



## DXI-SERIES CORE CONCEPTS EXPLAINED for systems with DXi 2.3 and later software

Email us at [Learning.Admin@quantum.com](mailto:Learning.Admin@quantum.com).

Last update: April 2, 2018



### **NOTICE**

This document may contain proprietary information protected by copyright. Information in this document is subject to change without notice and does not represent a commitment on the part of Quantum. Although using sources deemed to be reliable, Quantum assumes no liability for any inaccuracies that may be contained in this document. Quantum makes no commitment to update or keep current the information in this document, and reserves the right to make changes to or discontinue this document and/or products without notice. No part of this document may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any person other than the purchaser's personal use, without the express written permission of Quantum.

<b>Overview</b>	<b>3</b>
<b>Edge to Core Solutions</b>	<b>4</b>
DXi6000 Series	4
DXi8500	4
DXi V-Series	4
<b>Writing and Accessing Data</b>	<b>5</b>
Presentation Type	5
Network Attached Storage (NAS)	5
Virtual Tape Library (VTL)	5
OpenStorage (OST)	6
Oracle Recovery Manager (RMAN)	6
The Concept of Deduplication	7
Deduplication Terminology	7
How Data Is Stored to Disk	7
Duplicate Blocks	7
Additional OST Support	8
DXi Accent Software	8
Granular Restore Technology (GRT)	8
<b>Replicating Data</b>	<b>9</b>
Replication Requirements	10
Using Replication with NAS and VTL Presentations	11
Replication Types	12
Continuous/Namespace Replication	12
Why Configure Continuous/Namespace Replication?	12
How Do You Enable Continuous Replication?	12
How Do You Schedule Namespace Replication?	12
How Do You Manually Run Namespace Replication?	13
How Do You Access Replicated Data?	13
When Do You Need to Fail Back Data?	13
When Do You Need to Recover Data?	13
File/Cartridge Based Replication	14
When Should You Manually Synchronize Data?	14
Why Configure File/Cartridge Based Replication?	14
How Do You Enable File/Cartridge Based Replication?	14
Replicating Data with the OST Presentation	15
Optimized Duplication	15
Replicating Unique Data and Metadata	16
Accessing Replicated Data	16
Automated Image Replication (AIR)	17
Concurrent Optimized Duplication	17
<b>Moving Data to Tape</b>	<b>18</b>
Backup Application Specific Path to Tape	18
OST Direct To Tape	19
<b>Managing Disk Space</b>	<b>20</b>
Space Reclamation	20
What Is It?	20
When Does It Occur?	21
<b>Encryption and data security</b>	<b>22</b>
Data-in-Flight Encryption	22
What Is It?	22
Data-at-Rest Encryption	23
Secure File Shred	25
Low Capacity Management	26
What Is Low Capacity Management?	26
DXi4000 Series Example	26
<b>Application Environments</b>	<b>27</b>
What is Veeam?	27
Example Veeam Integration with DXi4700 and DXi6900	27
What Is the Dynamic Application Environment (DAE)?	28
Example DAE Integration with DXi4700 and DXi6900	28
<b>Finding Additional Information</b>	<b>29</b>

# OVERVIEW

This document provides core concept information, along with examples, to help you better understand your DXi disk backup system. The core concepts covered in this document are:

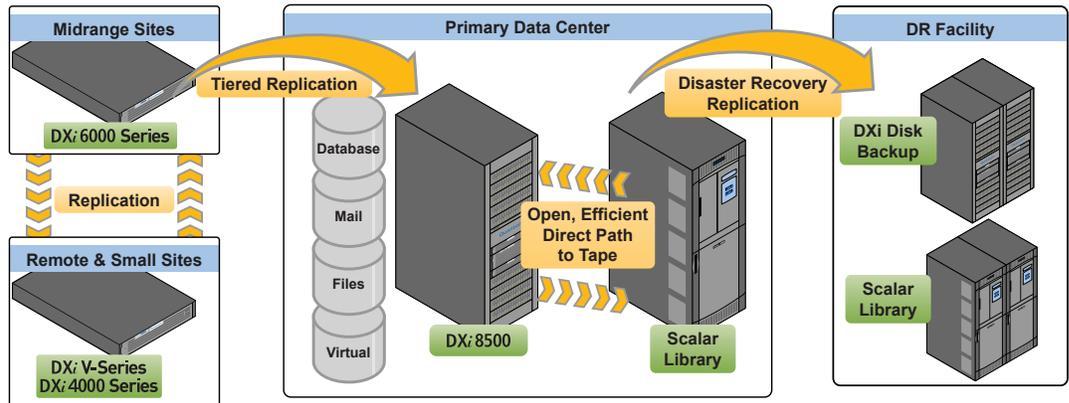
- Quantum's edge-to-core solutions
- Writing and accessing data
- Replicating data
- Moving data to tape
- Managing disk space

The core concepts information in this document applies to DXi V-Series, DXi4000 Series, DXi6000 Series, and DXi8500 systems with DXi 2.3 and later software.



# EDGE TO CORE SOLUTIONS

DXi systems are critical components of Quantum's edge-to-core solutions. The following graphic illustrates where each DXi system fits into these solutions.



## DXi4000 Series

The DXi4000 series systems are designed for small sites and remote offices where customers want to provide disaster recovery (DR) by replicating data to a larger data center.

## DXi6000 Series

DXi6000 Series systems are designed to provide simple, affordable protection for midrange and enterprise sites. DXi6700, DXi6800, and DXi6900 systems can be DR sites for other DXi6700, DXi6800, and DXi6900 installations, or replication sources for a data center/DR site with a DXi8500 system.

## DXi8500

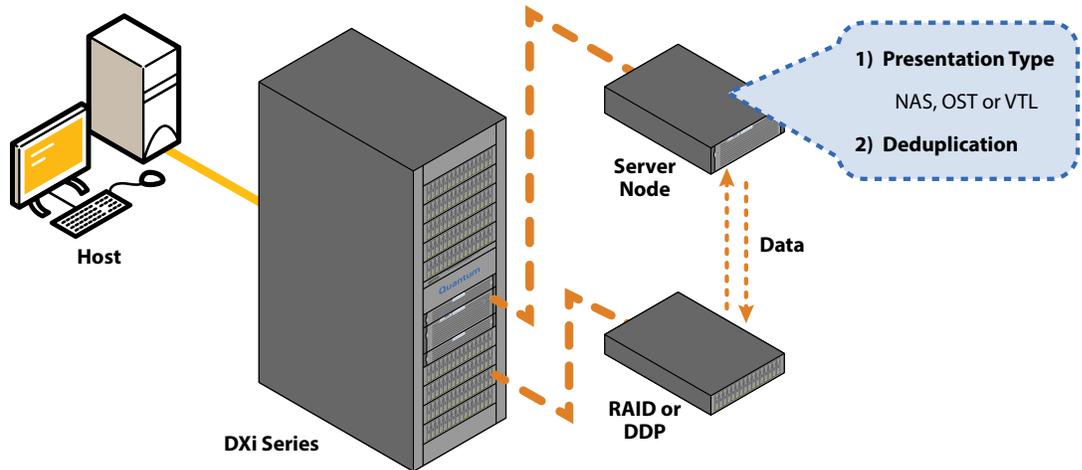
The DXi8500 system offers flexible configuration and scalability. DXi8500 systems are designed to be installed in a corporate data center, providing deduplication at the core.

## DXi V-Series

DXi V-Series systems (DXi V1000 and DXi V4000) are virtual appliances that combine deduplication functionality with the simplicity and flexibility of virtual machines to provide affordable backup and DR protection for any physical or virtual data. In addition, for DR protection, they can replicate to any other DXi appliance.

# WRITING AND ACCESSING DATA

DXi systems let you choose how data will be stored to disk to meet your business storage needs. From the DXi graphical user interface (GUI), you can configure the **presentation type** and choose to enable data **deduplication**. Both of these choices affect system performance, the use of disk space during the write/ingest process, and how the data will be seen from a media server. The following image shows a typical configuration.

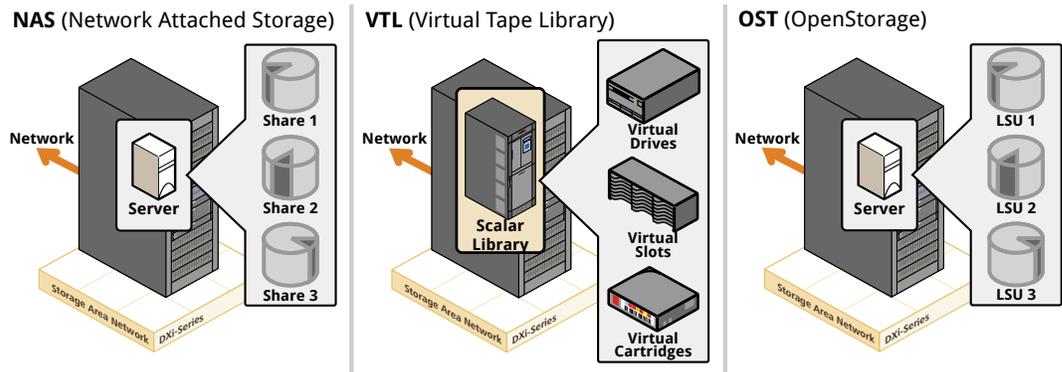


## Presentation Type

As shown in the image above, you can configure any or all of the presentation types that a particular DXi system supports:

- Network Attached Storage (NAS)
- Virtual Tape Library (VTL)
- Open Storage (OST)
- Oracle Recovery Manager (RMAN), a NAS configuration option - [see the RMAN discussion](#)

You can divide a DXi system so that it presents multiple presentations at the same time, as shown below.



## Network Attached Storage (NAS)

The NAS presentation lets you connect the DXi system directly on a LAN as a network resource, with its own network address. If you select NAS and select either the Common Internet File System (CIFS) or Network File System (NFS) as the protocol, you must configure a NAS share on the DXi system so that the DXi can be used as a NAS appliance for backup.

## Virtual Tape Library (VTL)

The VTL presentation allows you to present the DXi system's disk to the backup software to look like one or more virtual tape libraries. A virtual library has the same components as a physical library: virtual drives, bins (slots), media (with barcodes), and changer (robot). All of the components on the DXi system exist in virtual form because they are simulated by software. However, the backup application sees the virtual components as physical tape library components.

