

EMC DL1500 and DL3000 with EMC NetWorker

Best Practices Planning

Abstract

This white paper contains a compilation of specific configuration and best practices information for the EMC[®] DL1500 and DL3000 with NetWorker[®].

July 2009

Copyright © 2008, 2009 EMC Corporation. All rights reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com

All other trademarks used herein are the property of their respective owners.

Part Number h5713.4

Table of Contents

Executive summary	5
Introduction	5
Audience	5
Terminology	5
DL overview.....	7
The deduplication process	8
Key factors affecting the deduplication ratio	8
Sizing	10
Space management.....	10
Native format data	11
Deduplicated data	11
Phasing deduplication into the environment.....	11
General NetWorker settings and considerations	12
SAN vs. NAS environments	12
Using the DL with NetWorker in a SAN environment	14
DL SAN attach	14
Device discovery (persistent binding)	14
Device drivers.....	15
Scanning devices	15
Configuring the backup server	15
Simultaneous backups.....	16
Single library – multiple hosts	16
Virtual tape libraries	17
Emulations.....	17
Number of slots	17
Assigning and unassigning virtual libraries	17
Virtual tape cartridges	17
Reclaiming space.....	17
Coping with previously multiplexed data.....	18
Using the DL with NetWorker in a NAS environment	20
General recommendations.....	20
Network connectivity	20
Backup-to-disk configuration.....	22
Mounting NFS shares	23
AIX 5.2.....	23
HP-UX 11i.....	23
Linux	24
Solaris.....	24
Configuring CIFS shares.....	24
Moving data from emulated tape to physical tape	26
Backup application specific Path-to-Tape.....	26
NetWorker site requirements for PTT	27
Auto Archive.....	27
Auto Archive site requirements	28

NetWorker cloning	28
Replicating deduplicated data	29
Replication considerations	29
Seeding (pre-populating the target system).....	31
Recovery/Failback	32
Conclusion	32
References	32
Appendix: Using replication with NetWorker	33
DL setup.....	33
NetWorker setup	33
Authentication setup	33
Create the keys	34
Install the public key on the DL	34
Test the keys	34
NetWorker post-backup script example.....	34
Recovery from a replicated NAS share	36

Executive summary

The EMC® DL1500 and DL3000 appliances provide simple and reliable disk-based backup and recovery systems with the added feature of deduplication. The DL products with deduplication integrate seamlessly into your current backup environment to provide leading-edge backup and restore operations. The DL was designed and implemented for optimal performance and ease of use, and easily scales to meet your storage needs.

Introduction

This white paper provides specific configuration and best practices recommendations for the DL1500 and DL3000 appliances when used with EMC NetWorker® in NAS and SAN environments.

Audience

This white paper is intended for EMC customers, system engineers, and members of the EMC and partners professional services community who are interested in configuration and best practices information when using the DL appliances with NetWorker.

Terminology

- **Backup-to-disk (B2D)** – A backup solution where data is written to hard disk instead of tape.
- **Balance-rr mode** – Method for controlling the bundling of several physical Ethernet ports together to form a single logical channel. This mode is also referred to as Mode 0 (see also *Link Aggregation Control Protocol*).
- **Cartridge-level replication** – Process of verifying that the contents of an individual cartridge in a source VTL is present in a target VTL and the sending of the metadata description of the source cartridge.
- **CIFS** – Protocol for a Windows-based network fileshare.
- **Cloned storage node** – A computer with NetWorker software installed whose primary purpose is to make copies of save sets,
- **Data Pool** – Deduplicated data stored on the DL back-end array.
- **Data Protection Advisor (DPA)** – An EMC data protection management solution that provides automatic and continuous data collection, conditional analysis that triggers alerts, and a single, consistent interface for reporting.
- **Deduplication** – Process of detecting and identifying the redundant variable-length blocks (or data segments) within a given set of data to eliminate redundancy.
- **Directory/file-level replication** – Process of verifying that the contents of a directory or an individual file in a NAS share on the source system is present in a NAS share on the target system and the sending of the metadata description of the source directory or file.
- **Failback operation** - The replicated NAS share or VTL is transferred from the target system to the source system.
- **Immediate deduplication** – Process of deduplicating the backup data stream as it is ingested by the DL appliance (see also *scheduled deduplication*).
- **Ingest** – Process of receiving backup data from the backup server.
- **Link Aggregation Control Protocol (LACP)** – Method for controlling the bundling of several physical Ethernet ports together to form a single logical channel. LACP allows a network device to negotiate an automatic bundling of links by sending LACP packets to the peer (a directly connected device that also implements LACP). This mode is also referred to as Mode 4 (see also *Balance-rr mode*).
- **Logical Unit Number (LUN)** – Identifying number of a SCSI object that processes SCSI commands.

-
- **Multiplexing** – The interleaving of data on a target, required to make concurrency/parallelism technically realizable on anything but an Advanced File Type Device (AFTD). See *parallelism*.
 - **Locked** – State in which a target NAS share or VTL is unavailable for updating with the latest replicated data. Any data replicated to a directory or file within a locked share or to a cartridge within a locked VTL will be placed in an update queue on the target system. (see also *unlocked*)
 - **Metadata description** – Sequence of unique blocks and pointers that represent the original data. The receipt of this metadata description by the target system (and that data that the metadata describes) results in the generation of a recoverable image.
 - **Namespace replication** – Process of verifying that all data associated with the source is present at the target followed by replication of the metadata description of the source.
 - **NAS** – Network-attached storage.
 - **Native format data** – Uncompressed backup data that has not been deduplicated.
 - **NFS** – Protocol for cross-platform network fileshares.
 - **Parallelism** – The generic term for concurrency in NetWorker, from the name of the client, server, and group settings controlling their respective concurrency. See *multiplexing*.
 - **Partial replication** – The recovery point representing that a VTL or NAS share or readable image representing a cartridge or directory/file on the target system is missing data present in its source equivalent.
 - **Policy-based deduplication** – The choice of immediate, scheduled, or no deduplication.
 - **Recovery operation** - The replicated NAS share or VTL (recovery point) is made available for use on the target DL system.
 - **Recovery point** – A replicated NAS share or VTL on the target system whose metadata description has also been replicated. A replicated virtual tape or directory/file on the target system.
 - **Remote replication** – Backup data residing on a DL is copied over a LAN or WAN to another DL for disaster recovery protection.
 - **Save set** – Group of files or a file system that has been backed up on storage media.
 - **Scheduled deduplication** – Deduplication of data occurs after the backup stream has been ingested by the DL (see also *immediate deduplication*).
 - **Shoeshining** – An interruption in the data stream to physical tape drives that requires the tape to be repositioned on the drive's head in a manner that causes a stop-start or shoeshine motion of the tape device. This behavior is due to an inadequate data flow to the tape drive and slows overall performance.
 - **Source-based deduplication** – Deduplication of backup data occurs at the client before it is sent to the primary storage system.
 - **Staging** – The movement (or migration) of a save set vs. the copying of a save set. Staging is typically associated with the more generic backup-to-disk approach using Advanced File Type Devices in NetWorker – where once a save set is staged, the original copy is deleted.
 - **Storage node** – A computer with NetWorker software installed whose primary purpose is to receive a backup stream from one or more backup clients, and write to save sets on storage media.
 - **Sync ID** – Tag by which a source-target destination is identified in a DL replication configuration for cartridge-level or directory/file-level namespace replication.
 - **Target-based deduplication** – Deduplication of backup data occurs at the target storage system.
 - **Truncation** – DL space management function that removes the oldest blocks of data from storage having a deduplication equivalent when used capacity exceeds 70 percent.
 - **Unlocked** – State in which a target NAS share or VTL is available for updating with the latest replicated data. If a locked NAS share or VTL is unlocked, any directory/file or cartridge with queued replication data will be updated at that time. (see also *locked*)
 - **Virtual tape library (VTL)** – Software emulation of a physical tape library system.

DL overview

The DL1500 and DL3000 are stand-alone backup-to-disk (B2D) appliances offering a network-attached storage (NAS) front end (over IP) and an optional virtual tape library (VTL) front end (over FC). Both configurations of DL include data deduplication to efficiently store data written to the appliance. These appliances can be configured to present an NFS share or CIFS share, in NAS personality, or a virtual tape library, in VTL personality, to the backup storage node. When enabled, both NAS backup to disk (B2D) and VTL functionality can be used simultaneously. With policy-based deduplication, the data can be deduplicated as soon as the backup stream is sent to the appliance or postponed until the backup window has closed.

The DL appliance with its data deduplication capability does the following:

- Eliminates redundant data from backups to reduce storage, enabling longer onsite retention and reduced replication costs
- Performs sub-file comparisons on variable length data blocks to capture small block inserts and overstrikes in unstructured data
- Provides policy-based deduplication that is customer-tunable to adjust to changing backup environments and optimize backup performance
- Understands backup software metadata to optimize the backup process
- Includes built-in data compression that is additive to deduplication in the data reduction process

The DL appliance's replication option leverages the deduplication capability of the DL, substantially reducing the amount of backup data that needs to be migrated to a remote site. DL replication provides rapid local and remote restore with the following benefits:

- Supports replication between any combination of DL1500/3000 appliances
- Permits bi-directional replication between DL appliances
- Replicates deduplicated NAS shares and/or virtual tape libraries to reduce bandwidth
- Checks for the presence of each data block before replication so that only unique data blocks are sent to the target DL to further reduce network traffic
- Supports up to 10 source DLs to one target DL
- Maintains a common data deduplication repository at the target DL for maximum storage savings
- Replication of entire VTLs or NAS shares or individual tape cartridges (for VTL) and files (for NAS shares)
- Automatic recovery on the target system of replicated tape cartridges and files
- Deferment of tape cartridge/file replication during a backup window when a scheduled deduplication is used
- Detailed replication reporting
- Provides optional 128-bit AES data encryption for replicated data

The DL appliance's Path-to-Tape (PTT) option provides two means for copying virtual tapes to physical tape:

- Backup application-specific path-to-tape that maintains backup application catalog consistency of all tape copies
- Auto Archive for the export/import of virtual tape contents with physical tape under the control of the DL Console Manager

The deduplication process

The deduplication process divides the native data into variable length blocks and calculates a signature for each block. For each signature, the deduplication process determines if any other identical signature already exists, and if not, hardware-compresses the unique block and writes it to its data pool. If it finds any matches to these signatures, it replaces that redundant block with a pointer to the unique block so that only one copy of that data block will exist on disk. The metadata description (signatures and pointers) of the data is stored separately from the backup data on a RAID 6 LUN on the disk array in the DL3000 and on a RAID 1 LUN in the DL1500. A metadata description is stored for all data deduplicated within the DL.

