

Storage Consolidation and Archive Solutions Using DiskXtender and NearStore

TR 3177 by Legato Systems and Network Appliance, Inc.
[07/23/2002]

1. Audience
2. The Purpose of This Document
3. Solution Overview
4. NetApp NearStore Appliance
5. DiskXtender 2000
 - 5.1. Overview
 - 5.2. How DX2000 Integrates with Network Appliance
 - 5.3. Typical Deployment of DX2000 and NetApp NearStore Appliances
 - 5.4. Interoperability with NetApp SnapMirror Feature
6. DiskXtender UNIX/LINUX (DXUL)
 - 6.1. Overview
 - 6.2. DXUL Storage Manager Integration with Network Appliance
 - 6.3. Typical Deployment of DXUL and NetApp NearStore Appliances
7. Conclusion
8. Additional Information

[TR3177]

1. Audience

The information in this paper is intended for system administrators responsible for installing software and maintaining servers, devices, and clients on the network. This information will also be pertinent to administrators of data-generating applications responsible for properly and efficiently accommodating the growing requirements for application storage.

2. The Purpose of This Document

The purpose of this paper is to provide information on joint solutions offered by Legato Systems and Network Appliance that deliver expanded storage solutions to accommodate the growing requirements of storage consolidation and archiving with Network Appliance's NearStore disk storage systems.

The information on the following pages is intended to provide solution design and implementation detail for the system administrator beyond that of marketing collateral but short of the product documentation.

3. Solution Overview

Storage consolidation and archival demands introduce several issues for network administrators and their clients. The primary issue is an inability for the users and administrators to Logically View the distributed storage resources. They are forced to manually navigate through the added storage resources to access their data. As the amount of data in the distributed enterprise explodes, storage administrators face increasing challenges in terms of storage consolidation.

A superior solution to this problem uses Legato Systems DiskXtender and Network Appliance™

NearStore devices. The simple and easily installed solution provides users and applications with transparent access to the added storage systems. This solution solves the following problems:

- Too many network connections: Users and applications continue to utilize their existing drive network mappings, yet the available storage resources are effectively expanded when NearStore is added to the network. DX2000 provides transparent access to the added storage pool so no network configuration changes are required.
- Disorganized stored data: The inability to locate data on the network is minimized as users maintain the same access points. DX2000 and DiskXtender UNIX®/Linux® (DXUL) provide a single logical view, enabling clients and applications to access data through their familiar drive mappings. Data will remain accessible from the same location where they originally saved it. Automated policies relocate the data to the added storage devices, leaving behind a transparent pointer to access the file.
- Inefficient use of direct-attached Storage (DAS): DX2000 and DXUL group multiple NAS devices into a storage pool, allowing the capacity of the devices to be more efficiently utilized. DX2000 and DXUL policy management provides automated control of DAS to control and optimize the location of data resources.

This document provides an overview of the DiskXtender 2000 for Windows NT®/2000 and DiskXtender for UNIX/Linux solutions and how they are deployed to better manage the consolidation and archiving of data to NAS storage. Legato and Network Appliance have collaborated on the development of this solution and have jointly participated in testing its functionality.

4. NetApp NearStore Appliance

NearStore is the storage industry's first centralized rapid recovery, backup, and archive consolidation solution for the enterprise. The NearStore product family provides ready-access data retention and retrieval solutions that are affordable, inherently scalable, and easy to deploy and manage.

The NearStore product family leverages Network Appliance Data ONTAP™ technology and takes full advantage of value-added software from Network Appliance, including SnapMirror®, SnapRestore®, SnapVault™, and DataFabric® Manager, allowing products in the family to scale up to nearly 100TB. The family is designed to serve a new market space for business information that has previously been too expensive to keep on disk. At a low cost per megabyte of storage, NearStore systems are ideal complements to primary storage solutions (such as NetApp filers), making online archives of infrequently accessed data affordable. The large capacity of NearStore solutions also helps storage administrators consolidate management of multiple archive solutions onto far fewer systems, reducing precious administration time and allowing IT departments to focus energies elsewhere.