## OpenStorage (OST)

The OST presentation lets a DXi system present storage servers to a NetBackup and Backup Exec media server through the OST API. A storage server consists of logical storage units (LSUs), which are similar to directories in a NAS file system, or to tape cartridges in a VTL partition.

The OST presentation requires the Veritas NetBackup (6.5.3 or later) or Veritas Backup Exec 2010 or later host application, and the OST plug-in client, on the media server. Plug-in clients are host-OS dependent and are supplied by Quantum. To use a DXi system in OST mode, you must configure an OST storage server and LSUs on the DXi system. You must also map the LSUs on the media server so that they can perform backups and data can be restored from them. Additionally, you may need to set policies for OST replication (optimized duplication) and OST Direct-to-Tape on the media server.

## Oracle Recovery Manager (RMAN)

This NAS presentation option is available for DXi4700 and DXi6900 systems running DXi 3.2.5 and later software.

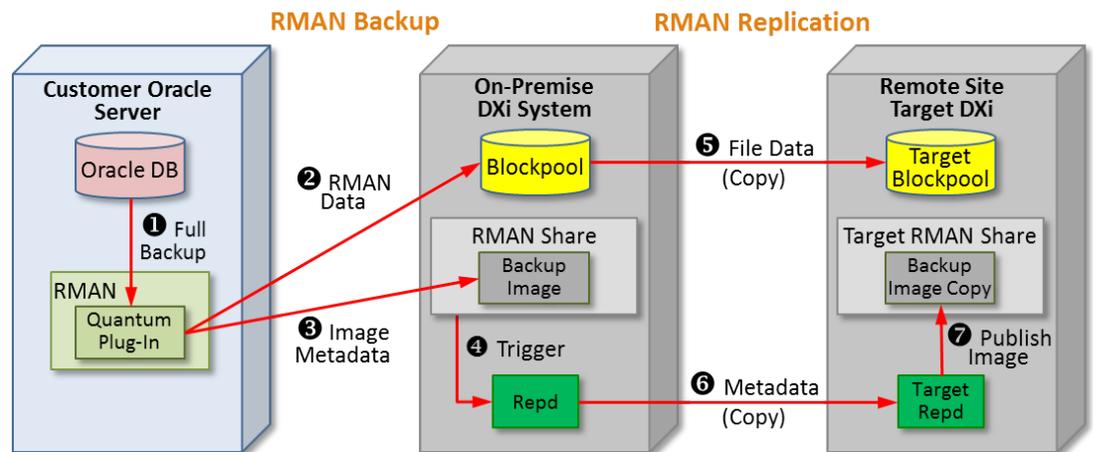
The Oracle RMAN presentation facilitates integration between DXi systems and Oracle servers for backup and recovery of Oracle data. Enabling Oracle RMAN requires a Quantum plug-in to be installed on the Oracle server. NAS shares on the DXi that will be used for Oracle backup must be configured as application-specific RMAN shares.

The Quantum RMAN plug-in provides client-side deduplication, which reduces bandwidth requirements between the Oracle server and the DXi. The plug-in also supports incremental backup and recovery, and replication from a DXi source RMAN share to a DXi target RMAN share.

A simplified representation and explanation of an RMAN implementation is shown below.

### RMAN Backup

On a **1 Full Backup**, the plug-in performs client-side deduplication and sends **2 RMAN Data** to the **Blockpool** and **3 Image Metadata** to the **RMAN Share** in the **On-Premise DXi System**.



### RMAN Replication

After the RMAN backup, if replication is enabled, **Repl** (replication daemon) is **4 Triggered** to replicate the backed-up, deduplicated RMAN data and metadata to a **Remote Site Target DXi**. The **5 File Data** and the **6 Metadata** are copied to the **Target RMAN Share**. Finally, the **Target Repl** **7 Publishes** the replicated **Image** to the **Target RMAN Share**.

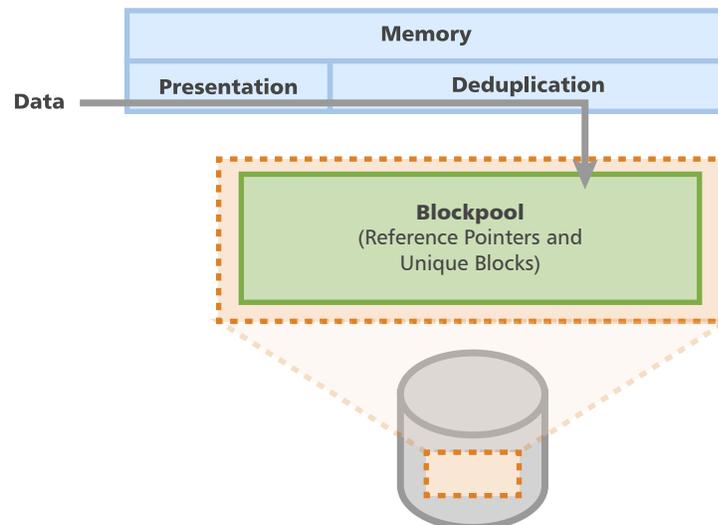
## The Concept of Deduplication

Data deduplication is really a simple concept with very smart technology behind it. With deduplication, a unique block of data is stored only once—for example, in the first file that contains that data. If the same block of data appears again (in the same file or in another file), a pointer to the first occurrence is stored, which takes up less space than storing the entire block of data again.

### Deduplication Terminology

To better understand deduplication, it's important to clearly understand the data path for deduplication, as well as the related terminology. Here are some of the important terms:

- **Presentation:** How the DXi system presents itself to the host application. Supported presentations are NAS, VTL, and OST.
- **Blockpool:** The area of the DXi system storage space that contains the unique blocks of data, along with the reference pointers for deduplicated data written to the DXi system.
- **Reference Pointers:** Pointers that reference unique data written to the DXi system. The pointers take the place of storing the redundant data multiple times and provide the ability to read deduplicated data.
- **Unique Block:** A unique instance of data, which is retained in the blockpool. Redundant data (another occurrence of the same data) is replaced with a reference pointer to the unique data block.



### Data Flow

DXi systems deduplicate data as it is ingested into the DXi appliance. All data is read directly out of the blockpool.

### How Data Is Stored to Disk

Data deduplication recognizes differences at the block level, within files and between files. Quantum's patented deduplication technology segments a stream of data into variable-length blocks and writes those blocks to disk. Along the way, it creates a fingerprint (hash code) for each data segment, and an index of the fingerprints it has seen. The index, which can be recreated from the stored data segments, lets the system know when it sees a new block and when it sees a previously stored block (which already has code in the index), so that it won't store another copy of that block.

### Duplicate Blocks

When the deduplication software sees a duplicate block, it inserts a pointer to the original block in the dataset's metadata, rather than storing the block again. In addition, the blocks that the system finds in the index are not stored again. Only the count of the block usage is incremented. If the same block shows up more than once, multiple pointers are created, which can save large amounts of disk space.

Even if the software sees a repeated block a month or a year after it first encountered the original block (if the original block has not been expired), it recognizes that it has already stored the data and doesn't have to store it again.

### **Additional OST Support**

DXi systems support the DXi Accent and Granular Restore Technology (GRT) features in conjunction with OST.

## **DXi Accent Software**

Quantum's DXi Accent software accelerates backups and reduces network bandwidth requirements. When DXi Accent software is enabled, part of the deduplication process is distributed to the backup server, so that only unique blocks are moved to the DXi system during backup. DXi Accent is invoked during the backup process and is currently used in conjunction with OST.

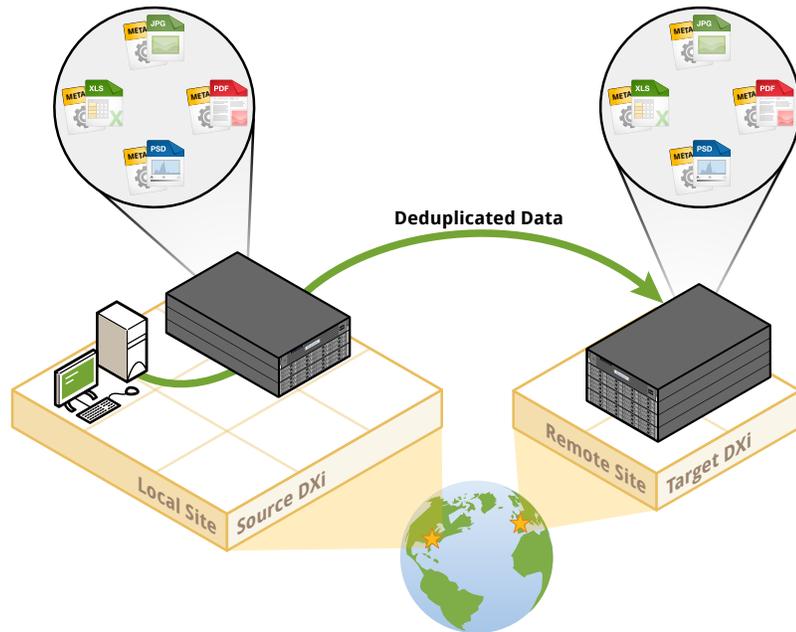
DXi Accent also supports AES (Advanced Encryption Standard) encryption methods. When AES encryption is enabled, all data sent from the media server to the DXi is encrypted. DXi Accent encryption uses the AES cryptographic algorithm with 128-bit or 256-bit length encryption keys. If you desire, you can supply your own digital certificates for use with DXi Accent encryption.

