reviewing the same objects every time because once objects are captured, there is no need to review them again. The Reset Incremental Replication option will force a rescan from the beginning, and all objects will be verified against the Apollo database. This option is time consuming and should be avoided.

Job Scheduling

In the menu bar, select Tools and select System Back-up. Choose the destination directory to where you'd like to back-up the StorFirst Apollo database, and choose the number of copies to make. Select the job scheduling option that is most suitable to your needs. The Now option runs the job once until completion. Jobs that are run once are in the Pending job folder during execution and then move to the Completed job folder. The simplest option for scheduling a recurring job is Daily, starting at a set time. Jobs on a recurring schedule never complete and remain in the Pending job folder.



Advanced Scheduling

Think about the HCP object storage and retrieval load patterns. For example:

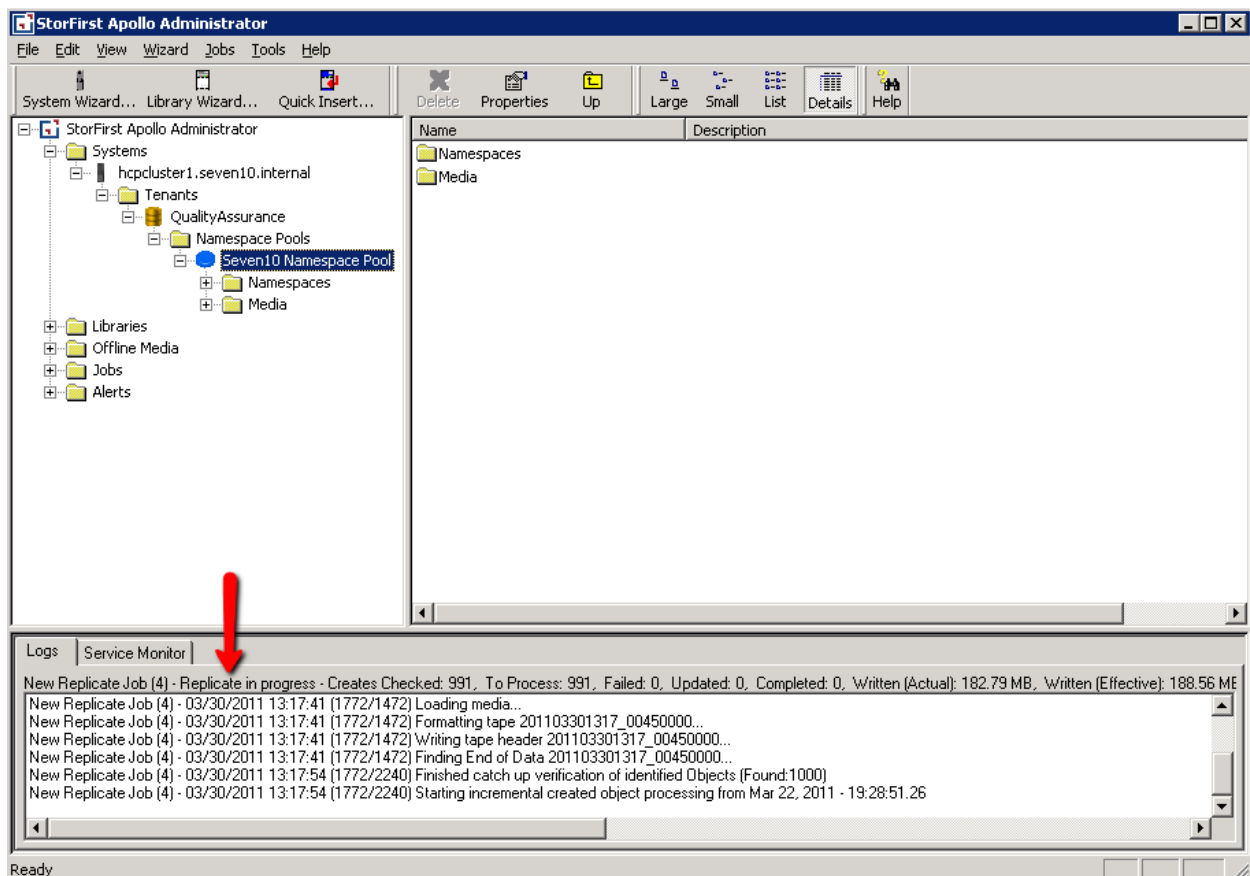
- Daytime has light retrieval and low tolerance for long latency
- Early evenings have high archival rates
- Night and weekend activity is unusual and low priority

To match these patterns, you can create a weekly job that is running late nights during the work week and all weekend long to catch up, just in case.

Replicate Job Execution

The replicate job will first replicate any objects that were lost by purging tapes, then replicate newer objects, and finally scan the entire HCP namespace pool for deleted objects and mark the database accordingly. Note that there is a lag time between deletion in HCP and discovery by Apollo because Apollo and HCP clients are not synchronized.

The replicate job provides up-to-date status information in the status line at the top of the Logs page.

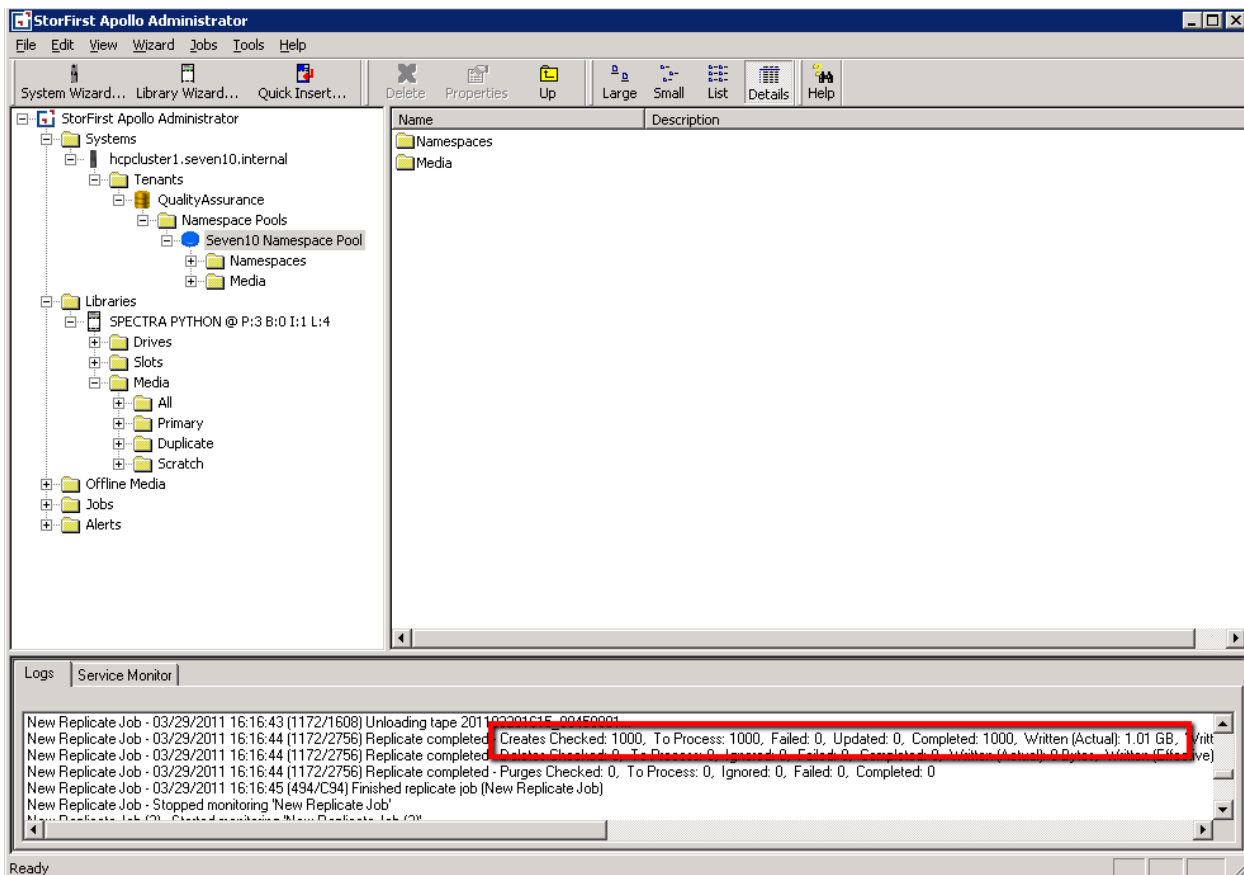


Viewing the image below, you'll find a description of each type of object in the red highlighted box:

Creates Checked: Shows the progress of scanning for objects in the namespace in chronological order, starting where the last job ended

To Process: Represents the number of objects that were placed in the job queue for replication. (Replication is a separate set of threads from scanning.)

