Dell HPC NSS Tiered Storage Solution

A Dell Technical White Paper

Quy Ta, Onur Celebioglu

Dell HPC Engineering

November 2011| Rev 1.0

Contents

Figures

Executive summary

This solution guide describes the Dell HPC Tiered Storage Solution (TSS). The Dell Tiered Storage Solution integrates CommVault Simpana 9 Data & Information Management Software with the Dell NFS Storage Solution (NSS) and PowerVault ML6000 Tape Libraries to provide a multi-tiered Hierarchical Storage Management system. The goal is to provide a data storage technique that offers Hierarchical Storage Management (HSM) and archive solutions that automatically move data between storage tiers through the use of defined storage policies. This document describes the architecture, use case scenarios, performance, and best practices for such solutions.

Introduction

Clusters have become one of the most popular architectures for High Performance Computing (HPC) today.[1] Along with the increasing popularity in HPC architectures, the HPC storage market has witnessed an explosion in data growth, presenting an increasing concern for data management. Enterprise IT departments are searching for solutions to address the data growth and regulatory/compliance requirements, while limiting the effect on IT budgets and resources. Data archiving and management are key emergent areas to address the rising data storage needs effectively.

Customers are seeking effective solutions to:

- Organize data storage and retrieval into separate tiers for cost management and storage space efficiency.
- Simplify data storage processes.
- Focus on shrinking expensive primary storage requirements and migrating data to cheaper secondary storage tiers.

HPC Customers typically have three kinds of needs for an HPC storage solution.

Scratch Space (tier1): high throughput and scalable cluster working space. This solution is provided by the Dell | Terascala HPC Storage Solution (DT-HSS), which offers a high throughput scale-out storage appliance based on the Lustre file system and Dell PowerVault storage arrays.[2]

Primary Storage (tier2): reliable, cost effective and good performance storage for user data. This solution is provided by the Dell NFS Storage Solution (NSS) that uses the NFS file system on top of the Red Hat Scalable File System (XFS) with Dell PowerVault storage. It provides an easy to manage, reliable, and cost-effective solution for unstructured data.[3]

Long-Term Storage (tier3): low cost, high capacity storage for long-term retention. Long-term storage is disk based, tape based, or a combination of both.

The Dell Tiered Storage Solution (TSS) enables Dell customers to manage the archival and movement of their data between these tiers.

More specifically, the TSS leverages sophisticated, policy-based data management automation processes to move data from tier to tier, leaving behind a small stub file containing the file metadata and moving the data blocks to a lower tier (commonly referred to as HSM). Additionally, you can use it for archiving purposes where not file stub is left behind and a new method of accessing the data is required. Typically, frequently used data can be moved to a faster tier, while infrequently used data can be moved to slower media at lower tiers. Correspondingly, as files are used more often, they can reside on the first or fastest tier and then moved down to the slowest or lowest tier as they fall out of use, but must be retained to meet regulatory or compliance requirements.

The following sections describe the TSS architecture in detail with focus on the CommVault Simpana 9 Data Management Software [4] in a Hierarchical Storage Management (HSM) implementation. Subsequent sections discuss a selection of typical use cases scenarios of an HSM solution focusing on performance and best practices of implementation for the solution. An extensive appendix covers detailed steps on configuring a TSS.

The Dell Tiered Storage Solution is delivered as an all-inclusive storage data management solution and is available with deployment services and full hardware and software support from Dell and CommVault. The solution design uses the following principles:

- Standards-based HPC components
- Ease of deployment
- Ease of use

Dell tiered storage solution technical overview

This section provides a quick summary of the technical details of the Dell Tiered Storage Solution (TSS) offering.

The Dell TSS offering consists of a Dell PowerVault DL2200 Disk-Based Backup Appliance and a PowerVault ML6000 Modular Tape Library; you can use more than one. The DL2200 appliance serves as the HSM data manager and utilizes CommVault Simpana 9 Data & Information Management Software, but it is not in the data path for the actual HSM or archive process. This is the central point of the HSM network, where storage resources and policies are created and administered. The PowerVault ML6000 Modular Tape Library serves as the designated ***long-term storage*** tier in this solution. This Long-Term Storage or LTS tier allows for low cost/high capacity storage and for long term retention of data. You can consult the accompanying documentation [(7)] or the Dell deployment service team for proper rack and stack procedures as well as loading and initializing of the drives and tape cartridges.

Dell TSS supports the Dell NFS Storage Solution[(3)] (NSS) as the designated ***primary*** storage tier in both standalone (NSS) and high availability (NSS-HA) configurations. The NSS and NSS-HA solutions use the NFS file system on top of the Red Hat Scalable File System Add-on (based on XFS) with Dell MD PowerVault as back end storage, to provide an easy to manage, reliable, and cost-effective solution for unstructured data.

Dell also offers the Dell | Terascala HPC Storage Solution (DT-HSS) that you can be designate as the ***scratch*** storage tier. Though not actively supported in the HSM network, manual movement of data between the DT-HSS scratch storage space and NSS primary storage space can be achieved through use of manual move commands or by scripts and even in job scheduling scripts. The data can then be HSM or archive between the NSS primary tier and ML6000 Tape library long-term storage tier.

Figure 1 and 2 illustrate the TSS reference architecture with Dell NSS and NSS-HA configurations.

Figure 1. TSS reference architecture with NSS configuration

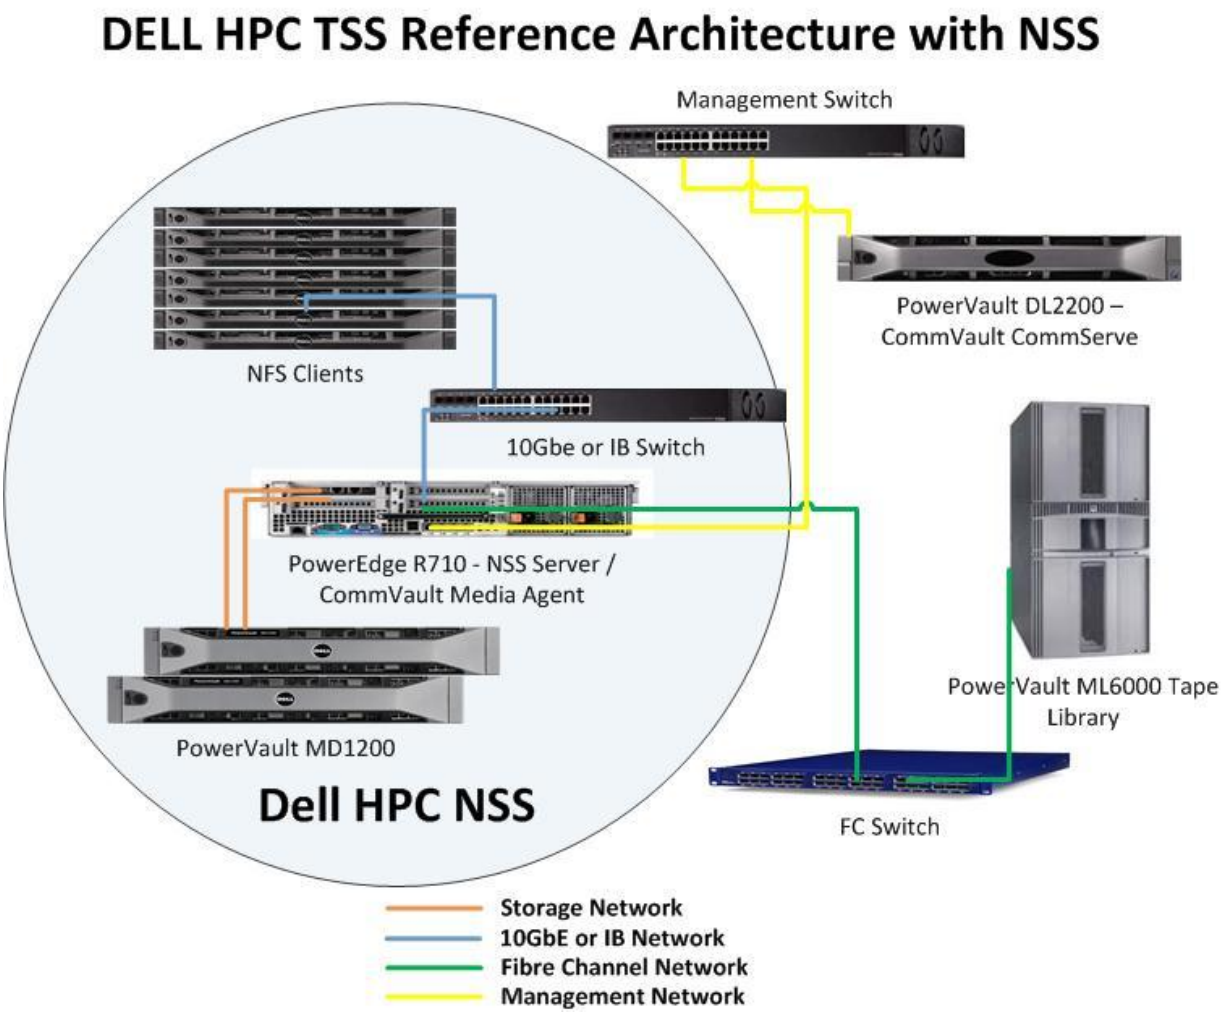

Figure 2. TSS reference architecture with NSS-HA configuration

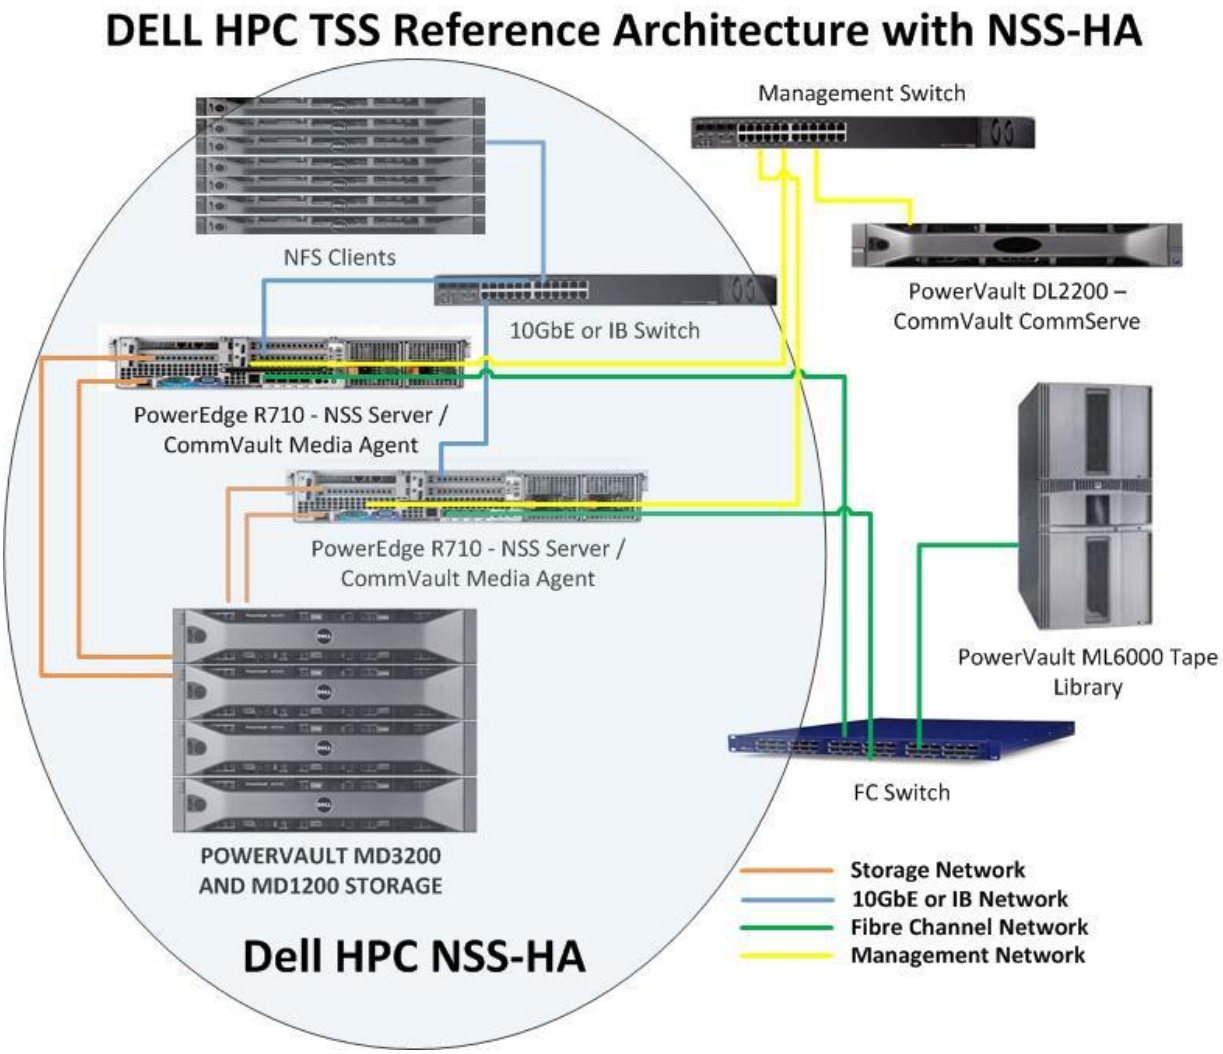

TSS data paths

Figure 3 and 4 illustrates the typical data paths that the Dell TSS solution supports in a 3 tier configuration (fast scratch [Dell DT-HSS], primary [Dell NSS], long term). Figure 3 illustrates the HSM Data Flow, where manual or script assisted data management occurs between the fast scratch tier and the primary storage tier, while HSM policy driven data management occurs transparently between primary storage tier and long-term storage tier. Figure 4 illustrates the Archive Data Flow, where data migration or Archival process is not transparent and requires admin provided interface for retrieval.

NOTE: In both illustrations, the Dell DT-HSS support is limited to manual or automated (via scripting) migration between “Fast Scratch” and “Primary Storage” tiers and is not part of the TSS HSM or Archive architecture.

Figure 3. TSS HSM data flow path

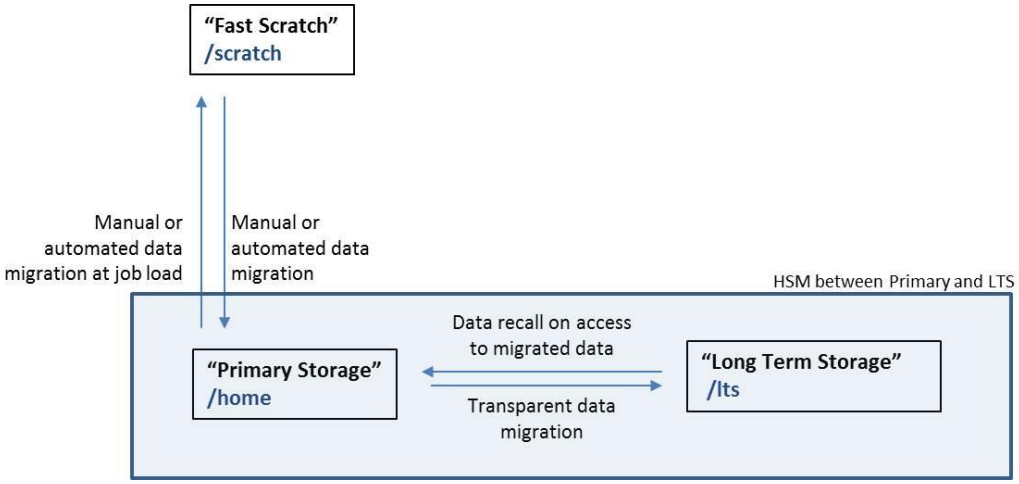

Figure 4. TSS archive data flow path

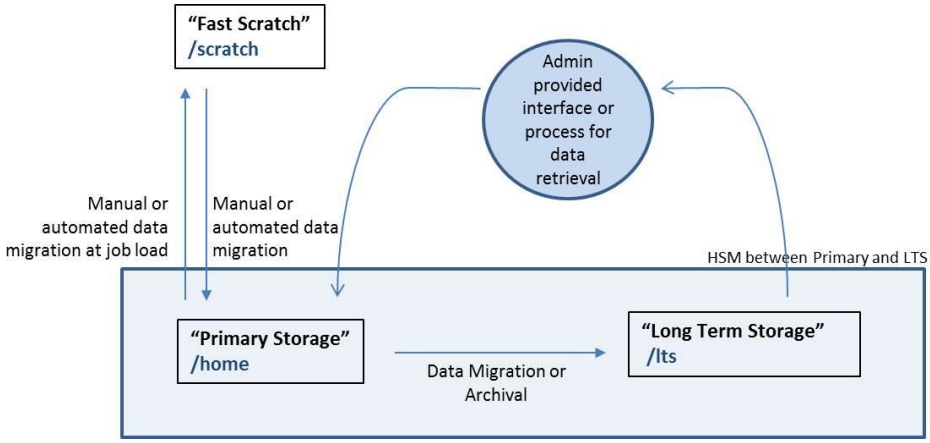

CommVault Simpana 9 Data & Information Management software

The backbone of a Hierarchical Storage Management solution is the software. The HSM system monitors the way data is used, then through defined storage policies, moves less frequently used or lower priority data safely to slower tiers, reclaiming storage space on faster tiers. Very detailed logic is required to catalog the data and watch for frequently used files that should reside on the upper tier. The software is also responsible for managing the requests and ensuring those requests happen in a timely manner.

The Dell Tiered Storage Solution deploys a PowerVault DL2200 Appliance that is powered by CommVault Simpana 9 software to serve as the central point for the HSM network. CommVault Simpana 9 is a common platform; modular data management software solution that provides a powerful set of storage management tools that help enterprises move and manage critical data.[4] The DL2200 Appliance comes factory pre-installed with Windows 2008 R2 and CommVault Simpana 9 Software. Upon initial boot, the administrator is presented with a setup wizard that provides guidance through configuring the appliance and network settings, as well as the CommVault CommCell setup.

Key software features of CommVault Simpana 9 in HSM

There are two key features of CommVault Simpana 9 in HSM: Policy Based Migration and Transparent Access.

Policy based migration

Rules definable by thresholds, owners, age, size, filters, folders or other search criteria allow added flexibility. Data can move from its existing tier space to a different tier space, or copy to another tier, thereby enabling data protection within the same solution.

Transparent access

End users maintain their access as before the archiving or HSM occurred. Also, the user can search and recall the information from the archive or stub file in the originating tier in the HSM scenario, minimizing the help needed from administrators.

CommVault Simpana 9 Components

The following section goes into more detail on the various components that make up the Simpana 9 package.

CommCell

The CommCell consists of Common Technology Engine (CTE) components including one CommServe, one or more Media Agents, and one or more Client Agents. The CommCell is managed by a CommCell Console, which includes tabs for Security, Storage Resources (including Libraries), Policies (including Replication, Schedule and Storage Policies) and Reporting. The TSS will have the PowerVault DL2200 serve as the central management for the CommCell via the CommCell Console.

CommServe

The CommServe Server ties the CommCell components together; it is the coordinator and administrator of the CommCell component. The CommServe Server communicates with all agents in the CommCell to initiate data migration, protection, management and recovery operations. Similarly, it communicates with MediaAgents when the media subsystem requires management. The CommServe Server maintains a database, the CommServe Database Engine, which contains all of the information related to the CommCell configuration. The TSS includes a PowerVault DL2200 Appliance to be used as the CommServe Server.

MediaAgent(s)

The MediaAgent transfers data between the client computer(s) and the storage media. Each MediaAgent communicates locally or remotely to one or more storage devices. The TSS includes the CommVault Linux Media Agent on the NSS server node(s).

Client Archive Agents

Migration Archiver Agents are the software modules responsible for periodically moving unused or infrequently used data on their host computers to secondary storage, thereby reducing the size of data on the primary storage. The TSS includes the CommVault File Archiver Agent on the NSS server node(s).

User scenarios for tiered storage solution

This section describes some key user scenarios that are verified and supported with the Dell TSS solution. This section is not a comprehensive list, but rather a list of key user scenarios that Dell TSS supports.

Migration policies

The storage administrator can set policies to migrate data from a user home directory on the ***primary tier*** to ***long-term storage (LTS) tier***. In the case of an HSM operation, Dell verified that a stub file for persistent recall of migrated data was created after the migration of the data blocks. In the case of an Archive operation, Dell verified that migrated data is no longer available and a stub file for data was not generated on the originating tier. Some policies that were verified include, but not limited to:

- Policy to migrate files that are 7 days or older from last modify or last access date.
- Policy to migrate files that are larger than 1GB in size.
- Policy to migrate files that are larger than 500MB to tape based LTS while files that are less than 500MB to disk based LTS.
- Policy to trigger migration jobs when primary file system utilization reached 80% peak.