As additional data is written to the DL, its data is also broken down into variable length blocks whose signatures are compared against those for existing unique blocks. Only new unique blocks are stored along with pointers to existing unique blocks in the appropriate sequence that can be used to reconstitute the data.

Key factors affecting the deduplication ratio

The ratio of the amount of storage space required to store the total number of save sets in a conventional disk storage system to the storage capacity used in a deduplication system is referred to as the deduplication ratio. There are many factors that affect deduplication ratios. Some key factors include:

Backup policies

Retaining data for longer periods of time improves the chance that common data will already exist in storage, resulting in greater storage savings and better deduplication ratios.

Whether performing backups of similar data according to a daily full or weekly full/incremental model, the amount of data storage required is essentially the same as only unique data is stored using either model. Since the amount of data sent to the DL is substantially greater if doing daily full backups, the deduplication ratio will be considerably higher for a policy of daily fulls rather than a policy of performing a weekly full/incremental model. Of course, the amount of data moved by the client would be much, much higher.

Change rate

When the first data stream is deduplicated, the number of unique blocks within it can vary widely depending on both the data type and backup application. The deduplication process may find little or no data redundancy to 50 percent or more data redundancy.

When multiple backups of the same policy are written to the DL, however, storage savings often increase significantly with each new save set as only those data blocks that are unique to each backup need to be written to disk. In conventional business operations, this unique data may represent only 1-2 percent of the data present in each additional backup set. The remainder of the backup consists of pointers to data already present on the DL.

Data types

Backups of user data such as text documents, PowerPoint presentations, spreadsheets, some database types, source code, and Exchange are known to contain redundant data and are good deduplication candidates.

Other data sources such as audio, video, and scanned images consist of unique data and are not good deduplication candidates unless backed up multiple times.

Compression

Compression does not affect deduplication of files that do not change. However, compression can cause even a small change in a file to ripple throughout the compressed file, greatly reducing the commonality that can be found between current and prior versions, or similar but not identical documents on different machines. As a result, software client-side compression prevents optimal deduplication from taking place since data that is already compressed cannot be efficiently deduplicated once it reaches target storage.

Data that changes frequently should be uncompressed when transmitted to the DL to increase the matches the deduplication engine can find in the data stream. Compressing frequently changing data prior to sending it to the DL will most likely result in negligible redundancy and poor deduplication ratios. The DL, in addition to deduplication, hardware-compresses the deduplicated data when writing to storage.

EMC recommends avoiding use of standard NetWorker directives “Standard Windows with compression” or “UNIX with compression”.

Encryption

Encryption does not affect deduplication of files that do not change between backups. However, encryption will cause even a small change in a file to ripple throughout a file when it is encrypted, greatly reducing the commonality that can be found between current and prior versions, or similar but not identical documents on different machines. Changing the data zone encryption passphrase will also change the encrypted form of every document, preventing deduplication from finding commonality between current and prior versions. As a result, software client-side encryption prevents optimal deduplication from taking place since data that is already encrypted cannot be efficiently deduplicated once it reaches target storage.

However, encryption is an important part of securing a company's data from theft or accidental loss. EMC NetWorker encrypts at the client to distribute the encryption workload and ensure all copies of the backup data are encrypted. NetWorker does not have a mechanism to encrypt just a clone to physical tape sent offsite.

To achieve encryption of backup data, consider the use of an encryption appliance deployed between the DL and the physical tape library.

If encryption is required without a physical encryption appliance, avoid changing the data zone passphrase frequently, so that the current version of an unchanged file is more likely to be identical to the prior version. If poor deduplication ratios are achieved, it may be appropriate to create a dedicated NAS share or VTL with deduplication turned off, to receive those backups.

Multiplexing and parallelism

EMC NetWorker customers typically use parallelism between 4x and 12x, for many reasons:

- To avoid shoeshining on a physical tape drive
- To increase concurrency for a given number of devices
- To reduce complexity and user management effort for a given level of concurrency
- To shorten the backup window by writing save sets from multiple clients to a single drive simultaneously
- To reduce the load on the storage nodes and server

While shoeshining does not apply to a DL, the other reasons do.

When parallelism is applied to a tape or tape-emulating device like a VTL, parallelism often means multiplexing. Multiplexing interleaves backup streams, writing a little of save set 1, then a little of save set 2, and so on, so that none of the clients sending save sets need to wait for the other clients to finish. Unfortunately, this interleaving of save sets has a significant impact on deduplication efficiency when the DL is used as a virtual tape library (VTL). Multiplexed streams hinder the deduplication process from efficiently identifying blocks of common data because of the additional header information added to the data with parallelism. In order to realize the full benefit of deduplication, EMC recommends multiplexing be turned off when using the DL as a VTL with NetWorker by setting target sessions and max sessions to a value of 1.

IMPORTANT: When using the DL as a VTL with NetWorker, ensure that the DL is running version 1.0.1 build 2021-2034 or later. Upgrading the DL version to this level or reducing or eliminating parallelism will enhance the deduplication of subsequent backups. It will not affect any backups that the DL has already processed.

To take advantage of the new code revision or elimination of parallelism on existing backups, they must be cloned to new virtual tapes and the original tapes removed to free space.

Note that use of parallelism is not an issue if the DL is used as a NAS backup-to-disk target (by creating an Advanced File Type Device (AFTD) on CIFS or NFS shares presented by the DL) as each save set is always stored as a separate file regardless of parallelism settings.

IMPORTANT: Backup environments designed for backup to NAS are architected very differently from backup environments designed for physical or virtual tape. Engage EMC Solution Architects in a rearchitecture if you are interested in moving from an environment designed for physical or virtual tape to NAS.

More details and recommendations on dealing with multiplexing and the DL appear in “Coping with previously multiplexed data” on page 18.

Sizing

Storage capacity needs to be sized to adequately handle the amount of data anticipated to be retained in both native and deduplicated format. Backups that are larger than expected or contain data that deduplicates poorly can require much more storage space.

Experience and analysis have shown that there are three distinct use cases for the DL deduplication products. Each use case has its own storage capacity needs. Some space must always be available for new data. That amount of space, however, will depend on how the system is used.

- Immediate deduplication. With an immediate deduplication policy, deduplication may not keep up with ingest. Systems using an immediate deduplication policy ideally should include additional storage capacity of 50 percent of the largest amount of data ingested on a single day during the week.
- Immediate deduplication with fast read/restore. To ensure that data required for rapid restore is available in native format, these systems include additional storage capacity that accounts for the truncation threshold.
- Scheduled deduplication. These systems include additional storage capacity that is sized for the largest backup in native format on a single day during the week.

EMC strongly recommends performing a sizing assessment when including replication in the backup environment. The sizing assessment can help determine if replication can occur in the desired timeframe based on the replication network bandwidth, latency of the connection between the source and target systems, and estimated amount of data to replicate on each day.

Refer to the *EMC Disk Library and EMC DL Sizing Tools Practitioner's Guide* for more details on the information required to assess sizing requirements and the sizing tools available for this purpose.

The latest available DL sizing tools and the *EMC Disk Library and EMC DL Sizing Tools Practitioner's Guide* are available on Powerlink at (restricted audiences only):

Products > Hardware/Platforms > Disk Library Family > Sales Tools.

Space management

When a backup stream is ingested, the native data is written to the share or tape cartridge. When deduplication occurs, the unique blocks contained in that backup are written to storage in a different location and the metadata index of signatures and pointers representing deduplicated backup is stored. The native format data and unique blocks present on the system represent the used storage space.

The DL retains the native format data for as long as possible to facilitate quick recoveries. If this native data has been deduplicated for longer-term retention, data recovery will be from the deduplicated data once the native data is removed.

EMC recommends the following practices:

- Configure NetWorker to use expired tapes before scratch tapes as expired tapes still consume space on the DL1500 or DL3000 until space reclamation has been run.
- Expired media is not subject to space reclamation until the tape is also relabeled. Relabeling the expired tape places it in a state that allows the space reclamation process to dereference and subsequently delete the unique blocks associated with the backups on that tape.
- Either NetWorker's Auto Media Management feature needs to be configured on the jukebox, a backup script using NetWorker CLI commands needs to be created to force relabeling of tapes when they are

expired, or the backup administrator must properly size the number of tapes required so tapes are relabeled as they are expired. NetWorker will always use the least recently used tapes and if there are a lot of unnecessary tapes, space reclamation will be inefficient.

Used storage space is reclaimed for reuse in several ways. The following describes the processes employed for native format data and deduplicated data.

Native format data

Recovery of the storage space occupied by native format data is primarily under the control of NetWorker. When NetWorker expires a volume, there is no direct communication of the event to the DL. This has an impact upon DL space management when the DL is configured as a VTL. As a result, the DL's Console Manager will still indicate that logical tape cartridge as containing data and therefore still occupying storage space on the DL. For the DL to reuse this space, NetWorker must also relabel the logical tape cartridge by writing a data block to the beginning of the scratch volume. Relabeling is a task that is common to the management of physical tapes and may be performed via scripting or with NetWorker's Auto Media Management feature (see "Reclaiming space" on page 17).