## **Granular Restore Technology (GRT)**

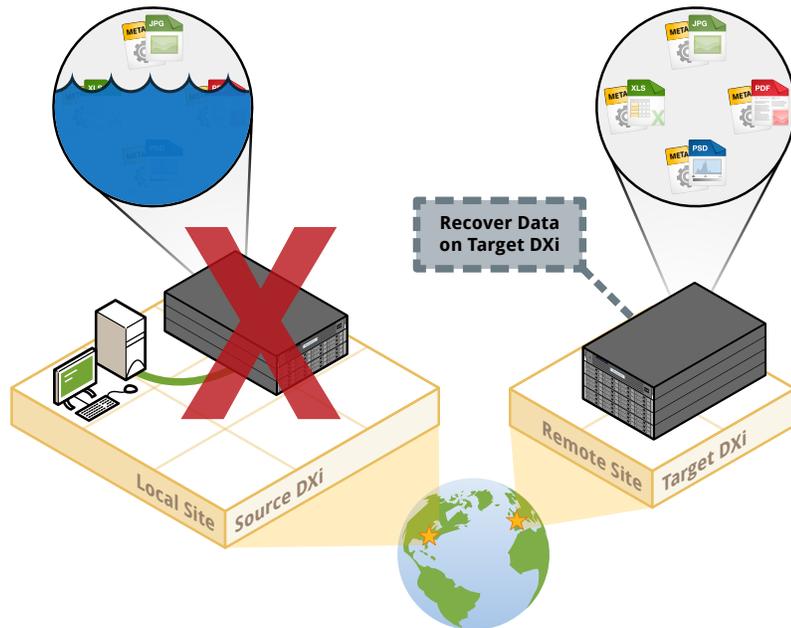
DXi systems support Veritas NetBackup and Backup Exec Granular Restore Technology (GRT). The combination of Quantum's DXi and Veritas GRT provides you with the ability to create a single backup image of a Microsoft application (for example, Exchange, SharePoint, SQL). With GRT, you can restore either the entire application/database or a single component (for example, a single mail message) for rapid restores without requiring additional copies or the use of a staging area, while leveraging the efficiency of DXi deduplication technology.

# REPLICATING DATA

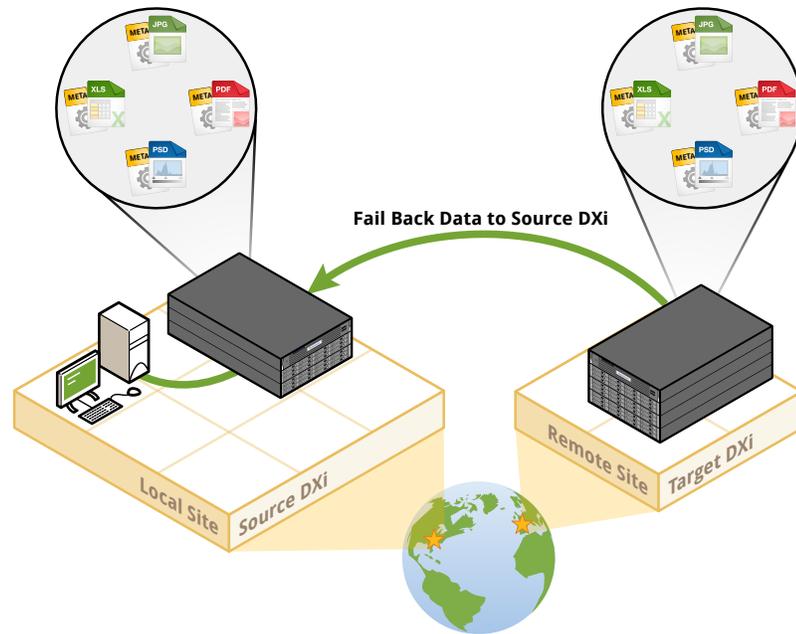
*Data replication* is the process of copying deduplicated data from one DXi system to another for purposes of disaster recovery (DR).



In the event of a disaster in which data on a source DXi system is lost, you can quickly *recover* the replicated data on the remote target system, allowing your business to resume normal operations.



After the source DXi is back in operation at its original location, you can *fail back* data from the target DXi to the source DXi, which restores all data on the source.



## Replication Requirements

Replication requires at least two DXi systems, each with an installed Replication license. One DXi system must serve as a *source* system and the other as a *target*. A DXi system can act as both a source and target system, and you can configure two systems to cross replicate.

On the target DXi system, the source DXi system must be added as a replication source. The target DXi system can be configured to receive replication data from up to 10 different source DXi systems.

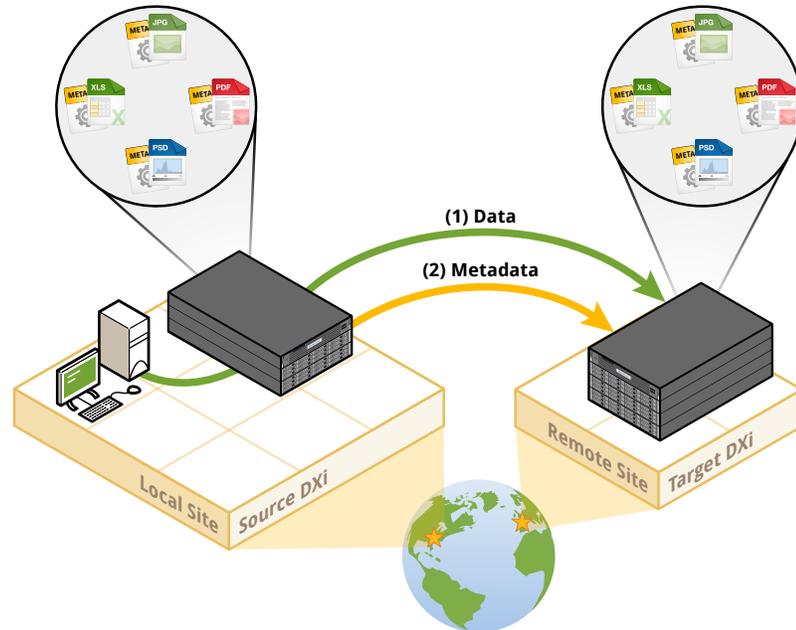
On the source DXi system, the target DXi system must be added as a replication target. A source DXi system can send replicated data to up to two different target DXi target systems.

## Using Replication with NAS and VTL Presentations

To be replicated, VTL partitions and NAS shares must have deduplication enabled. Enabling deduplication optimizes replication speed, since it reduces the size of the data to be replicated.

Replication occurs in two stages:

1. Unique blocks of data are sent from the source DXi to the target DXi.
2. Metadata is sent from the source DXi system to the target DXi system. This allows the target DXi to reflect the current data state of the NAS share or VTL partition on the source DXi.



## Replication Types

Two types of replication can be configured:

- Continuous/Namespace Replication (also referred to as simply *replication*)
- File/Cartridge Based Replication

These are described in the following sections.

## Continuous/Namespace Replication

*Continuous/namespace replication* occurs when replication is enabled for a deduplicated NAS share or VTL partition and a replication schedule is configured (or manual replication is performed on a regular basis). For Continuous/namespace replication to occur, the source DXi system must be configured to send data to the target DXi system. Similarly, the target DXi system must be configured to accept data from the source DXi system.

To optimize the process, deduplicated data is *continuously* sent in the background from the source DXi to the target DXi. However, a snapshot that preserves the file structure (*namespace*) of your data is sent to the target system only when a scheduled or manual replication job occurs. A snapshot contains all of the metadata that is necessary to recreate a share or partition, just as it was at the point in time when the snapshot was created.

A saved snapshot is necessary to recover your data at a later time. For this reason, it is not enough to simply enable replication for a share or partition. You must also configure a replication schedule (recommended) or perform manual replication on a regular basis to send snapshots of the share or partition to the target DXi.

## Why Configure Continuous/Namespace Replication?

Configure Continuous/Namespace Replication for disaster recovery purposes. If a disaster occurs and your source DXi system is no longer available, the point-in-time snapshots that you created will allow you to recover a NAS share or VTL partition to a previous state from the replicated data on the target DXi.

After your source DXi system is available again, you can fail back a snapshot to your source DXi system and recover the data.

## How Do You Enable Continuous Replication?

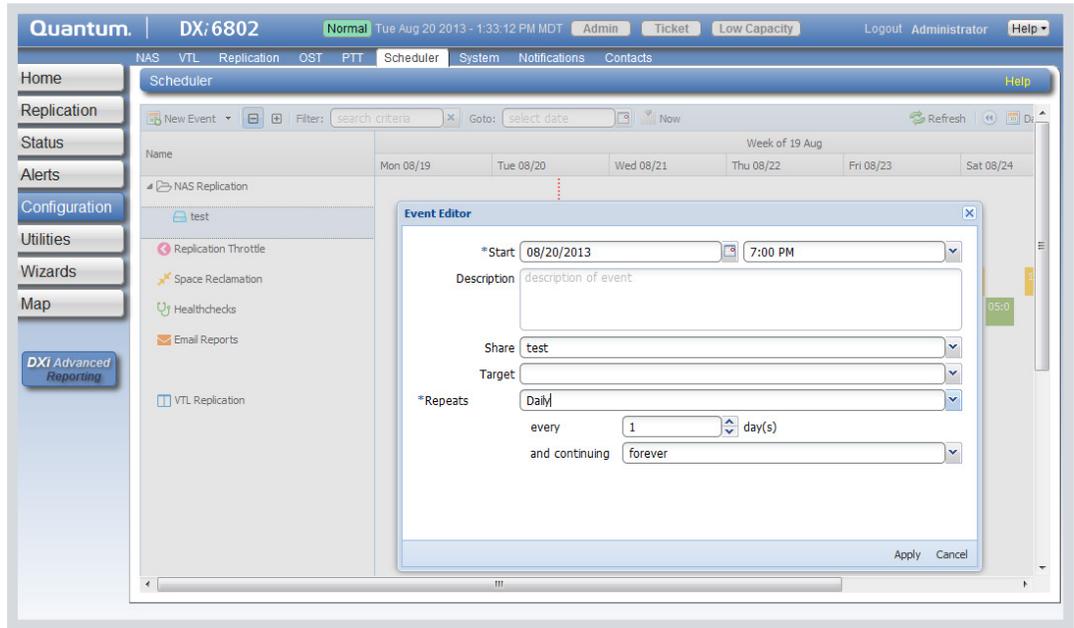
You can enable Continuous Replication when you add or edit a NAS share or VTL partition. The following example shows how to enable Continuous Replication for a NAS share on the **Edit NAS Share & Replication Settings** page (**Configuration > NAS > Summary**). The process is similar for a VTL partition.