The storage administrator can set policies to migrate or archive data according to the user, group or parent folders of the data. This includes creating different policies for different groups.

The storage administrator can avoid overloading on the data management system (CommVault) if multiple users try to recall large amounts of data. You can accomplish this by setting recall throttling parameters when creating sub-client(s) and associated policies.

- Maximum Stub Recovery (limit)
- Time between Recall to Count as Successive in Seconds (interval)
- Time to Wait after Maximum Successive Recalls Limit is Reached in Seconds (cool-down)

Data protection levels

The storage administrator, using policies, can make multiple copies of user data. The policy can determine which directories or files are copied to multiple locations, providing a way for users to protect their data accordingly. You can achieve this task through use of the ***Auxiliary Copy*** feature of Simpana 9 and can specify recovery from specific ***LTS tier*** through the **Copy Precedence** settings in the **Storage Policy Properties**. [5] For example, a policy can:

- *Migrate data in* ***/home/userX/redundant-copy*** *on* ***Primary tier*** (using either HSM or Archive) and move copies of the data on two separate ***LTS tier***s.
- *While* data outside of ***/home/userX/redundant-copy*** on ***Primary tier*** can be migrated to only a single instance on ***LTS tier***.

Data backup and recovery

The storage administrator can create policy to generate a full backup of the ***primary*** storage tier. You can use this backup to recover user data in the event of a disaster. To accomplish this task, perform a **full** system backup of the ***primary*** storage tier data and then a **full** system restore of the data.

In the event of a CommServe disaster, you can perform a **Disaster Recovery Backup** from the CommCell and then use the **CommServe Disaster Recovery Tool** to restore the database metadata and windows registry. [5]

Data partitioning

The storage administrator can create a policy to store user data for a given user or group of users on a designated set of tape cartridges in the ***LTS*** library. Conversely, the administrator can physically partition the ***LTS*** library tier to store data for different group of users. To accomplish this task, configure the ML6000 into several **logical libraries** and present them as multiple **storage resource libraries**.

Data sharing

The storage administrator can provide a way for users to share their data with other system users or groups by using the **User Administration and Security** feature of Simpana 9.

- System users can share data with other users who have access to the primary storage tier.
- Shared data is still subject to migration policies as defined by the storage administrator.

Performance and capacity scalability

Expansion of the LTS tier should not cause any major interruption to the existing system. The expansion includes expanding for capacity, performance, or both. This can be achieved with the capacity on demand feature of the ML6000 library along with slot increment software license key allowing for non-disruptive scalability.

Data encryption

The Advanced Encryption Standard (AES) in 128, 192, 256 bit encryption is desired to be available for data that resides on LTS tier. Simpana 9 supports many data encryption algorithms including Blowfish, AES, Serpent, Twofish and 3-DES. Data is encrypted according to the method you select when you **Configure the Client for Data Encryption** (client-level encryption) or **Configure a Storage Policy Copy for Data Encryption** (auxiliary copy-level encryption). You can select from several algorithms and key lengths.

Data compression

The storage administrator can compress data on **LTS** tier. You can achieve this by means of software compression on the **subclient** level, which includes options to compress the data in the Client/MediaAgent or hardware compression for libraries with tape media. This is supported by ML6000 library / LTO5 drives in a **2:1** ratio.

Reporting

The storage administrator needs a reporting capability to study overall status of the storage system and health. The administrator needs the ability to generate reports on ***primary*** tier utilization. You can achieve this through generating the **Library and Drive Report** and **Storage Information Report**.

Primary to LTS tier data movement and policies

Migration occurs transparently to the ***LTS*** tier based on policies defined by the storage administrator. A stub is left behind for users to see and use through **persistent recovery**.
A Recall is initiated if the user opens the file for view or edit. The file is recalled from the ***LTS*** tier back to the ***primary*** tier.

If the recall for a file in ***LTS*** fails, the Archive agent returns I/O error on failure to recall. If the file is only partially restored, the stub information is preserved. The next read on the file triggers the recall again.

All file attributes, including Linux permissions, are maintained on the files that have been migrated to the ***LTS*** tier.

The cluster user stores his/her data on ***primary*** tier; the user can then force the data to move from ***primary*** to ***LTS*** tier. The user is setup with permissions to do so in the CommCell Console. The user can then select an archive **subclient** that they have access to and initiate the archive job.

The user can see the file stub in the originating directory structure and can retrieve the file when required through the use of **persistent recovery**.

If the user accidently deletes a file that has been migrated to ***LTS*** tier and if a copy of the file remains on ***LTS***, the user can utilize the end user **web search console** on the **Web Server** installed on IIS server to view and recover their specific archived data back to the original or a staging location.

Primary to scratch to LTS tier

The user can manually or via script, copy/move data from Primary tier /home/user1/jobX, to Scratch tier /scratch/user1/jobX, as part of the job; and manually or via script copy/move job results back from /scratch/user1/jobX to /home/user1/jobX. The user can integrate data copy (across scratch and primary) as part of the job through the use of either the manual operation or the custom scripted operation.

Performance and scaling

This section describes the performance results of job process times for HSM migration and persistent data recall events. The test methodology utilizes data sets of various sizes as well as the number of file(s) in job, to illustrate the effect of overall data set size vs. the number of files within a migration job. Dell explored data compression states to illustrate performance gains achieved from compression settings of a 2:1 ratio when compared to no/zero data compression settings (native). In each of the migration and recall jobs, a storage policy was used to trigger job start and was also configured to perform data verification to assure data integrity of migrated data. [(5)]

The first test, summarized in Table 5.1, tested a single file of various sizes to measure how long it took to migrate the data to the LTS and leave behind a file stub as well as verify the transaction.

Single large file

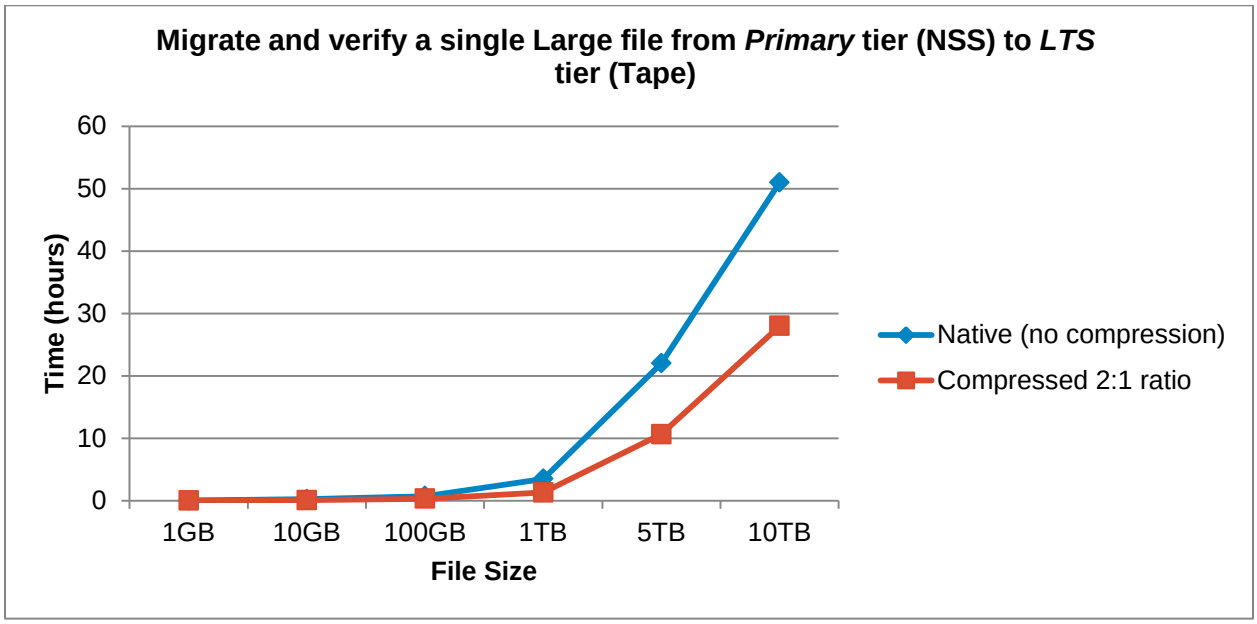

The second test uses a large number of files in an HSM migration for a larger total data size. The results are tabulated in Table 5.2

Many small files for large total size data set

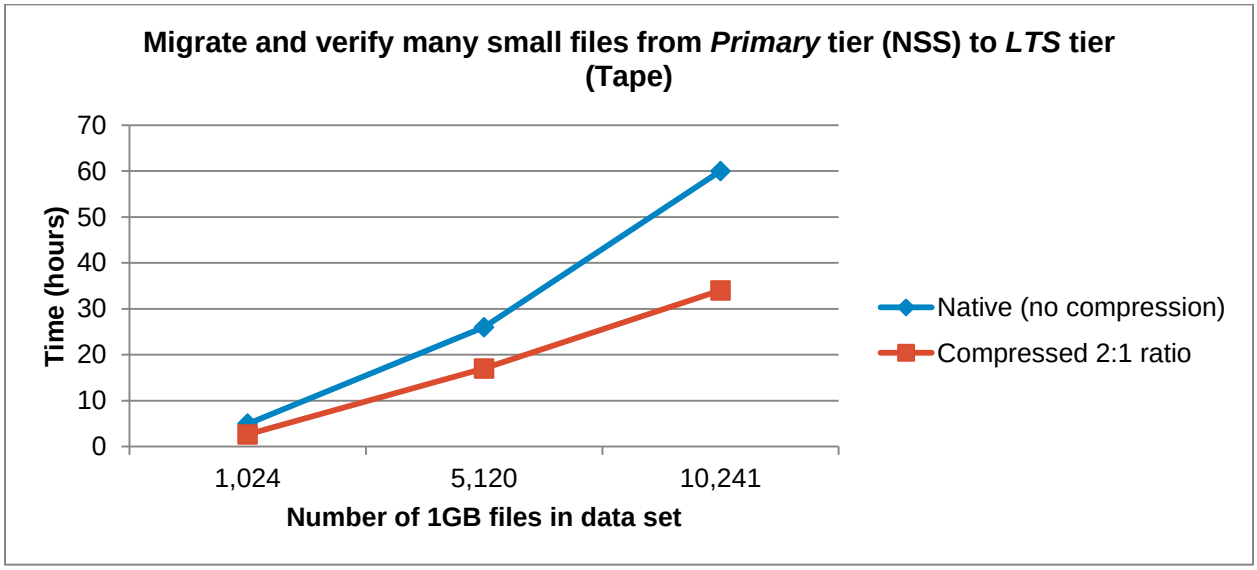

The third test measured the amount of time to migrate small files. Because the files are small, the migration is fairly quick, but there is additional time to load the tape in the drive and write the data. The results are summarized in Table 5.3

Single small file

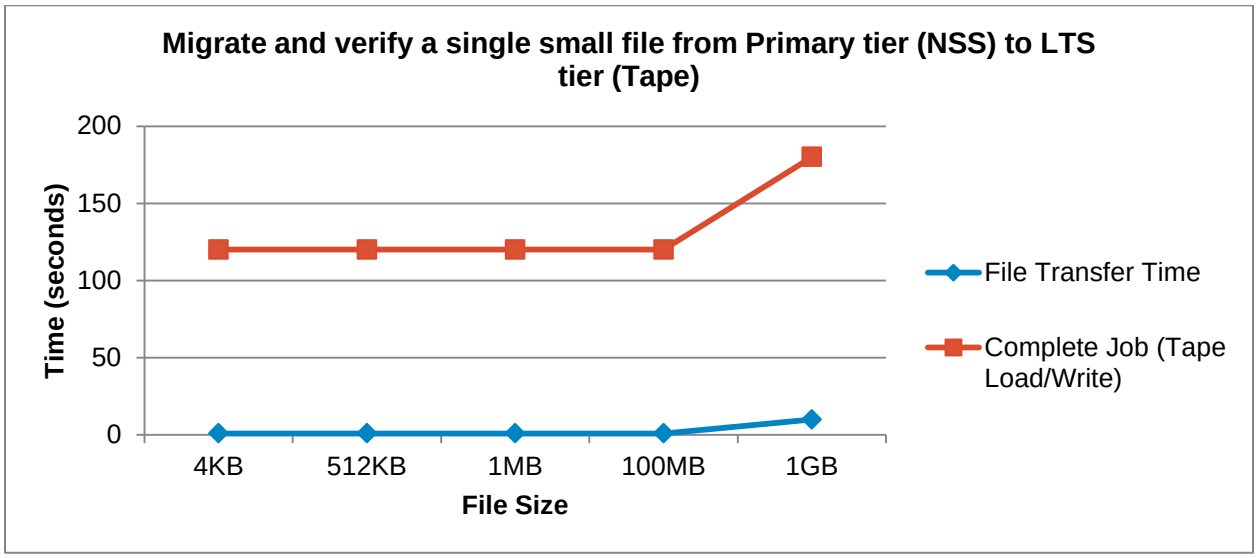

The fourth test that was done is actually the inverse of the first test. That is, it measures the amount of time to recall the data from the LTS to the primary storage. The results are summarized in Table 5.4 below.

Single file recall

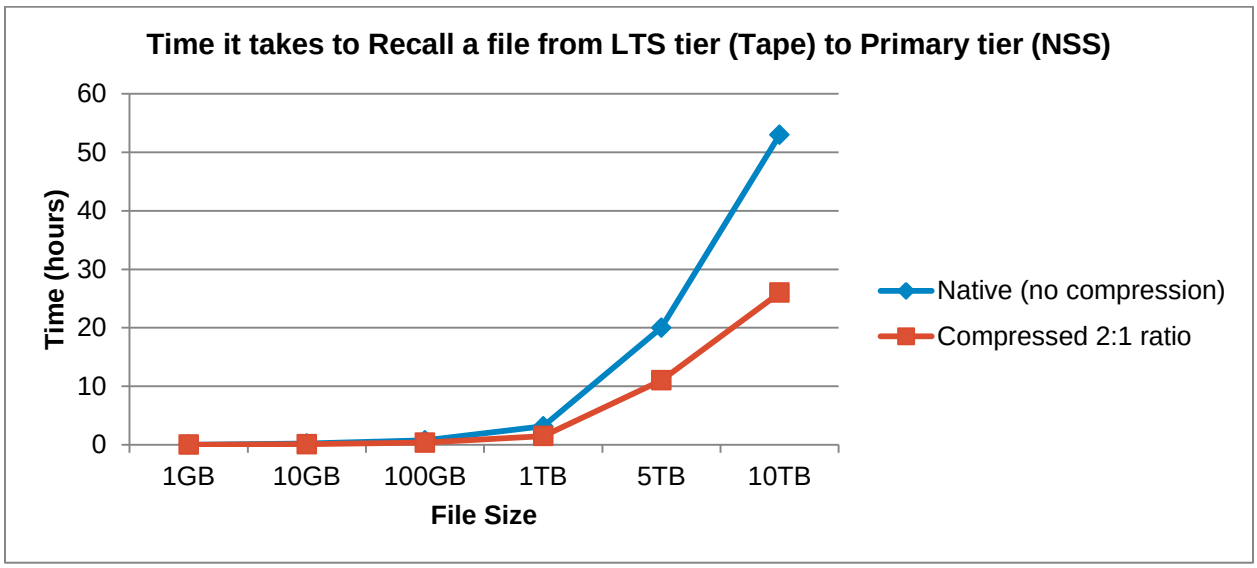

The final test examines the impact of migrating data using multiple tape drives. One to six drives were used in the testing with the resulting throughput in MB/s listed in Table 5.5 below.

ML6000 LTO5 Scalability

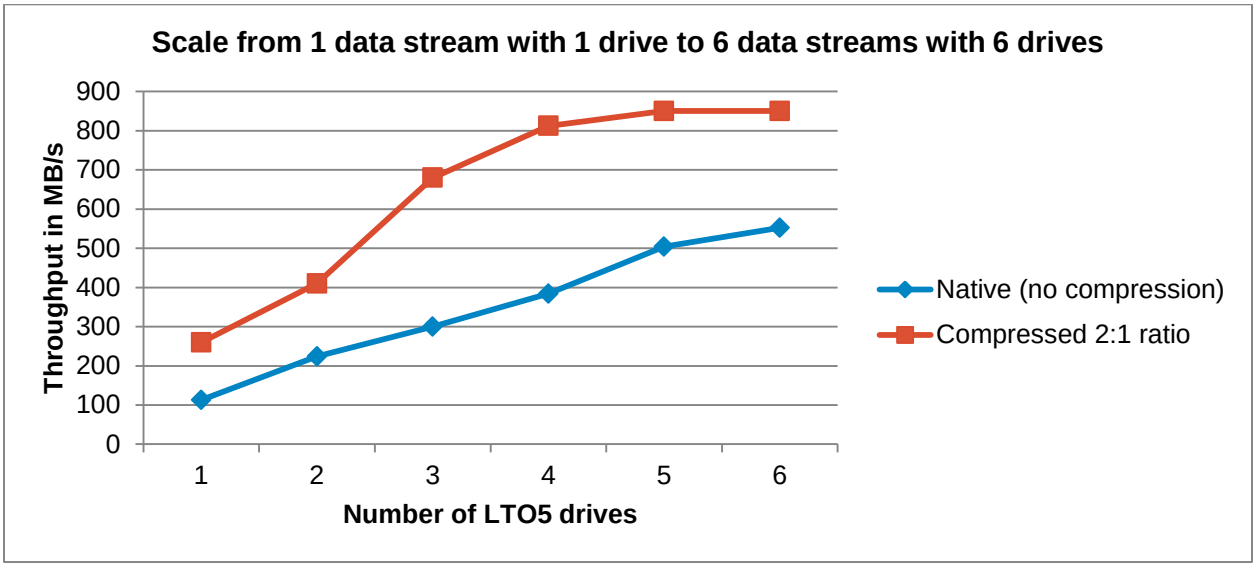

Optimizing and best practices

CommVault Simpana 9 Software has many options and parameters that can be tuned to improve performance for a particular implementation.[5] This section reviews a few of those options and parameters. This is not meant to be a complete list but rather a general guide or introduction to areas that can you can research for possible performance improvements.

Improving throughput to storage media

The throughput to a storage medium depends on the speed at which chunks are written to that medium and the number of entries per chunk in the Index Cache. You can increase the CHUNKSIZE to improve the throughput to the storage medium. However, the disadvantage is that granular restores (for example, single file restore) are slower. On the other hand, large restores, like a full machine rebuild are generally faster.

By default, the chunk size is configured to get the optimal throughput to the storage medium. Dell recommends that you modify the chunk size when performing data movement to tape. You cannot improve the performance of data movement to a disk by modifying the chunk size.

The following table gives default the chunk size for tape and disk backups and the recommended range of chunk size:

STORAGE MEDIA	BACKUP TYPE	DEFAULT CHUNK SIZE	RECOMMENDED RANGE OF CHUNK SIZE
Tape	Granular Backup	4 GB	8 GB, 16 GB, 32 GB
	Database Backup	16 GB	8 GB, 16 GB, 32 GB
Disk and Optical	All Backups	2 GB	----
Direct Attached NDMP	All backups	4 GB	----

Increasing block size

You can increase the block size for faster write operations to tape media. The default block size for write operations is 64 KB.

For more information about modifying the block size, refer to **Increasing Block Size** in the CommVault documentation. You can also use the **setFlushBlockSize tool** available in the ResourcePack to modify the block size.

Before increasing the block size, make sure that the following criteria are satisfied:

- The new block size is supported by the Host Bus Adapter Driver installed in the MediaAgent and the tape device.
- All the MediaAgents associated with a Storage Policy support the block size configured for that storage policy.
- If different MediaAgents are used for backup and restore operations and the backup MediaAgent has a larger block size, then make sure that the restore MediaAgent is configured with Host Bus Adapters and tape drives that are able to read the data written with larger block sizes.

Increasing job manager update interval

The Job Controller window displays all the current jobs in the CommCell. A status bar at the bottom of the job controller shows the total amount of jobs; the number of jobs that are running, pending, waiting, queued and suspended; and the high and low watermarks. This information is updated at the close of each chunk, or within 5-minute intervals; whichever occurs sooner (see the screenshot below as an example).

When larger CHUNKS are configured for data movement operations, the amount of time between Job Manager Updates is automatically extended by that configuration.