Failed: Counts the objects that had a replication failure during read, implying some unexpected error transferring the object. This count should normally remain at zero. Failed objects will be logged by the service log.



Updated: Represents duplicate HCP query results when replicating.

Completed: Counts the objects successfully transferred to tape.

Written (Actual): Amount of data that the completed objects represent.

Written (Effective): Amount of data that the completed objects represent on tape but differs from Actual due to block size settings on tape.

Deletes Checked and Purges Checked: Represents objects which were deleted or purged in the namespace before or after replication by Apollo.

When the job finishes, the starting time stamp for the next run of the replicate job will be saved in the database so that, for the next job, it will resume replicating from that point. This function can be overridden and reset to the beginning of time by the "Reset Incremental Replication" attribute of the replicate job.

Multiple Replicate/Restore Jobs

Apollo can have multiple replicate and restore jobs. Replicate and restore jobs can run simultaneously on namespaces in different HCP systems, as well as on multiple namespaces of the same HCP system. If you want to replicate or restore multiple namespaces on the same HCP system, see the following requirements for each type of job.

Replication Job

One requirement for a replication job is that there must be at least one drive freely available for the job to use when started. If the job specifies to use more drives than are available, the number of drives will be scaled back to the number not in use. If there are no drives available, the job will quit.

Another requirement is that there must be enough threads from the free thread pool for the specified number of transfer threads. The replication job will not scale back the number of threads automatically. Therefore, you need to be careful to configure overlapping scheduled replication and restore jobs appropriately. Please consult your StorFirst Apollo support representative for recommendations on setting up multiple overlapping jobs.

Restore Job

The requirements for the restore job are that there must be at least one drive available. If there are no drives available, the job will quit.

Another requirement is that there must be enough threads from the free thread pool for the specified number of transfer threads. The restore job will not scale back the number of threads automatically. Therefore, you need to be careful to configure overlapping scheduled replication and restore jobs appropriately. Please consult your StorFirst Apollo support representative for recommendations in setting up multiple overlapping jobs.

The restore job will automatically mount any and all available (that is, not in-use) drives for its own use. Therefore, you need to be careful to configure overlapping scheduled replication and restore jobs appropriately.

Tape and Drive Use

When a scratch tape is provisioned from the scratch pool, it is named using the Apollo timestamp, the name of the HCP namespace using it is written to the tape, and, from then on, the tape is reserved for use by that namespace. Objects on tape are broken up in blocks and interspersed with blocks of other objects. All blocks making up a complete HCP object will be on a single tape. Objects do not span multiple tapes.

Tapes are first selected from the media pool belonging to the specific namespace and then from the scratch media pool. Tapes with the least remaining available capacity are selected first. If there are no available tapes, the job session will terminate with an alert.

Each job uses all available tape drives up to the limit set by the job attributes. If there are no available tape drives, the job waits until a drive becomes available or until its time slot expires. The total number of partially written tapes in a namespace media folder will not exceed the maximum number of drives used by its replicate jobs.

Objects are acquired in parts by a set of parallel threads and queued in memory. For the highest possible performance, the data is streamed to tape using a cyclic buffer in memory.

When the replicate job reaches the end of a tape, the tape is closed. All together, the tapes used by a single HCP namespace constitute a useful set.

Replicate Verification

The replicate job is constantly scanning HCP for new objects. The process reviews all objects in the namespace and will replicate objects that are not found in the Apollo database. The Apollo database is used to ensure that each object is captured exactly once. Objects that were deleted in HCP are identified and marked in the Apollo database.

Auditing Replication

Apollo offers the option of external audits of its replication process. Right-click on an HCP tenant, namespace, or namespace pool and select Report all Objects. This will start a reporting job that will produce a text log file of all objects in Apollo. The text file can be found in the Jobs Log folder of the Apollo installation directory. The file name will be the same as the job name except that the extension will be TXT. The format of the spreadsheet is Object ID, HCP Name, Namespace Name, Status, Create Time (GMT), Write Time (GMT), Object Size, and Physical Media Name. The screen capture below shows an object report for an HCP system in Microsoft Excel format. Apollo object reports can be easily imported into Excel by doing the following:

1. Open Excel.
2. In Excel, open the object report .txt file.
3. Choose the Delimited option in the first screen of import. Do not choose the fixed-width option.

4. Click Next.
5. Make sure only the Comma option is checked.
6. Click Finish.
7. For readability, expand the header columns (double-click on the right edge).

The screenshot shows an Excel spreadsheet with the following data:

Record Number	Version	Tenant Name	Namespace Name	Status	Ingest Time	Change Time	Retention Time	Hold	Object Size
6	8.32526E+13	QualityAssurance	bobversdisabled	Created	3/21/2011 21:11	3/22/2011 19:25	0	False	821344
70	8.32526E+13	QualityAssurance	bobversdisabled	Created	3/21/2011 21:12	3/22/2011 19:26	0	False	2042576
134	8.32526E+13	QualityAssurance	bobversdisabled	Created	3/21/2011 21:12	3/22/2011 19:26	0	False	597744
198	8.32526E+13	QualityAssurance	bobversdisabled	Created	3/21/2011 21:12	3/22/2011 19:26	0	False	170528
262	8.32526E+13	QualityAssurance	bobversdisabled	Created	3/21/2011 21:13	3/22/2011 19:26	0	False	1576096
326	8.32526E+13	QualityAssurance	bobversdisabled	Created	3/21/2011 21:13	3/22/2011 19:26	0	False	288592
390	8.32526E+13	QualityAssurance	bobversdisabled	Created	3/21/2011 21:13	3/22/2011 19:27	0	False	419888
454	8.32526E+13	QualityAssurance	bobversdisabled	Created	3/21/2011 21:12	3/22/2011 19:27	0	False	1082544
518	8.32526E+13	QualityAssurance	bobversdisabled	Created	3/21/2011 21:13	3/22/2011 19:27	0	False	2034128
582	8.32526E+13	QualityAssurance	bobversdisabled	Created	3/21/2011 21:13	3/22/2011 19:27	0	False	900032
646	8.32526E+13	QualityAssurance	bobversdisabled	Created	3/21/2011 21:13	3/22/2011 19:27	0	False	625024
710	8.32526E+13	QualityAssurance	bobversdisabled	Created	3/21/2011 21:12	3/22/2011 19:27	0	False	1528480
774	8.32526E+13	QualityAssurance	bobversdisabled	Created	3/21/2011 21:12	3/22/2011 19:28	0	False	1170944
838	8.32526E+13	QualityAssurance	bobversdisabled	Created	3/21/2011 21:11	3/22/2011 19:28	0	False	1072592
902	8.32526E+13	QualityAssurance	bobversdisabled	Created	3/21/2011 21:11	3/22/2011 19:28	0	False	1845376
966	8.32526E+13	QualityAssurance	bobversdisabled	Created	3/21/2011 21:11	3/22/2011 19:28	0	False	1286992
7	8.32526E+13	QualityAssurance	bobversdisabled	Created	3/21/2011 21:11	3/22/2011 19:25	0	False	1154432
71	8.32526E+13	QualityAssurance	bobversdisabled	Created	3/21/2011 21:12	3/22/2011 19:26	0	False	1873408
135	8.32526E+13	QualityAssurance	bobversdisabled	Created	3/21/2011 21:12	3/22/2011 19:26	0	False	1074608
199	8.32526E+13	QualityAssurance	bobversdisabled	Created	3/21/2011 21:13	3/22/2011 19:26	0	False	884336

CHAPTER 7

TAPE DUPLICATION

Tape Duplication

The purpose of tape duplication is to create full or incremental daily offlined backups of data. The driving force behind this feature is to provide reassurance that an object is located in more than one physical tape location.