To further self-manage its storage capacity, the DL monitors its available storage for user data and has a truncation feature that begins executing automatically when used storage capacity reaches 70 percent. The truncation feature removes the oldest native format data blocks from storage until the used storage capacity reaches 65 percent. This truncation feature operates only on fully hydrated blocks of data that have a deduplicated block counterpart. It does not remove any native format data that does not undergo deduplication (deduplication policy of never) or has not yet been deduplicated.

Deduplicated data

After NetWorker relabels a volume for reuse, the metadata index that described that volume is no longer valid. The unique blocks and pointers associated with that volume need to be dereferenced. When unique blocks are no longer referenced by any volumes, they will be removed from the block pool and the storage space they occupy will be recovered. The DL's space reclamation feature performs this dereferencing and unique block removal. The space reclamation feature is set to run once per week by default. This feature is resource-intensive and running it infrequently or after expiring and relabelling a large number of tapes can substantially lengthen the time space reclamation will take to complete. This can negatively impact other DL processes such as ingest, deduplication, and replication.

Best practices

EMC recommends running space reclamation daily. At a minimum, it should be run whenever tapes are expired *and* relabeled. This feature can be scheduled or run manually.

If possible, schedule space reclamation to occur outside ingest windows. This will reduce the competition for DL resources and minimize any impact on ingest, deduplication, or replication.

Phasing deduplication into the environment

EMC recommends that backup clients are phased in over time to a new DL, starting with an initial full backup followed by the other backup types. Add similar clients gradually so that over time the device has its block pool populated in a controlled manner.

For example, if the DL is installed during a normal incremental schedule, the initial backups to the DL are tied to the existing technology's last full backup. If a restore is needed, data would be required both from the previous technology as well as the DL. If the customer does not want to perform a full backup outside the normal policy schedule, they need to wait until the full backup is to occur to begin sending data to the DL.

General NetWorker settings and considerations

The following describes general EMC NetWorker settings and best practices for optimizing the backup environment. Specific recommendations for NAS and VTL environments appear later in this paper:

- Avoid running disk-intensive applications such as virus scanning on the backup client when it is backing up or restoring files.
- If the source data for the backup is located on a single, non-RAID, physical disk and multiple streams are running in parallel, the source physical disk could become a performance bottleneck because of parallel reading. Therefore, only run a single backup stream on NetWorker clients when the data to be backed up is on a single physical disk. On the other hand, use parallelism on the client when backing up multiple drives.
- Balance when backups start, rather than schedule hundreds of backups to begin at 8 P.M. For example, schedule 50 to start at 8 P.M., another 50 to start at 8:30 P.M., etc. Look at savegroup and client completion times, or drive activity, to balance the load.
- To ensure steady-state load, examine the drive target sessions, and try to keep 10 percent more sessions running throughout the backup window — fewer than this risks stalling target devices, more than this places unnecessary load on the infrastructure.
- On large systems with more than several hundred gigabytes to protect, eliminate data travel through the network by configuring the client as a storage node or dedicated storage node. If the client doubles as a storage node, data is transmitted directly to the DL backup device, thus providing a more efficient use of network resources.
- Depending on the characteristics of the data, a small number of large files could render better performance than a large number of small files.
- Increase the number of storage nodes and devices for better performance, within reasonable limits.
- Every device requires memory and CPU resources on the storage node and on the NetWorker server. Carefully build any environment with more than 100 devices. After reaching 100 devices, add devices gradually, perhaps 20 percent of the desired number at a time, to be assured the infrastructure can cope with the load.
- If considering switching a SAN environment to a NAS environment or vice versa, carefully check the LAN infrastructure. A dedicated backup LAN may be appropriate.

Note: Using Advanced File Type Devices is very different from using VTLs and will require an understanding of the benefits and limitations before deployment. Briefly, AFTD can simultaneously receive many save sessions. They can simultaneously handle many recover sessions. However, only one clone can be performed at a time, so ensure that a single stream can deliver the performance needed to clone data before growing an AFTD too large.

Information on licensing NetWorker for use with the DL1500 and DL3000 can be found in the article *Licensing EMC Disk Library 3D 1500 and 3000 with NetWorker and NetWorker Fast Start on Powerlink*[®], EMC's password-protected customer- and partner-only extranet.

SAN vs. NAS environments

The DL can operate either as a VTL or as a NAS device. NetWorker's multiplexing feature, while positively impacting backup performance, has a negative impact on deduplication performance in the SAN-based VTL environment. It is possible to reduce or eliminate multiplexing while achieving the benefits of deduplication with adequate performance in this environment through careful planning and management.

In NAS environments, EMC NetWorker's Advanced File Type Device accepts concurrent streams, writing them into separate files in the directory structure of the AFTD. This makes it possible for DLs to find and remove common data, achieving optimal deduplication performance.

Careful consideration to the type of backup environment to use must be given prior to DL integration. Consult the EMC Solution Architect team to assist in determining the best environment to meet objectives and involve them in rearchitecting if moving from SAN to NAS or NAS to SAN.

Some considerations to keep in mind:

- An environment with an investment in tape, either physical or virtual like the EMC Disk Library line, is often architected around SAN, not LAN and Ethernet. Administrators and operators know how to manage it, monitor it, and configure it. AFTDs on NAS require different planning, different infrastructure, different configuration, and different monitoring.
- Customers already using AFTDs can adopt NAS on DL without much of an infrastructure or mindset change, but those well versed in tape libraries will see significant differences.

It is critical to size the data storage area holding an AFTD correctly. A too-small AFTD may fill up and bring backups to a halt. Sizing a DL NAS share is done by ordering the correct amount of disk storage in the DL so there is sufficient storage capacity for backups.

It is very important to understand and monitor when using staging, which copies save sets to another target before deleting them from the source. AFTDs do not span, and so the operator must monitor usage to ensure the device does not fill up.

With a VTL, space problems can sometimes be solved by adding more slots. However, using AFTD requires growing the file system.

Using the DL with NetWorker in a SAN environment

The DL, when used with its optional virtual tape library (VTL) feature, can be configured to look like a number of tape libraries with associated tape drives to the backup servers on the SAN.

The VTL feature is a licensed option of the DL1500 and DL3000 and is used only in SAN environments.

This section covers DL behaviors to expect as well as configuration recommendations on to how to achieve optimal performance when using it with NetWorker in a Fibre Channel SAN environment.

For the NetWorker versions and supported operating systems, refer to the *EMC Support Matrix*.

Refer to “Key factors affecting the deduplication ratio” on page 8 for additional recommendations and settings.

IMPORTANT: When using the DL as a VTL with NetWorker, ensure that the DL is running version 1.0.1 build 2021-2034 or later.

DL SAN attach

The DL1500 has two 4 gigabit Fibre Channel ports and the DL3000 has four 4 gigabit front-end Fibre Channel ports for target mode FC attach. All connections to these ports are made via a Fibre Channel switch. Direct attachment of any device to these front-end ports is not supported. The following recommendations apply when connecting the DL to a backup server via a Fibre Channel switch:

- When emulating an EMC Disk Library, use a Fibre Channel switch listed on the *EMC Support Matrix*.
- Avoid using hard (port) zoning schemes with backup devices.
- Always use a single initiator single target zoning scheme by world wide port name (WWPN).
- Switch encryption solutions are not supported by the DL. The DL provides an inline encryption feature with the optional replication feature to meet data-in-flight encryption requirements.
- Limit FC extended fabric (ISL link) configurations to three hops between the backup server/storage node and the DL.

Device discovery (persistent binding)

When an operating system scans its storage buses (Fibre Channel or SCSI) for devices, it may assign a particular device to a different device name from the one it used during a previous scan. Since backup software saves in its database information about the device names of the various tape drives it sees during the configuration process, device reassignment during a scan will impact the availability of backup devices to the application.

Persistent binding is the ability to fix a particular device to a particular device name. In general, this is done by associating a unique device characteristic (such as a WWPN or serial number) with a device name to maintain persistence of a particular device to a specific device name. The method for doing this varies for different operating systems and Fibre Channel disk and tape device drivers.

If persistent binding is improperly configured, the operating system may provide a device with multiple names — giving a device one device name after a boot, but giving it a different name on a subsequent boot.

IMPORTANT: EMC recommends the use of persistent binding at the operating system level. This will prevent interruption of backup operations to the DL and difficult-to-diagnose problems after a system reboot.

Device drivers

EMC NetWorker does not require special tape drivers or driver settings. Use the tape drivers present in the backup software's operating system for the virtual tape drives in the virtual libraries.

Scanning devices

The operating system can scan for new devices. If all devices are not detected, a gap may exist in the LUN numbering. NetWorker's **inquire** command can be used to see a list of available devices, and visually ensure that all emulated tape libraries and their drives are visible to the operating system.

Run **inquire -l** to force LUN detection across gaps in LUN numbering.

If there is a gap in the LUN numbering, run **jbconfig -l** to configure the environment.

Configuring the backup server

NetWorker operates in a client/server model consisting of NetWorker servers, clients, storage nodes, and storage devices. When backing up to a virtual tape library on the DL using a storage node, a backup client sends the data over the LAN to the storage node and the storage node sends data across the SAN to the data to the DL. If the backup client is also a storage node, the client sends the data directly across the SAN to the DL. The backup metadata associated with the backup goes via IP over LAN/WAN from the backup client to the backup server where it is stored in the client file index and media database. The metadata is instrumental to NetWorker's ability to quickly locate and restore backup files.

1. Create a virtual tape library on the DL, specifying the emulation type, number of slots, number of tape drives, and their emulation type. This VTL will be a backup device on the NetWorker storage node. If desired, multiple VTLs can be created to service different backup servers/storage nodes.
2. Create virtual tape cartridges on the DL for the VTL, specifying the cartridge type, quantity, and capacity. If save sets will be cloned to a physical library, use the same cartridge type and blocksize as the physical library — cartridge capacities may be different between physical and virtual, however.