Figure 5. Job Controller

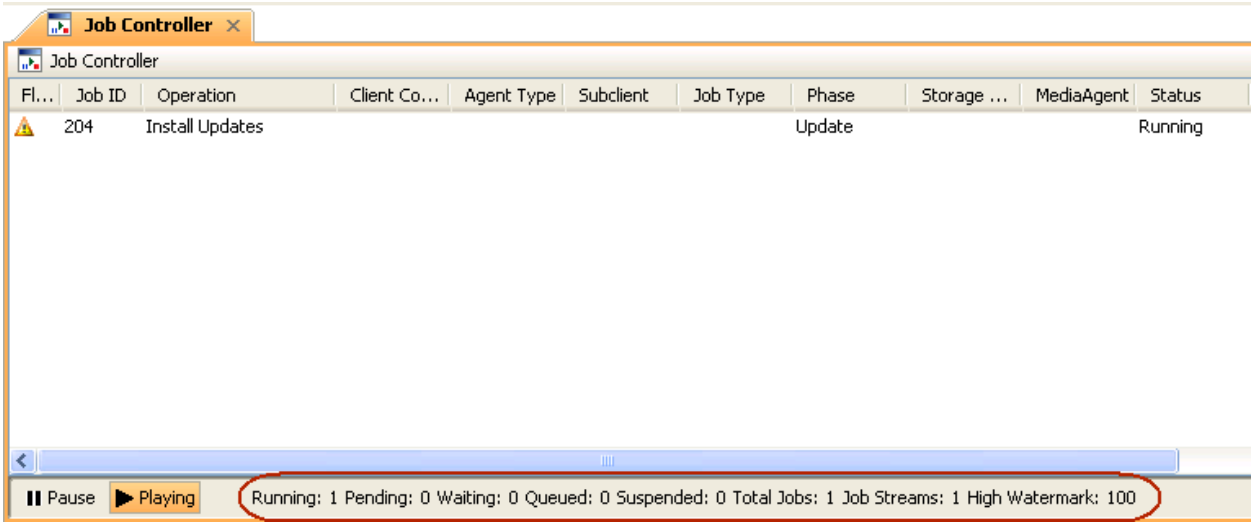

You can modify this interval to increase the performance of data movement operations. To modify the interval:

1. From the CommCell Browser, right-click the CommServe icon, and click **Control Panel**.
2. Double-click **Job Management**.
3. Click the **Job Updates** tab.
4. Click the integer in the **Protection (Mins)** column to modify the update interval for data protection jobs.

   Click the integer in **Recovery (Mins)** column to modify the update interval for data recovery jobs.

   In the **State update interval for ContinuousDataReplicator** box, click the integer in the box to change the time.
5. Click **OK** to save your changes.

Figure 6. Job Management

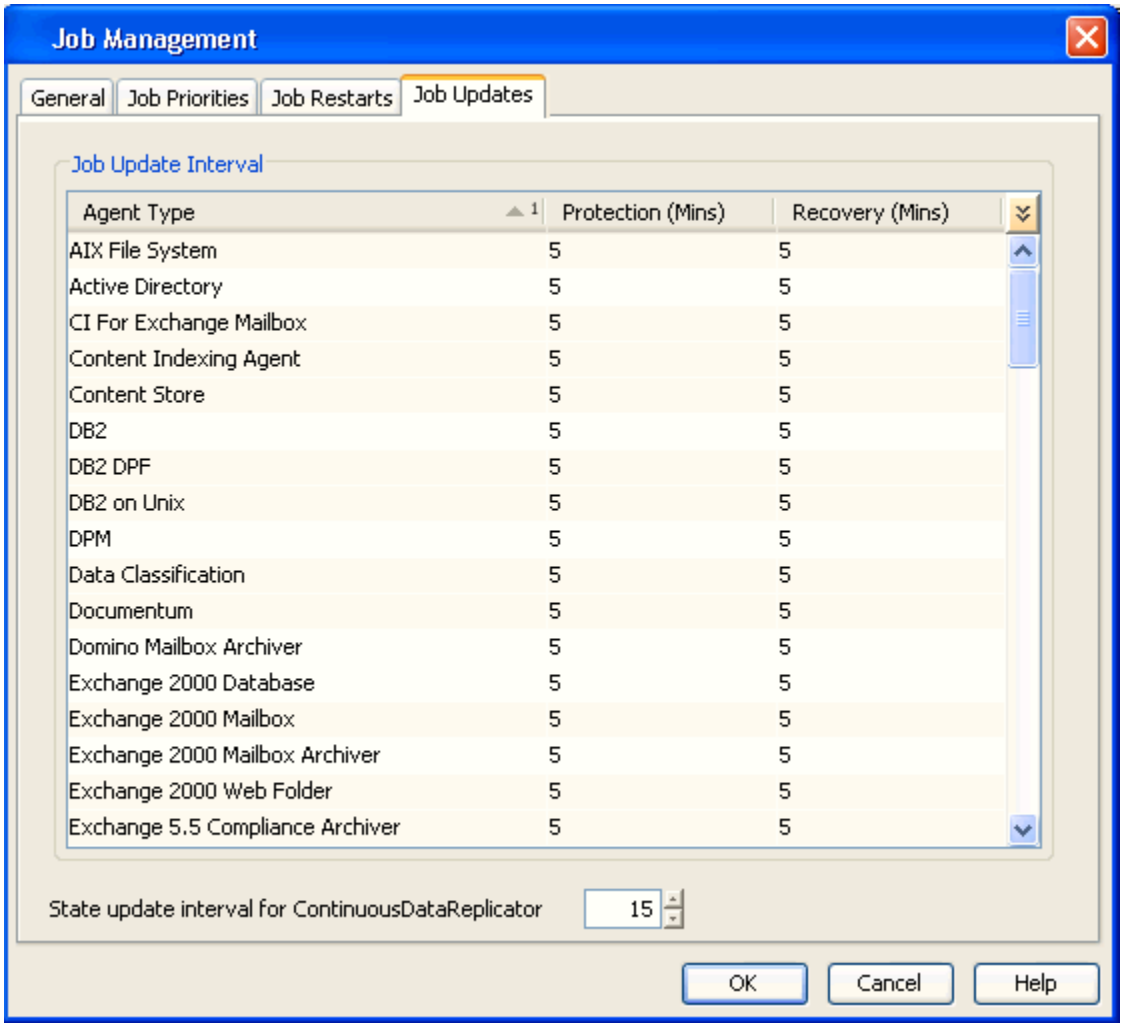

Increasing data transfer throughput from the client

You can increase the data transfer from the client to the MediaAgent by increasing the number of Network Agents. Network Agents establish data pipes to transfer data from the client to the MediaAgent. Therefore, increasing the number of Network Agents increases the data transfer throughput from the client.

Follow the steps given below to increase the data transfer throughput:

1. From the CommCell Browser, navigate to **<*Client*>** | **<Agent>** | **defaultBackupSet**.
2. Right-click the **<*subclient*>** or **<*instance*>** and select **Properties**.
3. Click the **Storage Devices** tab and then click the **Data Transfer Options** tab.

4. Enter the number of network agents in the **Network Agents** box.

   The default value is 2. You can set the value to 2 or 4.

5. Click **OK** to save your changes.

Figure 7. Data Transfer Option

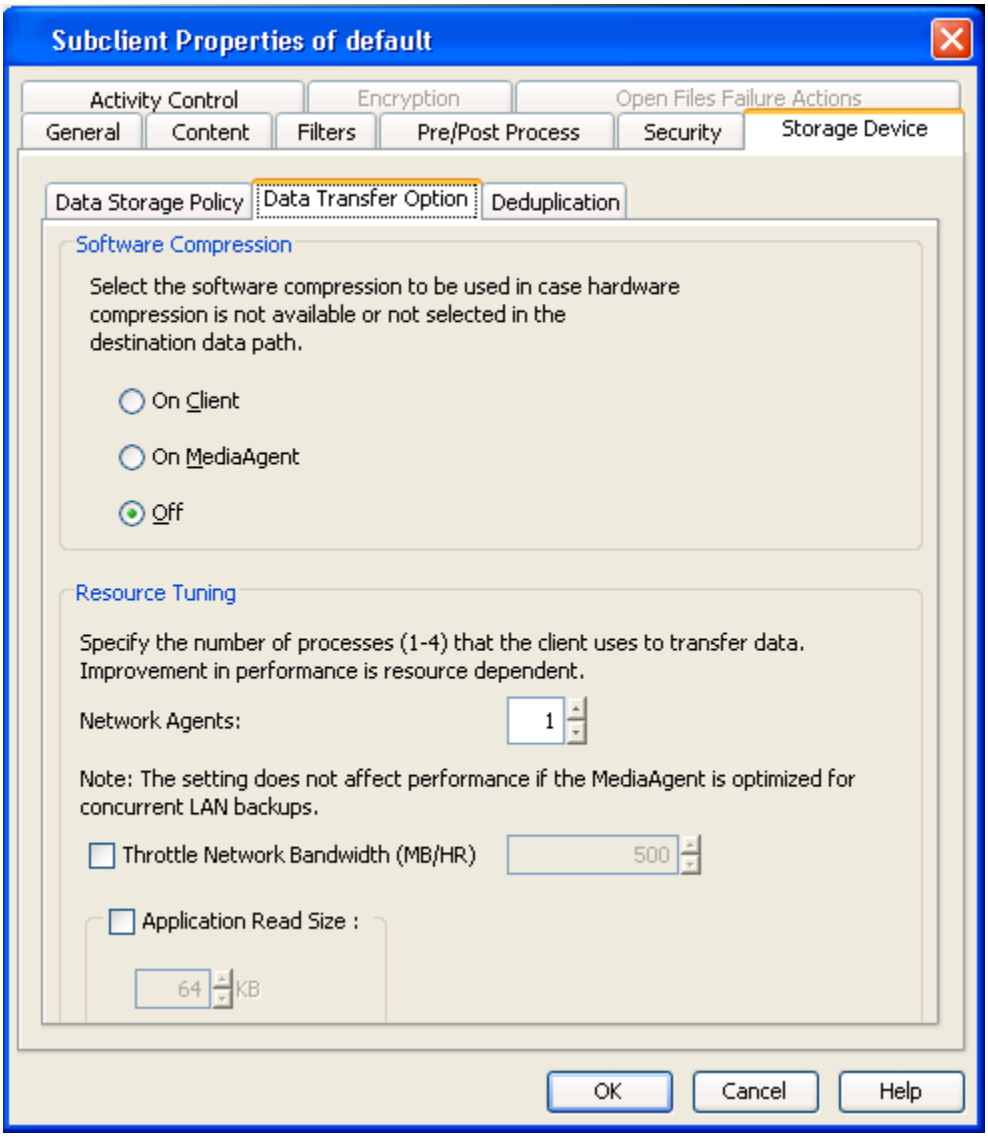

Increasing the pipeline buffers

The Data pipe buffers determine the amount of shared memory allocated on each computer for data pipes. The size of each buffer is 64K. By default, 30 data pipe buffers are established on each server for data movement operations. You can increase the data transfer throughput from the client by increasing the number of data pipe buffers.

Use the following registry key to modify the data pipe buffers:

REGISTRY KEY	LOCATION	SUPPORTED VALUES
nNumPipelineBuffers	• Windows - HKEY_LOCAL_MACHINE\SOFTWARE\CommVault Systems\Galaxy\Instance<xxx>\CVD • Unix - /etc/CommVaultRegistry/Galaxy/Instance<xxx>/MediaAgent	60, 120, 150, 300

Conclusion

This solution guide completes and enhances Dell's HPC storage solutions strategy with the introduction of an HSM offering. The Dell Tiered Storage Solution is available with deployment services and full hardware and software support from Dell and CommVault Systems, Inc. This document provides complete information on architecting, deploying, and tuning such a solution. The guidelines include hardware and software information along with detailed configuration steps and best practices make it easy to deploy and manage.

References

1) Clusters have become the most popular architecture for HPC today.
http://www.top500.org/overtime/list/36/archtype

2) Dell | Terascala HPC Storage Solution (DT-HSS)
http://content.dell.com/us/en/enterprise/d/hpcc/Storage_Lustre.aspx

3) Dell NFS Storage Solution for HPC (NSS)
http://content.dell.com/us/en/enterprise/d/hpcc/storage-dell-nss

4) CommVault Simpana 9 Data & Information Management Software
http://www.commvault.com/simpana.html

5) CommVault Simpana 9 Books Online
http://documentation.commvault.com/commvault/release_9_0_0/books_online_1/default.htm

6) Dell PowerVault DL2200 System Documentation
http://support.dell.com/support/edocs/stor-sys/pvdl2200/en/index.htm

7) Dell PowerVault ML6000 Tape Library System Documentation
http://support.dell.com/support/edocs/stor-sys/ml6000/en/index.htm

Appendix A: TSS Recipe

Pre-install preparation

SELinux

If you have SELinux enabled on the client computer, disable it. However, if it is not an option to disable it, then you must create the SELinux policy module as the user root before installing the CommVault Agent software or performing an archive operation. The SELinux Development package must be installed on the client.

To create SELinux policy module, perform the following steps as user "root":

1. Create the following files in the **/usr/share/selinux/devel** directory:

   ```
   <directory>/<file_name>.te
   ```

   Where <directory> is /usr/share/selinux/devel and <file_name> is the name of the file created to save the policy module statement. Use the same name for the policy module and the file.
   For example: When you are creating a policy module for the backup_IDA application, you can use the following file name: backup_IDA.te

   The content of the file is as follows:

   policy_module(<name>,<version>)
   ##############################

   Where <name> is the name of the policy module. You can give any unique name to the policy module, such as a process or application name. The version of the policy module is specified in <version>. It can be any number, such as 1.0.0. For Example: While creating a policy module for the backup_IDA application, you can use the following content. policy_module(backup_IDA,1.0.0)

   Create the file:
   ```
   <directory>/<file_name>.fc
   ```

   Where <directory> is /usr/share/selinux/devel and <file_name> is the name of the file created to save the policy module statement. Use the same name for the policy module and the file.
   For example: When you are creating a policy module for the backup_IDA application, you can use the following file name: backup_IDA.fc

   The content of the file should be as follows:

   The following list of files is not exhaustive. If the process fails to launch, check /var/log/messages; if required, add any files you need to the following list of files.

   ```
   /opt/<software installation
   <directory>/Base/libCTreeWrapper.so -- gen_context
   (system_u:object_r:texrel_shlib_t,s0)
   /opt/<software installation
   <directory>/Base/libCVMAGuiImplgso -- gen_context
   (system_u:object_r:texrel_shlib_t,s0)
   /opt/<software installation
   <directory>/Base/libdb2locale.so.1 -- gen_context
   (system_u:object_r:texrel_shlib_t,s0)
   /opt/<software installation
   ```

```
<directory>/Base/libdb2osse.so.1 -- gen_context
(system_u:object_r:texrel_shlib_t,s0)
/opt/<software installation
<directory>/Base/libDb2Sbt.so -- gen_context
(system_u:object_r:texrel_shlib_t,s0)
/opt/<software installation
<directory>/Base/libdb2trcapi.so.1 -- gen_context
(system_u:object_r:texrel_shlib_t,s0)
/opt/<software installation
<directory>/Base/libDrDatabase.so -- gen_context
(system_u:object_r:texrel_shlib_t,s0)
/opt/<software installation
<directory>/Base/libIndexing.so -- gen_context
(system_u:object_r:texrel_shlib_t,s0)
/opt/<software installation
<directory>/Base/libSnooper.so -- gen_context
(system_u:object_r:texrel_shlib_t,s0)
```

2. Create the policy file from command line, using the following commands. Make sure that you execute the following commands from the **/usr/share/selinux/devel** directory.

```
[root]# make backup_IDA.pp
Compiling targeted backup_IDA module
/usr/bin/checkmodule: loading policy configuration from tmp/backup_IDA.tmp
/usr/bin/checkmodule: policy configuration loaded
/usr/bin/checkmodule: writing binary representation (version 6) to
tmp/backup_IDA.mod
Creating targeted backup_IDA.pp policy package
rm tmp/backup_IDA.mod tmp/backup_IDA.mod.fc
[root]# semodule -i backup_IDA.pp
[root]#
```

3. Execute the policy module using the following command:

```
[root]# restorecon -R /opt/<software installation directory>
```

SELinux is now configured to work with this CommVault Linux Agent Software.

Installing CommVault agent: Interactive install

1. Log on to the NSS / NSS-HA node as **root**.
2. Run the following command from the CommVault Simpana 9 Software Installation Disc 3 (DellEnterprise_900_UnixLinuxMac_DVD3).

   ./cvpkgadd

   The product banner and other information are displayed.
3. Press **Enter**.
4. Read the license agreement. Type **y** and press **Enter**.
5. If your computer is 64-bit, type **2** and press **Enter.**

```
Please select a setup task you want to perform from the list
below: Advance options provide extra setup features such as
creating custom package, recording/replaying user selections and
installing External Data Connector software.

1) Install data protection agents on this computer

2) Advance options

3) Exit this menu

Your choice: [1]
```

6. Press **Enter**.

```
Certain Simpana packages can be associated with a virtual IP, or
in other words, installed on a "virtual machine" belonging to
some cluster. At any given time the virtual machine's services
and IP address are active on only one of the cluster's servers.
The virtual machine can "fail-over" from one server to another,
which includes stopping services and deactivating IP address on
the first server and activating the IP address/services on the
other server. You now have a choice of performing a regular
Simpana install on the physical host or installing Simpana on a
virtual machine for operation within a cluster.

Most users should select "Install on a physical machine" here.

1) Install on a physical machine

2) Install on a virtual machine

3) Exit

Your choice: [1]
```

7. If you only have one network interface, press **Enter** to accept the default network interface name and continue. If you have multiple network interfaces, enter the interface name that you want to use as default, and then press **Enter**.

```
We found one network interface available on your machine. We will
associate it with the physical machine being installed, and it
will also be used by the CommServe to connect to the physical
machine. Note that you will be able to additionally customize
Data pipe Interface Pairs used for the backup data traffic later
in the Simpana Java GUI. Please check the interface name below,
and make connections if necessary:

Physical Machine Host Name: [angel.company.com]
```

8. Press **Enter**.

```
Please specify the client name for this machine. It does not have
to be the network host name: you can enter any word here without
spaces. The only requirement is that it must be unique on the
CommServe.

Physical Machine Client name: [angel]
```

9. Type the appropriate number to install **File Archiver for UNIX Agent and the MediaAgent**. A confirmation screen marks your choice with an "**X**". Type **"d"** for Done, and press **Enter**.

```
Install Simpana on physical machine angel Please select the
Simpana module(s) that you would like to install.

[ ] 1) MediaAgent [1301] [CVGxMA]

[ ] 2) UNIX File System iDataAgent [1101] [CVGxIDA]

[a=all n=none r=reverse q=quit d=done >=next <=previous ?=help]

Enter number(s)/one of "a,n,r,q,d,>,<,?" here:2
```

10. Press **Enter**.

```
Do you want to use the agents for restore only without consuming
licenses? [no]
```

11. Type the appropriate number to install the latest software scripts and press **Enter**.

```
Installation Scripts Pack provides extra functions and latest
support and fix performed during setup time. Please specify how
you want to get this pack. If you choose to download it from the
website now, please make sure you have internet connectivity at
this time. This process may take some time depending on the
internet connectivity.

1) Download from the software provider website.

2) Use the one in the installation media

3) Use the copy I already have by entering its Unix path

Your choice: [1] 2
```

12. Press **Enter**.

```
Keep Your Install Up to Date - Latest Service Pack provides extra
functions and latest support and fix for the packages you are
going to install. You can download the latest service pack from
software provider website.

If you decide to download it from the website now, please make
sure you have internet connectivity at this time. This process
may take some time depending on the internet connectivity.
```

```
Do you want to download the latest service pack now? [no]
```

13. Press **Enter** to accept the default path.

```
Please specify where you want us to install Simpana binaries. It
must be a local directory and there should be at least 176MB of
free space available. All files will be installed in a "simpana"
subdirectory, so if you enter "/opt", the files will actually be
placed into "/opt/simpana".

Installation Directory: [/opt]
```

14. Press **Enter** to accept the default location.

```
Please specify where you want to keep Simpana log files. It must
be a local directory and there should be at least 100MB of free
space available. All log files will be created in a
"simpana/Log_Files" subdirectory, so if you enter "/var/log", the
logs will actually be placed into "/var/log/simpana/Log_Files".

Log Directory: [/var/log]
```

15. Type **“no”**.

```
Most of Software processes run with root privileges, but some are
launched by databases and inherit database access rights. To make
sure that registry and log files can be written to by both kinds
of processes we can either make such files world-writeable or we
can grant write access only to processes belonging to a
particular group, e.g. a "simpana" or a "dba" group. We highly
recommend now that you create a new user group and enter its name
in the next setup screen. If you choose not to assign a dedicated
group to Software processes, you will need to specify the access
permissions later.

If you're planning to backup Oracle DB you should use "dba"
group.

Would you like to assign a specific group to Software? [no]
```

16. Type **“d”** for done.

```
Access Permissions for Other Users

Installer will assign full access rights to root user and its
belonging group
for all installed Software files and its processes.

For any other users, you can specify the access permissions now.

However, since you chose not to assign a dedicated group in
previous step,
```

```
make sure you specify sufficient access rights for other users if
you are also
planning to install Software agents involving third party
software protection.

[X] 1) Allow read permission to other users
[X] 2) Allow write permission to other users
[X] 3) Allow execute permission to other users

[a=all n=none r=reverse q=quit d=done >=next <=previous ?=help]
Enter number(s)/one of "a,n,r,q,d,>,<,?" here:
```