NetApp NearStore streamlines backups to tape and enables near-instantaneous data recovery. It provides much faster access to data for costs comparable to high-end tape or optical systems. Built with renowned NetApp simplicity and reliability, NearStore offers vast archive capacity with minimal ongoing management cost. NearStore is an especially good solution for archiving e-mail, medical and satellite images, MCAD drawings, integrated chip designs, home directories, etc. When it's combined with DX2000 it allows you to consolidate all your distributed data to a centralized storage system to reduce management cost and increase access performance.

5. DiskXtender 2000

5.1. Overview

DiskXtender 2000 (DX2000) is the premier solution for automated storage management and data access for enterprise data. Supporting many storage architectures that leverage a variety of storage media, DX2000 has been designed to provide Windows 2000 and Windows NT servers with:

- Infinite storage capability
- Transparent access to stored information
- Automated data management
- Data storage consolidation
- Logical view of all stored data

5.2. How DX2000 Integrates with Network Appliance

The distributed architecture of DX2000 enables the deployment of DX Data Managers to an unlimited number of Windows 2000 or Windows NT servers. DX Data Managers can be configured to utilize one or more supported media services.

In this case, the network-attached storage media service would be defined in the DX Administrator, enabling the storage of one or more NetApp NearStore appliances to be accessible from an extended drive on their Windows 2000 or Windows NT servers.

DX2000 manages data residing on Windows 2000 and Windows NT NTFS partitions/file systems. The architecture of the Legato Systems and Network Appliance solution, exploits the capability of DX2000 to migrate data to NearStore. A Network Appliance NearStore presents a CIFS connection to Windows-based clients allowing DX2000 to transparently copy data from the managed file system into the NearStore. Policies defined in the DX2000 by the administrator control which files are migrated or copied to NearStore devices.

From the creation of the file to the migration of the data to NearStore, to the reduction of a file to a zero-byte pointer and optional deletion of files from the managed server, DX2000 policies enable the complete life cycle management of data residing on direct-attached server storage. The policies that can be defined for an installation of DX2000 include move, purge, and deletion rules. A move rule defines when a file is copied from primary storage to NearStore. A purge rule defines when to truncate the copied file to reclaim primary storage space, but keep it accessible from NearStore storage. A delete rule removes the file entirely from both primary and NearStore appliances.


All rules utilize the same criteria for defining files that will be candidates for action. The attributes considered are:

- Directory name
- File name
- File extension
- File size
- File attributes (hidden, system, compressed, read-only, archive)
- File age based on create time, modified time, or last access time

5.3. Typical Deployment of DX2000 and NetApp NearStore Appliances

A simple implementation of DX2000 is illustrated in the following diagram. By installing DX2000 on the NT/2000 servers, data resident on the NT/2000 servers can be copied/archived to the single NearStore appliance connected over a LAN or WAN.

Windows Data Migration to NearStore



Windows Server

DiskXtender 2000 installed on each server. Data migrates to NearStore based on administrator-defined policies.