Note: The DL requires blocksizes of 512 KB or smaller.

3. Assign the VTL to a NetWorker storage node on the SAN to which the DL is attached.
4. Configure persistent binding on the NetWorker storage node.
5. Scan for devices at the NetWorker storage node. The VTL and its tape drives will appear.
6. Apply NetWorker licensing to enable the logical devices. NetWorker requires the VTL Autochanger bundle license for version 1.0 of the DL. This license enables the DL and provides entitlement for 10 TB of physical disk capacity within the DL.

NOTE: A change in DL version 1.1.3 and later enhances NetWorker 7.4 and 7.5's virtualization capabilities with the DL1500/DL3000, including the configuration of additional VTLs and an unlimited number of virtual cartridge slots within the library. Because this change affects which NetWorker licensing model is used with DL, customers currently using NetWorker 7.4 or 7.5 with DL version 1.1.3 will be required to apply a new enabler and authorization code in order to operate the Disk Library in VTL mode. **Without the enabler and authorization code, NetWorker will lose access to the DL for recoveries or backups (no data resident in the DL is lost however; only the ability to access that data).** Contact EMC licensing renewal for more information.

Customers using NetWorker 7.3 or earlier are not affected when upgrading to DL version 1.1.3. The VTL Autochanger bundle license will be required. Upgrading NetWorker to version 7.4 or later will require a new enabler and authorization code.

7. Perform the normal NetWorker setup for backups. See the *EMC NetWorker Administration Guide*.

A single client with multiple disk spindles, or multiple clients, can concurrently back up multiple save sets as parallel backup streams. In general, multiple parallel backup streams will increase the aggregate throughput of the backup process and will shorten the backup window of the client, for a given number of devices.

To use the DL deduplication feature, EMC recommends turning off multiplexing in NetWorker so that backup data from multiple sources is not interleaved on the virtual tapes. Keep in mind, however, that turning off multiplexing will adversely affect performance unless the environment is adjusted to compensate as described next.

1. Edit the properties of each virtual device in NetWorker to set the **target sessions** and **max sessions** to 1.
2. Reduce parallelism and re-examine schedules so that the number of active sessions in NetWorker is always approximately 10 percent over the aggregate target sessions.
3. Add more virtual devices gradually to the existing environment. Do not assume that the storage node and server infrastructure can support significantly more devices.

Simultaneous backups

The DL1500 and DL3000 emulate a number of industry-standard library and tape drive types. These appliances support multiple virtual libraries and tape drives simultaneously.

These aggregate throughput rates assume all available Fibre Channel target ports are used.

To address backup bottlenecks commonly seen with shared devices, two solutions can be deployed:

- Create a library and tape drives for each backup host's exclusive use to ensure the best possible performance.
- Create a single library with many tape drives, but only assign a few tape drives to each backup server. In this configuration, each backup server has its own dedicated resources. This also simplifies the initial installation and avoids any driver-related issues you could run into when sharing a tape device between unlike OS platforms.

With either solution, it is a good practice to limit the number of virtual tape drives so when all are servicing I/O at the same time each drive can sustain at least 20 MB/s. You may need to stagger backups to support fast and slow data streams to obtain maximum utilization of the available bandwidth.

Note: A NetWorker storage node is limited to either 16 devices (NetWorker Edition) or 32 devices (Power Edition). The devices can be logical, physical or a combination of logical/physical.

Single library – multiple hosts

For multiple hosts to use the same devices, the DL requires you to create different groups for each host. A group consists of exactly one host, one target, and one or more devices. The DL does not permit multiple hosts to access the same group.

This environment will require DDS licensing and configuration in NetWorker.

Virtual tape libraries

Emulations

The VTL permits you to create combinations of any emulated library with any emulated tape drive. NetWorker is generally insensitive to these combinations.

Other backup software applications are sensitive to combinations that do not exist in the real world, and will not work properly when unsupported combinations of library and tape drive types are specified. When they fail, it may not be obvious that it is an invalid combination.

EMC recommends creating virtual library and drive combinations that are formally supported by the actual physical libraries/drives.

Number of slots

The number of slots for a particular virtual tape library is not limited to the amount supported by the physical tape library it is emulating, but the number defined must be licensed appropriately in NetWorker.

Note: EMC recommends the use of the VTL Autochanger bundle license. This license supports a DL configuration with up to three (3) virtual libraries and an unlimited number of virtual cartridge slots. The license agreement calls for these three virtual libraries to all correspond to the same physical DL. Do not violate license terms by creating VTLs on different DLs.

Assigning and unassigning virtual libraries

Until a virtual library is assigned to a SAN client, it is not enabled for use by the backup servers. To ensure correct library behavior when assigning or unassigning virtual libraries, ensure that device numbering on the bus starts with LUN 0 and that there are no LUN numbering “holes.” LUN numbering holes can occur when devices are unassigned.

To change the LUN values, use the Console Manager as follows:

1. Go to Configuration.
2. Select **SAN Clients**.
3. Select **Edit**.
4. Change the LUN values for the device list as necessary.

Virtual tape cartridges

When you create virtual tape cartridges, they are spread across the storage LUNs in a round-robin fashion. This maximizes performance by spreading the I/O load across multiple LUNs.

The cartridge size can be customized but cannot exceed the capacity of its physical equivalent. This ensures that the compressed data stored on a virtual tape cartridge will fit on a physical tape once it has been decompressed and then recompressed by the physical tape drive. The DL1500 and DL3000 will not allow you to create tape cartridges larger than their physical equivalents. A size of 100 GB has been found to be a good minimum virtual cartridge size.

Reclaiming space

Much like physical tape, the bytes from expired save sets still reside on the virtual cartridge until the cartridge has been relabeled, truncating them. When a virtual tape becomes recyclable in NetWorker, relabel it to free up space in the DL.

By default, NetWorker will not automatically label or relabel tapes. But automatic tape labeling will be done if the Auto Media Management feature is set on the jukebox resource:

The following is from the *EMC NetWorker Administration Guide*:

Auto Media Management gives the NetWorker server automatic control over media loaded in the storage device. When Auto Media Management is enabled during device configuration, the NetWorker server automatically:

- Labels the volume (recognizes EDM labels and does not overwrite them).
- Mounts the volume.
- Overwrites volumes it considers to be unlabeled.
- Recycles volumes eligible for reuse that are loaded into the device.

The NetWorker server considers a volume to be unlabeled under the following conditions:

- Has no internal label.
- Is labeled with information other than a NetWorker label.
- Is labeled with a NetWorker label, but the density indicated on the internal label differs from that of the device where the volume is mounted.

Note: The NetWorker server considers volumes that were labeled by a different application to be valid relabel candidates if Auto Media Management is enabled. Once the NetWorker server relabels the volume, the previously stored data is lost.

If NetWorker automatically recycles [=reuses] a volume after all of the save sets on the volume have passed the customer-set useful lifetime, the two 32 KB labels at the beginning of the tape are *not* overwritten; the tape is written beginning at the tape position just after the second tape label. In other words, the tape usage size does not return to zero.

Note: When the NetWorker server recycles a volume, the volume label does not change if the volume remains in the same media pool. That is, if a volume labeled Dev.006 is recycled, it will retain the volume label Dev.006 and will not receive a new label with the next sequence number. The original data on the volume, however, will be overwritten by the new data

Coping with previously multiplexed data

If a DL has too much space already consumed by existing multiplexed data, and that data needs to be retained, the data can be un-multiplexed by careful cloning. The cloned tapes will deduplicate when rewritten to the DL. Follow this by relabeling the original multiplexed tapes to free the non-deduplicated space. (Be sure everything is off of the originals before relabeling!)

Note: This process will require 1) scripting; 2) extra storage capacity; and 3) processing through a storage node.

There is an additional benefit from this extra clone step beyond deduplication. The data is now in an un-multiplexed format for faster recovery.

Look to the hardware and time requirements for existing cloning to judge what is needed for an extra clone step. Be conservative when estimating as single save set cloning will need at least N reads of the tape, where N is the amount of multiplexing. (Normal cloning of multiplexed data only requires one read of the tape.)

Demultiplexing through cloning requires sufficient space in the DL to accommodate the multiplexed and demultiplexed cloned copies of the same data. This will not be practical if the DL is already full.

However, if sufficient capacity exists, cloning can occur. The cloning operation is performed from the command line and is on a scale that will likely require scripting. *Save sets must be cloned individually.*

For example:

```
# pseudocode: EMC does not support scripts
for vol in $volume_names
```

```
do
  ssids=`mminfo -q volume="$vol" -r ssid`
  for ss in $ssids
  do
    nsrclone -S $ssid
  done
done
```

The script must not try to clone all the save sets on a volume with a single nsrclone command. A performance optimization in NetWorker will leave data from those save sets multiplexed on the clone volume, which defeats the purpose here.

Using the DL with NetWorker in a NAS environment

This section covers DL behaviors to expect as well as configuration recommendations on to how to achieve optimal performance when using it with NetWorker in a NAS environment.

General recommendations

NAS B2D shares created on the DL1500 and DL3000 can be hidden from network browsing when they are created using the **Hide this share from network browsing** option.

Refer to “Key factors affecting the deduplication ratio” on page 8 for additional recommendations and settings.

Network connectivity

The DL contains six or eight Gigabit Ethernet ports for the DL1500 and DL3000, respectively, configured as a bonded network interface. All six or eight ports service replication traffic, appliance management traffic, and NAS backup data traffic. Therefore careful planning and scheduling can optimize the performance of each particular operation. One way to achieve this is by scheduling these operations to be performed at particular times.

The bonding method used for the DL ports can be either the round robin policy (Mode 0), which transmits packets sequentially across all ports for load balancing and fault tolerance, or LACP (Mode 4), which utilizes IEEE 802.3ad dynamic link aggregation. The underlying DL appliance software and operating system manage the I/O distribution between each of the bonded ports in an attempt to optimize the aggregate bandwidth possible across the links.