In HCP, the data for an object never changes. For this reason, the replication job in Apollo is essentially always incremental. Once an object is successfully put onto a physical tape, there is no need to check in HCP to see whether that object has changed.

Two important terms used to describe this feature are *primary tape* and *duplicate tape*. A primary tape is the tape a replication job uses to transfer objects from HCP onto a tape. A duplicate tape is used by a duplication job to transfer objects from a primary tape to another tape.

HCP backup differs from traditional file system backup. The total set of data on a tape in Apollo is unique. Therefore, the goal of the duplication feature in Apollo is to have one full tape and one full duplicate.

The most practical idea is to keep one tape, which simply fills up constantly, in the library at all times. Each day, the data on that tape is replicated to another tape. This other tape either could be kept in the library to be used again for more duplicate data or could be offlined and placed in a secure location.

An incremental and offline scenario may require three tapes for the data on any given day until a full second copy can be offlined permanently. These tapes include:

- The primary tape used every day for the first copy of replicated data
- A full duplicate tape that keeps a second copy of all the replicated data
- One (or more) incremental tapes for duplicate data that has yet to make it out to the full duplicate tape, most likely because that full duplicate was offlined for a time

Tape Duplication Details

The two ways to create a duplicate of data are manual and automatic.

Manual Type

A manual duplicate can be created by right-clicking on any primary tape and selecting Duplicate Tape.

The screenshot shows a dialog box titled "Options" with a close button (X) in the top right corner. The dialog is divided into four sections:

- Duplication Mode:** Contains three radio button options:
 - Full** Duplicate entire contents of primary tape
 - Append** Duplicate new objects since last duplication to single tape
 - Incremental** Duplicate new objects since last duplication using multiple tapes
- Target Tape Partition:** Includes the instruction "Choose a partition to use as a duplicate. Only scratch tapes can be used." and a dropdown menu that is currently empty.
- Eject Notification:** Includes the instruction "Enabling this option makes this job notify the administrator daily, via email, which tapes should be ejected or inserted each day." and a checkbox labeled "Enable Eject Notification" which is currently unchecked.
- Tape Rotation:** Includes the instruction "Select a day to always perform a full duplicate, regardless of tape rotation. Saturday and Sunday are available for 7 tape sets only." and two dropdown menus:
 - The first dropdown is set to "Friday" and is labeled "Day to always create a full duplicate".
 - The second dropdown is set to "5" and is labeled "Number of incremental tapes to use per primary tape".

At the bottom of the dialog, there are three buttons: "< Back", "Next >", and "Cancel".

With a manual duplication, the tape to use as a duplicate must be specified by the user. Only empty scratch tapes are available in the dropdown box.

Modes

The two modes for a manual duplication job are Full and Append:

- If you select Full, the job will format or reformat the selected tape and copy the entire contents of the primary tape to the duplicate tape.

⚠ Note: Because this job reformats the duplicate tape every time, there is the possibility of exposure to lost data should the primary tape become corrupt before the duplicate is created. Therefore, it is not recommended to rerun a manual duplication job using the same tape as a target.

- If you select Append, the job will format the selected tape once and copy the entire data content of the primary tape to the duplicate tape. If this job is run a second time, it will copy, by appending to the duplicate tape, only the data content which was new to the primary tape since the last duplication.

Automatic Type

One or more duplicates can be created automatically by right-clicking on a namespace and selecting Duplicate Tape(s) in Namespace.

With an automatic duplication, the job will automatically decide which tapes to use as duplicates. The job will search for eligible tapes from the namespace first and then from the scratch pool, if required.

Requirements of eligible tapes are:

- Must have a block size less than or equal to the block size of the primary tape
- Must have a total size greater than or equal to the current amount of data on the primary tape
- Must be enabled
- Must be online
- Must be blank or already part of the incremental set for this primary tape

Modes

The two modes for an automatic duplication job are Append and Incremental:

- The purpose of the Append mode is to incrementally write to one duplicate tape until the primary tape is full, at which point both tapes are finalized and become read-only. At this point, one or both of the tapes can be offlined.

If this option is selected, the job will format the selected tape once and copy the entire data content of the primary tape to the duplicate tape. When this job is run a second time, it will copy, by appending to the duplicate tape, only the data content which was new to the primary tape since the last duplication.

- The purpose of the Incremental mode is to incrementally write to two or more duplicate tapes until the primary tape is full. These two or more duplicate tapes can be offlined after each duplicate job session, normally each day, in order to ensure a two-site data disaster recovery scenario.

Eject Notification

If the Enable Eject Notification option is checked, a duplication alert will be issued by the system if you need to take action. Please see “Alerts and Notifications” section (p. 25) that deals with alert notifications and types. The default notification method for this duplication alert is by email.

For example, if incremental mode is selected, after a duplication job is run, an instruction will be issued to eject the duplicate tape from the library. If a tape is to be used as the next tape for a duplication job and that tape is offlined, the instructions will also indicate which tape to bring back into the library and insert.

Whether or not this option is checked, you can view these instructions by looking at the job instruction log. To view this log, right-click on the job and select View Instruction Logs.

Tape Rotation: Number of Tapes in Set

This “Number of incremental tapes to use per primary tape” option specifies how many tapes to use in an incremental set *per primary tape*. The minimum number of tapes is two, which will create a full duplicate per primary tape each day using two tapes.

If no new data is written to the primary tape during any given day, the duplication job will detect this when run and will not advance the rotation unless the rotation or day specifies that a full duplicate is scheduled.

Note: *The rotation is advanced each time the duplication job is run if and only if the duplication job actually transferred objects to a duplicate tape.*

Tape Rotation: Day to Always Create a Full Duplicate

The “Day to always create a full duplicate” option can be set to Rotation Only or to a certain day of the week. With rotation only, the rotation strictly decides on which days of the week to create a full duplicate. This can best be described through an example. If a three-tape set is selected:

1. On the first day the job runs, a full duplicate is created.
2. On the second day the job runs, an incremental duplicate is created.
3. On the third day the job runs a full duplicate is created using a third tape.
 - *Note: At this point, the data on the first and second tapes is considered unnecessarily redundant and the duplication job will overwrite these as necessary.*
 - *Note: If necessary, these tapes can be uninventoried and provisioned for use elsewhere.*
4. On the fourth day, the job will create an incremental copy of the primary tape since the last duplication, and it will use the first tape in the set (the one used in step 1 above) to do so.
5. On the fifth day, the job will create a full copy of the primary tape using the second tape (the one used in step 2 above).

When selecting a day to create a full duplicate, if the duplication job is run on that day, a full duplicate will be created, regardless of whether the rotation specifies to do so.

This rotation mechanism ensures that no duplicate data is ever overwritten unless it is redundant (that is, there are *more* than two copies of that data).

✎ Note: Saturday and Sunday are available only if you're using a seven-tape set to avoid duplicate creation on weekends, when user intervention may be required.

✎ Note: If the duplication job is run multiple times in the same day on the full-day override, it will create a full duplicate each time.

Duplication Job

At the completion of a duplication job, the duplicate tape will be fully part of the set of duplicates for the primary tape.

Once the primary tape becomes full, it gets finalized. This can be observed by noting that the free space on the primary tape is zero. If an incremental job encounters a primary tape that is finalized, the job detects whether a finalized duplicate exists and if not creates one. After completion of the duplication job, the job will ignore the primary tape. If you want to create another duplicate of the primary tape, this duplicate will need to be created manually.

For any manual duplication jobs, if the primary tape is finalized, the duplicate tape will also be finalized.

Uninventory

- If a primary tape is uninventoried because it became missing or was damaged, Apollo will promote the duplicate tape to a primary tape by assigning all the objects' primary locations located on the Apollo database to the duplicate tapes. Any objects that did not make it out to the duplicate tapes will be flagged for a re-replicate from HCP only if the objects were not deleted.