The screenshot shows the Quantum DXi 6802 web interface. The top navigation bar includes 'NAS', 'VTL', 'Replication', 'OST', 'PTT', 'Scheduler', 'System', 'Notifications', and 'Contacts'. The main content area is titled 'Edit NAS Share & Replication Settings'. On the left, there is a sidebar with navigation options: Home, Replication, Status, Alerts, Configuration (selected), Utilities, Wizards, and Map. The 'Configuration' section is expanded to show 'NAS Share Settings' and 'Replication Settings'. The 'NAS Share Settings' section includes fields for Name (test), Description, and Options (Enforce read-only access, Hide from network browsing). The 'Replication Settings' section is divided into 'Send' and 'Receive' sections. The 'Send' section has a checked checkbox for 'Enable replication to the replication target DXi', which is highlighted by a red arrow and labeled 'Enables Continuous Replication'. Below this, there is a table for 'Replication Target DXIs' with columns for 'Replication Target DXIs', 'Status', and 'Encryption'. The table shows one entry: '10.20.86.25' with 'Ready' status and 'AES 256-bit' encryption. The 'Receive' section has a checkbox for 'Enable Directory/File Based replication to this DXi' which is unchecked, and an 'Access' dropdown set to 'Unlocked'.

## How Do You Schedule Namespace Replication?

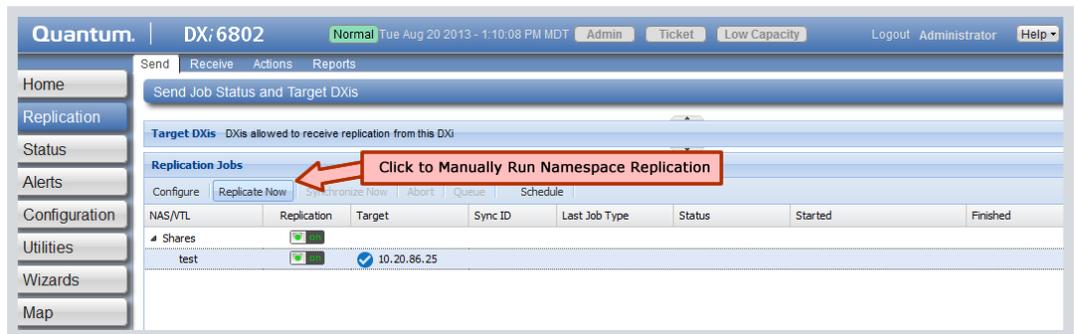
To schedule Namespace Replication, go to the **Scheduler** page (**Configuration > Scheduler**), expand either **NAS Replication** or **VTL Replication**, and select your share or partition.

In the **Event Editor** screen, enter a **Description** and select options for the other fields.



### How Do You Manually Run Namespace Replication?

To manually run Namespace Replication, go to the **Send Job Status and Target DXis** page (**Replication > Send**), select the NAS share or VTL partition, and click **Replicate Now**.



### How Do You Access Replicated Data?

When Continuous/Namespaces Replication is enabled and running, two copies of the data exist. The original data resides on the source, and a replicated copy resides on the target. You can recover the data by using the Failback and Recover processes.

### When Do You Need to Fail Back Data?

If the replicated data resides on the target, and if some data on the source becomes corrupted, destroyed, or modified in an undesirable way, you can use the Failback option to copy the replicated data on the target back to the source.

**The Process:** Failing back data copies the data, and the metadata, from a target system to a source system. However, after the data and metadata have been copied, they are not accessible until the Recover process is run on the failed back NAS share or VTL partition.

### When Do You Need to Recover Data?

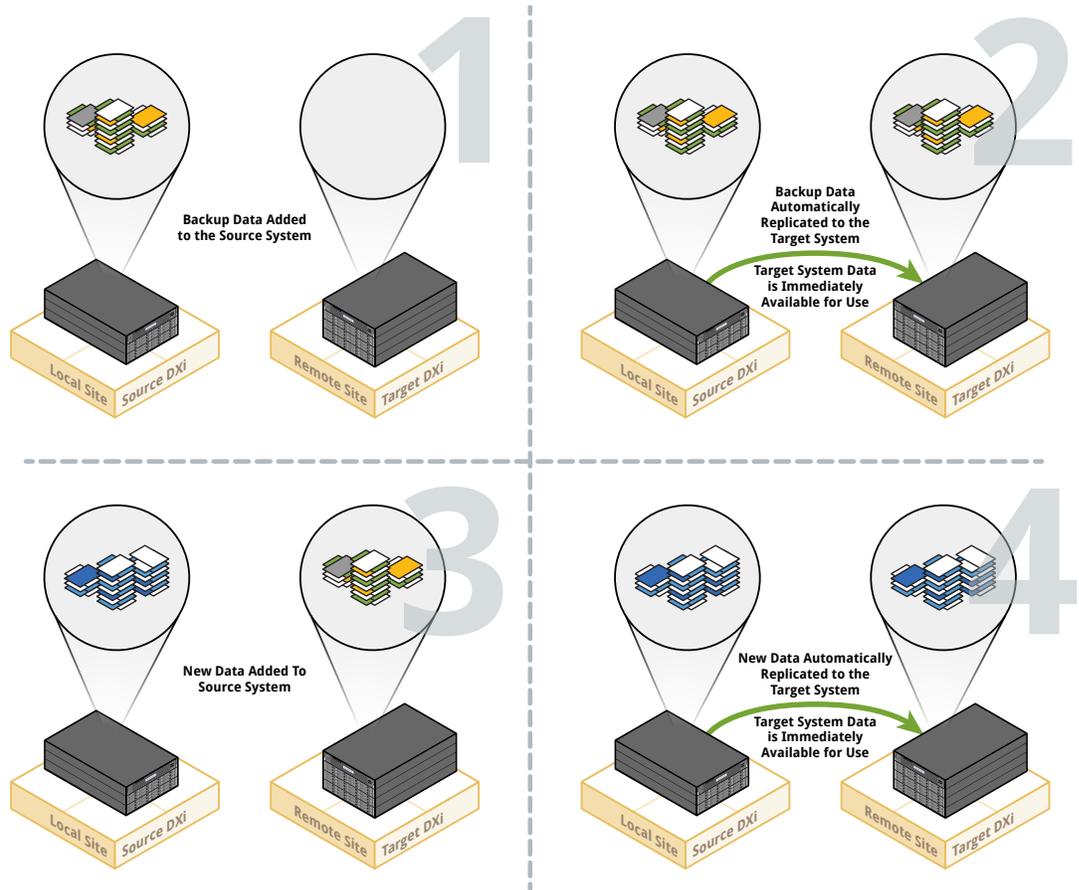
If you want to make the replicated data on the target accessible, you would recover data. For example, if your source becomes unavailable and you have a replicated copy on a target system, you could recover the data that's on the target. You would also need to recover data after failing back that data to the source, in order for that data to be accessible.

**The Process:** Recovering the data creates a copy of the metadata in a new share or partition, so that the data is accessible.

## File/Cartridge Based Replication

Like Continuous/Namespacing Replication, *File/Cartridge Based Replication* sends data from a NAS share or VTL partition on a source DXi system to a target DXi system, where it can be accessed. However, with File/Cartridge based replication, replication and recovery occur automatically after a cartridge is unmounted in a VTL partition, after a file is closed in a CIFS share, or after a certain period of time in an NFS share. The key is that this occurs *automatically*, meaning that scheduling or manual intervention is *not* needed.

For example, if files are deleted on the source DXi system, they will automatically be removed on the target DXi system. When new files are added to the source DXi system, they will automatically replicate to the target system. After automatic recovery, the data is immediately available for use on the target DXi system.



### When Should You Manually Synchronize Data?

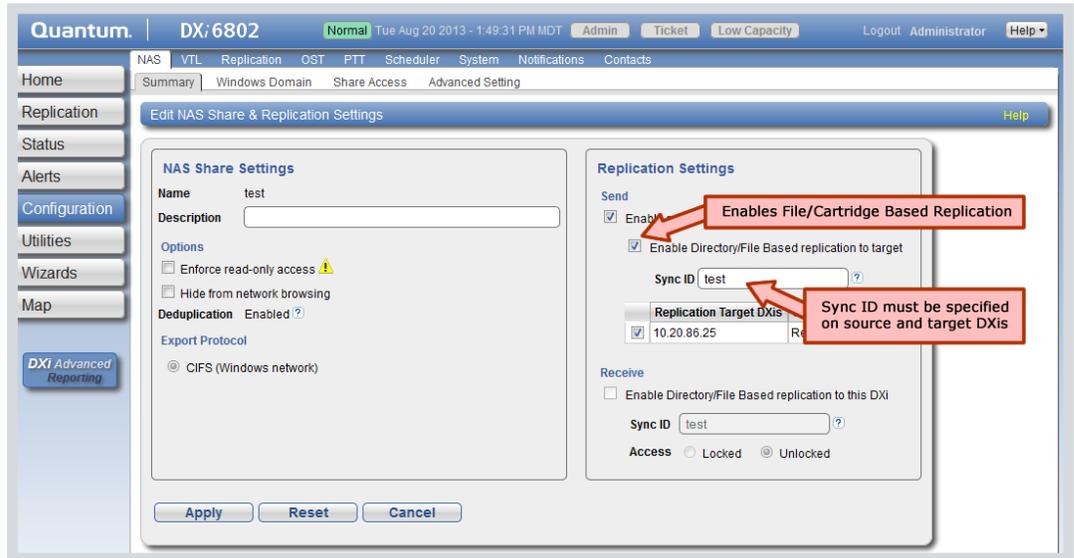
You can manually synchronize data between the source DXi system and target DXi system, if needed. If a File/Cartridge Based Replication job fails or is disabled for a period of time, manually synchronizing data brings the source DXi system and target DXi system into *agreement*.

### Why Configure File/Cartridge Based Replication?

With File/Cartridge Based Replication, data from your source DXi is automatically replicated and recovered on your target DXi. This means that data on the source and target is always in sync. The replicated data on the target DXi is immediately available for use. You don't have to take the extra step of recovering your data from a snapshot.

### How Do You Enable File/Cartridge Based Replication?

You can enable File/Cartridge Based Replication when you add or edit a NAS share or VTL partition. The following example shows where you can enable File/Cartridge Based Replication for a NAS share on the **Edit NAS Share & Replication Settings** page (**Configuration > NAS > Summary**).