17. If you indicated "**Yes**" in Step 15, you are prompted for the group name that you must use to launch processes. Enter the group name and then press **Enter**.

```
Please enter the name of the group which will be assigned to all
Software files and on behalf of which all Software processes will
run. In most of the cases it's a good idea to create a dedicated
"simpana" group. However, if you're planning to use Oracle
iDataAgent or SAP Agent, you should enter Oracle's "dba" group
here.

Group name: sky1

REMINDER

If you are planning to install Simpana Informix, DB2, PostgreSQL,
Sybase or Lotus Notes iDataAgent, please make sure to include
Informix, DB2, etc. users into group "sky1".

Press <ENTER> to continue ...
```

18. Type a network TCP port number for the Communications Service (CVD) and press **Enter**. Type a network TCP port number for the Client Event Manager Service (EvMgrC) and press **Enter**.

```
Every instance of Simpana should use a unique set of network
ports to avoid interfering with other instances running on the
same machine. The port numbers selected must be from the reserved
port number range and have not been registered by another
application on this machine.

Please enter the port numbers.

Port Number for CVD : [8600]

Port Number for EvMgrC: [8602]
```

19. If you do not want to configure the firewall services, press **Enter**.

```
Is there a firewall between this client and the CommServe? [no]
```

20. Type the fully qualified CommServe host name and press **Enter**.

```
Please specify hostname of the CommServe below. Make sure the
hostname is fully qualified, resolvable by the name services
configured on this machine. CommServe Host Name:
mycommserve.company.com
```

21. Press **Enter**.

```
Commcell Level Global Filters are set through Simpana GUI's
Control Panel in order to filter out certain directories or files
from backup Commcell-widely. If you turn on the Global filters,
they will be effective to the default subclient. There are three
options you can choose to set the filters.

1) Use Cell level policy

2) Always use Global filters

3) Do not use Global filters

Please select how to set the Global Filters for the default
subclient? [1]
```

22. Type the appropriate number to select the **Client Group** and press **Enter**.

```
Client Group(s) is currently configured on CommServe
mycommserve.company.com. Please choose the group(s) that you want
to add this client angel.company.com to. The selected group(s)
will be marked (X) and can be deselected if you enter the same
number again. After you are finished with the selection, select
"Done with the Selection".

[ ] 1) Unix

[ ] 2) DR

[a=all n=none r=reverse q=quit d=done >=next <=previous ?=help]s

Enter number(s)/one of "a,n,r,q,d,>,<,?" here: 2
```

23. Enter the number corresponding to the storage policy through which you want to back up the External Data Connector and press **Enter**.

```
Please select one storage policy for this IDA from the list
below:

1) SP_StandAloneLibrary2_2

2) SP_Library3_3

3) SP_MagLibrary4_4

Storage Policy: [1]
```

24. Type "**3**" for the **Exit** option and press **Enter**. The installation is now complete.

```
Certain Simpana packages can be associated with a virtual IP, or
in other words, installed on a "virtual machine" belonging to
some cluster. At any given time the virtual machine's services
and IP address are active on only one of the cluster's servers.
The virtual machine can "fail-over" from one server to another,
which includes stopping services and deactivating IP address on
the first server and activating the IP address/services on the
other server. Currently you have Simpana installed on physical
node angel.company.com.

Now you have a choice of either adding another package to the
existing installation or configure Simpana on a virtual machine
for use in a cluster.

1) Add another package to angel.company.com

2) Install Simpana on a virtual machine

3) Exit

Your choice: [3]
```

Installing CommVault agent: Clustered environment (NSS-HA)

Verify that NSS-HA stack is installed, configured, and operational. Follow the previous procedure "A.2. Install CommVault Agent: Interactive Install", making sure to install on the ACTIVE node FIRST and selecting to install the File System *i*DataAgent as virtual on each node of the cluster (Step 6 above). Each node within a cluster, whether active or passive, must have the necessary software components installed. Installing these software components to be cluster-aware involves both the active and passive nodes, whether the software component is installed directly to a particular physical node or not. During the cluster server installation, software binaries are installed on all physical nodes in the cluster, and the configuration files on all physical nodes are edited to include information about all installed cluster servers.

Note the following guidelines:

- Dell recommends that you install the MediaAgent or Agent software for the cluster server on the active node first.
- You can perform software installation and configuration for all available preferred nodes automatically when you install the software from a physical node to the cluster server.
- You can subsequently install the software on any passive node that was unavailable during the automated installation.
- Install the MediaAgent software on the physical node. When you install the MediaAgent software on the physical node, you can use GridStor® to provide the necessary failover capabilities. See CommVault documentation on Alternate Data Paths (GridStor) - Clustered Environments for more information.
- Install the MediaAgent software from the active node to the cluster server to take advantage of failover protection provided by the clustering software. You should perform this installation with all passive nodes available, so the software installs on all nodes in the cluster at once. If any passive node is not available during the installation, perform a separate passive node installation.

- Configure the index cache either on the shared disk or on a network share accessible to all the nodes on the cluster. Do not configure the index cache on the physical nodes.

Installing CommVault agent: Commcell console

To install CommVault Agent: ComCell Console:

1. From the CommCell Browser, select **Tools** | **Add/Remove Software** |**Install Software.**

Figure 8. CommCell Console

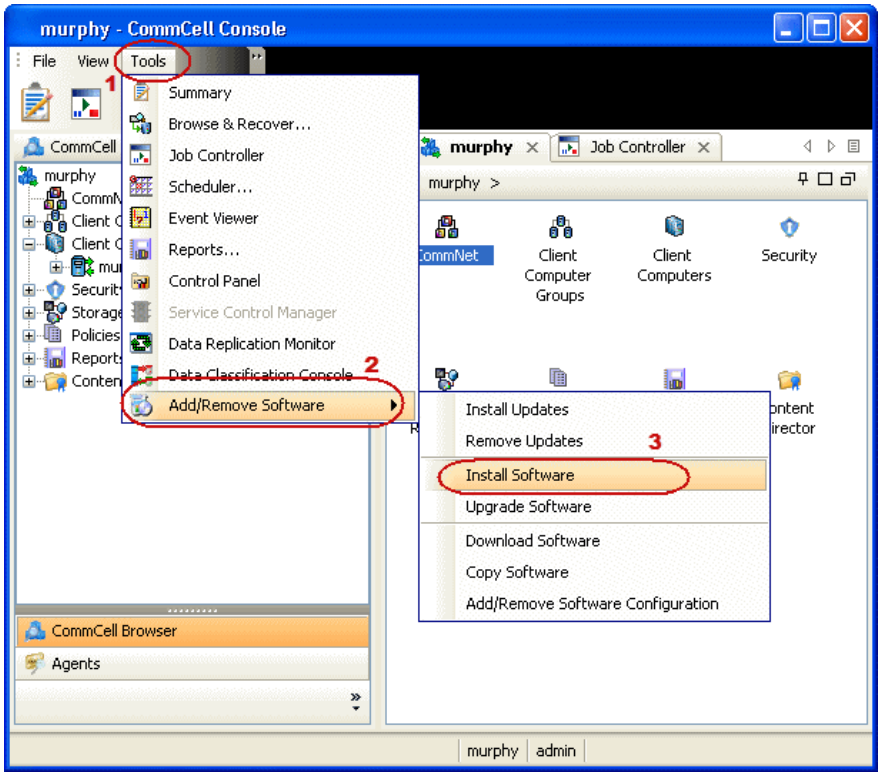

2. Click **Next.**

Figure 9. Installer

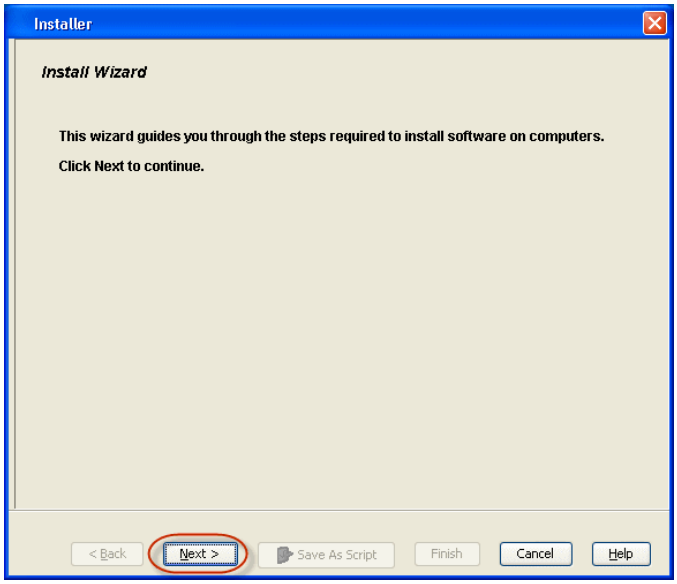

3. Select **Unix and Linux**. Click **Next.**

Figure 10. Installer – Select OS

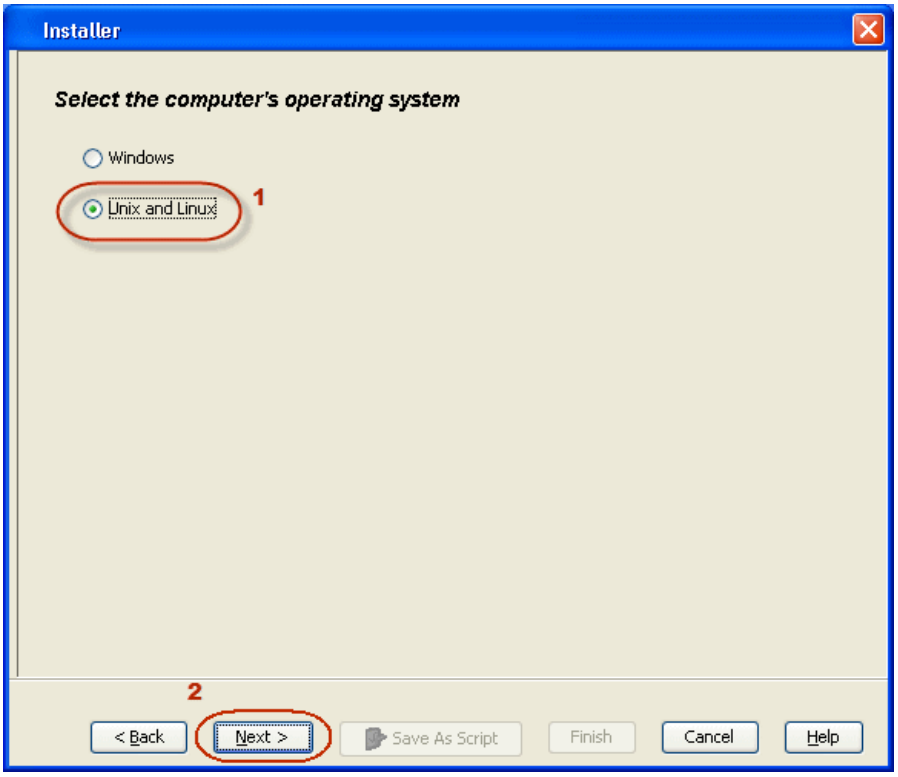

4. Select **Manually Select Computers**. Click **Next.**

Figure 11. Installer – Discover computers

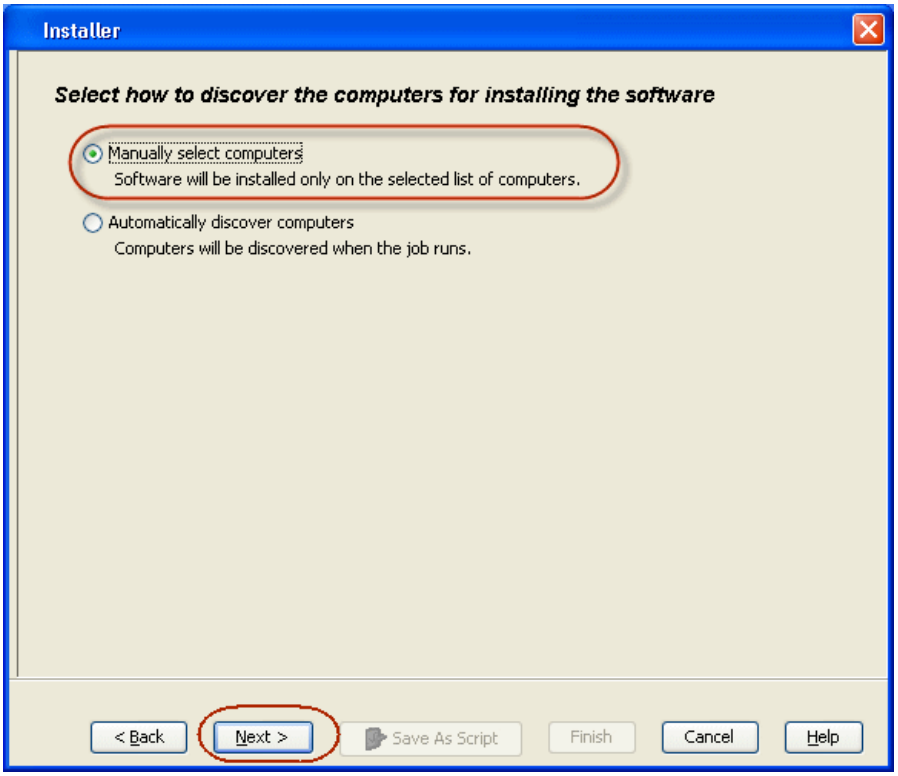

5. Enter the fully qualified domain name of the computer in which you want to install. The **File Archiver for UNIX Agent** will be installed on this client computer.
6. Click **Next**.