When possible, Apollo assigns the objects to tapes that are completely full, up-to-date duplicates of primary tapes. In this case, the full duplicate becomes a primary tape, and all incremental tapes become uninventoried and can return to the scratch pool to be used for other purposes. If no full up-to-date duplicate exists for the primary tape, all tapes, full or incremental, become primary tapes to be used for replication jobs. These tapes then become eligible to have duplicates created of themselves.

- If a duplicate tape is uninventoried, all locations of objects on that tape are removed from the database and those objects become eligible to be duplicated to another tape.

✎ Note: An uninventory should not be performed unless either the administrator wants to offline tapes or primary or duplicate tapes are actually lost or defective.

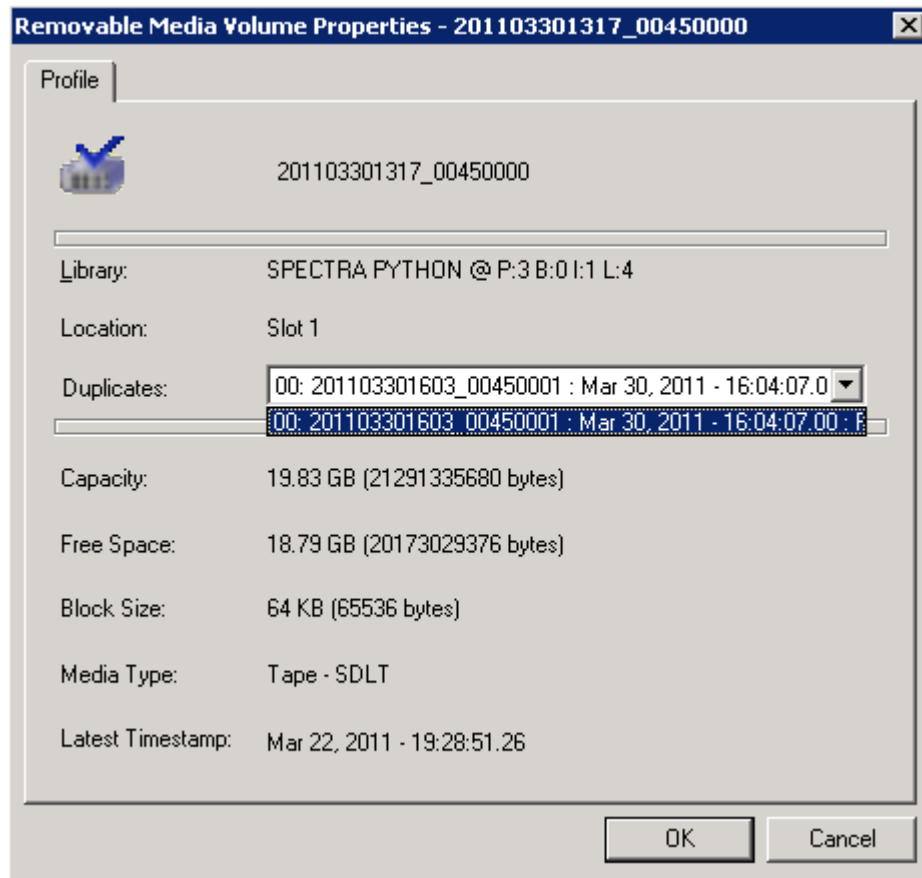
Inventory

- Inventory of primary tapes is not affected by the duplication feature.
- Inventory of duplicate tapes requires that the primary tape be inventoried first. If it is not, the job will not allow a duplicate tape to be inventoried. If, however, the primary tape was lost or defective, the duplicate tape can be inventoried as a primary tape by selecting the Missing Primary Tape Override checkbox as part of the inventory job.

User Interface

You can determine whether a tape is a primary tape or a duplicate tape in a few ways. The first is through the tape properties, which can be accessed by right-clicking on the primary tape in the StorFirst Apollo Administrator window. If the tape is a primary tape, it will list its duplicates, if they exist, in a dropdown box. The format of entries in the box is:

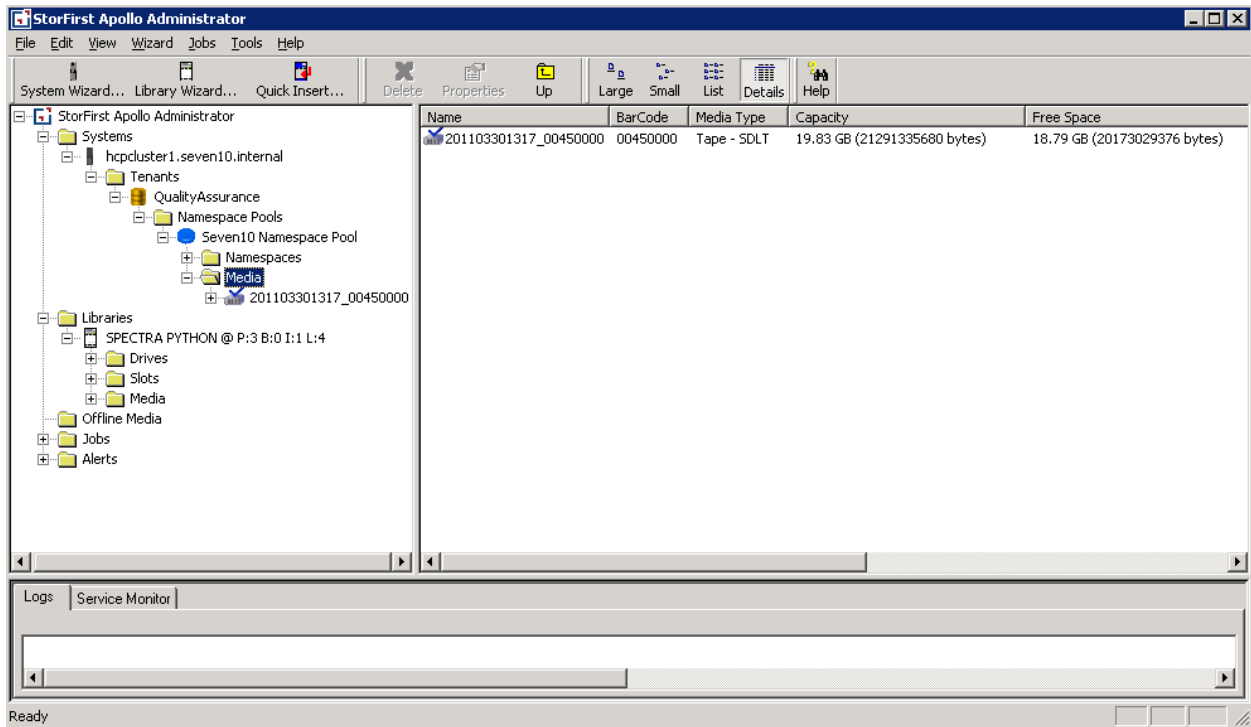
Number in Incremental Set: Duplicate Tape Label: Date Duplicate Written To : Full or Incremental.