## Replicating Data with the OST Presentation

DXi systems can replicate (duplicate) OST data to another DXi using the following methods:

- Optimized Duplication
- Automatic Image Replication (AIR)
- Concurrent Optimized Duplication

These methods are covered below.

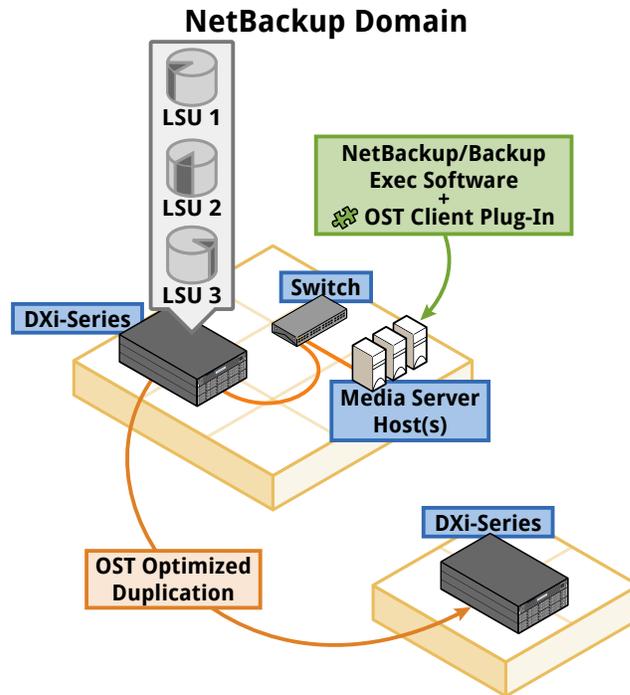
### Optimized Duplication

OST replication (referred to by Veritas as *Optimized Duplication*) is used with OST presentations. This replication mode is specific to the Veritas OST API supported by NetBackup 6.5.3 and later and Backup Exec 2010 and later—it does not work with other backup applications.

Veritas backup applications use this functionality to initiate optimized duplication backup images between systems in the same domain. The backup applications can duplicate/replicate to multiple targets. In this case, the element that is replicated is a logical storage unit (LSU) defined by OST and the DXi system. The backup applications manage the replication/duplication process. In other words, the movement of unique blocks from one LSU to another is initiated by the backup application.

Replication occurs in the background and uses Quantum's deduplication capabilities (by sending only unique data blocks) to reduce the bandwidth requirements. Replication to the target DXi system is still initiated, managed, and controlled by the media server, while the actual data copy is offloaded to get the maximum benefits from the DXi system's replication capabilities.

Data is replicated at the backup image level and the image copies are tracked by the NetBackup or BackupExec catalog. This allows data to be recovered from any available copy. Policies can be set on the master server, which initiates and automates the duplication.



#### Replicating Unique Data and Metadata

Blocks are not replicated until the DXi system receives a cue from the NetBackup or Backup Exec server. The source and target (or targets, since one source can be set to replicate to more than one target) are set up in policies and are not defined inside the DXi system. The namespace is sent, along with the unique blocks, on a file-by-file basis.

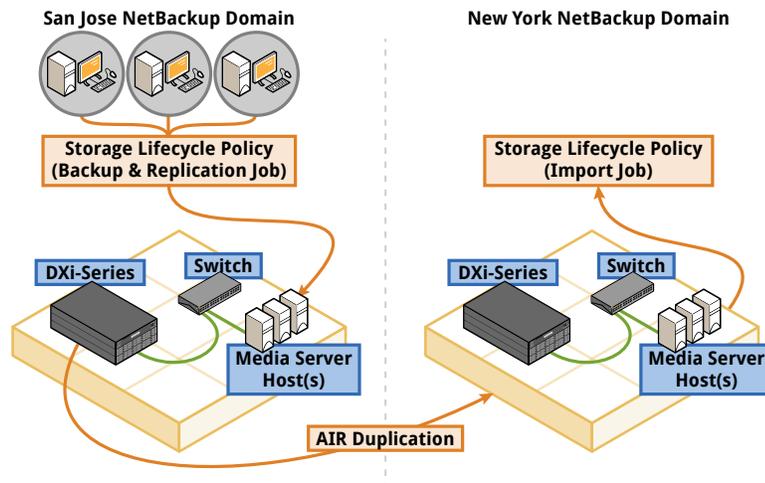
#### Accessing Replicated Data

Data is accessible from the NetBackup or Backup Exec local media server and master server. The catalog identifies all instances as holding the same data, so the data can be read directly by NetBackup or Backup Exec from any location.

## Automated Image Replication (AIR)

If you are using Veritas NetBackup 7.1 or higher, with DXi 2.2.x or later software, you can configure an LSU for Automatic Image Replication. When Automatic Image Replication is enabled, data on an LSU is automatically replicated to a remote LSU that resides on a DXi in a different NetBackup domain. The timing of the duplication, and the backup images that are duplicated, are determined by the storage lifecycle policies (SLPs) that you configure in NetBackup.

It is important to remember that with Automatic Image Replication, the local and remote LSUs reside in different NetBackup domains. This differs from optimized duplication, which occurs between two LSUs that reside within the same NetBackup domain.

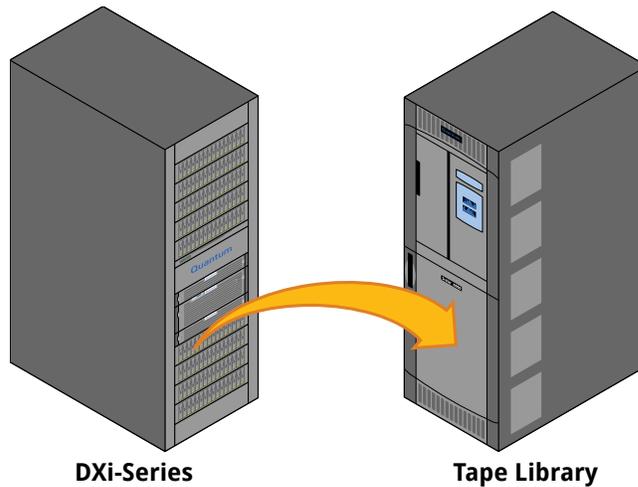


## Concurrent Optimized Duplication

For both optimized duplication and Automatic Image Replication, you can optionally enable Concurrent Optimized Duplication. When this feature is enabled, as data is written to the storage server, it is simultaneously replicated to the target DXi system. Then, when optimized duplication or Automatic Image Replication subsequently occurs, the operation is more efficient, because a portion of the required data has already been replicated to the target server.

# MOVING DATA TO TAPE

DXi systems can move data directly from the DXi disk to physical tape cartridges in an attached tape library without sending data through a backup application media server. The tape cartridges can then be stored offsite as part of your DR plan. This is called *Path to Tape (PTT)* or *Direct to Tape*.



DXi systems support two types of Path to Tape functionality:

- Backup Application Specific Path to Tape: Supported on DXi systems with the VTL presentation type configured.
- OST Direct to Tape: Supported on DXi systems with OST presentation type configured. This option is specific to Veritas and the OpenStorage API.

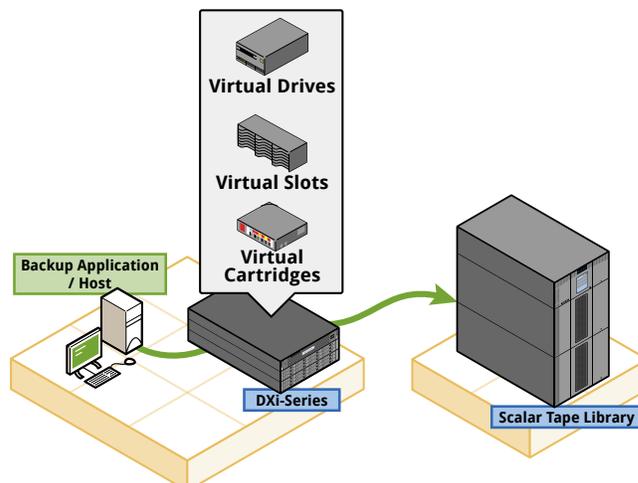
## Backup Application Specific Path to Tape

In Backup Application Specific PTT, data on the DXi system's virtual cartridges is written to physical media in a physical library that is directly connected (through NDMP), bypassing the media server. In this case, the backup application controls the process of copying the data, and it keeps track of where the data resides. In other words, the backup application manages the copies of data.

In the following Backup Application Specific PTT example, the backup server is backing up data to a VTL partition on the DXi system. A Quantum Scalar library is directly attached to the DXi and is configured for Backup Application Specific PTT.

With this configuration, the backup server can direct the DXi to duplicate the backup images stored via VTL partition to physical media in the Scalar library.

The backup application is aware of both copies of the backup images and can recover data from either location.



In Backup Application Specific PTT, the physical tapes contain a copy of the data from the disk copy, but they do not hold images of the virtual cartridges. This means that the virtual cartridges and physical cartridges will have different barcodes, and that they could differ in media type and cartridge count (for example, data held on one virtual DLT cartridge might be written to two physical LTO cartridges). The backup application tracks the data in both locations across the different media.

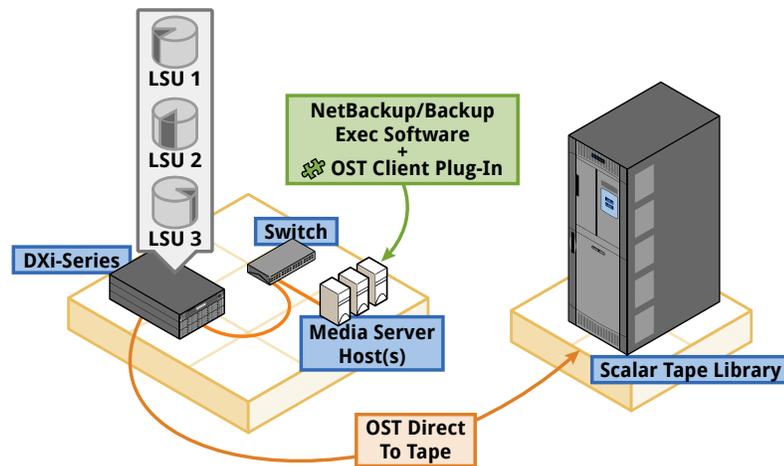
## OST Direct To Tape

For systems using the OST API with Veritas NetBackup, a different direct PTT option is available. Quantum refers to this option as *OST Direct to Tape*. This option works much like Backup Application Specific PTT, except that the source data is the data in an OST LSU on the source DXi system, not in a VTL partition.