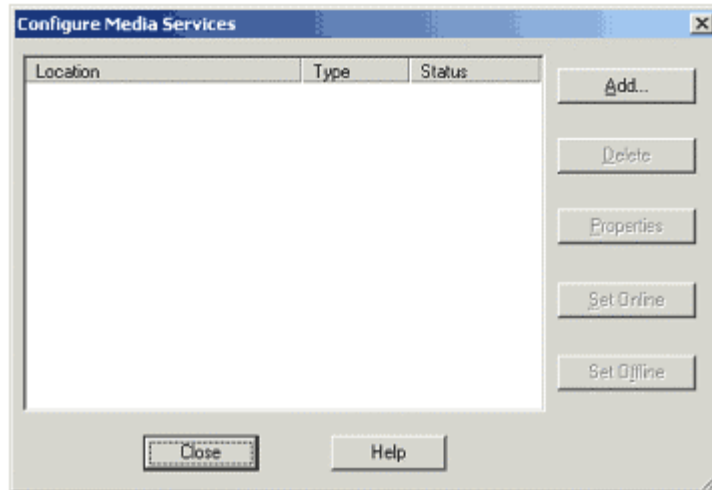
NearStore

Data Migration Flow

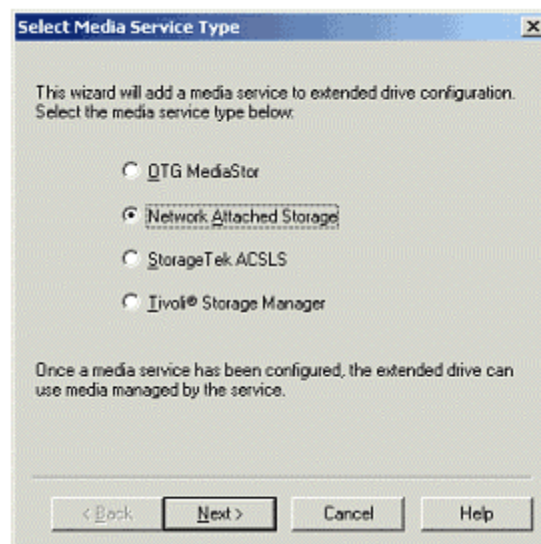
When DX2000 is utilized in conjunction with a NetApp NearStore appliance, the NearStore appliance now becomes a shared resource to the NT/2000 servers. As the storage needs of these servers grow, the space is dynamically made available from the NearStore appliance. If more NT/2000 servers require use of the NearStore capacity, they can be easily configured with the DX2000 software. Likewise, as the NearStore appliance reaches its capacity, additional NearStore appliances can be added to the DX2000 storage pool easily and transparently to the users and applications.

Following these simple steps will configure DX2000 to work with the NearStore appliance.

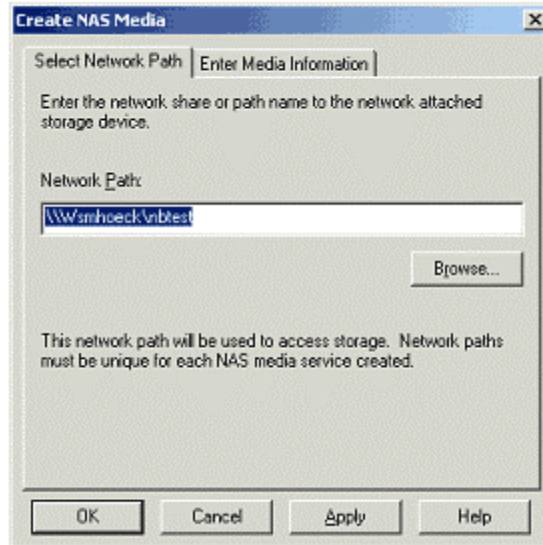
1. Create a share on the NetApp NearStore appliance for each installation of DX2000 to access. Each share should be completely independent of other shares used by DX2000 or any other applications. DX2000 should be the only application accessing the share and any subfolders. Assign permissions to the share allowing the service account running DX2000 to have full read/write/change capabilities.
2. Install DX2000 on the server where data is to be managed. The install can be performed on the machine, or you can utilize our ability to deploy our software to other servers from the single install.
3. Start the DX Administrator to configure the server.
4. Select the Service drop-down menu and select Configure Media Services, displaying the following:



5. Click Add to select the appropriate Media Service. Select Network-Attached Storage and click Next, then Finish.

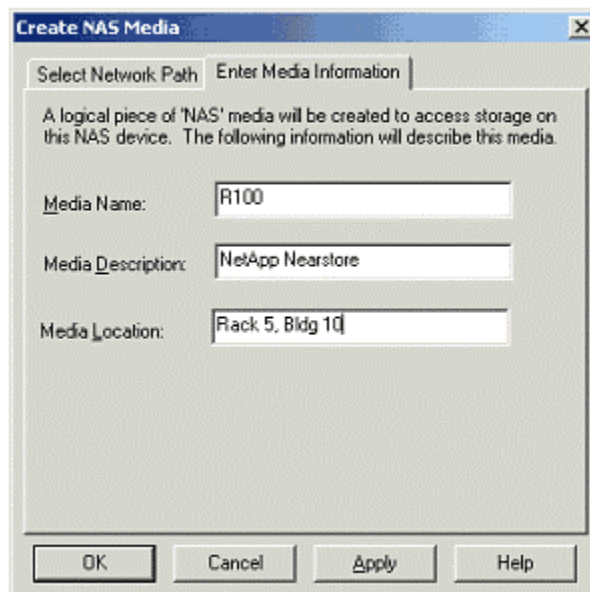


6. Once created, the NAS Media Service must be edited to add the media (i.e., NearStore appliance) that will be used. Click Properties to display the media. Click Add and the following dialog will display.