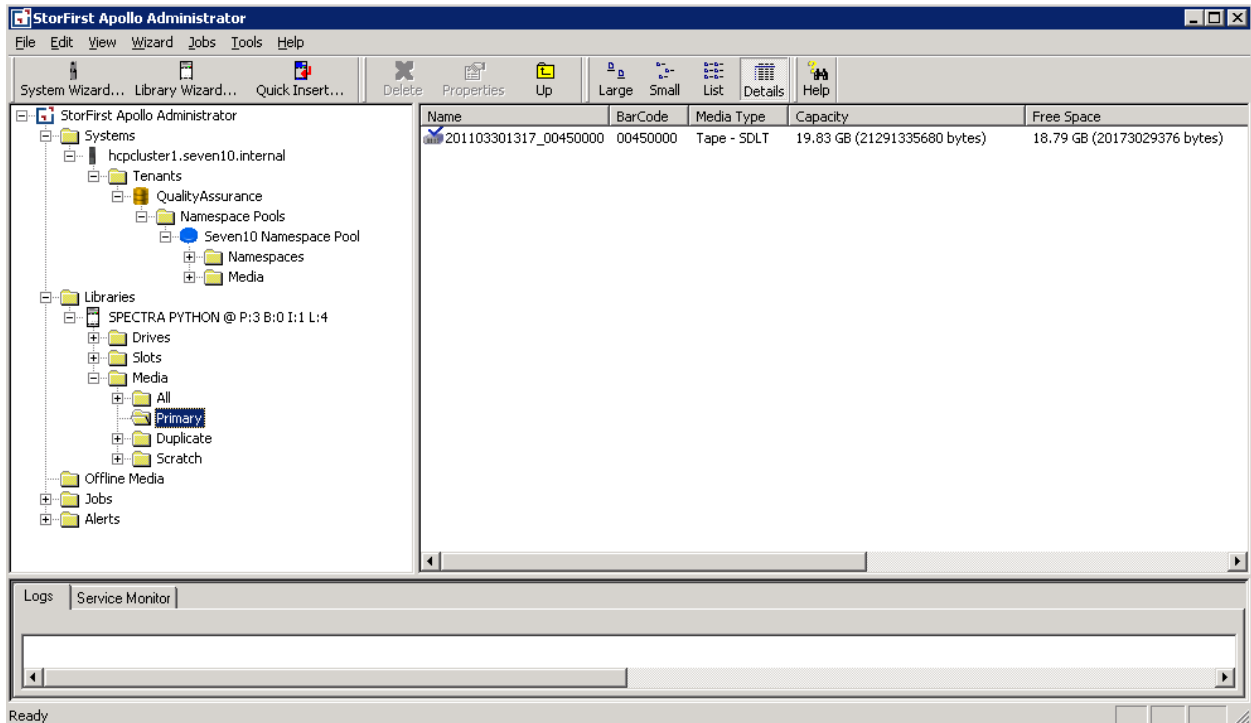
If the tape is a duplicate, it will indicate its primary tape in the dropdown box. The format of entries in the box is as follows:

Number in Incremental Set: Primary Tape Label: Date Duplicate Written To : Full or Incremental.

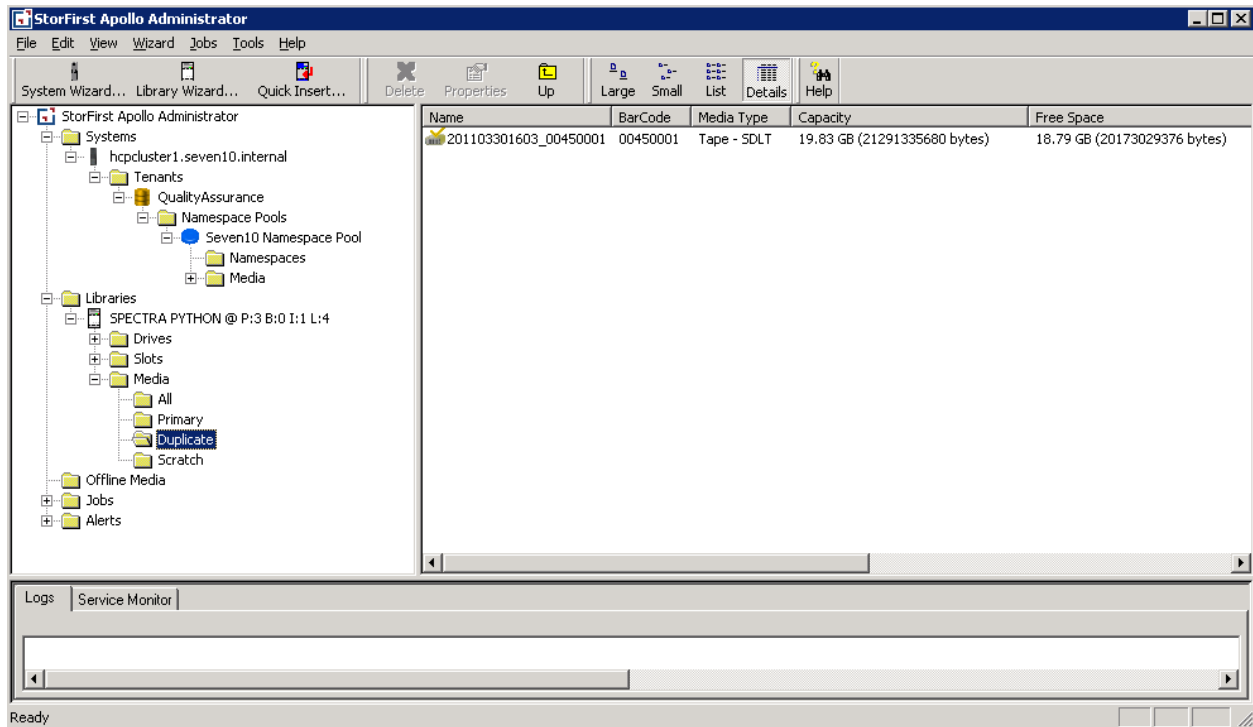
Primary tapes are listed under the Media folder for each namespace pool:



Primary tapes are also listed in the Primary folder within the Media folder for each library:



Duplicate tapes are located only in the Duplicate folder within the Media folder for each library:



CHAPTER 8

THE RESTORE JOB

This chapter describes restoration of objects from tapes to HCP. Restoring data from Apollo to HCP would be considered a disaster recovery scenario if it resulted from a HCP failure; however, restoration may also simply be needed to restore objects that were deleted by mistake, or it may be needed as a result of a change in configuration or equipment deployment.

Note: *Apollo allows only one active replicate or restore job per HCP namespace pool. Therefore, make sure to stop replication from the target before starting restoration.*

Target Selection

Restore Type

Choose the "Restore All" option if you would like to restore all objects from all tapes associated with the selected device.

Choose the "Restore by Object" option if you would like to restore specific objects from all tapes associated with the selected device.

Choose the "Restore by Date Range" option if you would like to restore all objects with an ingest time (GMT) within the specified date and time range from all tapes associated with the selected device.

Restore All

Restore by Object

Restore by Date Range

between: 5/24/2011 16:11:07

and: 5/24/2011 16:11:07

< Back Next > Cancel

The target HCP namespace pool is selected by right-clicking on an HCP namespace pool in the StorFirst Apollo window and selecting Restore Objects in Namespace Pool.

Restore Type

Restoration is logically defined as the equivalent of a copy process from one lost HCP to another HCP. The target HCP was selected with the right-click that began the creation of this restore job. The logical source HCP, now possibly existing only in Apollo and not physically available, will be selected after you select a restore type in the menu. Source and target HCP namespace pools can be the same.

The Restore All option makes this job restore all objects from source to target.

The Restore by Object ID option lets you specify specific object IDs to the job through a file or by manually entering them into a field.

The Restore by Date Range option is similar to the Restore All type, except that it restores only the objects with a creation date in the specified range. The creation date an unchangeable attribute of all objects.

Restore Options

Restore Deleted Objects

In the Restore Attributes window, check the Restore Deleted Objects option if you want to restore objects that were previously deleted from the HCP namespace pool.

To be sure all deleted objects get restored to the HCP namespace pool, make sure that replication is fully up to date on this pool.

Leave this box unchecked if you do not want to restore previously deleted objects to the HCP namespace pool.

To be sure no deleted objects get restored to the HCP namespace pool, make sure that replication is fully up to date on this pool.

Open Verification Check

Check the Open Verification Check option if you don't want to restore objects that already exist in HCP. Apollo will issue an Open Object command to HCP to determine whether the object exists.

Restore by Object

When restoring by object, you must supply a fully qualified path for each object. You have a choice of entering the paths manually or importing them from a file. When entering the paths manually, there is a limit of 100 objects per restore job. When importing the paths from a file, there is no limit on the number of objects.

Restore

Objects are in all cases restored completely with their original HCP cryptographic hash values and metadata intact by using a restore function in the HCP SDK for exactly this purpose. Apollo delivers the complete object to HCP. HCP computes the hash value and checks the computed value against the original value. The restored metadata for each object includes the object creation time and retention setting.

Offline Tapes

All required tapes should be online for the restore job. Should any of the tapes required by the job be offline when the job starts, the job will immediately list the missing tapes in its log, one line per tape as follows:

```
Objects from OFFLINE tape 200810101533_000311 skipped - bring tape online and restart the job to complete restore.
```

Bring the tapes online and restore using the restore by individual source tape feature.