Figure 12. Installer – Enter Host Names

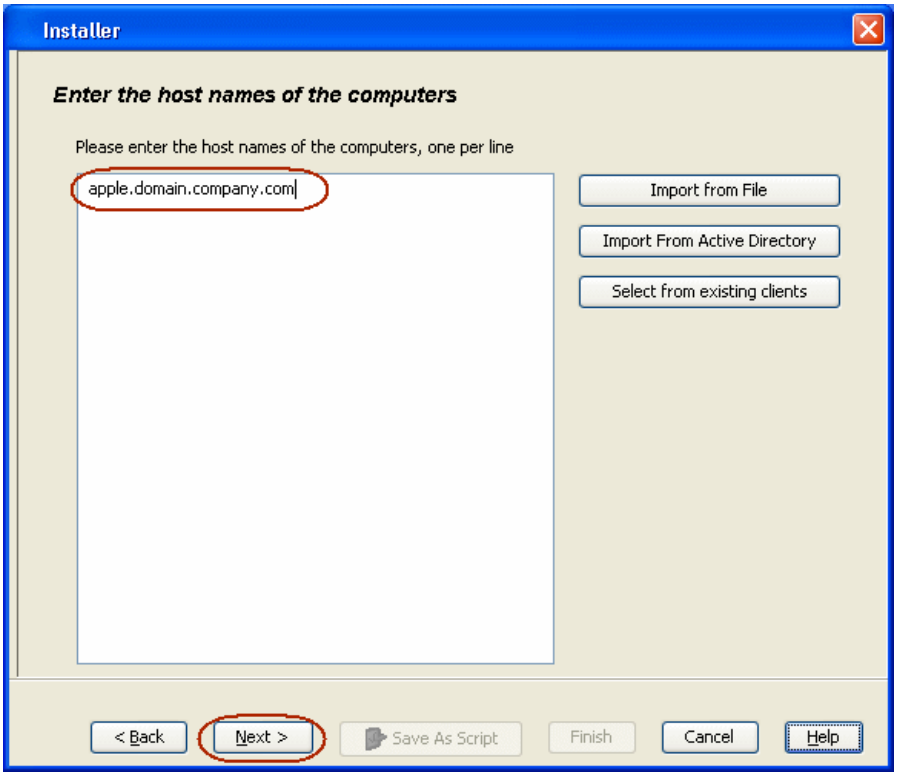

7. Click **Next**.

Figure 13. Installer – Select Software Cache

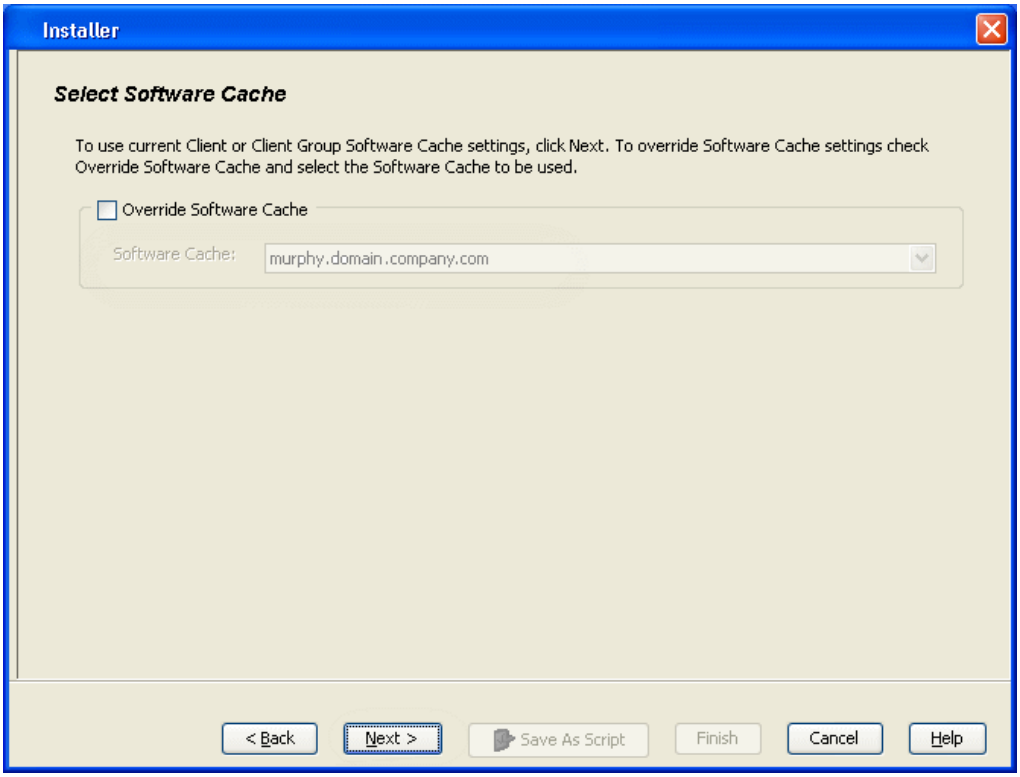

8. Specify **User Name** and **Password** of client computer. Click **Next**.

Figure 14. Installer – Enter Account Information

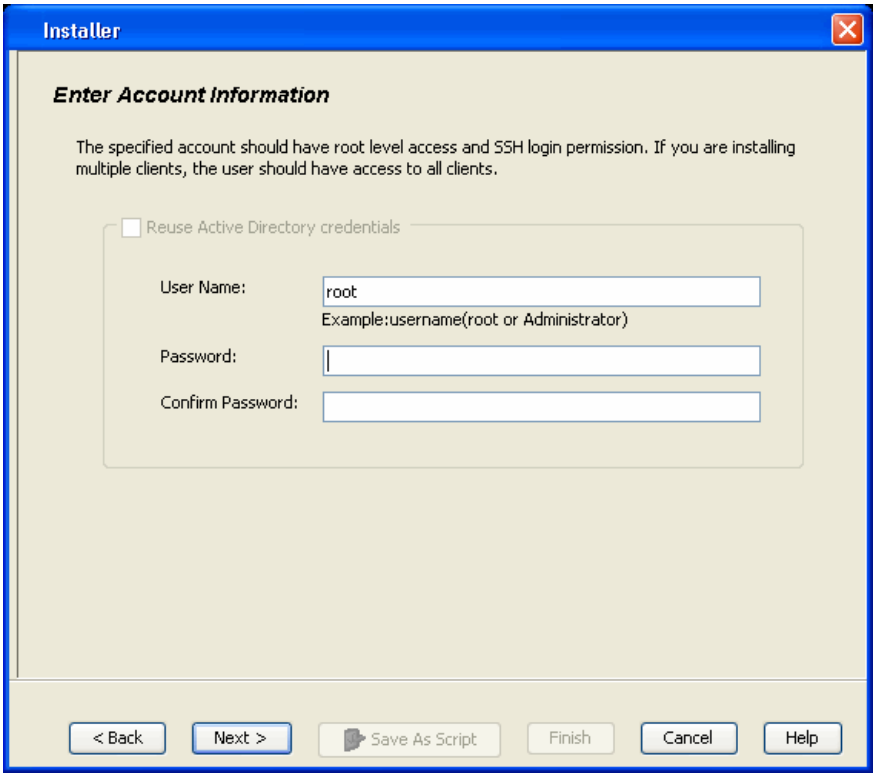

9. Select **File Archiver for UNIX Agent**. Click **Next**.

Figure 15. Installer – Select Package(s) to Install

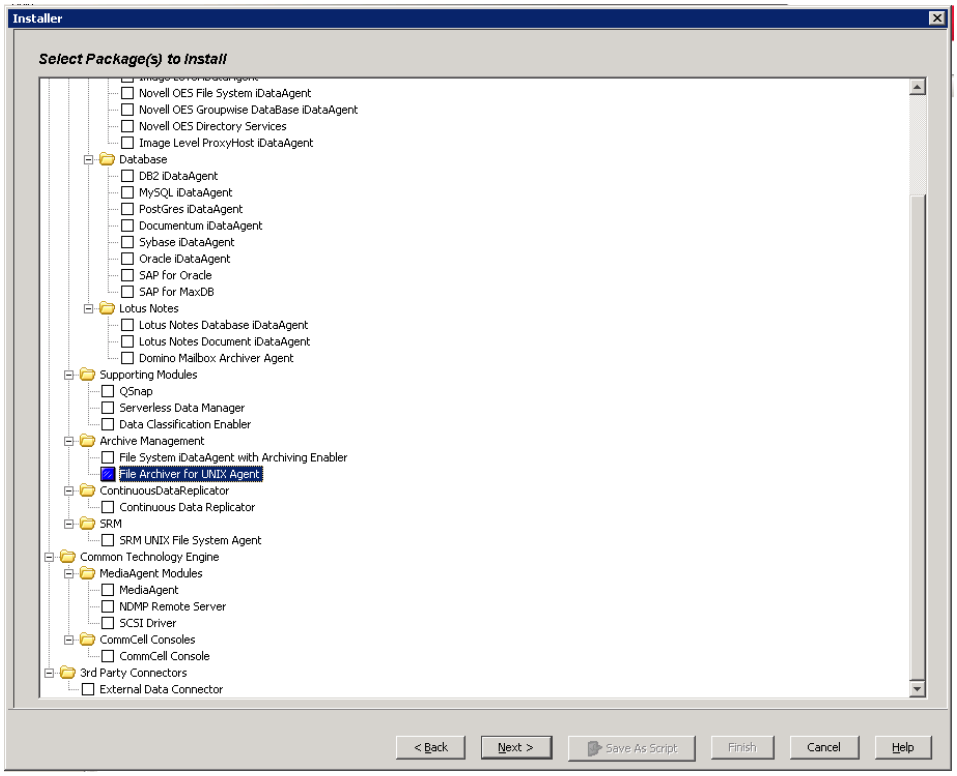

10. Click **Next**.

Figure 16. Installer – Enter settings for software

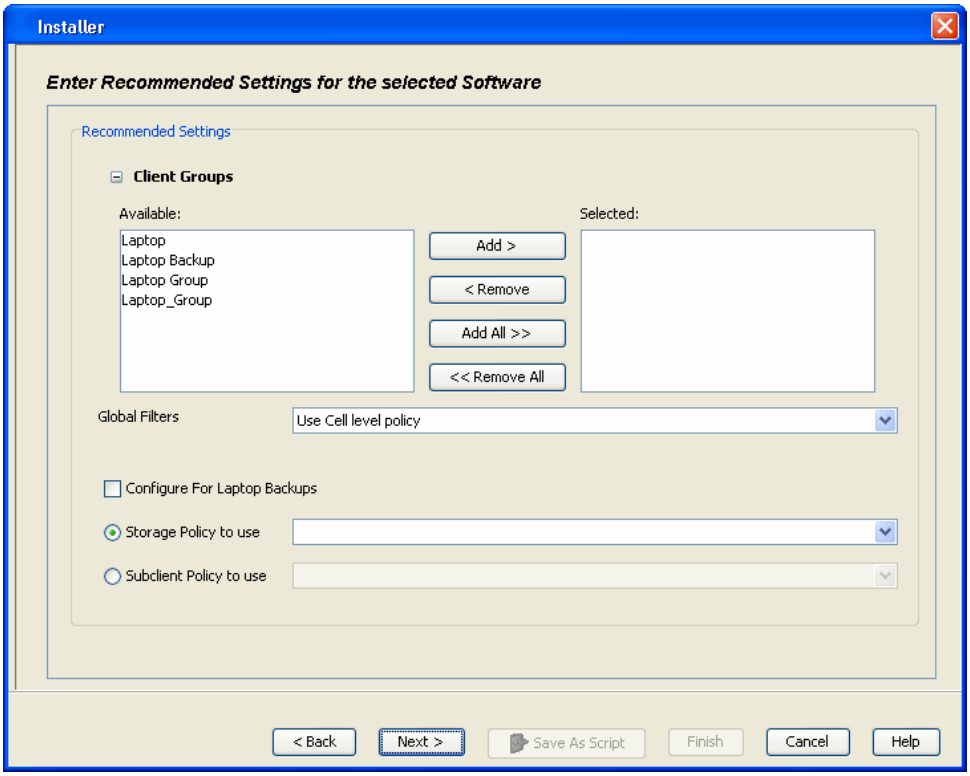

11. Click **Next**.

Figure 17. Installer – Enter Additional Install Options

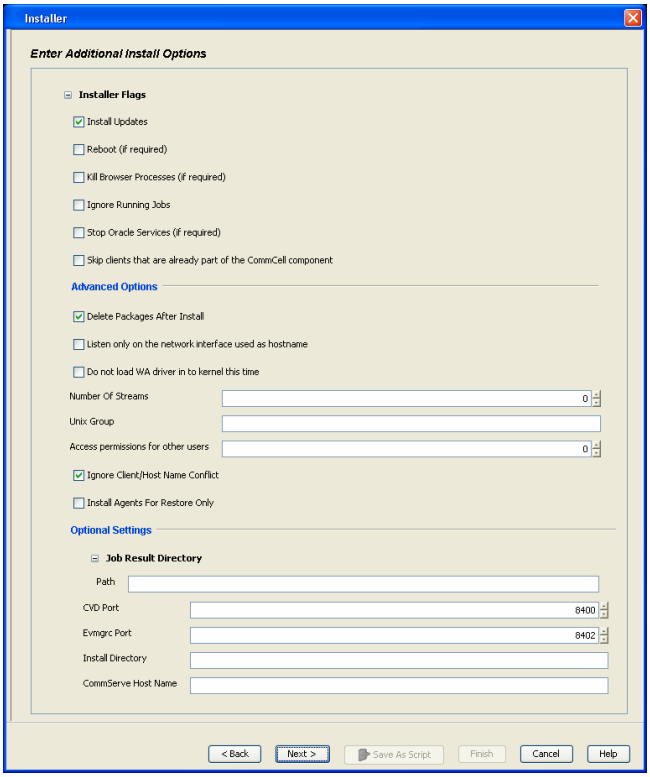

12. Select **Immediate**. Click **Next**.

Figure 18. Installer – When To Run The Job

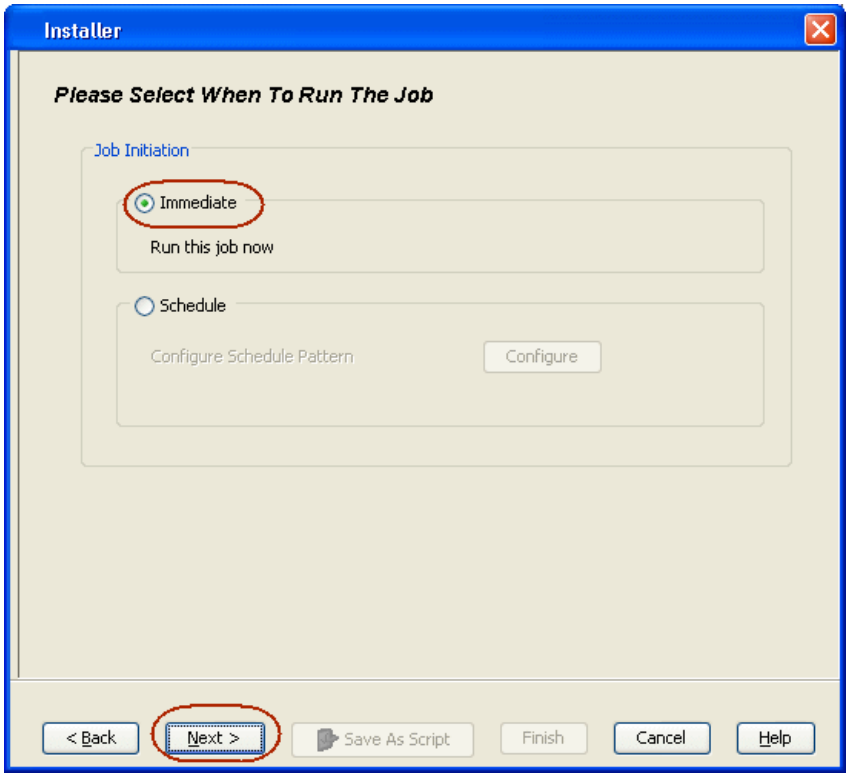

13. Click **Finish**.

Figure 19. Installer - Summary

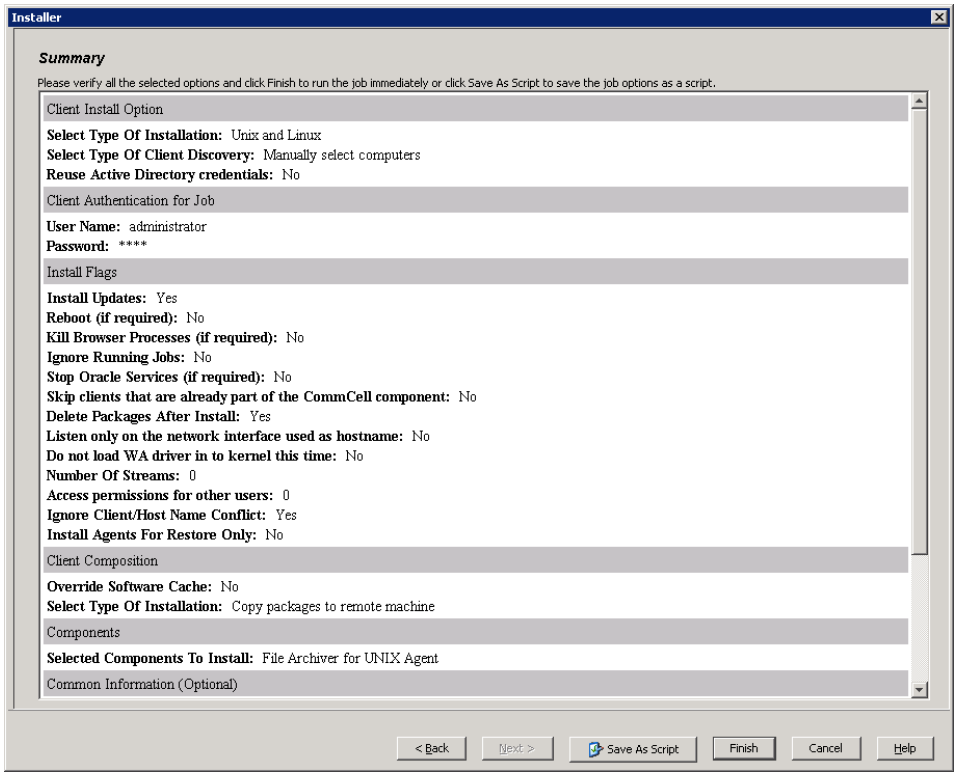

Configuring storage resources (example)

To configure a disk device (nss nfs storage tier):

1. Click the **Backup Target** button on **EZ Operations Wizard**. If the **EZ Operations Wizard** does not display, double-click the icon in the toolbar.

Figure 20. EZ Operations Wizard

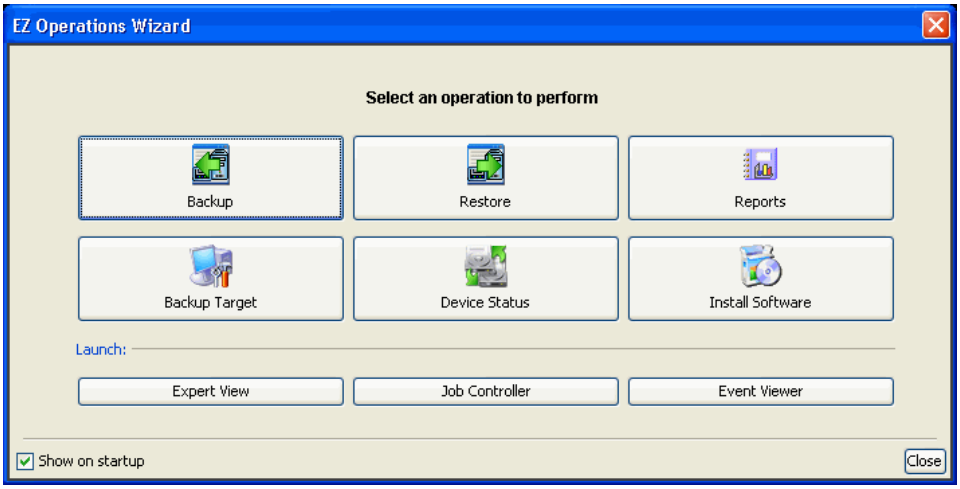

2. Click **Disk Library (For backup to disk)** and click **Next.**

Figure 21. Select Configuration Type

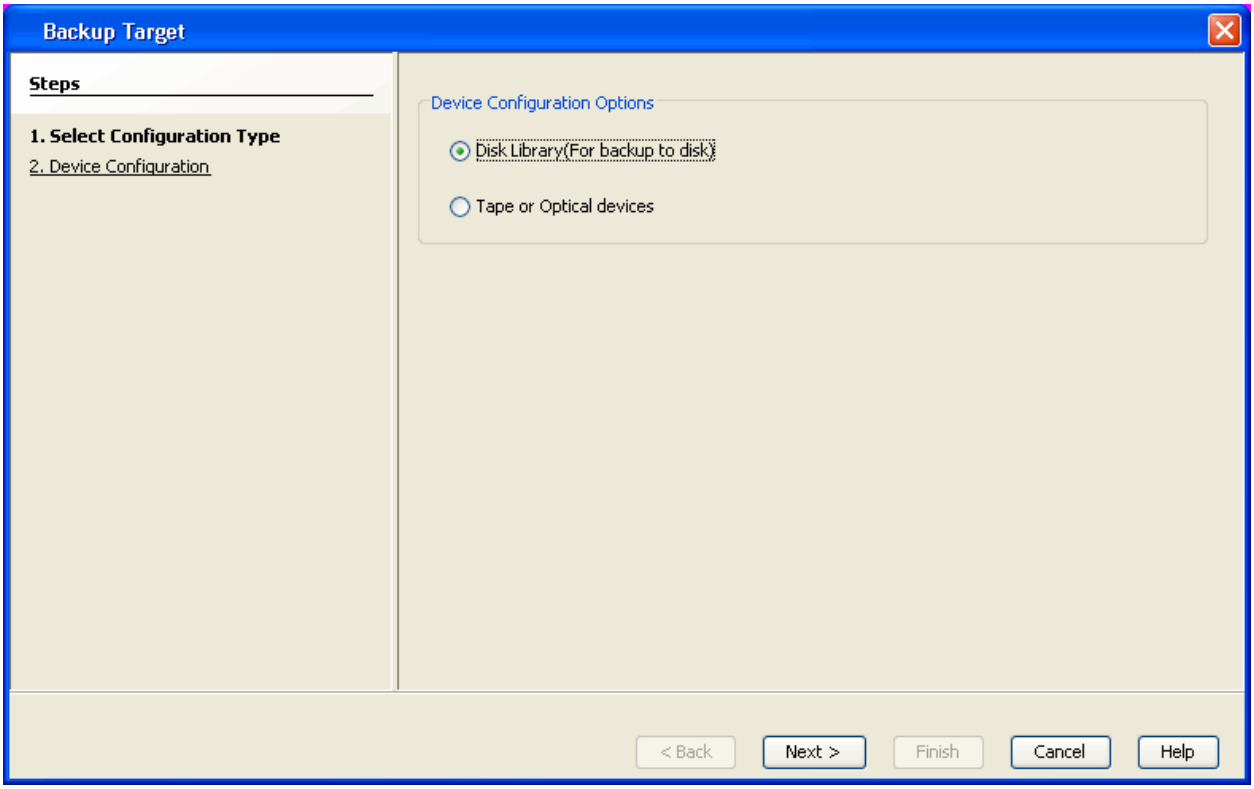

3. Click **Use Local Disk**. Type the name of the folder in which the disk library must be located in the **Enter backup destination folder** box or click the **Browse** button to select the folder. This will be the exported NFS volume from NSS stack if setting up NSS storage resource.
4. Click **Next**.

Figure 22. Disk Library Configuration

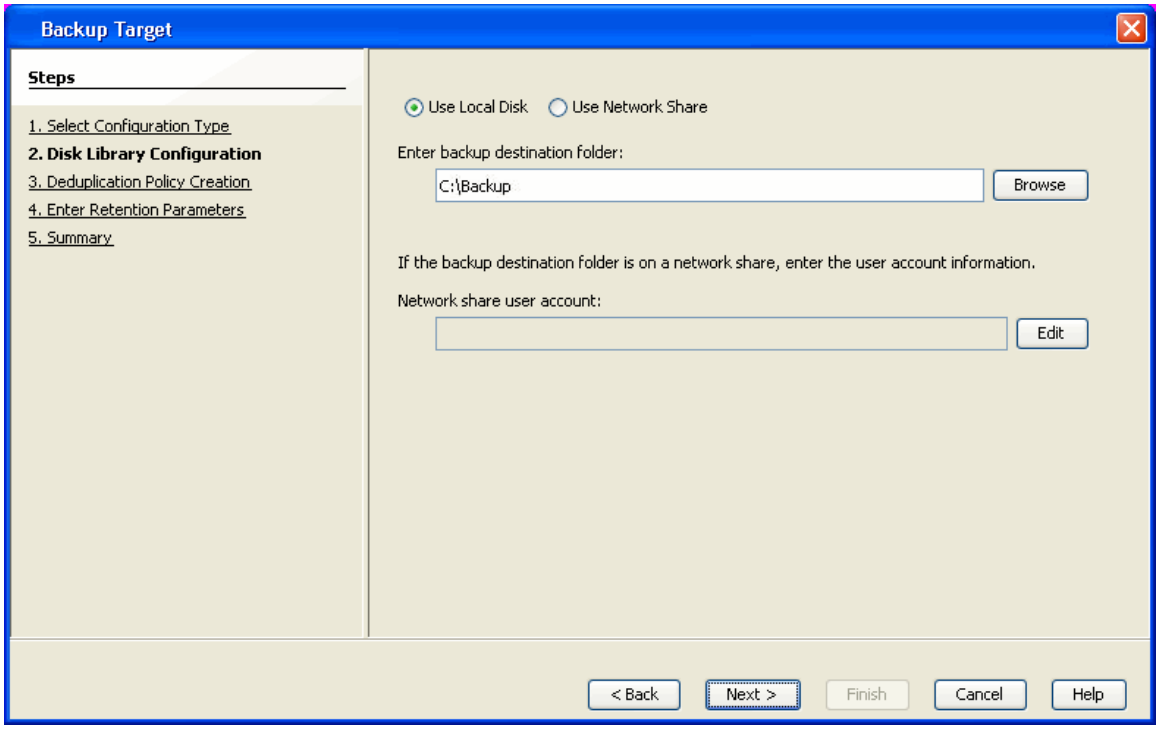

5. Select the **Enable Deduplication** option. This option saves disk space for storage. Type the name of the folder in which the deduplication database must be located in the **Deduplication Store location** box or click the **Browse** button to select the folder.
6. Click **Next**.