Mode 4 is available in DL version 1.1 and later. Mode 0 is available in DL version 1.0 and later.

Significant performance optimization will be achieved through configuration settings on managed Ethernet switches. For Mode 0 and Mode 4, EMC recommends using managed switches to group the physical port connections to which the DL appliance is connected.

For Mode 0, use a managed switch with the ports used by the DL configured for "trunking" or "link aggregation" or "EtherChannel" without LACP or pAgP to maximize performance potential. Refer to the vendor documentation for specific instructions when configuring the port group.

For example, to configure a Cisco 6500 Series switch for a common EtherChannel group, run the following command from the Cisco switch CLI:

```
set port channel mod/ports... [admin_group]

mod = switch module
ports = port numbers to be added to the port channel group.
admin_group = name of port channel group (optional)
```

For example:

```
set port channel 2/1-8
```

For Mode 4, switches that utilize IEEE 802.3ad dynamic link aggregation allow their ports to be automatically grouped together to form an ultra-high-bandwidth connection that greatly expands bandwidth capacity to the network. This aggregation of DL port connections potentially improves NAS data ingest or replication performance depending on the environment.

For example, to configure a Cisco 6500 Series switch for an EtherChannel group with LACP, run the following commands from the Cisco switch CLI:

```
set channelprotocol lacp mod
set port lacp-channel mod/ports
```

mod = switch module
ports = port numbers to be added to the port channel group

For example:

```
set channelprotocol lacp 2
set port lacp-channel 2/1-8
```

For more information on programming EtherChannel groups and the LACP mode when using Cisco 6500 Series switches, see:

<http://www.cisco.com/en/US/docs/switches/lan/catalyst6500/catos/8.x/configuration/guide/channel.html#wp1020899>

For Mode 0, all DL ports connect to the same Ethernet switch. Splitting bonded ports on a DL between multiple switches, however, is supported on *certain* switches when using Mode 4. Consult the switch vendor's documentation.

When using Mode 0, make sure LACP is disabled when configuring the EtherChannel port group on managed switches.

When using Mode 4, only connect to Ethernet switches that support IEEE 802.3ad dynamic link aggregation and are properly configured for LACP.

In general:

- Use a dedicated network by configuring a separate network or use QoS features that guarantee network bandwidth.
- Alternatively, use virtual networks (VLANs) to segregate backup from production network traffic.
- Set network interface cards (NICs) for servers and clients to full duplex. Set all routers to full duplex.
- Use only CAT 5e or CAT 6 cables (1 Gb/s rated cables).
- If using a DNS server, verify the DNS server configuration settings are correct.
- Use multiple DL ports when connecting to the network. The more DL ports used, the better the load balancing will be across the ports.
- For redundancy, connect at least two DL ports to an Ethernet switch.
- Set each switch port used by the DL to auto-negotiate/auto-sensing. The DL network interface cards are preset to auto/auto and cannot be changed.
- When connecting multiple DL Ethernet ports to the network when using Mode 4 (LACP), make sure the proper LACP settings on the switch are configured. Consult the switch vendor's documentation for the configuration steps.
- When using Mode 0 (balance-rr mode), use a managed switch configured for "trunking" or "link aggregation" or "EtherChannel" to maximize performance potential. Refer to the vendor documentation for the switch for specific instructions when configuring the port groups. Do *not* use the LACP settings on the switch.

Keep the following in mind:

- Concurrent operations on the DL *will* have significant impact on performance
- The DL does not support jumbo frames.
- Splitting bonded ports on a DL between multiple switches is *not* supported when using Mode 0.
- Splitting bonded ports on a DL between multiple switches is supported on *certain* switches when using Mode 4. Consult the switch vendor's documentation.

Backup-to-disk configuration

NetWorker operates in a three-tier client/storage node/server model consisting of NetWorker servers, clients, storage nodes, and storage devices. When backing up to disk using a storage node, a backup client sends the data over the LAN to the storage node and the storage node stores the data to the DL. If the backup client is also a storage node, the client sends the data directly to the DL. The backup metadata associated with the backup goes from the backup client to the backup server where it is stored in the client file index and in the media database. The metadata is instrumental to NetWorker's ability to quickly locate and restore backup files.

Best practice for AFTDs is to create one per pool on a storage node and not to place more than one on a file system. The AFTD should be the only thing on the file system. This is especially important for supporting replication in the future.

In the simplest LAN backup-to-disk configuration, there is a single NetWorker client, a backup server, and DL as the storage device all connected via an Ethernet LAN.

Setting up the backup-to-disk environment involves the following steps:

1. Create an NFS or CIFS share on the DL. This share will be mounted (or mapped) as a backup device on the NetWorker server/storage node. If desired, multiple shares can be created to store different collections of save sets.
2. Mount (or map) the share on the NetWorker server/storage node.
 - a. Create a mount point for the NFS share to be used by the backup application and mount it using the commands for your particular operating system as described in "Mounting NFS shares" on page 23.
 - b. Map the CIFS share to a drive in Windows Explorer and create a directory to be used by NetWorker as the location for the backup device as described in "Configuring CIFS shares" on page 24.
3. Apply the NetWorker Disk Backup Option license. This license enables the DL to be used as a disk target and provides entitlement for a selectable capacity of physical disk within the DL.
4. Create a NetWorker Advanced File Type Device that points to the share on the DL (or directory in the CIFS share).

- c. Open the NetWorker Management Console and then invoke Administration.
- d. Create a new device as follows:

Select Devices and under Devices tree, create a new device:

"Name" = to full path to the CIFS share created on the DL (example:
[\\10.40.166.100\pep_cifs1\nw](#))

"Media Type" = adv_file

Also specify "netadmin" and the password for "username" and "password" on the second tab.

5. Perform the normal NetWorker setup for backups. See the *EMC NetWorker Administration Guide*.

During a backup operation, the NFS or CIFS share on the DL designated as the backup device receives the save set directly from the client/storage node or backup server.

In large environments it is not likely that every single client will have storage node software installed, or storage node licenses available. When there are a large number of clients, it is much more likely that a small number of storage nodes will serve as an aggregation point for a large number of clients. The clients

will send their data to the storage node with which they are associated, and the storage node will back up the data to the share on the DL.

Mounting NFS shares

The DL can serve as a NAS appliance for backup purposes. If it is an NFS environment, access to the shares is restricted by hostname. Root access to a NFS share is not allowed and the access rights will be changed to “nfsnobody” as a security precaution.

DL version 1.0 does not support NFSv4.

NFS shares created on the DL are mounted on the NFS client in the following environments as described next:

AIX 5.2

Enter the following command:

```
nfsd -o nfs_use_reserved_ports=1
```

This mount command does not persist across AIX reboots. For AIX 5.2 or later, use the `-p` option to mount the share permanently. For releases of AIX prior to 5.2, the mount command must be included in a boot up script.

If you are using NFSv3, mount the NFS share created on the DL using this command:

```
mount -o timeo=600 {nfs_server}:/ {export path} / {mountpoint}
```

For example: `mount -o timeo=600 cosmowanda:/Q/shares/tl_nfs_1 /stoli1`

To mount the NFS share on another NFS version, use the command:

```
mount -o vers={1|2|3},timeo=600 {nfs_server}:/ {export path} / {mountpoint}
```

To show the list of file systems exported by the nfs server:

```
showmount -e {nfs_server}
```

and

```
showmount -e dfshares {nfs_server}
```

To show or clear the NFS statistics, respectively:

```
nfsstat -c [show NFS statistics]
nfsstat -z [clear NFS statistics]
```

HP-UX 11i

If you are using NFSv3, mount the NFS share created on the DL using this command:

```
mount {nfs_server}:/ {export path} / {mountpoint}
```

For example: `mount cosmowanda:/Q/shares/tl_nfs_1 /stoli1`

To mount the NFS share on another NFS version, use the command:

```
mount -o vers={1|2|3} {nfs_server}:/ {export path} / {mountpoint}
```

To show the list of file systems exported by the nfs server:

```
showmount -e {nfs_server}
```

and

```
showmount -e dfshares {nfs_server}
```

To show or clear the NFS statistics, respectively:

```
nfsstat -c [show NFS statistics]
nfsstat -z [clear NFS statistics]
```

Linux

If you are using NFSv3, mount the NFS share created on the DL using this command:

```
mount {nfs_server}:{export_path} /{mountpoint}
```

For example: `mount cosmowanda:/Q/shares/tl_nfs_1 /stoli1`

To mount the NFS share on another NFS version, use the command:

```
mount -o nfsvers={1|2|3} {nfs_server}:{export_path} /{mountpoint}
```

To show the list of file systems exported by the nfs server:

```
showmount -e {nfs_server}
```

To show the NFS statistics:

```
nfsstat -c [show NFS statistics]
```

Solaris

Mount the NFS share created on the DL using this command:

```
mount -o vers=3 {nfs_server}:{export_path} /{mountpoint}
```

For example: `mount -o vers=3 cosmowanda:/Q/shares/tl_nfs_1 /stoli1`

To mount the NFS share on another NFS version, use the command:

```
mount -o vers={1|2|3} {nfs_server}:{export_path} /{mountpoint}
```

The `mount {nfs_server}:{directory} /{mountpoint}` command defaults to NFSv4 in Solaris. Attempting to mount the share using NFSv4 results in a “No such file or directory error”. To avoid this error, specify the NFS version when using this command.

To show the list of file systems exported by the nfs server:

```
showmount -e {nfs_server}
```

and

```
showmount -e dfshares {nfs_server}
```

To show or clear the NFS statistics, respectively:

```
nfsstat -c [show NFS statistics]
nfsstat -z [clear NFS statistics]
```

Configuring CIFS shares

If you are operating in a Windows environment, you must specify the Windows domain information in the DL Configuration section of the DL Console Manager before you create the CIFS share. This involves adding the name of the workgroup or Active Directory domain used by the Windows system on which the CIFS share will be used. Once a CIFS share is created, access is allowed on a per-username basis.