Enter the CIFS/UNC path that will be used as the target for copying/archiving data by DX2000. DX2000 will ensure that appropriate permissions are set on the Network Path entered. The Browse button can be used to access Network Neighborhood.

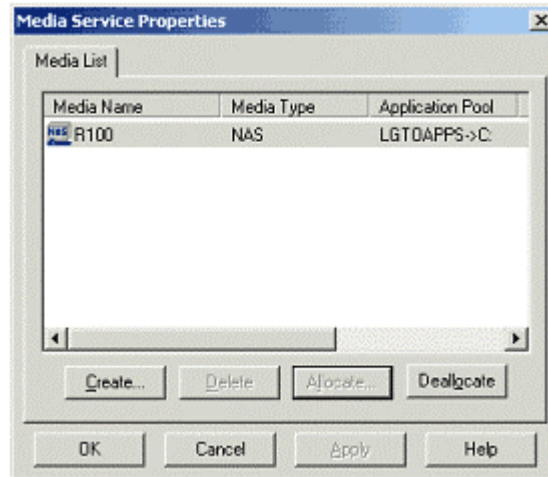
7. Click the Enter Media Information tab to complete the necessary information to define the media.



The Media Name will be used in the DX2000 Administrator GUI and reports. Media

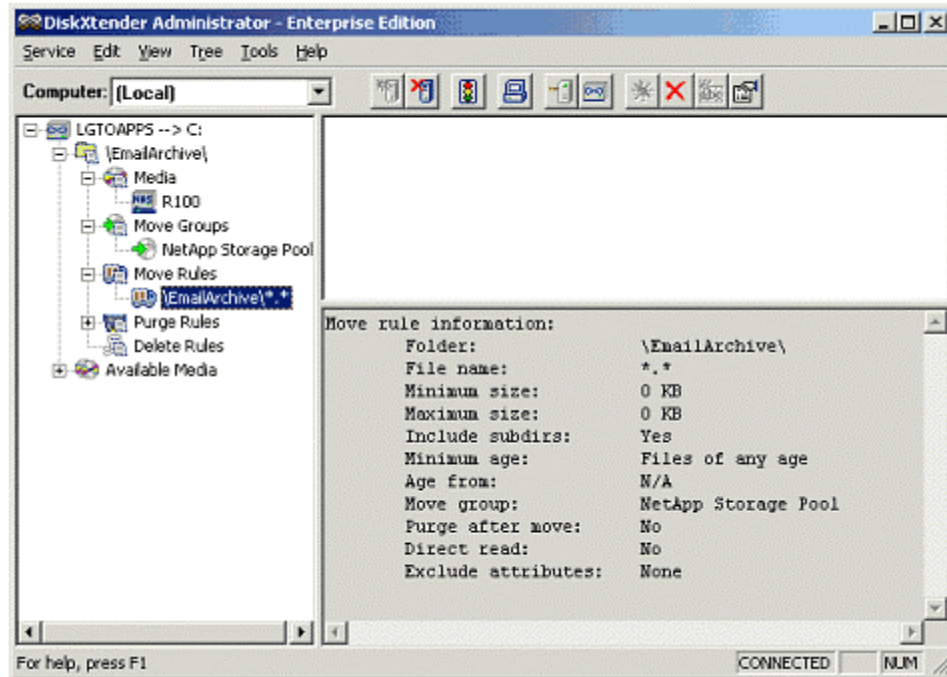
Description and Media Location are additional parameters that will display in the properties of the media.

8. Repeat the media creation process for all NearStore appliances that will be used with this installation of DX2000.
9. When you exit from Configure Media Services, DX2000 will prompt you to create the Extended Drive. While defining the Extended Drive the media defined in the Configure Media Services dialog will be displayed for you to allocate the media to the Extended Drive. You can perform this process at this time or return to Configure Media Services to Allocate the Media.