Figure 23. Deduplication Policy Creation

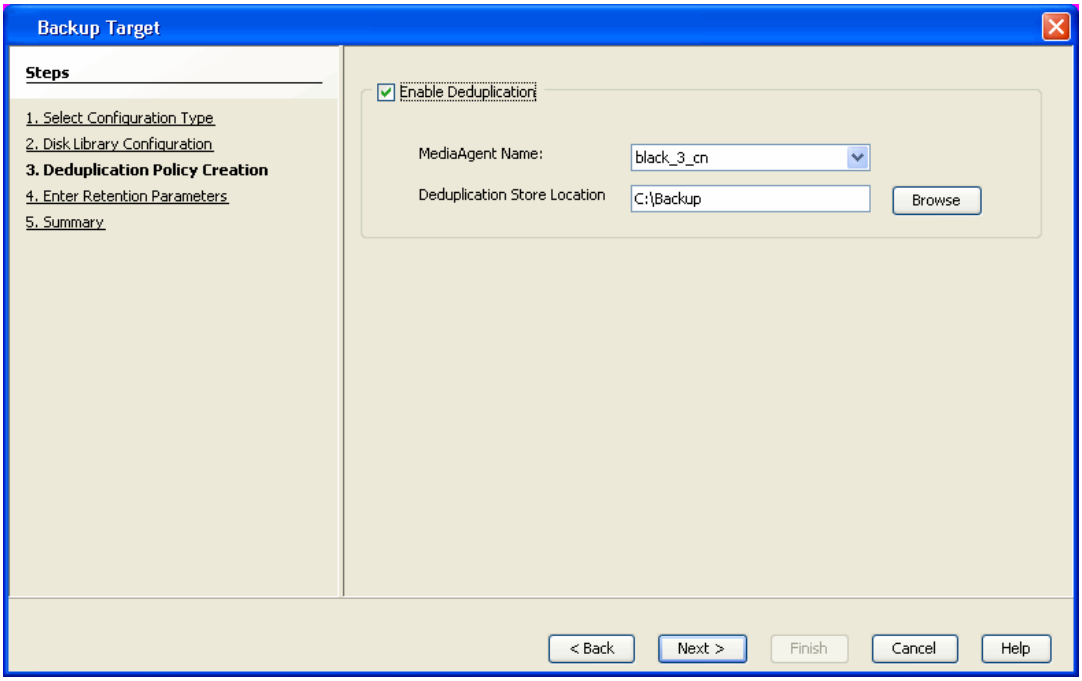

7. Click **Next**.

Figure 24. Enter Retention Parameters

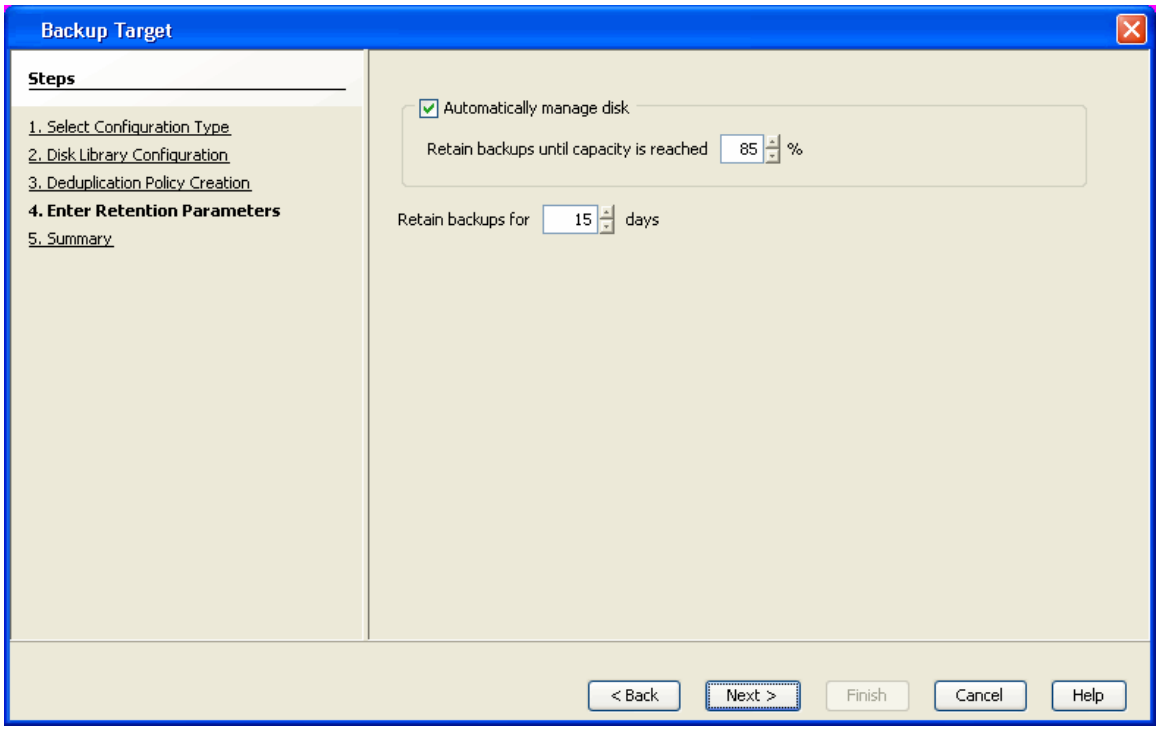

8. Click **Finish**.

Figure 25. Backup Target Summary

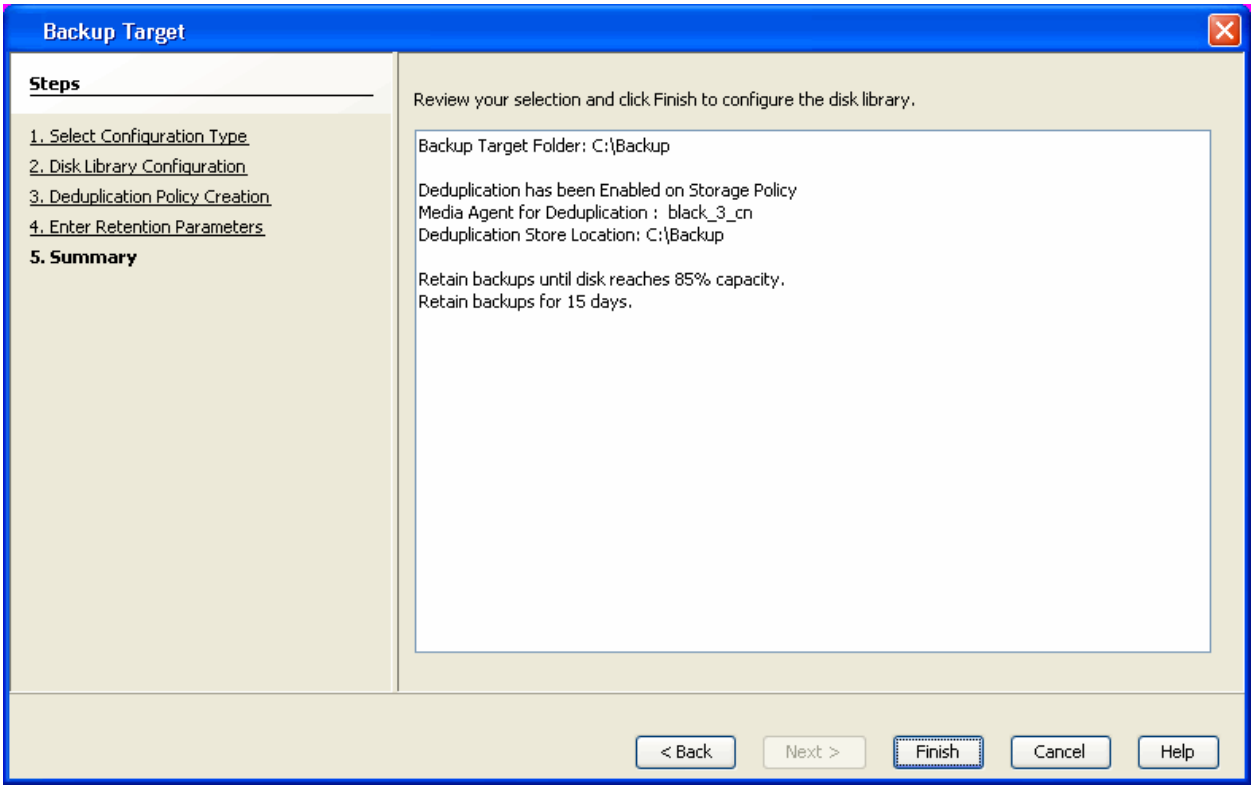

Configuring a tape device (ML6000 Tape Library)

To configure a tape device:

1. Click the **Backup Target** button on **EZ Operations Wizard**. If the **EZ Operations Wizard** does not display, double-click the icon in the toolbar.

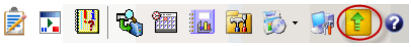

Figure 26. EZ Operations Wizard

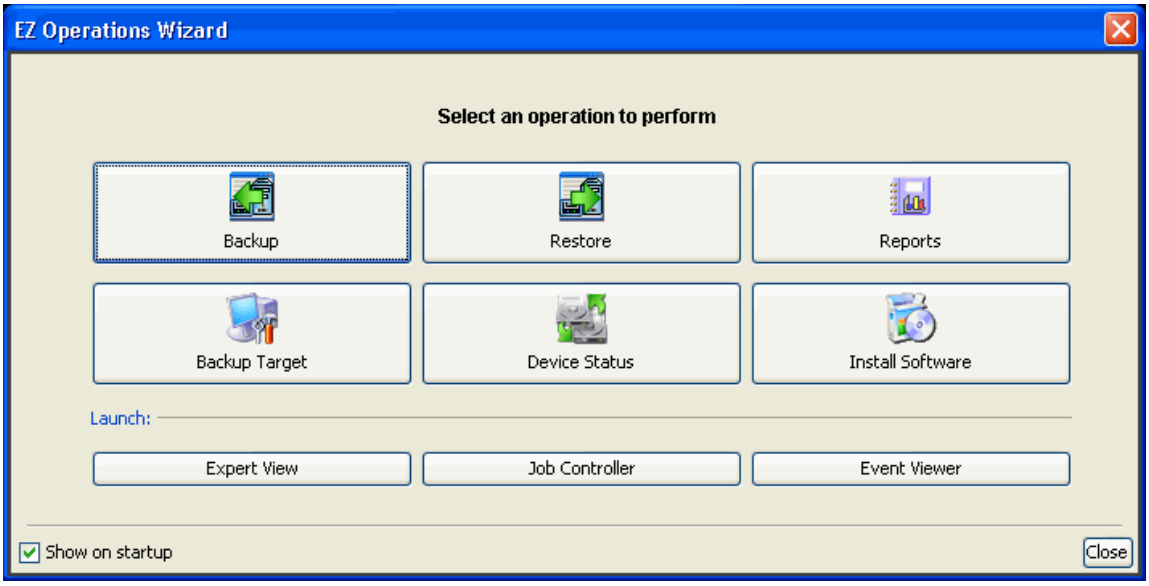

2. Select **Tape or Optical devices**. Click **Next**.

Figure 27. Select Configuration Type

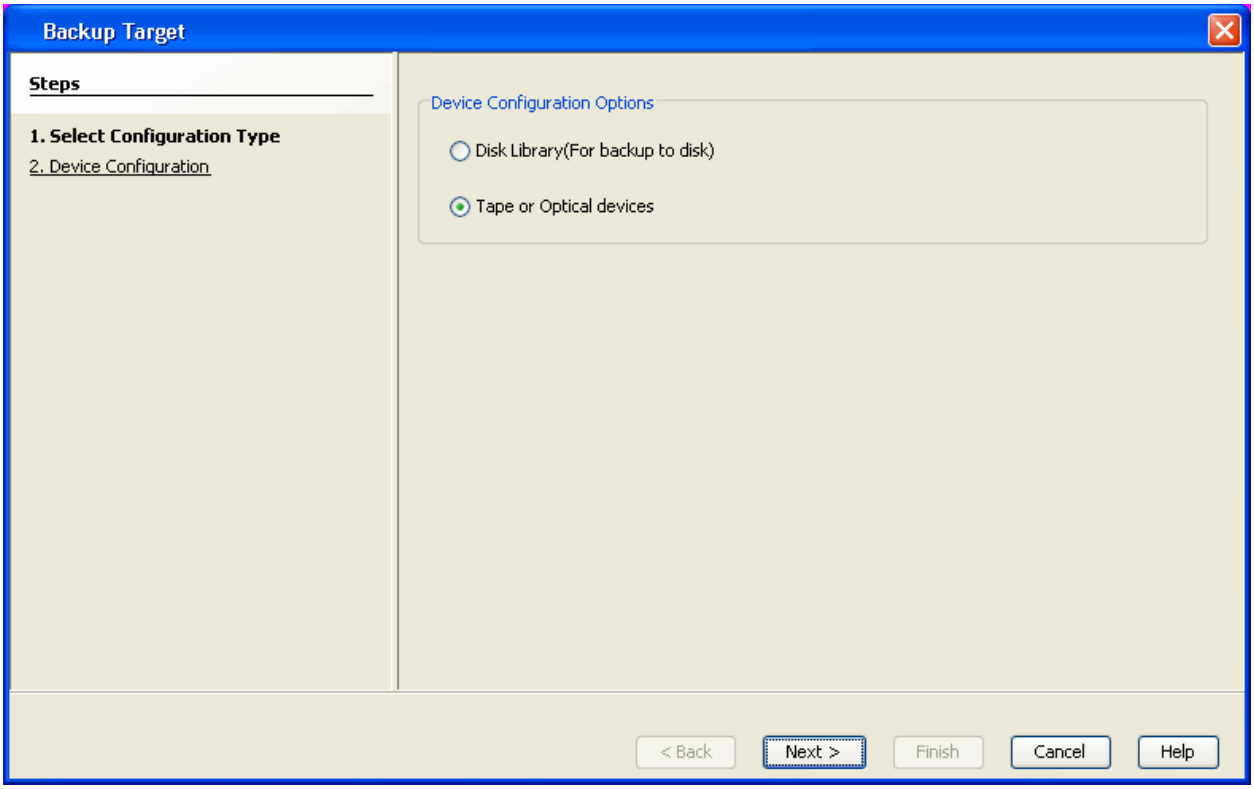

3. Click and select the library you wish to configure. Click **Next**.

Figure 28. Select Library

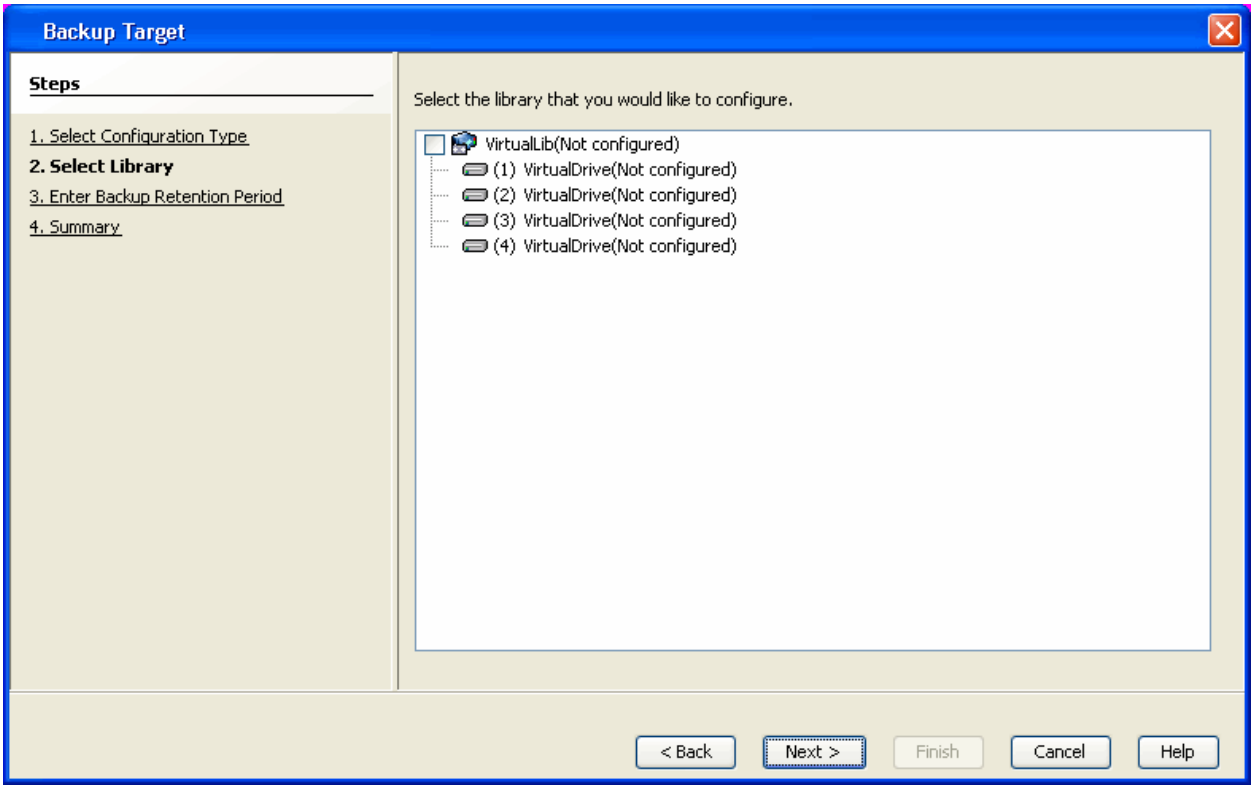

Click **Next.**

Figure 29. Enter Backup Retention Period

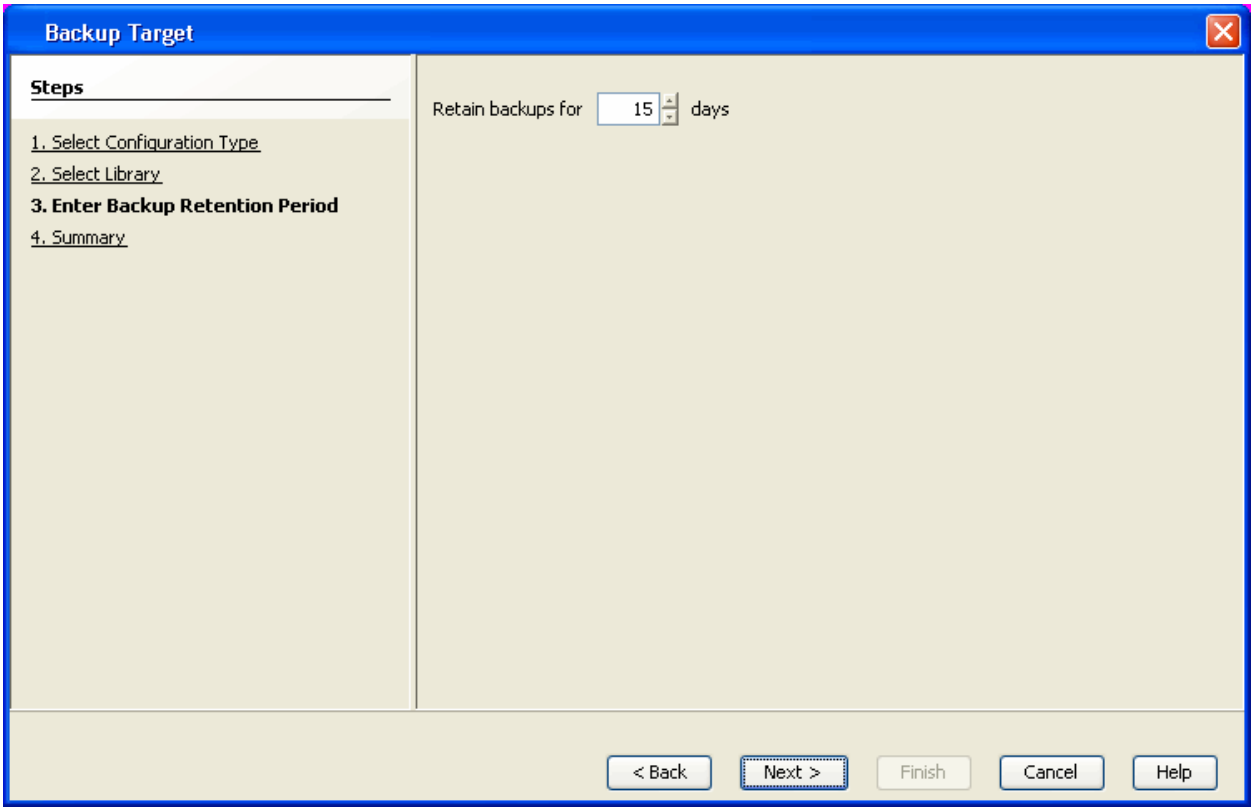

4. Click **Finish**.

Figure 30. Summary

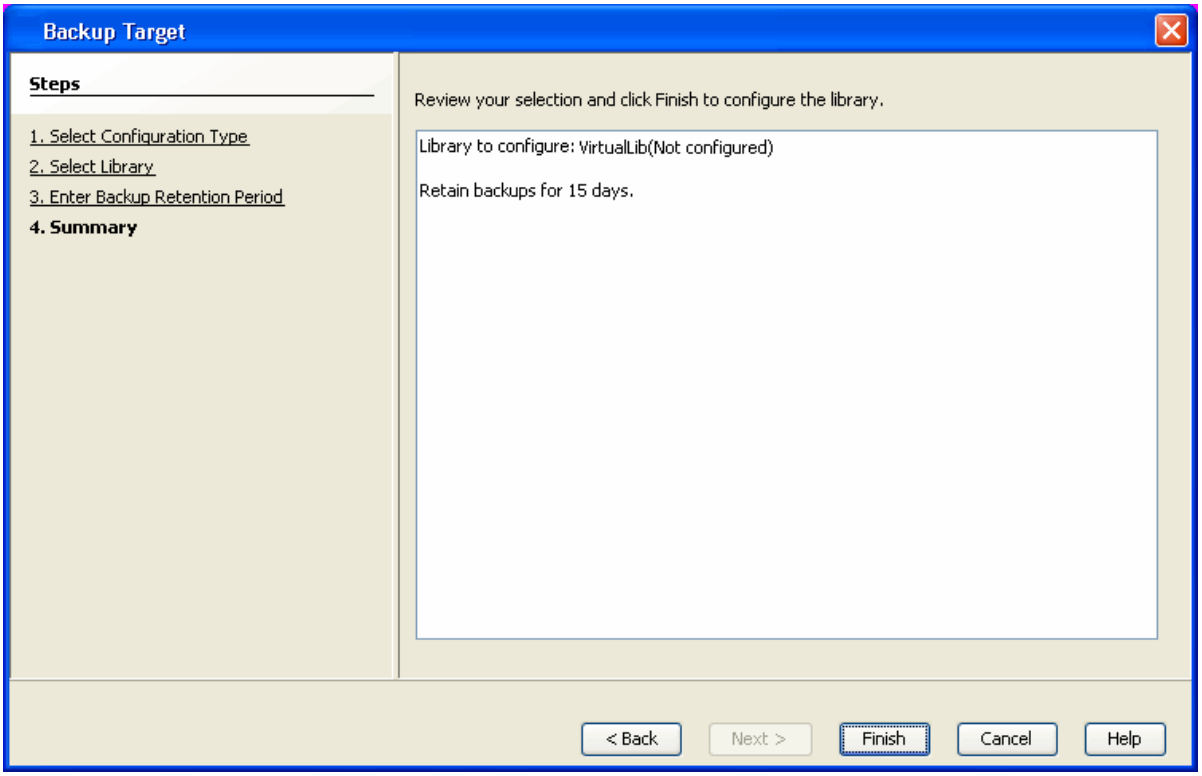

5. Click **OK**.

Figure 31. Success Message

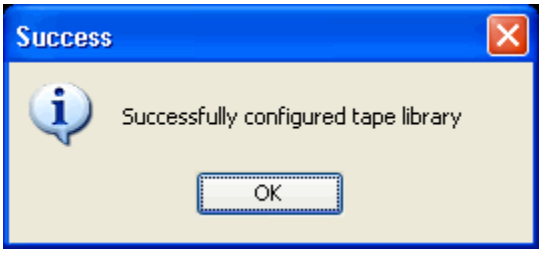

Configuring storage policies

A storage policy is automatically created when you configure a device. A Storage Policy acts as a channel through which data is transferred to the storage device. As the name indicates, a Storage Policy allows you to establish a comprehensive set of storage parameters - such as data retention, streams, deduplication, and so forth, for the data channeled through the storage policy.