Duplicate Tapes

Apollo restores data only from primary tapes. In order to restore from duplicate tapes, you must uninventorize the primary tapes. This will promote the duplicate tapes to primary tapes, and you can now restore from what were previously duplicate tapes.

CHAPTER 9

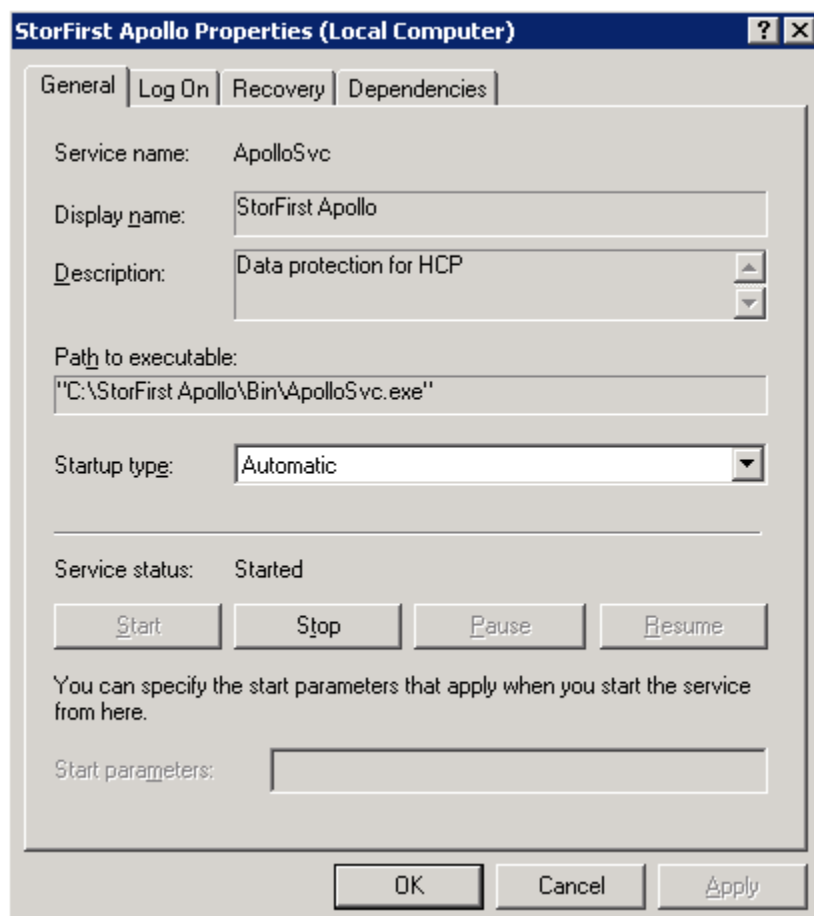
SUPPORT SERVICES

StorFirst in Windows

The StorFirst software is delivered and copied to a distribution folder of your choice, usually via a zip file downloaded from the Seven10 FTP site. The installation process creates the StorFirst Apollo destination folder with the software, database, and log file folders.

The running StorFirst Apollo system is in two parts on your server. The StorFirst Apollo service is in the background and can be manipulated by the Microsoft administrative tools in the Control Panel under Services. The service is configured to start automatically after the Windows boots. The StorFirst Apollo Administrator is run from the icon on your desktop. You can leave it running permanently or start it up as needed.

To view the StorFirst Apollo service Properties window, right-click the service in the Services window, which is accessible from the Windows Administrative Tools:



Support Call

Seven10 Storage Software Support can be contacted via:

- Phone – (978) 725-5525
- Email – support@seven10storage.com

Before initiating a service call or email, please be prepared to send the Apollo log files. Go to Tools → Zip Logs in the StorFirst Apollo Administrator window and click Zip Logs. This will zip the JobLogs folder, Service Logs folder, Seven10 Storage portion of the registry, application event Log, security event Log, and system event log.

Server Replacement

If you need to upgrade or replace your existing server, please contact Seven10 Support for assistance.

Software Uninstall

StorFirst Apollo can be removed by first stopping the StorFirst Apollo service and then using Windows Add or Remove Programs. This will remove all software binaries, shortcuts, and registry entries. The database files and logs will not be removed from the server. A complete removal of StorFirst Apollo is accomplished by deleting the StorFirst Apollo destination folder.

Changing SCSI Address

Server reconfiguration or server replacement can change the SCSI address used to communicate with a library and its drives. If StorFirst can no longer communicate as the result of such a change, it will mark the device as disconnected but retain all information about tapes and slot locations that may have been accumulated over time.

Tape Libraries

Simply using the Library Wizard again will create a new, but blank, representation of the library without that valuable knowledge. This transition can be implemented easily in two very different ways, both without loss of operational information. The first method makes all the tapes offline and then uses the Library Wizard for automated recognition of the library and its new SCSI addresses. This way is highly recommended. The second method is to manually change the SCSI addresses. This method is for expert support personnel and assumes that the new SCSI addresses are known.

The automated offline way follows these steps:

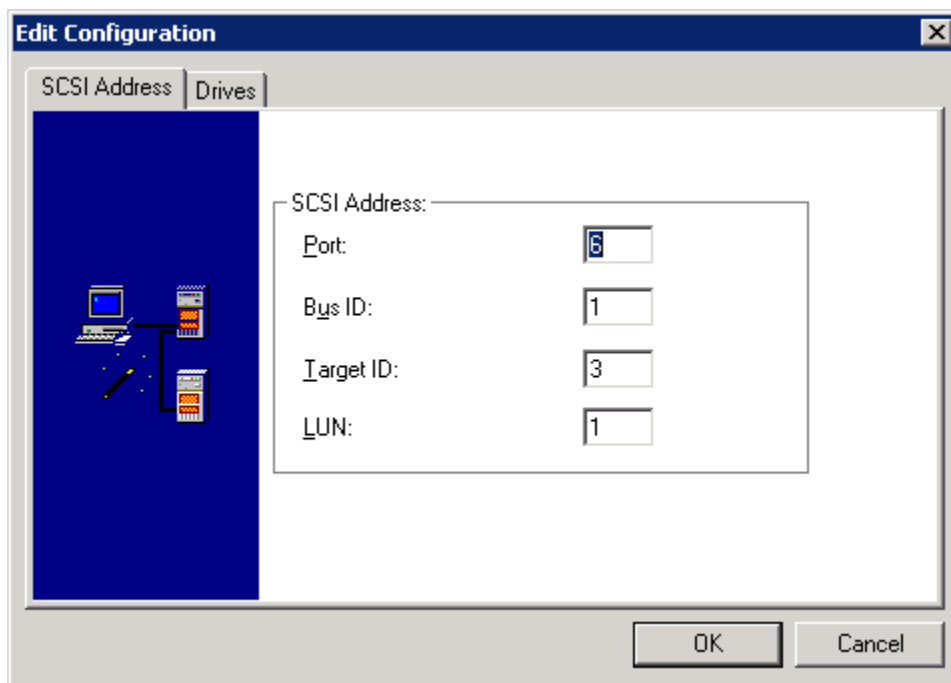
1. Disable and delete the library object in StorFirst Apollo; all its media are now offline.
2. Activate the Library Wizard and configure the library.
3. Select all the full slots, right-click, and select Inventory.

4. Select the Offline to Online option and run the job.

Step 1 can be done while the library is still operational or after the system changes while the library appears disconnected. Steps 2 through 4 must be done after the system changes.

The manual SCSI editor way follows these steps:

1. Implement the system changes and start StorFirst Apollo.
2. StorFirst will recognize the issue and mark the library with a red "X."
3. Right-click and select Edit Configuration.
4. In the Edit Configuration window, update the SCSI address of the library.
5. Click OK. StorFirst will verify the new addresses and confirm the identity of the library. The SCSI addresses of the drives are then updated automatically.



Replacing a Library

Libraries can be replaced for capacity, maintenance, and many other reasons. Media can be logically moved from one library to another if a new library is installed and the media were transferred with or without StorFirst participation. In any case, follow these steps:

1. Extract tapes in one of these ways:
 - If the library is operational, eject all media with StorFirst commands.
 - If the library is not operational, extract the media manually.
 - Disable and delete the library object in StorFirst.