Since only one value for a workgroup or Active Directory domain can be entered, either workgroup or Active Directory domain validation will be used but not both at the same time. Also, only one workgroup name or Active Directory domain value can be used by all Windows systems using this DL system.

1. Create a CIFS share on the DL and join a workgroup or domain type authentication domain (using the same type as the NetWorker server or storage node that will use the CIFS share).
2. Map the CIFS share to a drive in Windows Explorer and create a directory (for example, “nw”) to be used by NetWorker as the location for the Advanced File Type Device (which will be created in a later step).

Note: After this directory is created, you can dismount the CIFS share. You need not mount it again. NetWorker will mount the share directly using a UNC path.

3. Add a user **netadmin** with a password on the DL.
4. Add a Windows user called **netadmin** and make the user a member of the Administrators and Backup Operators group.
5. Under Device Manager/Services, edit the **NetWorker Backup and Recover Server** and **NetWorker Remote Exec Service**.
 - e. Under **Log on as**, choose **This account** and enter the netadmin username and password.
 - f. Restart these services.

Moving data from emulated tape to physical tape

There are many options for copying backups on virtual tape to physical tape for long-term data retention. Physical tapes can be created directly from virtual tape using EMC DL1500 and DL3000 copy-to-tape options or with the standard method of copying or *cloning* with backup software.

The DL has two copy-to-tape options:

- Backup application specific Path-to-Tape (PTT)
- Auto Archive

The DL copy-to-tape options require the connection of a physical tape library to port 3 of the DL's HBA. The respective feature licenses, when installed, convert this port from a SCSI target to a SCSI initiator. As a SCSI initiator port, port 3 of the HBA will no longer support backup host connectivity. Only one of these options can be licensed at a time on the DL.

Backup application specific Path-to-Tape

Path-to-Tape (PTT) is a licensed option of the DL1500 and DL3000 and is used only for copying backup images between virtual cartridges and physical tape or a virtual cartridge in an EDL VTL. This feature cannot be used to copy NAS backups – CIFS shares or NFS mount points – to tape.

The PTT feature is available with NetWorker 7.5.1 and later.

The PTT feature makes use of NDMP commands implemented internal to the DL to use the DL system's NDMP tape server to copy save sets or backup images on virtual tape cartridges under the direction of NetWorker. The PTT feature does not utilize the storage node for data movement for this copy operation, thus freeing up resources. NetWorker's catalog tracks all the tape copies performed in this fashion. Restores can be performed directly from either tape copy. Figure 1 illustrates the control and data flow that occur with the PTT feature.

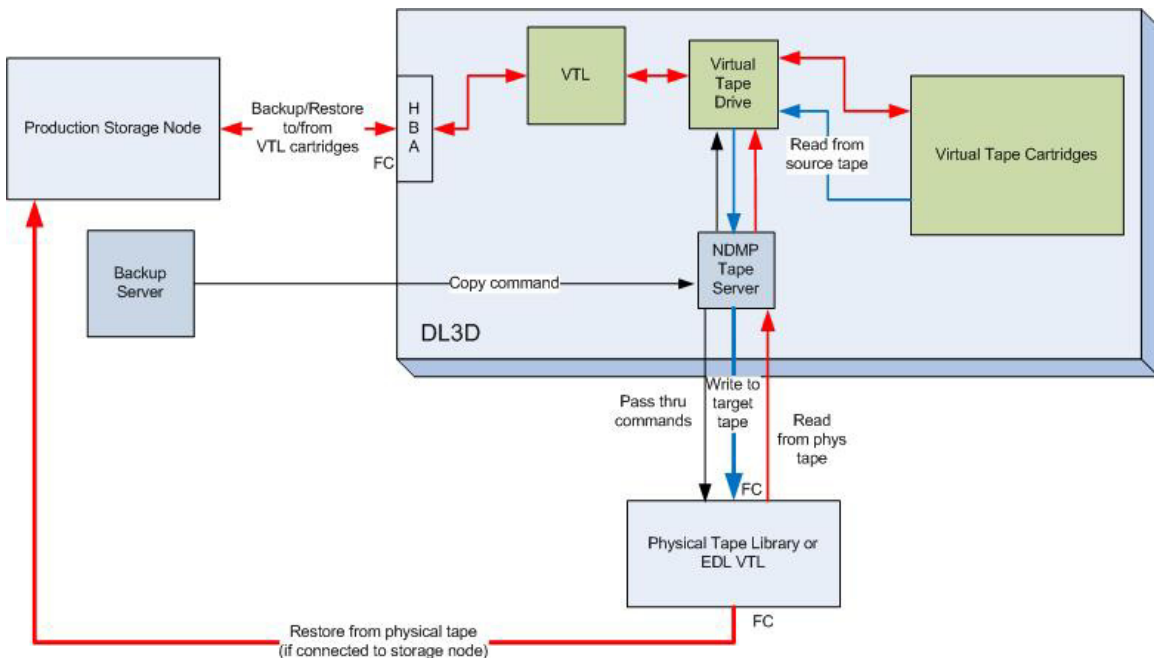


Figure 1. PTT control and data flow

NetWorker uses NDMP as a control protocol to tell the DL appliance to duplicate a save set on its virtual tape cartridge A to a physical tape or virtual cartridge B in an attached destination device. When

NetWorker initiates the copy, the virtual tape cartridge to be copied is loaded into a drive in the source virtual tape library and positioned to the start of the save set to be copied. The media to receive the copy is loaded into a drive in the destination physical or virtual library, labeled (or has its label verified), and positioned to be ready to accept data. Then the NDMP tape server software in the DL is instructed by NetWorker to copy from the source virtual tape cartridge to the destination media. The DL undeduplicates the data (if deduplicated) and copies it to the destination media, telling NetWorker when the copy process completes. If the destination tape becomes full during the copy operation, NetWorker will put that tape back into its slot in the library, load another tape into the drive, label it (or verify its label), and continue copying. NetWorker records which tape received the copy. Using this method NetWorker is able to catalog and track copies on both the source and destination media.

The copies generated by the PTT feature can be read without the use of a DL. In other words, the data on the tape is written in an unduplicated native tape backup format, with a different (BSP catalog aware) barcode.

NetWorker site requirements for PTT

- *DL version 1.1.2 or later* with both VTL and PTT licensed features installed
- Destination library (physical tape library or EDL VTL)
- *NetWorker 7.5.1 or later* with:
 - A NetWorker license used to enable the DL as a backup target
 - The appropriate VTL capacity licensing based upon physical capacity of the DL disk drives
 - An NDMP client connection license used to enable the DL as a data source in the PTT process
 - The appropriate NetWorker licensing for the destination library. This may be Autochanger licenses for a physical tape library or another VTL Autochanger bundle license for a second EDL VTL.

NOTE: If using DL version 1.1.3 and later, it includes an enhancement to NetWorker 7.4 and 7.5's virtualization capabilities, including the configuration of additional VTLs and an unlimited number of virtual cartridge slots within the library. Because this change affects which NetWorker licensing model is used with DL, customers currently using NetWorker 7.5 with DL version 1.1.3 will be required to apply a new enabler and authorization code in order to operate the Disk Library in VTL mode. **Without the enabler and authorization code, NetWorker will lose access to the DL for recoveries or backups (no data resident in the DL is lost however; only the ability to access that data).** Contact EMC licensing renewal for more information.

- Two storage node resources configured in NetWorker, one for the initial backup (the production storage node) and a second (for the DL) to perform the copy to tape, and appropriate licensing.

Auto Archive

DL version 1.1 supports connecting a physical tape library or EMC Series 4000 Disk Library to the DL and automatically exporting virtual tapes to corresponding physical tapes or virtual tapes as well as the importing of physical or virtual tapes to virtual tapes in the DL. The Auto Archive feature supports the following physical or virtual tape library models: Quantum iScalar 2000, iScalar 500, PX720, and PX500 only with IBM LTO-2, LTO-3, or LTO-4 tape drives.

Auto Archive is a licensed option of the DL1500 and DL3000. Unlike backup application specific PTT, these copy operations do not utilize NetWorker. Auto Archive operations are managed from the DL's Console Manager. All data is copied in native format.

Auto Archive provides the following capabilities:

- Auto archive – copies the virtual tape to a physical tape and retains the virtual tape
- Export – copies the virtual tape to physical tape and deletes the virtual tape
- Import – copies the physical tape to virtual tape

-
- Does not require integration with or support by backup software

The main advantages of Auto Archive are:

- No CPU cycles and no I/O bandwidth resources are consumed on the backup server.
- The backup server does not need to be connected to both the DL and the physical tape library.

When using Auto Archive, keep the following in mind:

- There is a possibility that the entire contents of the virtual tape *may not* fit on the physical tape. The default virtual tape size accounts for usable tape capacity, which may vary from manufacturer specifications.
- The tape format of the copy must match the format of the source tape. It is not possible to copy to a different tape type.
- Only the native tape format supported by the tape drive is supported. For example, with LTO-3 tape drives, use LTO-3 tapes.
- Manual record keeping is involved unless the backup software contains some type of vaulting software to track offsite cartridges.

When a virtual tape is archived or exported to a physical or virtual tape, the tape copy is identical to a tape that was written directly by NetWorker and can be used to restore files directly on a system (without the DL). The restore environment must conform to certain rules:

- NetWorker must be used to restore the tape.
- The same type of tape drive must be used as that which wrote the physical tape.

Auto Archive site requirements

- DL with both VTL and Auto Archive licensed features installed
- Destination library (physical tape library or EDL VTL) of the supported type (Quantum iScalar 2000, iScalar 500, PX720, and PX500 only with IBM LTO-2, LTO-3, or LTO-4 tape drives)
- VTL configured on the DL of the same emulation type as the physical tape library or EDL VTL being used

NetWorker cloning

When data backed up to a DL in VTL mode needs to be written onto physical tape, NetWorker can clone save sets written to emulated tape in the DL through a production storage node to a SAN-attached tape library. This operation is a standard NetWorker cloning procedure. This procedure uses emulated tape drives in the DL in conjunction with a tape device in another tape library to perform the cloning process. Cloning from a production storage node to a second storage node can also be performed over IP between the storage nodes.