1. From the CommCell Console, navigate to **Client** | **File System** | **Backup Set**. Right-click the **Subclient** and click **Properties**

Figure 32. CommCell Browser

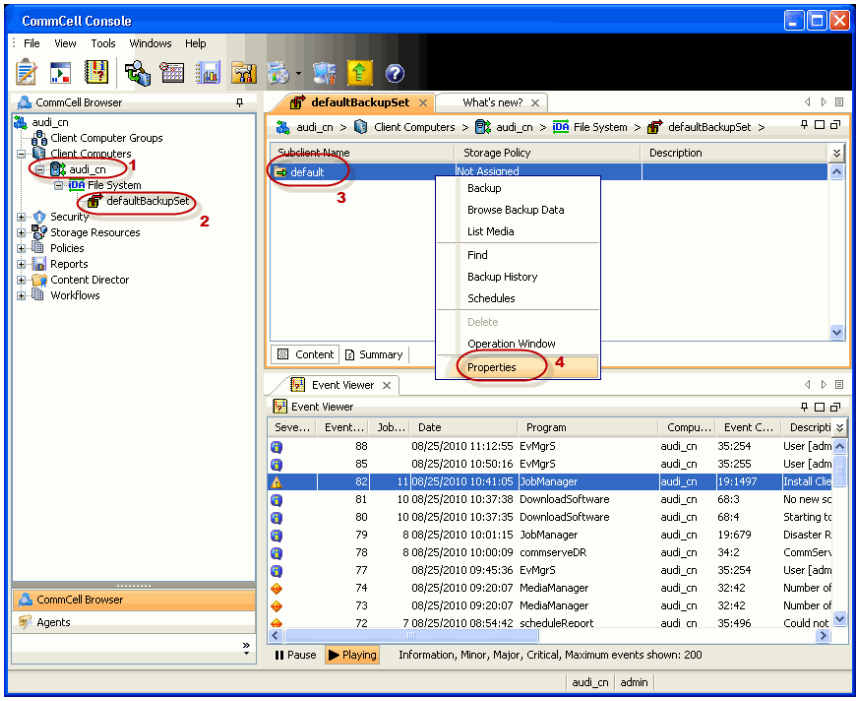

2. Click on the **Storage Device** tab.

Figure 33. Data Storage Policy

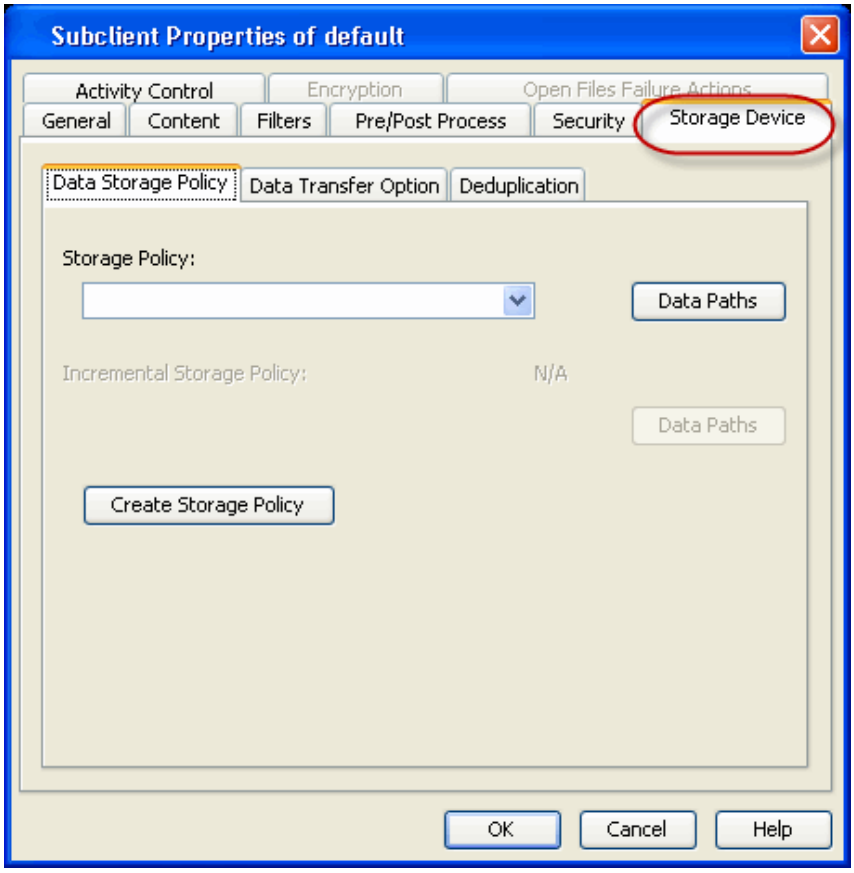

3. Click the **Data Storage Policy** tab.
4. From the **Storage Policy** list, click the name of a Storage Policy.
5. Click **OK**.

Figure 34. Storage Policy drop-down

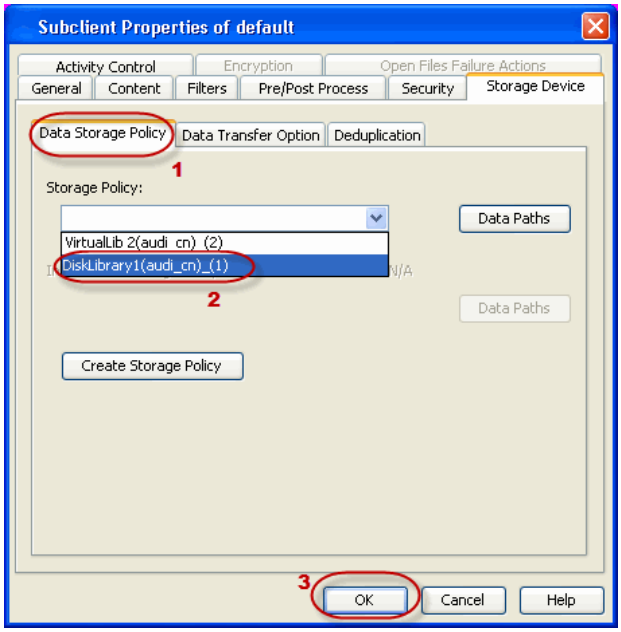

Configuring subclient agent (NSS/NSS-HA nodes)

To configure Subclient Agent:

1. From the CommCell Browser, navigate to **<*Client*> | File Archiver | Default Archive Set**.
2. Right-click **Default Archive Set | All Tasks** ; click **New Subclient**.

Figure 35. CommCell Browser

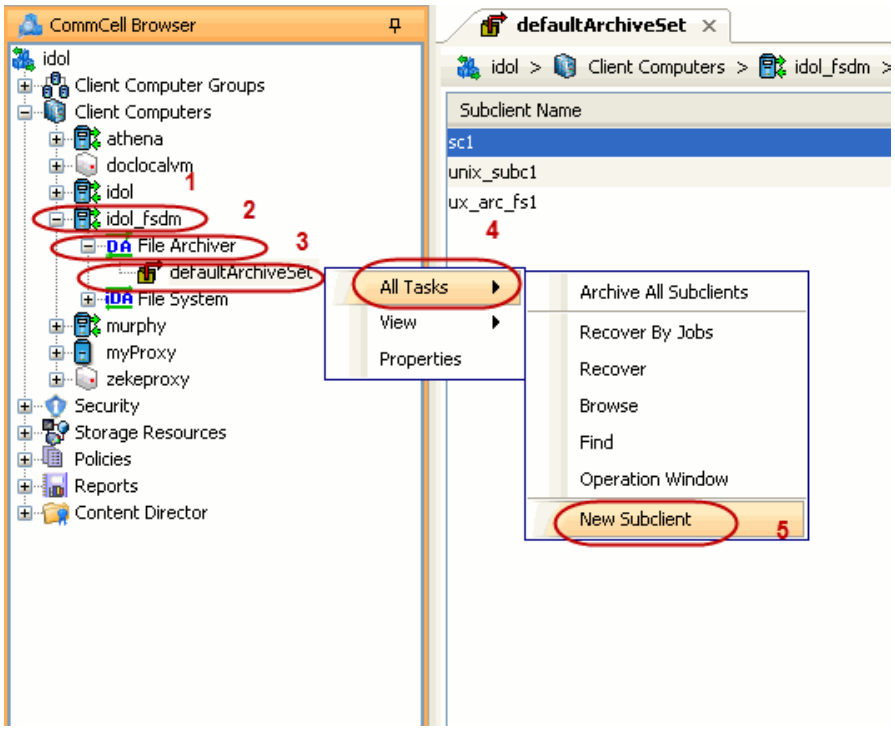

3. Enter **Subclient name**.

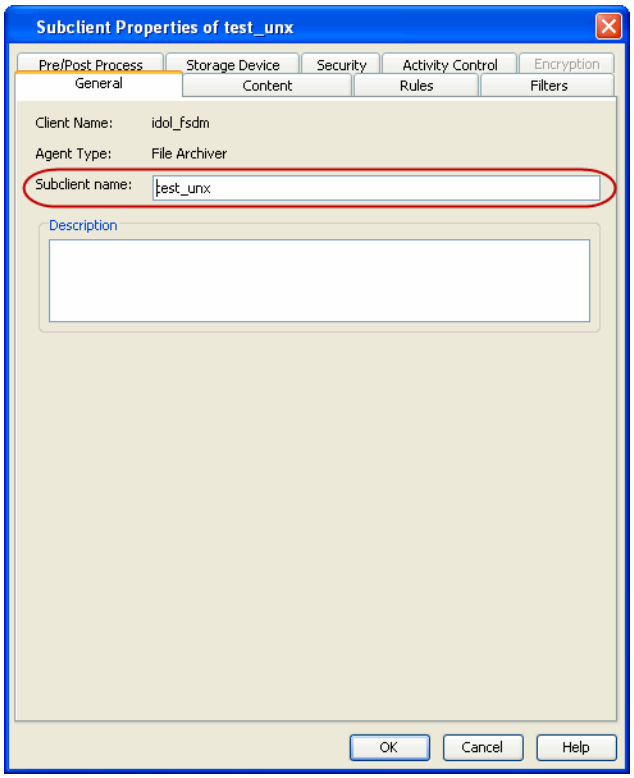

4. Click on the **Content** tab. **Browse** to the folder that contains files you want to archive or move. For NSS node, this is the NSS storage mount point.
5. Select the folder and click **Add.** Click **Close**.

Figure 36. Browsing Content

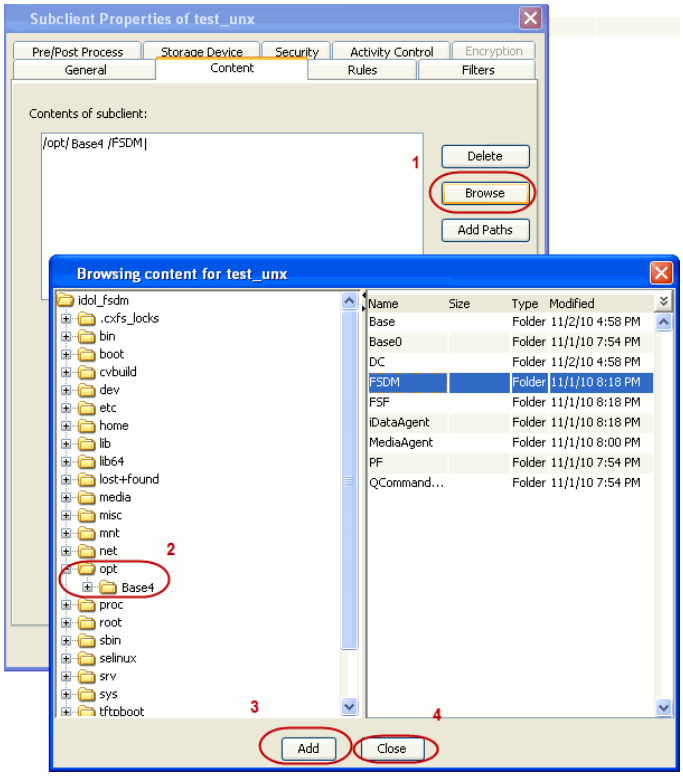

6. Click the **Storage Device** tab. Select the **Storage Policy** from the drop-down list.
7. Click **OK**. If you do not have a Storage Policy created, follow the steps given below to create a Storage Policy.

Figure 37. Storage Policy

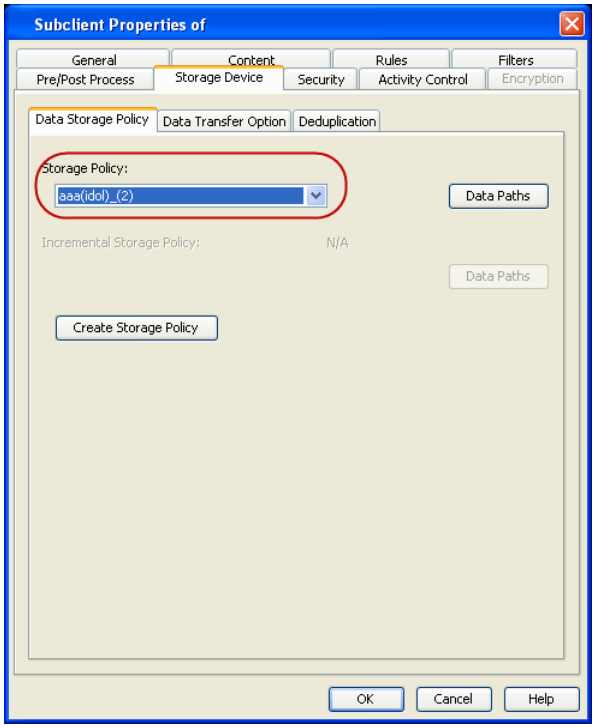

8. If the subclient is to participate in the HSM network, configure it to create stubs, by choosing the option, **Replace the File with a Stub**, under the **Rules** tab. Files that qualify the archiving rule will get stubbed.

Figure 38. Stub Management

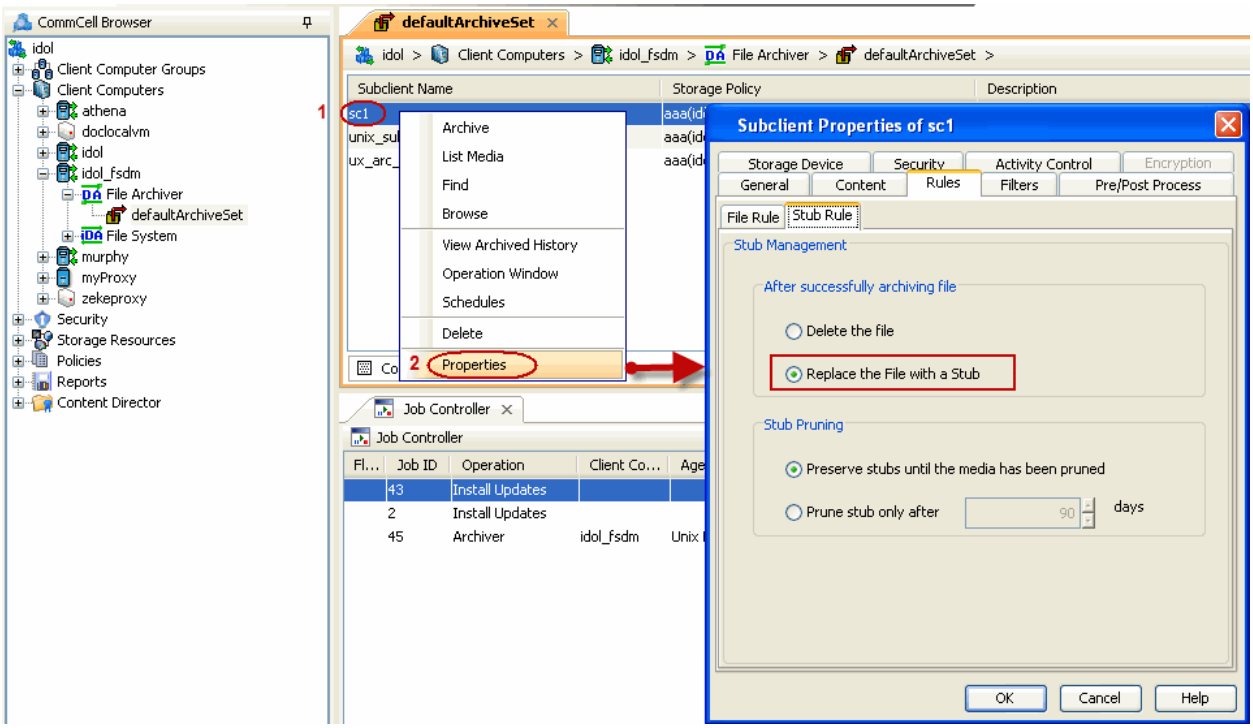

Creating storage policies manually

To create storage policies manually:

1. Click **Create Storage Policy**.
2. Follow the prompts displayed in the Storage Policy Wizard. The required options are as follows:
   - Select the Storage Policy type as Data Protection and Archiving.
   - Select No to allow legal hold to use the policy.
   - Specify the name of the Storage Policy.
   - Specify the name of the primary copy. The primary copy is automatically created along with the Storage Policy.
   - Select No for use of an existing global deduplication policy.
   - Specify name of the default library to which the Primary Copy should be associated. This must be a disk library. Make sure that you select a library attached to a MediaAgent operating in the current release.
   - Select the MediaAgent.
   - Verify the device streams and the retention criteria information.
   - Select if you want to enable deduplication for the primary copy. Also, select if you want to enable Client Side Deduplication option.
   - Name of the Deduplication Store, MediaAgent hosting the Deduplication Store, and the location of the Deduplication Store if you selected above.
3. Review the details and click **Finish** to create the Storage Policy. The primary copy is also created.