10. Now that the media has been allocated to the Extended Drive, all remaining steps to configure the system to migrate data to the NearStore appliance can be completed. The following steps describe how to create a move rule. The steps involved to create a purge or deletion rule are very similar.
 - o Create a media folder
 - o Assign the NearStore appliance media to the media folder
 - o Create a Move Group for the NearStore appliance media
 - o Create a Move Rule, which is set to migrate the data to the Move Group just created

The result of these steps is illustrated in this DX Administrator screen.



5.4. Interoperability with NetApp SnapMirror Feature

Enhanced configurations of this solution can leverage the capabilities of the SnapMirror feature provided by Network Appliance. Using SnapMirror provides fault tolerance of data that has been migrated from the DX2000 Extended Drive to the NearStore appliance. By mirroring the data, DX2000 can be reconfigured to access the migrated data from the mirrored volume or server in the event a NearStore volume becomes inaccessible.

If a NearStore volume is inaccessible it is necessary to reallocate media within DX2000 so that the CIFS share (DX2000 Media) is now pointing towards the mirrored volume.

To reallocate media use the DX2000 Administrator to perform the following steps:

1. Delete the failed media (i.e., CIFS share connection to failed volume) from the Extended Drive. Performing this function will remove all files from the Extended Drive that had been migrated to that CIFS connection.
2. Delete the hidden OtgXmsNas file from the mirrored NAS Media location. This file is found at the root of the share connection.
3. Create a new piece of NAS media defining the CIFS share to the mirrored location. Then allocate the media to the Extended Drive.
4. Add the new media to the extended drive. When the media is added to the Extended Drive, file pointers will be created for all files that are in the mirrored volume. This will reestablish access to the files that are in the mirrored volume or system location.
5. Add media back to Move Group.

Once this process is completed, access to the files is restored.

Administrator Tip:

It will be important to review the security settings for the files created during the media restore

procedure. When the failed media is removed from the system, all details of the file including security settings are deleted from the Extended Drive. When the new media is added to the Extended Drive, default permission settings are applied to the files. However, the folders that the files resided in will not be deleted and permission settings for the folders on the Extended Drive will not be impacted by this procedure.

6. DiskXtender UNIX/LINUX (DXUL)

6.1. Overview

DiskXtender for UNIX is the premier solution for automated storage management and access for enterprise data. With support for leading hardware and storage media, including Network Appliance NearStore, it provides:

- Support for industry-leading UNIX operating systems
- Virtual storage pool for heterogeneous storage devices
- Infinite storage capability
- Transparent access to stored information for users and enterprise applications
- Data protection and disaster recovery
- Massive/enterprise scalability

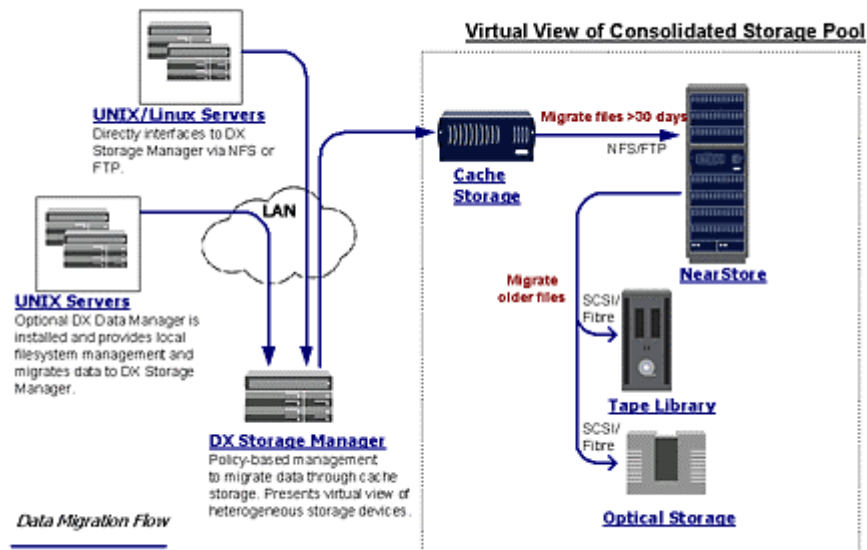
6.2. DXUL Storage Manager Integration with Network Appliance

The foundation of the DiskXtender UNIX/Linux is called the DX Storage Manager (DXUL-SM). It provides an extensive set of features to handle your UNIX data management needs. Through its implementation of a virtualized file system, DXUL-SM is able to manage and consolidate many storage resources into a single point of access for your applications and users.