The advantages and disadvantages of this strategy are as follows:

Advantages

- Enables cloning from the DL under NetWorker control, with standard NetWorker policy support; that is, multiple retention policies for different cloned copies of data
- Enables cloning at the save set level
- Enables copying from one emulated tape type to another tape type (physical or virtual). This is sometimes referred to as tape conversion
- Enables copying from multiple emulated tape cartridges to a single physical or virtual tape (tape stacking)

Disadvantages

- Requires storage node licenses
- Must maintain additional front-end SAN infrastructure for connections to a separate tape library as well as the emulated tape library
- Consumes SAN bandwidth – data must be copied from a virtual tape cartridge in the DL over the SAN to another physical device on the SAN
- Consumes storage node CPU, memory, and I/O bandwidth

Replicating deduplicated data

DL replication does not utilize the backup application. Replication is managed by the DL, and manual tracking of the replicated copies is required.

Replication is an optional licensed feature and is supported in the following configurations:

- DL1500 to DL1500
- DL3000 to DL3000
- DL1500 to DL3000
- DL3000 to DL1500

DL replication occurs in two stages. First, a copy of the unique data associated with each replication-enabled VTL/share that is not already present at the target DL transfers to the target DL as the data deduplicates. This is referred to as *continuous* replication. Second, a metadata description of the replicated VTL/share or cartridge or directory/file is sent to the target system. This second stage ensures that a recoverable image is available on the target DL. This second stage is referred to as *namespace* replication.

In version 1.1, there are two variations of namespace replication:

- VTL/NAS share replication (as is in version 1.0)
- Individual tape cartridge or directory/file replication

With VTL/NAS share replication, a VTL or NAS share is made available as a recovery point when its metadata description is replicated to the target system. This replication of the metadata description can be scheduled once every 24-hour period; it can also be done manually as many times as desired. The 10 most recent recovery points of a NAS share/VTL are available on the target DL for selection.

With individual tape cartridge or directory/file replication, the tape cartridge or directory/file becomes readable on the target DL as soon as its metadata description is replicated. This replication of the metadata description occurs when a tape is unmounted (when using cartridge-based replication) or when a CLI replication command is issued (when using directory/file-based replication). The VTLs and NAS shares on the target DL contain the most recent copy of each replicated tape or directory/file.

NOTE: Failure to perform a namespace replication of a replication-enabled VTL or NAS share will result in no recovery point or available cartridge or directory/file on the target DL.

See “Appendix: Using replication with NetWorker” on page 33 for an example of using NetWorker with the DL’s directory/file-based replication feature.

More information on replication can be found in the *EMC DL1500 and DL3000 Version 1.1 Replication – Best Practices Planning* white paper on Powerlink.

Replication considerations

Replication transmits data over Ethernet using TCP. The source DL can optionally encrypt the data using AES-128 before sending it across the TCP link to the remote DL where it is unencrypted before it is stored. The DL uses TCP port 1062 for replication of the unique data and port 80 for replication of the metadata

description. There is no physical Ethernet port on the DL reserved for the replication feature. See “Network connectivity” on page 20 for recommendations when connecting to the network.

When adding DL replication to the backup environment, keep the following in mind:

- Always assess the network between the source and target systems in *both* directions as part of the test and acceptance process and prior to putting the systems into production. This will ensure that sufficient bandwidth is available to support the maximum amount of data desired to be replicated within the replication window. This can also be done by performing a formal assessment if necessary.
 - Use “ping” or similar analyzer software to get an idea of latency between the two systems as well as determine if any network bottlenecks are occurring. Run this test multiple times on different days at different times to confirm the numbers.
 - Use the Console Manager’s built-in Performance Analyzer (Figure 2) to measure network performance between the two systems.

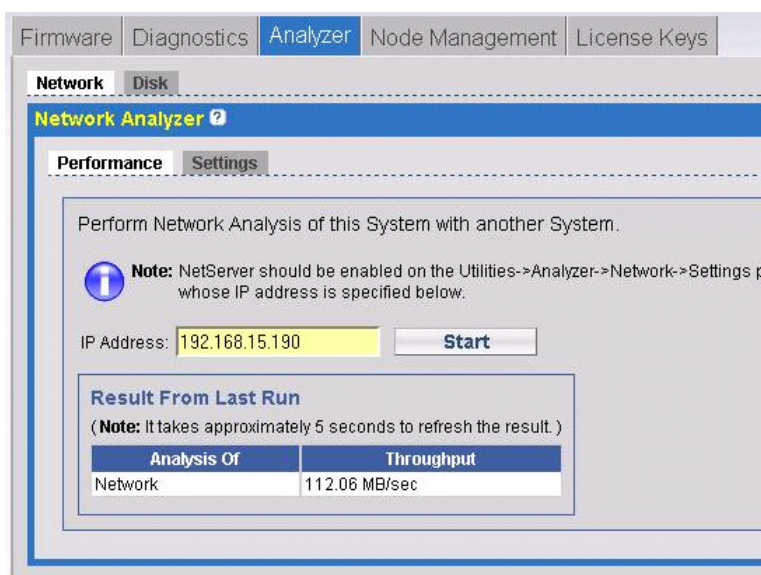


Figure 2. Network performance analysis

- Make sure both the source and target DLs are configured for replication before performing any backups to the source DL.
- Seeding of the target system is almost always required. See “Seeding (pre-populating the target system)” for recommended options.
- Do not place replication configurations into production until required replication windows are met. When replication falls behind (exceeds the replication window), it is hard to catch up.
- Consider using immediate deduplication to effectively extend the replication window and lower the chance of “falling behind.” Since replication is initiated when deduplication begins, this enlarges the replication window to maximize the amount of data that can be transferred should change rates increase unexpectedly, an unusually large amount of data is backed up, or excessive network traffic slows data movement.

Replication performance depends on daily change rates between new backup data and data residing on the target system.

- The higher the change rate, the more data that needs to be moved within the replication window.
- Lower than expected deduplication rates will add time to the replication and may exceed the required replication window.

-
- If deduplication is slowed, continuous replication will also be slow. Deduplication can be slowed by poor ingest rates, fragmentation, and other factors that affect the speed at which block level match detection can occur.
 - Network latency will have a significant impact on replication performance.
 - 128-bit AES encryption for replication is available. Encryption may affect replication performance. It should be disabled if the customer's WAN is already secured.

If you have opted to do VTL/NAS share-level namespace replication, the length of time it takes to verify that all data associated with that VTL/NAS share is present at the target DL will increase as the amount of data increases. This will impact the speed at which a recovery point is available on the target DL as the metadata description of that VTL/NAS share cannot replicate to the target DL until the verification completes. Consider the following options:

- Use smaller VTLs or NAS shares targeting specific data and performing rolling backups to these VTLs/shares to spread the replication out over time and allow the replication process to work with easily managed data sizes.
- Use cartridge-level or directory/file-level namespace replication instead. Because the data verification step occurs at the cartridge or directory/file level, the delay that can occur before the metadata description replicates is substantially reduced unless you append to large cartridges. In this case, always overwrite or use new cartridges or limit the size of the cartridge.

Seeding (pre-populating the target system)

Initial replication (first data copy) is a time-consuming operation and will usually take significantly more time than later replications because the data is typically new and unknown to the target DL.

For replication, the goal is to facilitate the transfer of the large amount of unique data produced by initial backups to the source shares/VTLs to the target DL. This is done by seeding the target DL with the same data as on the source DL so that only the minimum amount of data needs to be transferred between units on subsequent replications to maintain synchronization. The source and target systems should not be put into production until this seeding process is complete. Seeding is considered complete when the required replication window has been demonstrated to have been met.

There are a number of options that can be employed to accomplish this:

- **OPTION 1:** Locate the source and target DL so that they are on a dedicated Ethernet network and replicate locally. This will allow replication to proceed at the fastest possible network rate. EMC recommends replicating at least two full backups in this local configuration to demonstrate replication savings. If savings are not seen, further local replications may be needed to ensure the target system has sufficient data. (That is, the data transferred by subsequent replications is more representative of the quantity of data that will typically replicate within the window.) When the seeding completes, the target DL can be deployed to its intended location.
- **OPTION 2:** Locate the source and target DL so that they are both connected to the same storage node, and use the storage node to perform an inline copy or dual write of data from at least two full backups to the source and target systems, using a temporary NAS share or VTL on the target DL. Deploy the target DL to its target location and perform a namespace replication of the VTL/NAS share. Make sure DL space reclamation is turned off on the target DL. Delete the temporary NAS share or VTL on the target DL as well as references to the temporary copy in the backup application catalog.
- **OPTION 3:** Use tape to seed the data in the target DL. Using the tapes from at least two full backups present on the source DL, write the data to temporary NAS shares or a VTL on the target DL. Make sure DL space reclamation is turned off on the target DL. Perform a namespace replication. Delete the temporary NAS share or VTL on the target DL as well as references to the copied save sets/images in the backup application catalog if the tape copy was performed by a backup application.

Note: The goal of Options 2 and 3 above is to write native backup data directly to the source DL and target DL and allow each system to deduplicate the data.

If replication is not enabled for a NAS share or virtual tape library until multiple backup streams destined for that particular virtual device or NAS share are ingested, there will be a large number of unique blocks on the source DL that will need to be transferred over the network when the first namespace replication occurs. In this situation, the data verification performed as part of the namespace replication will detect that all the unique blocks on the source system must be sent over the network and there may be a significant delay before the data in the two systems is synchronized.

If there are disruptions during the replication process for any reason (for example, a network infrastructure problem), reseeding may be necessary to eliminate a large replication backlog that can occur as a result of the disruption. Large backlogs, when they exist, make it difficult for the systems to resynchronize.