Deleting a library object preserves all its media objects in the Offline folder. The net logical result is that all the useful media are in the Offline folder. The system has severed the

connection between the tape object and the library and slot. Physically, the tapes are stacked on a table.

2. Reconfigure:
 - a. Shut the system down. Disconnect the old library and connect the new.
 - b. Insert all the tapes in the library now, or do it later using StorFirst commands.
 - c. Turn the library on, wait for it to become ready, and then start the server. Note that hot-pluggable functionality is supported, but when in doubt, restart the server.
3. In StorFirst, use the Library Wizard to configure the new library.
4. Use the Insert Media command to load all media if not loaded already. Access this command by accessing Library, Slots, and then right-click on the desired slots. Select Insert Media. At this point, all the old media are in the new library slots marked full and known to StorFirst as Offline. It is time to connect the two.
5. Right-click the slots and select Inventory Slots. Select the Reinstate Media from Offline to Online option. This will scan all media, match them to the objects in the offline folder, and reconnect the tape and the slots. Now all media is back online.

When the procedure is done, the library has been replaced, and the system has been restored and is back in operation.

Notification of Alerts by Email

If your StorFirst Apollo server is behind a firewall, you can use a local SMTP server for notification of alerts by email. This can be done by entering your company email address and your company SMTP server in the Alert Configuration window. The username and password are optional.

System Restart

Sometimes power spikes or outages will adversely affect computers, I/O buses, or libraries. Should the system stop or hang under such circumstances, follow these steps to restore system operation:

1. Shut down the server and all managed storage units.
2. Power up the storage units and wait for ready.
3. Reboot the server.

Tape Block Size

The available block size may be limited by the tape drive, tape type, and/or the available BUS transfer sizes of the attached SCSI or FC card. If you select a block size that exceeds any of these values, Apollo will scale the value used back to the smallest maximum allowed by any of these items in the data path. If Apollo does this, a note in this format will be placed in the log:

5:TapeMgr SetBlockSize - resetting blksize to [Block Size in KB] (Host: [Block Size in KB], Tgt:[Block Size in KB], Min:[Block Size in KB], Max:[Block Size in KB])

Host and Tgt indicate the SCSI or FC card limits, and Min and Max indicate the device, drive, and tape limits.

APPENDIX A INSTALLING LIBATTACH


ACSL is a software solution from Sun Data Management Group providing library aggregation and sharing of libraries among library-management-enabled applications like backup and StorFirst. Windows-based applications use LibAttach, installed on their server, for access to ACSL over a LAN. To use ACSL with StorFirst, you must install LibAttach on the StorFirst server before using the Library Wizard.

LibAttach is not delivered with StorFirst and must be acquired independently from Sun Data Management Group. Installation will typically be available from Sun Data Management Group, and the following are only observations from the installation of LibAttach 1.4 at Seven10 labs and are offered here as advice. The official Sun Data Management Group documentation and their support organization is the only authority on the subject.

Step One: Modify the Hosts File

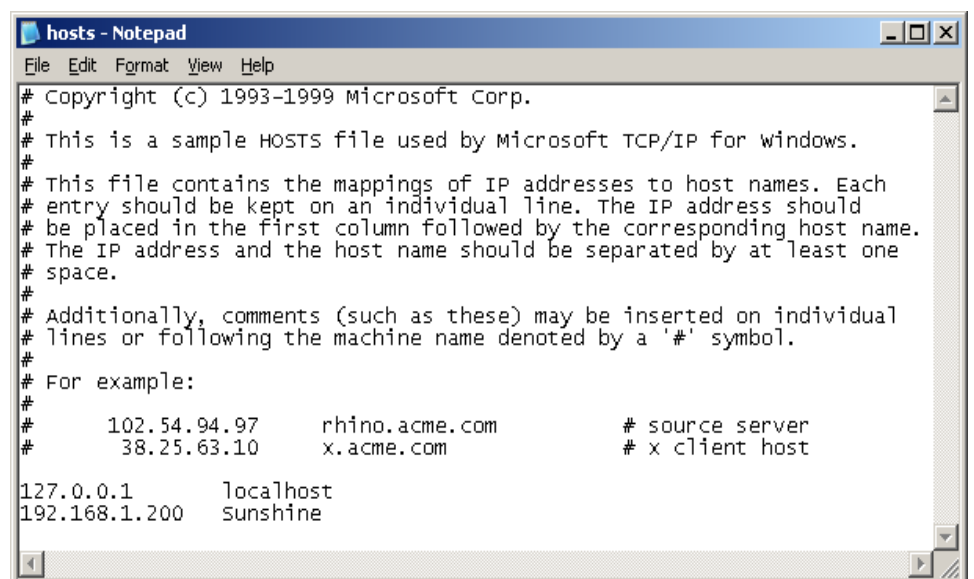
Add the ACSL server hostname and static IP address to the hosts file. The file can be found in the location shown on the right.

Open the hosts file with a text editor like Notepad, and add the ACSL server at the end of the file.



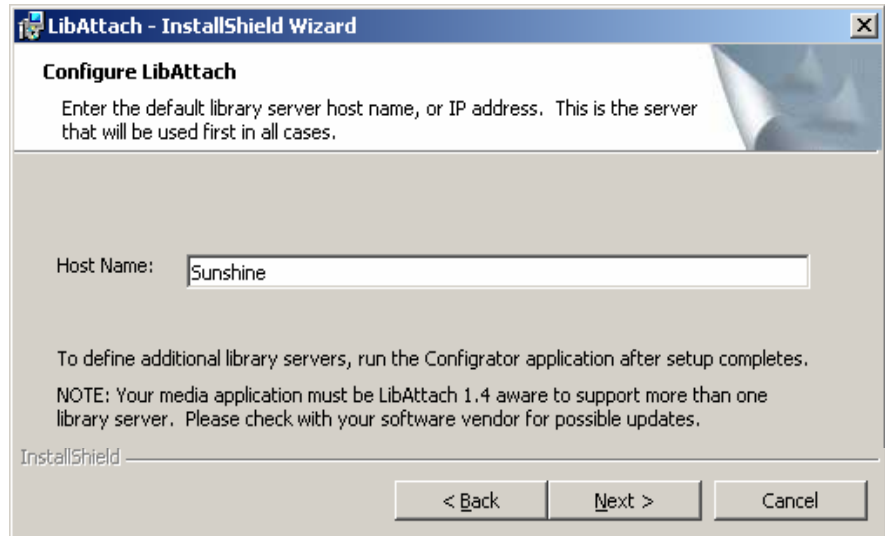
Step Two: Install LibAttach

Install LibAttach on the StorFirst server using the CD supplied by Sun Data Management Group.



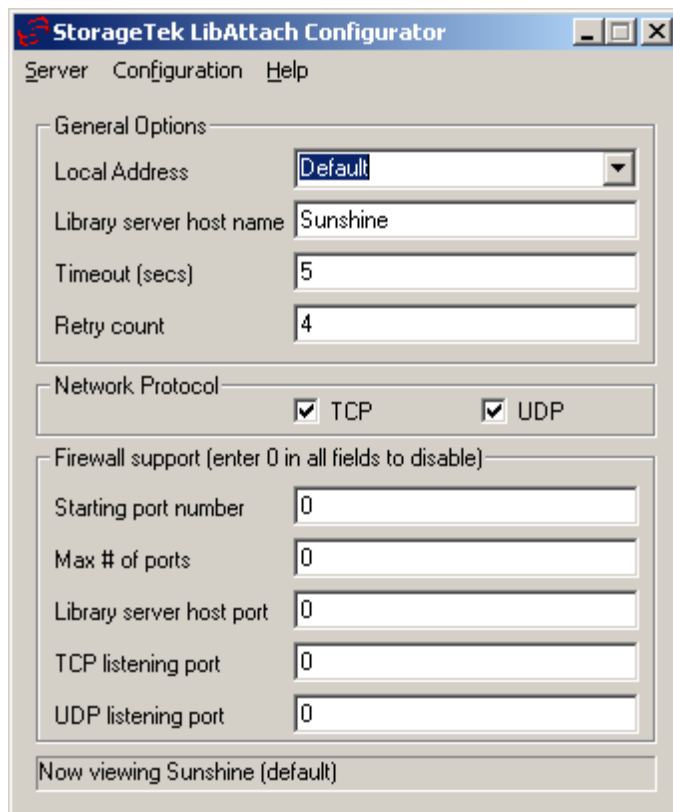
Step Three: ACSLS Server Name

LibAttach may automatically start the configuration after installation; otherwise, start the configuration from the Start menu. Enter the name of the particular ACSLS server to be used; only one such server is supported in StorFirst, and the LibAttach default is automatically selected. The ACSLS server IP address is retrieved from the hosts file automatically and transparently using the name as a key.