Figure 39. Create Storage Policy Wizard

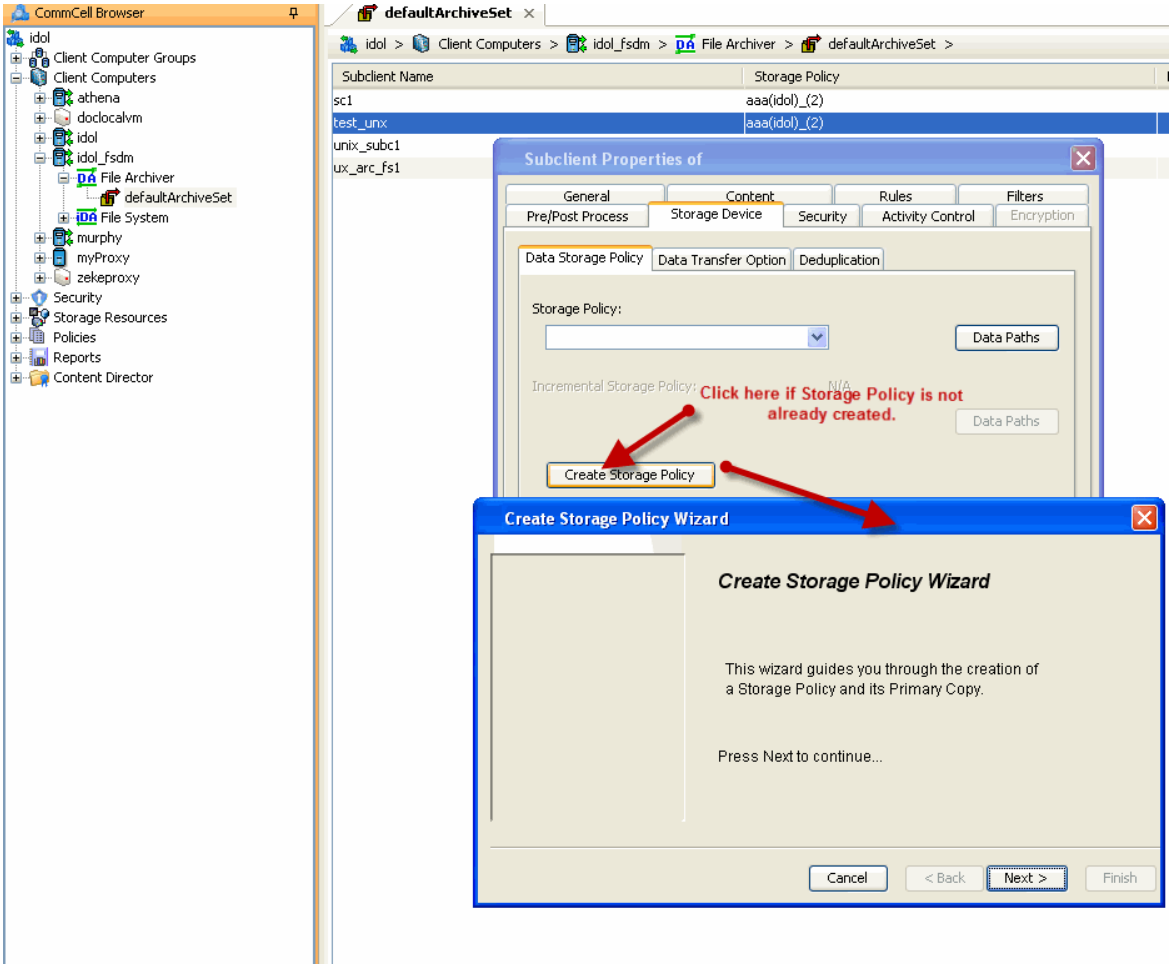

Testing install and data migration

To test install and data migration:

1. From the CommCell Console, navigate to **Client Computer | File Archiver | defaultArchiveSet**.

2. Right-click the **Subclient** and then click **Archive**.

Figure 40. defaultArchiveSet

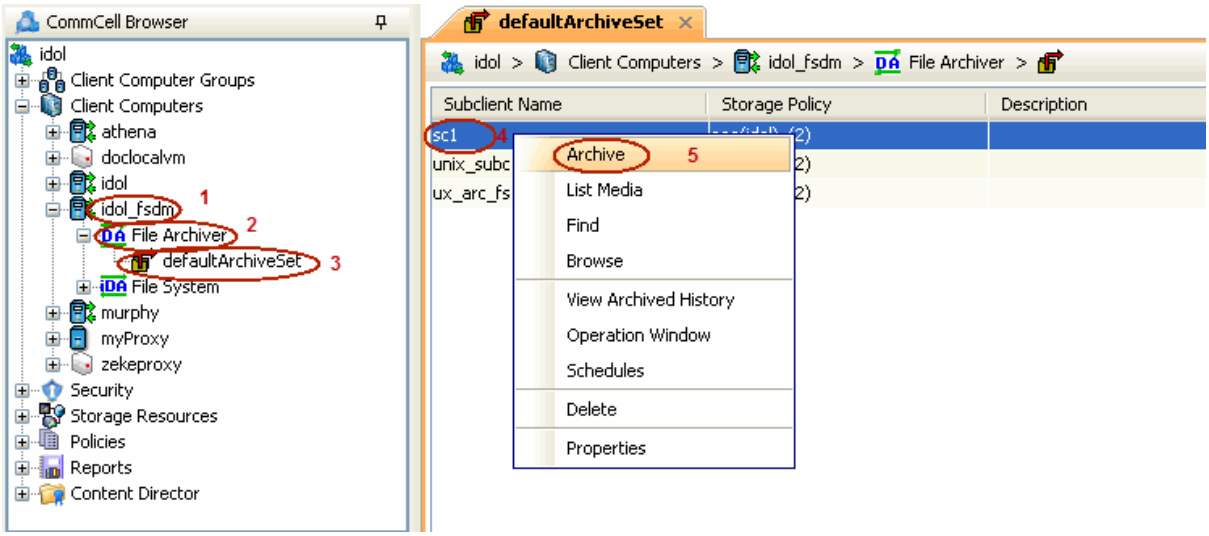

3. Select **Immediate** to run the job immediately and then click **OK**. You can track the progress of the job from the **Job Controller** or **Event Viewer** window of the CommCell console.

Figure 41. Archive Options for Subclient

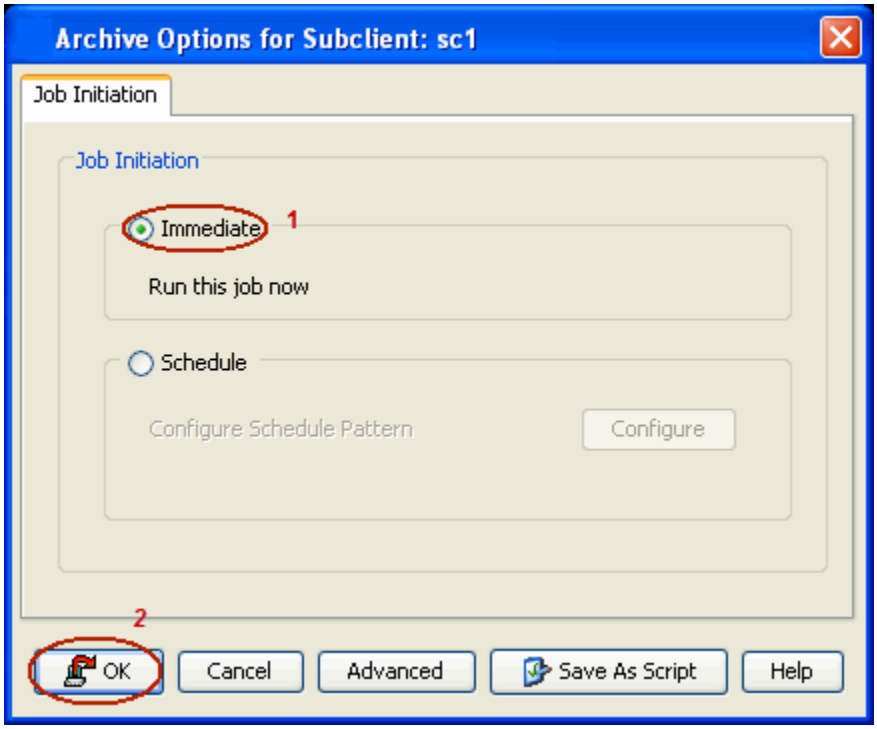

Figure 42. Job Controller

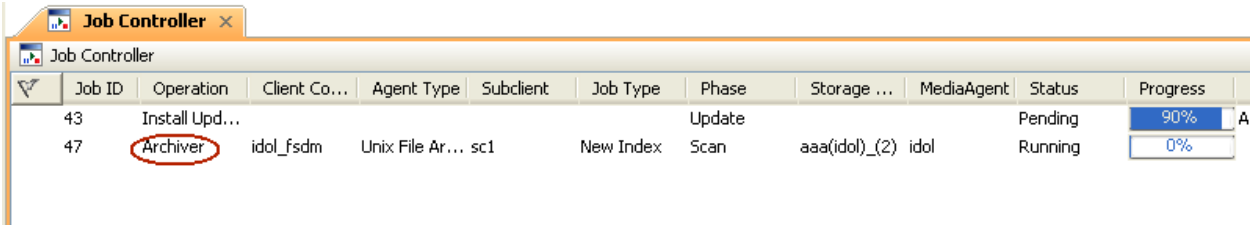

4. Once the archiving or migrate process is completed, files that meet the stubbing rules are stubbed. **Stubs** are placeholders of the original data after it has been migrated to the secondary storage tier. Stubs replace the original files in the location selected by the user. After stubbing, the size of the files on the disk reduces to a single block.

Figure 43. Files before stubbing

```
bose:/ext3 # df -h
Filesystem            Size  Used Avail Use% Mounted on
/dev/mapper/system-root
                      9.5G  5.8G  3.8G  61% /
udev                  126M  112K  126M   1% /dev
/dev/sda1              69M  8.9M   56M  14% /boot
/dev/mapper/savigrp-fsdm
                      2.0G   92M  2.0G   5% /fsdm
/dev/mapper/savigrp-ext3
                      485M  485M     0 100% /ext3
/dev/mapper/savigrp-cdrcache
                      4.0G  567M  3.2G  15% /cdrcache
/dev/mapper/savigrp-cdrcache
                      4.0G  567M  3.2G  15% /cdrdest
mesons:/vol/vol3/cvbuild
                      922G  767G  156G  84% /cvbuild
mesons:/vol/Qupdates_readonly/Qupdates-readonly
                      1.4T  931G  503G  65% /updates
oemdepot.commvault.com:/CV9OSP1
                      3.4G  3.4G     0 100% /90
oemdepot.commvault.com:/9OSP1
                      250G  151G  100G  61% /9OSP1
bose:/ext3 #
```

Figure 44. Files after stubbing

```
bose:/opt/simpana/Updates # df -h
Filesystem            Size  Used Avail Use% Mounted on
/dev/mapper/system-root
                      9.5G  3.0G  6.5G  32% /
udev                  126M  112K  126M   1% /dev
/dev/sda1              69M  8.9M   56M  14% /boot
/dev/mapper/savigrp-fsdm
                      2.0G   92M  2.0G   5% /fsdm
/dev/mapper/savigrp-ext3
                      485M  485M     0 100% /ext3
/dev/mapper/savigrp-cdrcache
                      4.0G  567M  3.2G  15% /cdrcache
/dev/mapper/savigrp-cdrcache
                      4.0G  567M  3.2G  15% /cdrdest
mesons:/vol/vol3/cvbuild
                      922G  767G  156G  84% /cvbuild
mesons:/vol/Qupdates_readonly/Qupdates-readonly
                      1.4T  931G  503G  65% /updates
oemdepot.commvault.com:/CV9OSP1
                      3.4G  3.4G     0 100% /90
oemdepot.commvault.com:/9OSP1
                      250G  151G  100G  61% /9OSP1
bose:/opt/simpana/Updates #
```

Testing data persistent recovery

To test data persistent recovery:

1. Open a Terminal or Console window.
2. Perform an action that executes an open and read on an archived file to initiate the recovery operation. For example, you can use the **vi** or **cat** commands to do this:

   vi *<path><filename>*

   cat *<path><filename>*

Figure 45. Console Window

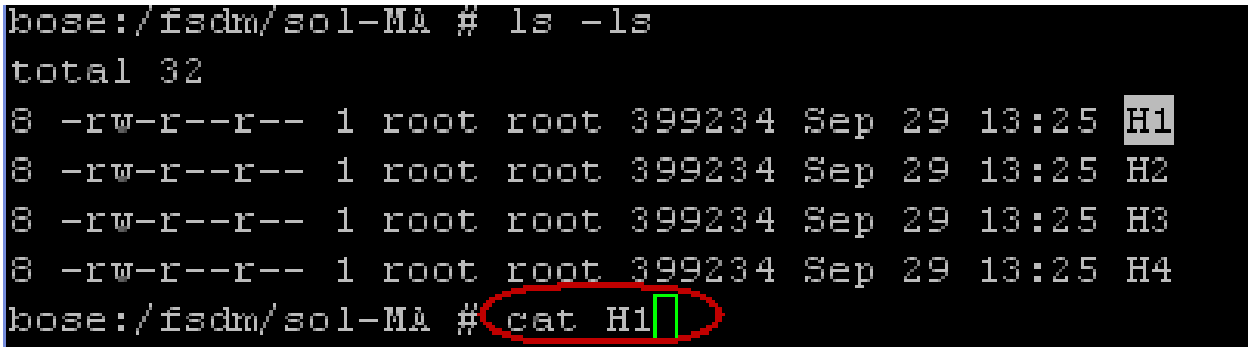

You can track the progress of the job from the **Job Controller** or **Event Viewer** window of the CommCell console.

Multiple stub recoveries (group or directories of files) are submitted to the Job Controller as one job called a *Persistent Recovery* job.

Figure 46. Job Controller

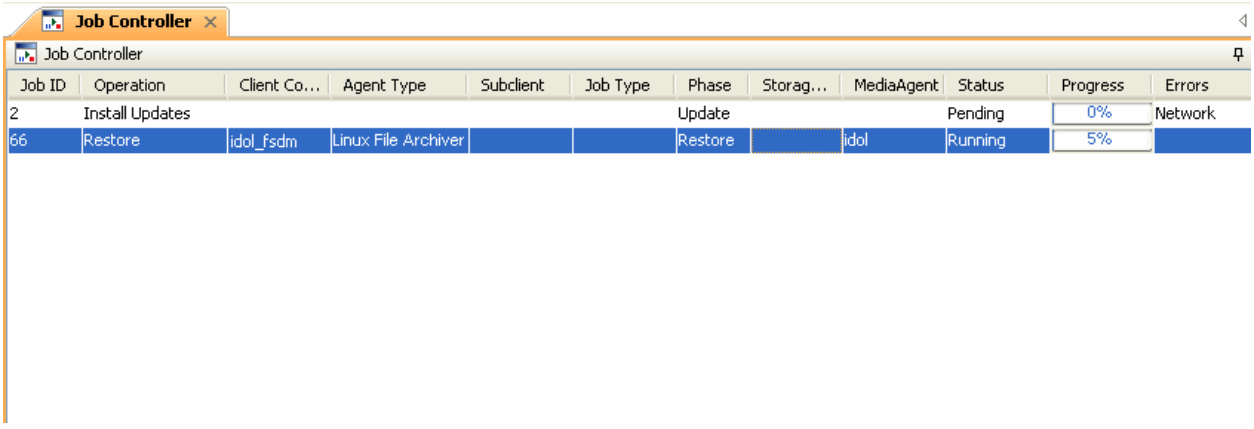

After recovery, stubs are replaced by files in the location specified.

Figure 47. Stubbed files

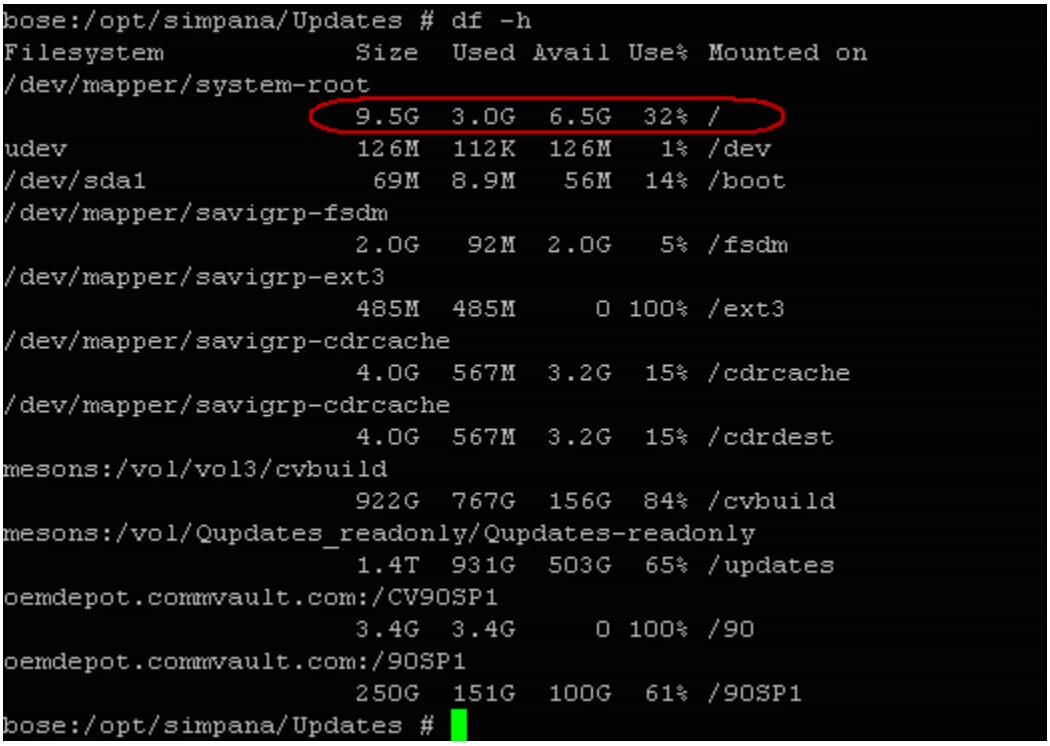

Figure 48. Recovered files

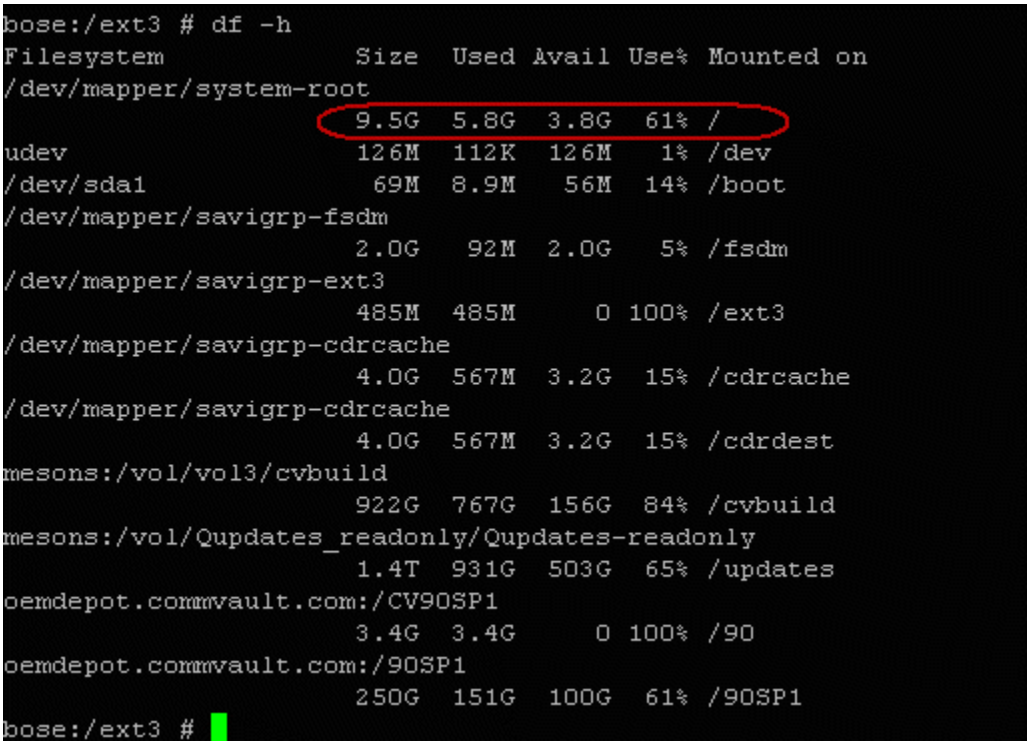

Appendix B: Tools used to test and generate data sets

IOzone

The IOzone tool was used to generate data sets used in verification of user case scenarios. It was also used to establish a baseline with the NSS/NSS-HA configurations. It measure sequential read and write throughput (MB/sec) as well as random read and write I/O operations per second (IOPS). You can download IOzone http://www.iozone.org. Version 3.353 was used for these tests and installed on both the NFS servers and compute nodes used.

dd

The Linux dd utility was used to generate data sets used in verification of user case scenarios as well as configuration baseline. DD is part of the coreutils package and the version native to RHEL6.1, dd (coreutils) 5.97, was used.