Deployed in the application/user space on a UNIX server, DXUL-SM exports its own managed file system through an NFS or FTP connection. A single DXUL-SM file system is designed to effectively maintain a massive quantity of files, in many cases surpassing the limitations of native UNIX file systems. This file system supports all UNIX security concepts of Users and Group definitions.

Files written to this file system are initially staged in disk cache (e.g., a NetApp filer) prior to being migrated to the configurable archive storage. To provide an online access to archive data, DXUL-SM has been certified to migrate/archive data to NetApp NearStore and filer appliances. Since DXUL-SM supports the creation of 15 copies of a file, configuration settings can be defined to write the copies to any one or combination of archive media. An illustration of this architecture is provided in the following figure.

UNIX/Linux Data Migration



As shown in the illustration, any network client can access the DXUL-SM file system using the exported FTP or NFS connection points available. In addition, DXUL Data Managers, located to the left of the clients in the figure above, can be configured to migrate data to the centralized storage pool.

The DiskXtender UNIX/Linux Data Manager (DXUL-DM) is available to customers as an add-on component to the DXUL-SM. DXUL-DM enables management of native file systems running on Sun™ Solaris™ and Compaq Tru64. Combining the two components leverages the performance characteristics of the native file system with the enhanced data management capabilities of DXUL-SM.

Utilizing an implementation of DMAPI, DXDM provides the ability to monitor file system activities, migrate data to/from the DXUL-SM, and purge files to reclaim disk space. Automated policies defined for the native file system control migration of data to the DXUL-SM based on high/low watermarks, time of last access, and size of the file.

6.3. Typical Deployment of DXUL and NetApp NearStore Appliances

The NearStore storage capacity is made available over NFS to an instance of DXUL Storage Manager. As displayed in the previous figure, several instances of DXUL-SM can share the same NearStore appliance for data migration.

Following these simple steps will configure DXUL-SM to work with the NearStore appliance.

1. Create container files on the NearStore mountpoint. Any utility (e.g., mkfile, dd, etc.) can be used for this operation. It is suggested to make these containers 10GB to 100GB in size. The appropriate number of container files to create can be determined by dividing the available storage capacity of the NearStore by the size of the container files.
2. Run 'tapelabel' on each of the container files from step 1 to write an internal volume label. Use "FileSystem" media type. For example,

3. `tapelabel - m<mediatype> <LTSVOLUMEPATHS>/000000 000000`

<mediatype> is 19 for Solaris, 44 for HP-UX, 69 for SGI, 94 for Tru64UNIX, and 119 for Linux.

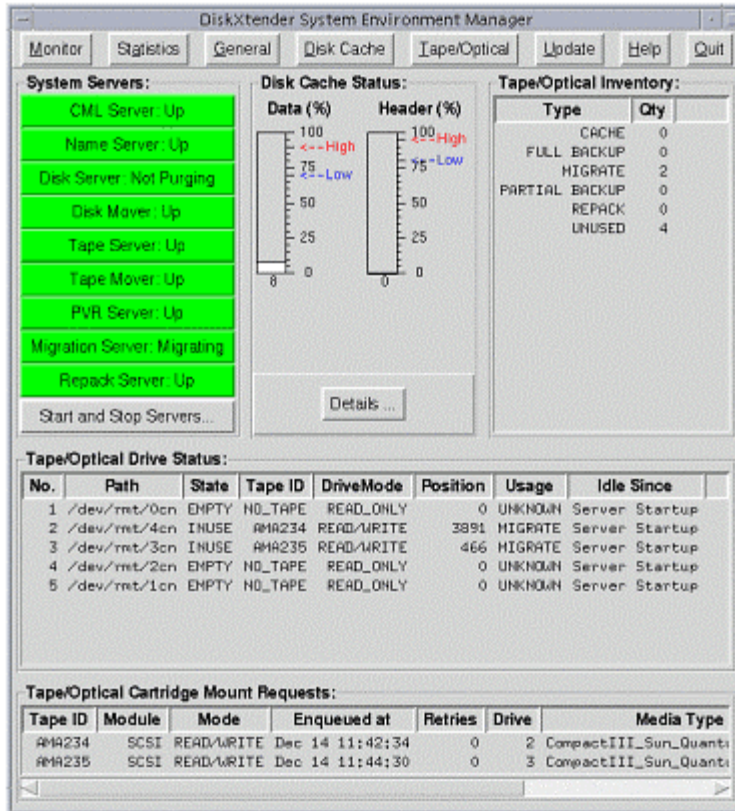
4. Configure the Logical Tape System (LTS) interface to provide mounting and dismounting service for the storage volumes.