Recovery/Failback

Should a VTL become unavailable on a source DL, two methods are available to access the data either at the target system (recovery) or by replicating the VTL to the source DL (failback).

In order for the replicated VTL to be accessed from the target DL by a secondary storage node, a recovery must be initiated. When a recovery on the target DL is performed, the virtual tape libraries appear on the target DL. In order for a virtual tape library to be accessed by a secondary storage node on the target side (visible to the SAN client ports), you will need to enter a name for the library, a virtual tape drive model number, and the number of drives on the DL.

When failing back to the source DL, delete the original virtual tape library on the source DL first. If it is not deleted, the BSP may end up seeing the same barcodes twice. Perform a failback on the target DL, and then perform a recovery on the source DL.

If the VTL and volume names are the same on the replicated side, disable the original library in NetWorker, create a new library, and issue reset and inventory commands once it is up. That will result in NetWorker updating the location field for all volumes so it will be able to use them automatically from that point on.

This can be done proactively with both VTLs (source and replicated one) configured in NetWorker all the time. Just make sure that only a single library is enabled in NetWorker at any given point in time – if both libraries are enabled, conflicts will occur.

If such a configuration is done beforehand, failover from the NetWorker perspective can be performed in less than one minute, but this is a manual process. It cannot be automated.

Conclusion

The processes and configuration recommendations described in this white paper are intended to assist you in maximizing the performance and usage of the DL in the majority of NetWorker backup environments. Since no two backup environments are exactly the same, it may be necessary to vary individual settings to meet each specific environment's requirements and goals.

References

The following EMC white papers may provide further information. They can be found on EMC.com or Powerlink, EMC's customer- and partner-only extranet.

- *EMC DL1500 and DL3000 - Best Practices Planning*
- *EMC DL1500 and DL3000 Version 1.1 New Features and Functions – A Detailed Review*
- *EMC DL1500 and DL3000 Version 1.1 Replication - Best Practices Planning*
- *EMC DL1500 and DL3000 Copying to Physical Tape – Best Practices Planning*

Appendix: Using replication with NetWorker

NetWorker can be configured to work collaboratively with replicated NAS shares. Recommendations for configuring the DL systems as well as NetWorker are outlined in the sections “DL setup” and “NetWorker setup,” respectively, in this appendix. It is important to note, however, that no scripting is needed for the replication process for the DL in VTL mode, because cartridge-based replication of a tape volume occurs automatically when it is dismounted on the source DL.

One of the premises of the replication process in the example below is that the NAS share must contain a single AFTD and nothing else. Anything else on the NAS share will suffer collateral damage from replication, because only NetWorker’s use of the share will be paused during replication.

EMC Solution Architects have a much richer script than that provided here to update the media database, delivered as part of a services engagement in a NetWorker 7.5 environment. That script does give the replicated copy a distinct volume ID, so that both AFTDs can be used simultaneously after replication for recovery or cloning. Talk to your EMC representative if interested in media database consistency.

DL setup

The source and target DLs must be correctly configured for replication for this script to work.

- Target DL is enabled to accept replication from the source DL.
- Source DL is enabled to replicate to the target DL.
- Source NAS share is created with deduplication enabled.
- Target DL is configured with a NAS share with deduplication enabled and having the *same* name as the source share.
- Target NAS share must be configured as a directory/file-based target, with a SyncID the same as the share name, and access must be unlocked.
- Source NAS share must be configured for directory/file-based replication and with a SyncID the same as the share name.
- Do *not* configure “replicate daily”.

Test that the share contents can be replicated between the two DL systems by performing a synchronization on that share through the source DL’s Console Manager (Data Services > Replication > Source Role > NAS). Make sure this works before continuing with the next steps.

NetWorker setup

Configure NetWorker as follows for replication:

- Mount the source share on the storage node. Create the AFTD.
- Label and mount the AFTD, making sure that the only thing on the NAS share is the AFTD.

Authentication setup

The server the replication script runs upon must have SSH configured to log in to the (source) DL, from the command line and without a password prompt. The script will rely on ssh to send the DL the commands to replicate. Using the SSH client program’s password authentication mechanism is not only insecure, but difficult to do via a script. As a best practice, use of public key authentication is preferred. With public key authentication, you can connect securely to the DL from the host server on which the script is run without having to enter a password.

For Windows servers, you may need to use cygwin for ssh communication and openssh (a package available from cygwin for key generation). If you do not have cygwin, you can obtain it from <http://www.cygwin.com>. It is not necessary to install an ssh application on Linux hosts.

Create the keys

This procedure takes you through the steps to create and copy ssh keys from the server from which you will run the post backup script for replication to the DL server.

For Windows servers only:

Perform the following steps on the host server from which you will be running the post backup script:

1. Open a command prompt and type the following:
`c:\cygwin\bin\ssh-keygen.exe -t rsa`
2. You will be prompted to enter a passphrase. Leave this blank.
3. Your public key will be saved as **id_rsa.pub** on the Windows server.

For Linux servers only:

1. Log in as root.
2. Enter the following at the command prompt to generate the public and private keys:
`ssh-keygen -t rsa`
3. Elect to put the keys into the default directory.
4. Press RETURN when prompted for a passphrase.
5. Go to the default directory containing the generated keys and verify that **id_rsa.pub** key exists.

Install the public key on the DL

1. Establish an ssh session with the source DL, logging in as **cliadmin**.
2. Append the `id_rsa.pub` key to the `authorized_keys` file in the `cliadmin` account on the source DL.

Test the keys

1. From a bash shell, test that the public key works by trying to ssh into the DL from your server as `cliadmin` without the password. This may prompt you to trust the host's credential, but not as for a password or a passphrase. For example:

```
$ ssh $DL3D -l cliadmin hostname
```

2. Run the remote hostname once more to confirm you get no prompts at all.

```
$ ssh $DL3D -l cliadmin hostname
```

If this does not work, you'll need to fix it before you can run the script.

NetWorker post-backup script example

The following is a lightweight sample script to drive DL trigger-based replication of a NetWorker Advanced File Type Device from a local (source) DL to a remote one, ensuring consistency of the device, and waiting for the replication to finish before reporting success or failure. It will not synchronize the NetWorker media database or update the database to know about the copy.

To use this script, run it once backups are completed. The script will replicate the whole backup device, and it is important for the device to be in a consistent state. (This example script will fail if the device cannot be unmounted, which requires it be idle.)

Moreover, this script should only be run once deduplication and replication have finished; until the replication is complete, the AFTD will remain unmounted. If this script is invoked too soon, the script will succeed — but the AFTD will be unusable for, perhaps, an extended period of time until it completes. To use the replicated AFTD, since this script does not provide catalog consistency, you must unmount the source AFTD and then mount the target AFTD (so that NetWorker only sees the volume in one place at a time).

The example script expects four command line arguments:

- The name of the NetWorker server
- The name of the (source) DL
- The name of the NAS share on the (source) DL
- The name of the (source) AFTD devices on the NetWorker system

This example script will report errors and info on stderr, and exit with a non-zero exit code on error.

This script has the following high level workflow:

- Unmount the source NetWorker AFTD for consistency
- Trigger a synchronization of the share
- Wait for synchronization to finish
- Remount the source AFTD

This script contains EMC intellectual property and is Copyright 2008-2009 EMC.

```
#!/bin/sh

#This script is an example of how to use EMC DL trigger-based replication in an EMC
NetWorker backup-to-disk environment.
#

SERVER=$1
AFTD=$2
AFTD_RO=$3
DL3D_NAME=$4
DL3D_SHARE_NAME=$5

errmsg_exit()
{
    echo "command failed: $1"
    exit 1
}

assert_set()
{
    if [ "$1" == "" ]
    then
        echo "error: must have $2 set in the environment"
        exit 1
    fi
}

assert_set "$SERVER" "SERVER, the NW server name,"
assert_set "$AFTD" "AFTD, the AFTD name,"
assert_set "$AFTD_RO" "AFTD_RO, the read-only AFTD name,"
assert_set "$DL3D_NAME" "DL3D_NAME, the name of the DL3D,"
assert_set "$DL3D_SHARE_NAME" "DL3D_SHARE_NAME, the name of the share on the DL3D,"

nsrmm -s "$SERVER" -u -f "$AFTD" || errmsg_exit "could not unmount $AFTD"
nsrmm -s "$SERVER" -u -f "$AFTD_RO" || errmsg_exit "could not unmount $AFTD_RO"
ssh -n "$DL3D_NAME" -l cliadmin syscli --sync nas --name "$DL3D_SHARE_NAME" \
```

```

    || errmsg_exit "could not sync $DL3D_SHARE_NAME on $DL3D_NAME"
status="InProgress"
while [ "$status" == "InProgress" ]
do
    sleep 120
    ssh -n "$DL3D_NAME" -l cliadmin syscli --getstatus sync --source > /tmp/status.$$ \
        || errmsg_exit "could not check sync $DL3D_NAME"
    status=`grep ",$SHARE_NAME," /tmp/status.$$ | cut -d, -f 4`
    rm /tmp/status.$$
done
if [ "$status" != "" ]
then
    echo "PANIC: unexpected status for DL3D replication: '$status'"
    exit 1
fi
nsrmm -s "$SERVER" -m -f "$AFTD" || errmsg_exit "could not mount $AFTD"
nsrmm -s "$SERVER" -m -f "$AFTD_RO" || errmsg_exit "could not mount $AFTD_RO"

```

Recovery from a replicated NAS share

When the target DL is configured with a NAS share, NetWorker may be able to recover data directly from the replicated DL. This process assumes that the DL systems have been configured according to the “DL setup” section above.

1. Pause replication on the source DL. Even though the source DL may be unavailable the state of replication may still be “active”.
2. Unmount the share on the source DL.
3. Mount the share on the target DL. NetWorker now recognizes the replicated data and can manage recovery from the target device.