To use this option, the OST plug-in must be installed on the media server, and a physical tape library must be directly connected to the DXi system through an NDMP connection.

After the backup data is stored on a DXi system, a direct-to-tape command is initiated using policies set in NetBackup. This command triggers the DXi to send the data over a Fibre Channel link to drives in a physical tape library, where NetBackup-formatted media is created.

NetBackup controls the media export, but the data is copied directly from the DXi system to tape. The location of the data (on disk and on tape) is maintained in the NetBackup catalog.



# MANAGING DISK SPACE

Do you know how much disk space you need to back up your conventional disk and tape systems? You probably have a general idea of your requirements, based on your experience with normal compression systems. Deduplication has similar results. Although customers' experience varies, some see a reduction in disk space needs of 90% or more.

The best way of predicting your storage needs is to consult with your Quantum sales engineer or authorized reseller. Their experience and product-specific sizing tools will help you make an informed decision on a deduplication solution.

As with any backup system, whether backing up to disk or tape, a DXi system requires normal capacity management. This consists of removing data that is no longer needed, and reclaiming space for new datasets.

In general, capacity management is coordinated through the backup application, as it would be for any other backup system. However, Quantum DXi systems are designed to provide *automated space management*. This means that the system alerts you when specific thresholds are exceeded, to let you know when normal space management actions are needed.

It's important to understand two key space management concepts:

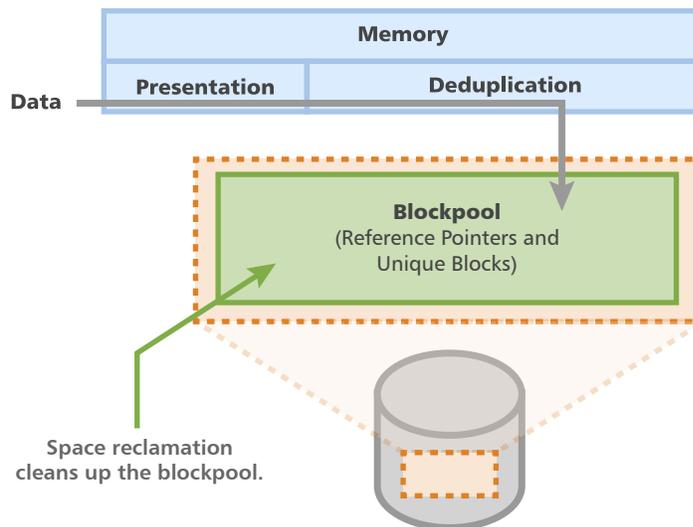
- Space reclamation
- Low space management

These are covered in the following sections.

## Space Reclamation

### What Is It?

*Space reclamation* is the process used to delete blocks that are no longer referenced by metadata, returning capacity to the free space pool (within the blockpool) for future reuse. In simple terms, space reclamation cleans up the blockpool.



Space reclamation can include up to four stages:

- **Stage 1: Reclaim Disk Space (Compaction).** The DXi checks if there is any unfinished compaction work from a previous space reclamation operation and performs it first. This is identical to Stage 4.
- **Stage 2: Calculating Deletion Candidates.** The DXi dumps the reference pointers that refer to the unique blocks and determines what work needs to be done.
- **Stage 3: Deleting New Candidates.** The DXi decrements the reference counts of the unique blocks in the blockpool for the data being deleted (NAS / OST) or overwritten (VTL).
- **Stage 4: Reclaim Disk Space (Compaction).** Unique blocks with zero reference pointers are removed from the blockpool and the remaining data is compacted, thereby creating space for new unique data.

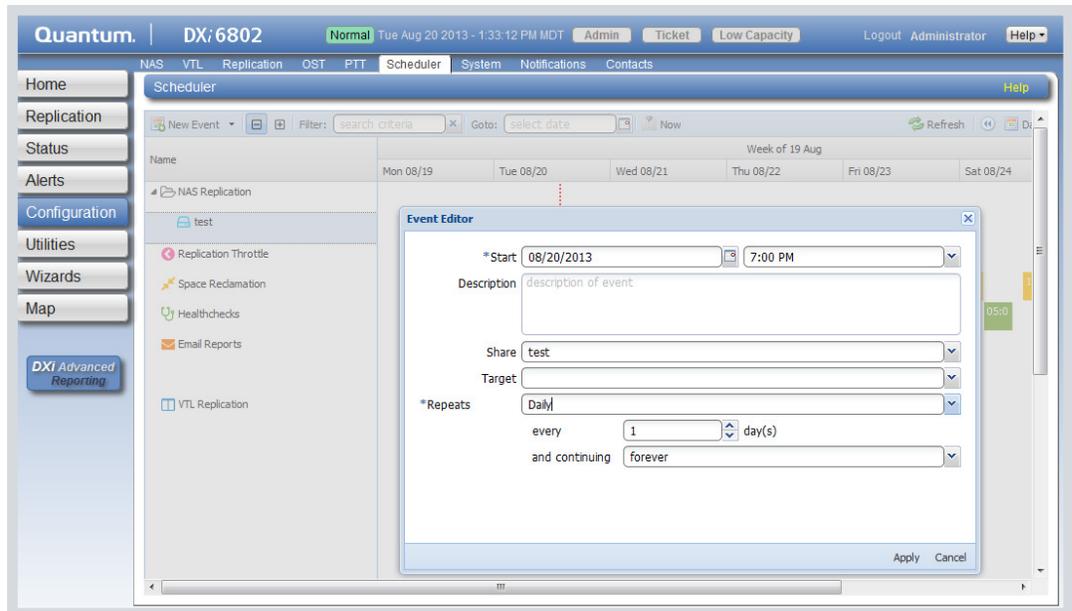
Depending on which space reclamation mode is initiated, not all space reclamation stages are run. The two space reclamation modes are described below:

- **New / Normal Mode:** To increase performance, when space reclamation is initiated manually on the **Space Reclamation** page or as a scheduled event, only Stage 2 (Calculating Deletion Candidates) and Stage 3 (Deleting New Candidates) are run. Stage 1 and 4 are not required in Normal mode because the DXi can automatically compact reclaimable space as needed and use it to store new deduplicated data.
- **Low Capacity / Legacy Mode:** When the DXi enters Low Capacity mode, space reclamation is automatically started, to free up disk space. In this case, all four stages of space reclamation are run.

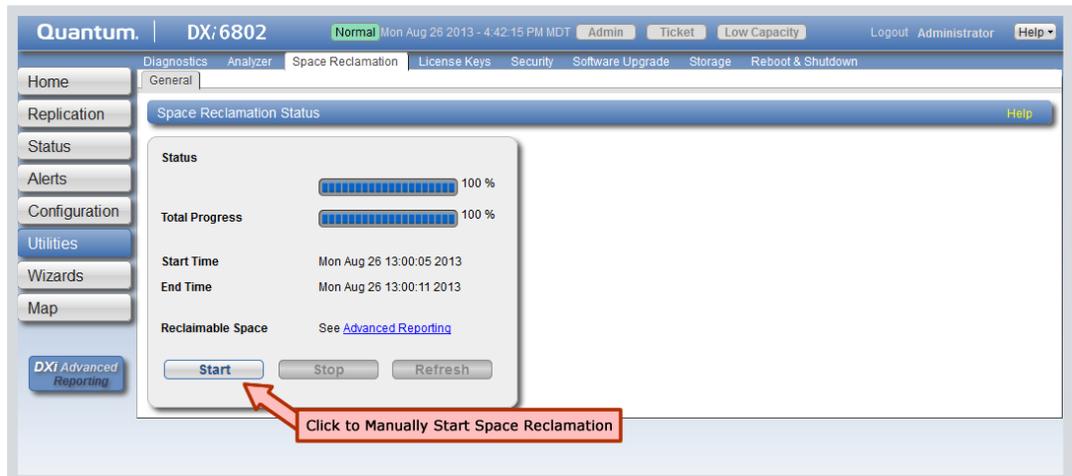
### When Does It Occur?

Space reclamation is scheduled to run daily by default. Running it on a daily basis is recommended, to maximize performance and capacity utilization. It is far more efficient to process a day's worth of new data than a week's worth.

You can modify the space reclamation schedule using the **Scheduler** page (**Configuration > Scheduler**) on the DXi GUI, as shown below. Schedule space reclamation to start approximately two hours after your backup job has completed.



Space reclamation can also be run on-demand from the DXi GUI on the **Space Reclamation Status** page (**Utilities > Space Reclamation**) by clicking **Start**, as shown below.