Step Four: Properties

The defaults work well if there is no firewall in the way.



Step Five: Test the ACSLS Connection

Make sure that LibAttach can successfully communicate with the ACSLS server. Select Test in the Configuration menu at the top of the StorageTek LibAttach Configurator window.



Troubleshooting

If the test fails, try to ping the ACSLS server from the LibAttach server. Check that the IP address was entered correctly into the hosts file. Make sure that a firewall is not the problem. If the network works and LibAttach connectivity still fails, contact Sun Data Management Group.

APPENDIX B

ACSLS

Configuring ACSLS

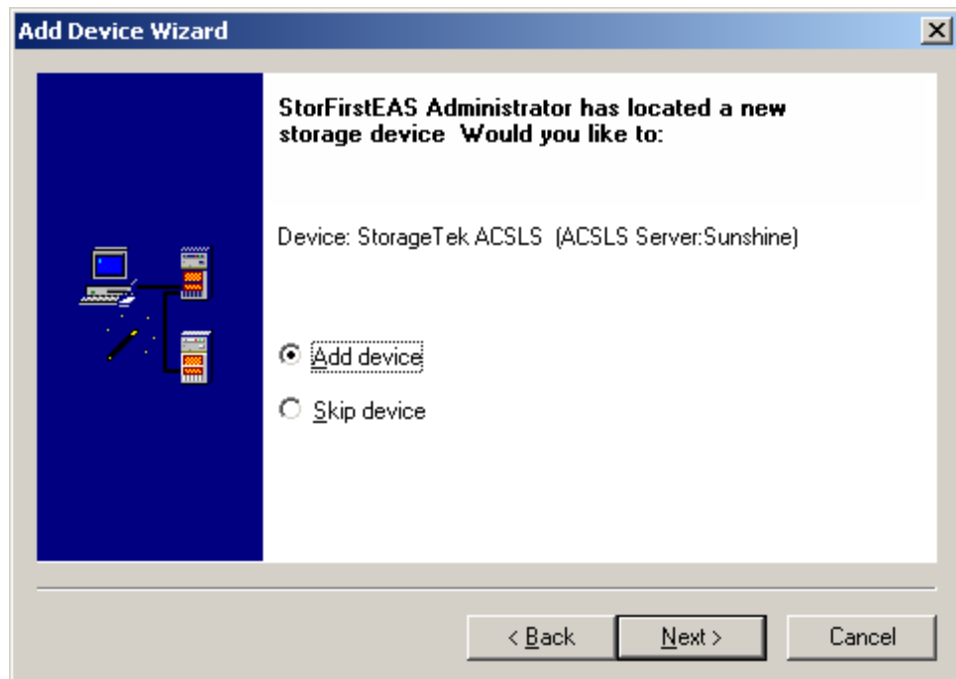
To use ACSLS, you must install LibAttach on the StorFirst server before using the Library Wizard. Once present, the ACSLS option can be added or skipped. ACSLS managed libraries should not be visible to StorFirst on a system SCSI bus because dual management is destructive. However, skipping the device in the Add Device Wizard will also work.

The Wizard will in turn present all un-configured libraries, managed by ACSLS or directly found on the system SCSI buses.

A drive can be excluded from use by StorFirst by selecting Unknown in the dropdown menu at this time or by disabling it later. This enables sharing of a library between multiple applications.

All drives used by StorFirst must be interchangeable and compatible with all media in the assigned slots. Two separate libraries can use different types of drives and types of tape.

With ACSLS, the drives that are offered will be the drives provided by ACSLS and accessible for I/O on the StorFirst systems SCSI buses, matched by serial number comparison. The drives are otherwise no different operationally from the non-ACSLs case.

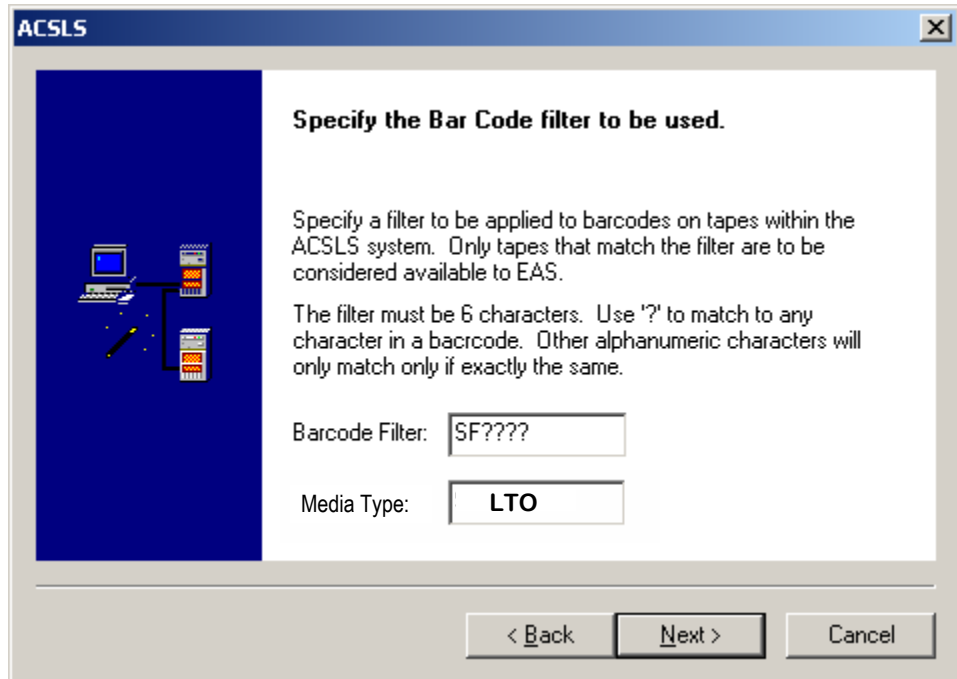


Media Selection with ACSLS

ACSLS supplies library aggregation and sharing but not partitioning. It may therefore present so large a virtual library that the usual StorFirst UI model is less useful. ACSLS operates on tapes and does not divulge slot locations. To preserve the common StorFirst user interface and its normal operation on slots, StorFirst automatically creates virtual slots for the tapes. To limit the number of tapes shown to those that are likely to be used by StorFirst, a barcode filter for selection is offered.

For example, if Apollo-designated tapes by convention have six-character bar codes starting with SF and followed by four digits, set the filter to "SF????". The set of media to be used can further be restricted by selection of media type. As StorFirst queries ACSLS for media, it will be informed about all

media in the aggregate virtual library managed by ACSLS, yet only the media passing the specified filter will be shown in the StorFirst Apollo Administrator window.



Specify the Bar Code filter to be used.

Specify a filter to be applied to barcodes on tapes within the ACSLS system. Only tapes that match the filter are to be considered available to EAS.

The filter must be 6 characters. Use '?' to match to any character in a barcode. Other alphanumeric characters will only match only if exactly the same.

Barcode Filter:

Media Type:

< Back Next > Cancel

Library Partitioning

Apollo takes complete control of the drives and slots assigned to it. In a single library, all assigned drives must be interchangeable and compatible with all media in the assigned slots. Two separate libraries can use different types of tape. StorFirst is capable of sharing a library with other Apollo instances or even other applications.

The shared resources, the robot, and the door may cause access collisions if a library is shared between applications. Should Apollo encounter a busy robot, it will simply wait and try again. This works well for sharing between StorFirst instances. Other applications may respond differently and with less tolerance.