Set up 64 NAS-type drives in `tpdevs`, enter a fully set LTS bit map line into `PVRINFO`, copy the `lts_mntpgm.NAS` and `lts_dismnt.NAS` from the `cs` directory into the `adm/bin` directory, and enter the `LTSVOLUMEPATH` line into the `DiskXtender.path` file. Add to the `DiskXtender.conf` file the following parameters: `tapeserver.unload_wait_time=5`
`pvrserver.lts_debug_flag=1`

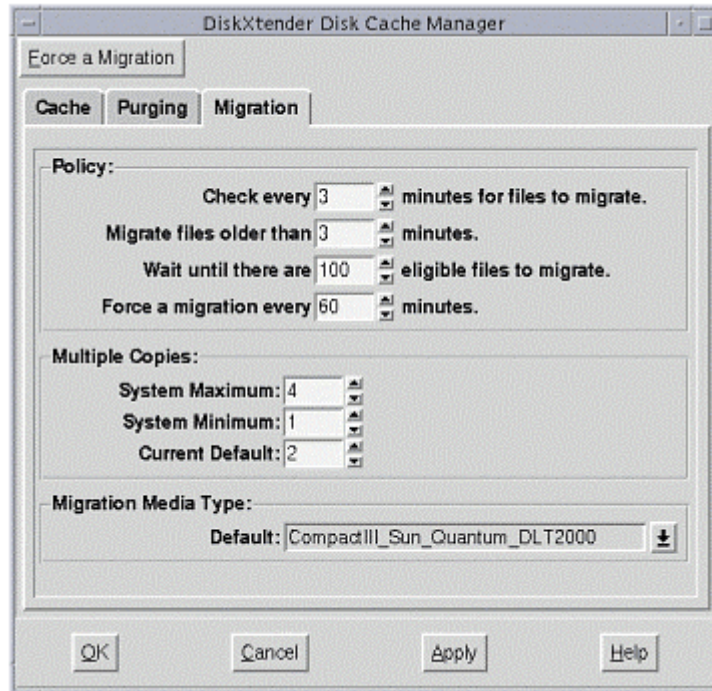
5. Make the storage volumes (labeled container files) known to the DXUL-SM tape inventory database by running 'settape' for each of the storage volumes created above. The tape server must be running for this command to run. For example,

```
settape -m <mediatype> -u 0 000000
```

Now that the configuration of using NearStore with DXUL-SM is complete, it is simple to apply migration policies by editing the cache properties. These properties can be edited using the DXUL-SM Graphical Interface. To start the DXUL-SM GUI, run the command: `dxsem` from `/usr/diskx` directory. This will present the following screen:



Click on the Details button to display settings for the disk cache, which include the ability to set migration and purge settings.



Choose the applicable tab to define and set policies for the disk cache, purging, and migration.

7. Conclusion

Through a shared vision, Legato Systems and Network Appliance have developed an integrated software and hardware solution that will optimize the process of storage consolidation and data archival. By using this solution, our joint customers can be confident that the storage technologies they're deploying deliver significant value while providing high-performance data protection, simplified storage management, and consistent high availability of critical information. Our technology partnership provides your IT organization with a robust storage consolidation and business continuance formula for the long term.

8. Additional Information

For additional information and documentation on the Legato DiskXtender Family, please visit:

<http://www.legato.com/products/diskxtender/>

For more information on the NearStore architecture, please visit:

http://www.netapp.com/tech_library/3187.html