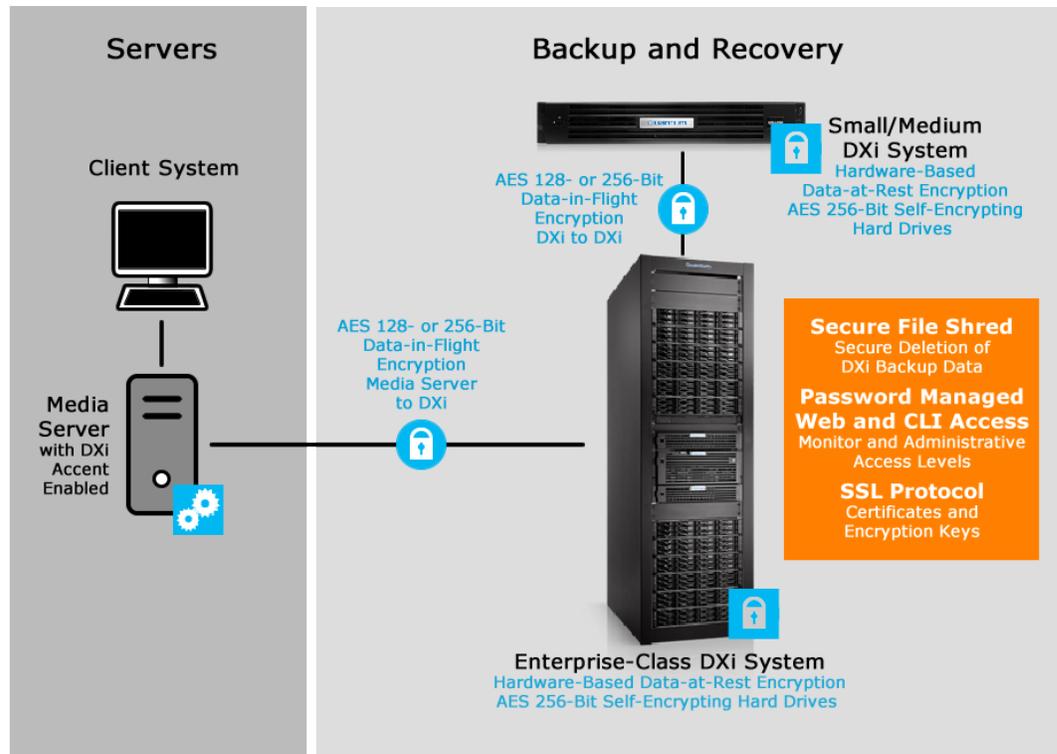


# ENCRYPTION AND DATA SECURITY

Data encryption and security features on DXi systems ensure that backup and replication data cannot be intercepted in transit, and that deduplicated data backed up to a DXi cannot be accessed in any way other than through the DXi on which it was originally stored.

DXi systems support the following encryption features:

- **Replication and Failback Encryption:** User-configurable 128- or 256-bit encryption.
- **Data-in-Flight Encryption:** DXi 3.0 and later software. Deduplicated data sent from the media server to a DXi system or replicated data sent between DXi systems can be secured with a license-enabled AES 128- or 256-bit encryption algorithm.
- **Data-at-Rest Encryption:** DXi 2.2 and later software. Self-encrypting hard drives can perform AES 256-bit encryption on deduplicated, compressed data with no system performance impact.



In addition to encryption, DXi systems support the following data security features:

- **Secure File Shred:** DXi 2.2 and later software. Securely delete backup images. See the [Secure File Shred](#) topic for more information.
- **Password Managed Web and Command Line Interface Access:** Set passwords for the Monitor and Administrative levels of access, and disable Command Line access.
- **SSL Protocol for Server Connections:** Configure certificates and encryption keys for data transmission.

## Data-in-Flight Encryption

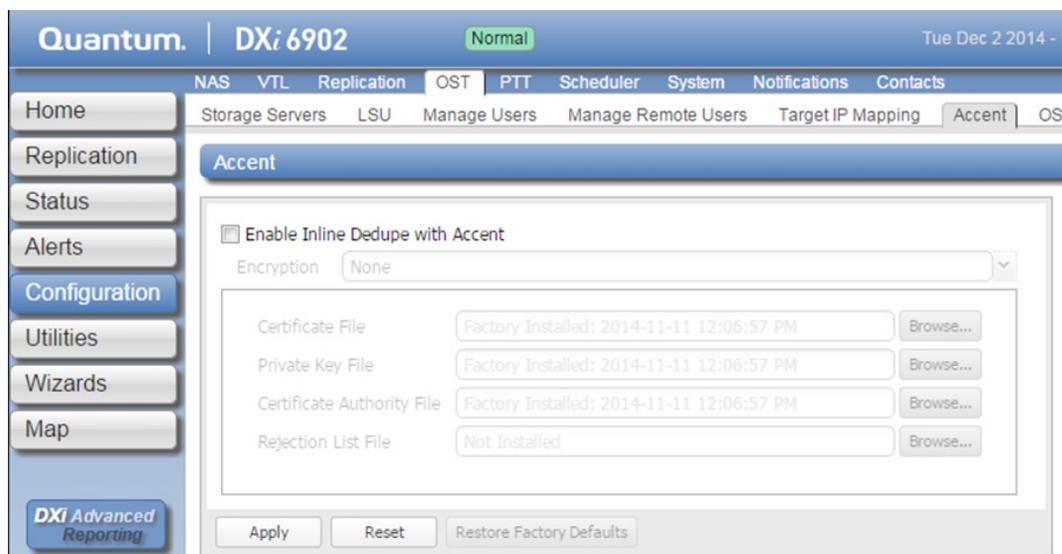
### What Is It?

Replication and Failback encryption have always been available on DXi systems. In DXi 3.0 software, Data-in-Flight Encryption became a license-enabled feature that uses 128-bit or 256-bit AES encryption to secure backup data when it is in transit between a media server and a DXi. It also secures replication data in transit between DXi systems. The AES encryption options are available only when the Data-in-Flight license is installed. This encryption is not available in restricted regions.

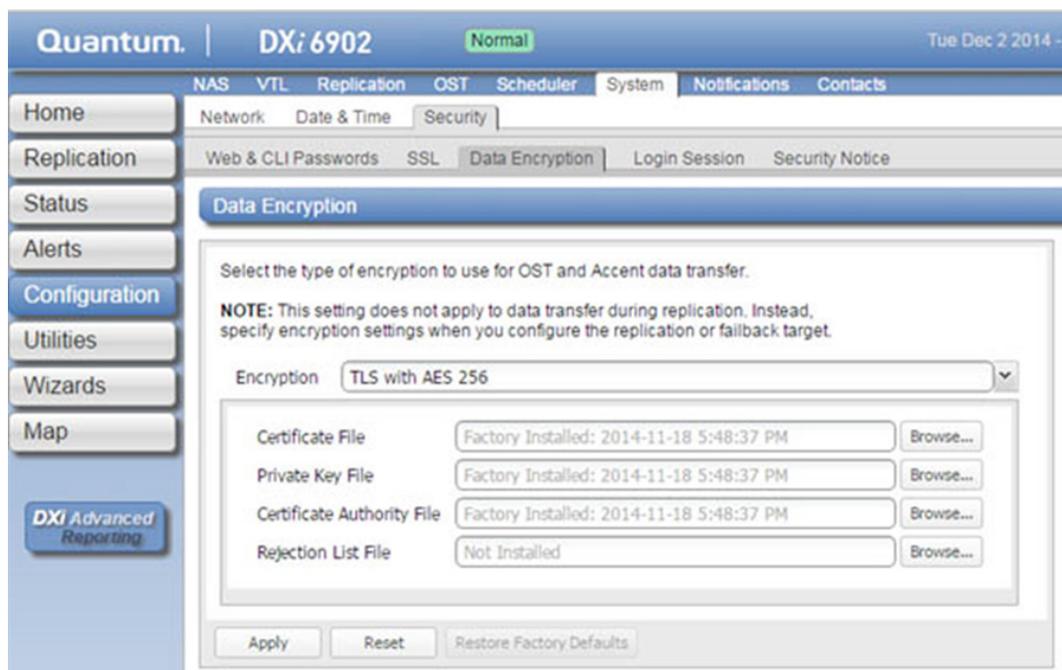
The Advanced Encryption Standard (AES) is a U.S. Government (NIST) standard for electronic data. The designations of 128-bit and 256-bit are the length of the encryption/decryption key.

After the Data-in-Flight Encryption license is installed, you can configure the OST DXi encryption settings from the appropriate GUI page for DXi 3.0 and earlier software or for DXi 3.1 and later software.

## Accent Encryption Settings in DXi 2.0.2 to DXi 3.0 Software



## OST and Accent Encryption Settings in DXi 3.1 and Later Software



## Data-at-Rest Encryption

Data-at-Rest Encryption uses Self Encrypting Drive (SED) technology to secure all hard drives in a DXi system so that, if they are removed from the DXi, they cannot be read using any other system or device. This encrypts all file data and metadata, configuration files, and the DXi software and operating system. Self-encrypting drives and Data-at-Rest Encryption are available on DXi4700, DXi6800, DXi6900, and DXi8500 with 3-TB drives. This encryption is not available in restricted regions.

To enable Data-at-Rest Encryption, all drive controllers and hard drives in the system must support SED technology. The Data-at-Rest license must also be installed, and you are asked to supply a passphrase that the DXi uses to generate an encryption key. The passphrase ensures that all physical disks are paired with their respective controllers, and that data can only be read back from the disk by the same controller that wrote it. If a controller must be replaced, the passphrase is required to enable the new controller to access the data on the physical disks.



+ License + Passphrase → Enable Data-at-Rest Encryption

### Self-Encrypting Drives

#### Important

After you enable Data-at-Rest Encryption, you cannot disable it or turn it off. Make sure to back up your passphrase and recovery files, because they may be required for future capacity expansion or for certain service scenarios.

#### Example Data-at-Rest Configuration

Quantum. | DXi 4700 | Normal | Fri Jan 24 2014 - 9:04:42 AM PST | Admin | Ticket | Low Capacity | Logout | Administrator | Help

NAS | Replication | OST | Scheduler | System | Notifications | Contacts

Home | Network | Date & Time | Security | Web & CLI Passwords | SSL | Login Session | Security Notice | Data-at-Rest | Help

Data-at-Rest

Data-at-Rest Encryption

Current Data-at-Rest Encryption **Disabled** | Check Status

Enable Data-at-Rest Encryption

\* Passphrase |

\* Confirm Passphrase |

Enable

\* Required field

Download/Send Data-at-Rest Recovery File

Download |  Email

\* Email Recipient |

\* Password |

\* Confirm Password |

Apply

\* Required field

## Secure File Shred

Secure File Shred is a DXi utility that securely and permanently erases sensitive data stored on a DXi system with 2.2 and later software. During secure shred, all residual data associated with deleted NAS files or VTL cartridges is securely erased from the DXi disk drives by performing a single-pass overwrite with zeros.

The secure shred process can take multiple days to complete, and while it is running, the DXi operates in limited mode. This stops backups, restores, and all scheduled jobs, such as replication. It is possible to cancel secure shred before it completes and return the DXi to normal operation.

### Example Secure File Shred Utility Page

The screenshot shows the Quantum DXi 4700 management interface. The top navigation bar includes 'Quantum. | DXi 4700', a status indicator 'Normal', the date and time 'Mon Jan 27 2014 - 8:20:38 AM PST', and user roles 'Admin' and 'Tick'. Below the navigation bar, a menu contains 'Diagnostics', 'Analyzer', 'Space Reclamation', 'License Keys', 'Secure Shred', and 'Software Upgr'. A left-hand sidebar lists navigation options: 'Home', 'Replication', 'Status', 'Alerts', 'Configuration', 'Utilities' (highlighted in blue), 'Wizards', and 'Map'. The main content area is titled 'Secure File Shredding' and contains a warning box with the following text: 'The Secure Shred process can take days to complete. The system must stop all jobs during this time. The system will reboot during this process.' Below the warning box, there are four rows of data: 'Last invoked' (None), 'Last successful' (None), 'Last result' (NA), and 'Last error string' (None). A 'Start Shred' button is located at the bottom of the main content area. At the bottom left of the page, there is a button labeled 'DXi Advanced Reporting'.

# Low Capacity Management

## What Is Low Capacity Management?

*Low capacity management* is the automated process initiated by the DXi system to continue operations as the system fills up, and to maintain continued access to stored data.

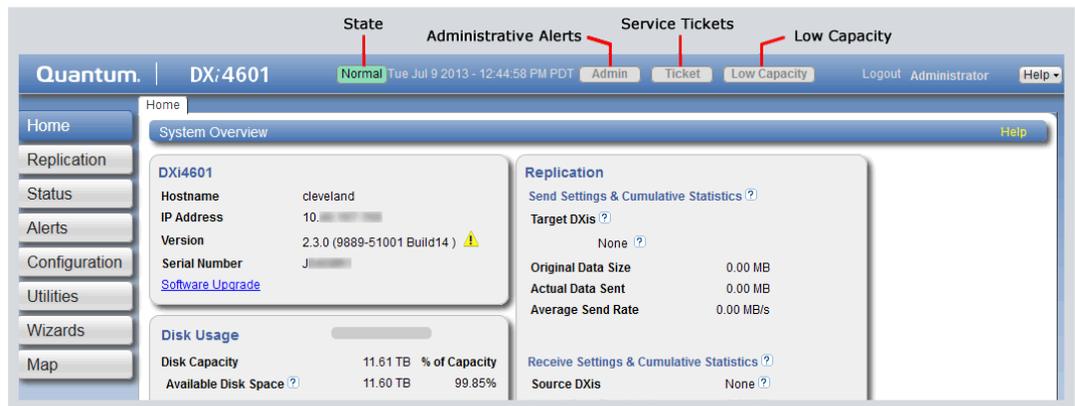
### What Is the Impact?

The thresholds that trigger the Low Capacity condition are different for each DXi model. Threshold examples for a DXi4000 Series system are provided below.

### DXi4000 Series Example

As disk capacity is used and free disk space approaches low levels on a DXi4000 system:

- The **Low Capacity** condition indicator turns on in the DXi GUI banner.
- An Administrative Alert and a Service Ticket are generated each time free disk space drops below one of the following threshold values:
  - o **850 GB** - The State indicator on the system banner displays **Attention** and the system enters the Low Threshold state. Ingest continues, and space reclamation starts.
  - o **250 GB** - The State indicator on the system banner displays **Low Space**. Ingest stops, and target replication to the system is paused, but data can still be read.
  - o **10 GB** - The State indicator on the system banner displays **No Space**. All data read/write activity is stopped.



# APPLICATION ENVIRONMENTS

## These concepts apply to the DXi4700-Series and the DXi6900-Series.

In DXi 3.4 and later software, you can install and license one of two optional virtual machine features.

These features expand the capability of the DXi4700-Series and DXi6900-Series systems by enabling them to run either the Veeam Agent or the Dynamic Application Environment (DAE).

The Veeam Agent is integrated into the DXi software and it facilitates the movement of virtual machines in the customer data center to the backup share in the DXi.

The Dynamic Application Environment (DAE) is a hypervisor that allows you to run applications in a virtual machine (VM) environment on DXi4700 and DXi6900 systems.

In this section, we will look at each of these concepts in a little more detail.

## What is Veeam?

### The Company

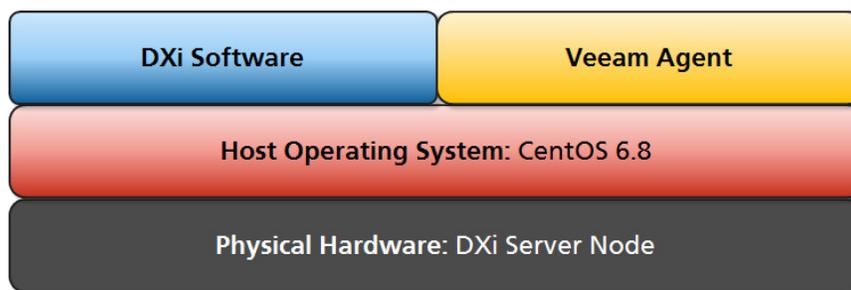
Veeam Software is a privately held information technology company that develops backup, disaster recovery, and virtualization management software for VMware and Hyper-V virtual environments.

### The Software

Veeam Backup and Replication is proprietary software that performs backup, replication, and disaster recovery for virtual machines (VMs). This typically runs on the customer's backup server to identify and protect VMs in the ecosystem.

Veeam Backup and Replication operates at the virtualization layer to back up VMs using the hypervisor's snapshots to retrieve VM data. Users can select full or incremental backups. Incremental backups are created using the built-in changed block tracking (CBT) mechanism. Also, there's an option to perform active full and synthetic full backups.

The Veeam Agent that is integrated into the DXi 3.4 and later software helps Veeam Backup and Replication move VMs from the ecosystem to the backup share in the DXi.

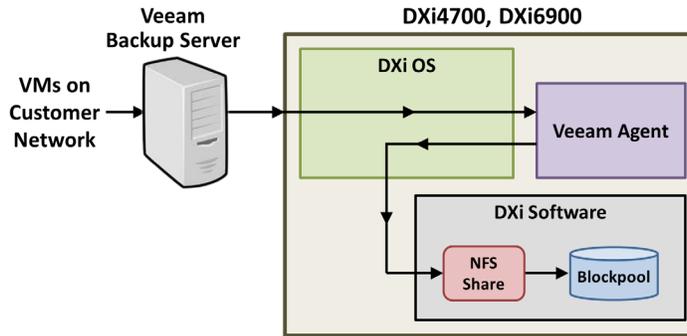


### Example Veeam Integration with DXi4700 and DXi6900

The Veeam Agent runs in the DXi process space with direct access to DXi resources. This provides better performance than running Veeam as a separate VM on the DXi.

The simplified block diagram on the next page shows the data flow in a typical Veeam installation. Management communication is omitted for clarity.

In the diagram below, the **VMs on the Customer Network** are identified by the **Veeam Backup Server**. The VM backups can be full, or incremental using the **Veeam Agent** changed block tracking. During a scheduled backup, the **Veeam Agent** works with the **Veeam Backup Server** to facilitate the data movement from the backup sources to the DXi. The backup data is sent by the **Veeam Agent** to the **NFS Share** for deduplication and storage in the **Blockpool**.



Veeam is supported on

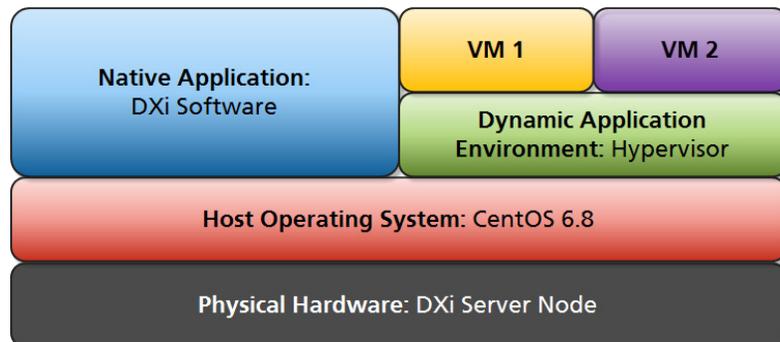
- DXi4700 G1 and DXi4700 G2 systems – NAS or Multi-Protocol systems.
  - o Veeam is NOT supported on Single Protocol VTL-only systems.
- All DXi6900 systems – DXi6900 G1, DXi6900 G2, and DXi6900-S.

## What Is the Dynamic Application Environment (DAE)?

The Dynamic application Environment (DAE) is a hypervisor that allows you to run applications in a virtual machine (VM) environment on DXi4700 and DXi6900 systems.

When DAE is licensed and installed on a DXi system with 3.4 and later software, you can install and run third-party applications in an environment completely separate from the DXi application. Depending on the configuration, DAE applications may or may not have direct access local DXi resources and shares.

Keep in mind that the DXi software may impose some limits on the resources available for the DAE VMs. This ensures that the DXi has enough resources to perform its required operations.



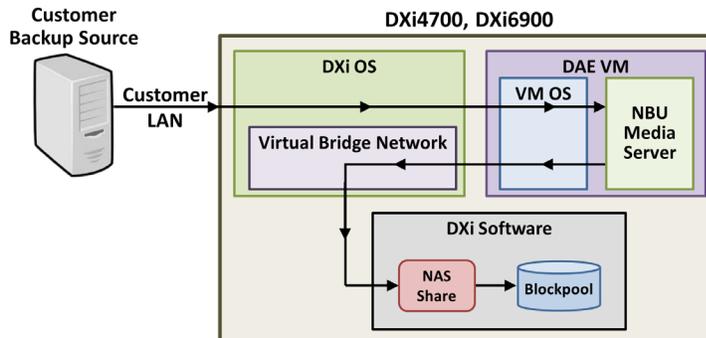
### Example DAE Integration with DXi4700 and DXi6900

An example use for DAE is installing the Veritas NetBackup (NBU) media server as a VM instead of purchasing and installing a dedicated server to run the NetBackup application.

The simplified block diagram on the next page shows the data flow when NBU is installed in a DXi as a DAE virtual machine. Management communication is omitted for clarity.

In the diagram below, backups of the customer data are scheduled by the virtualized **NBU Media Server**, and backup data is sent from the **Customer Backup Source** to the **DAE VM**, where it is routed through the **NBU Media Server**, then through a **Virtual Bridge Network** to a **NAS Share** for deduplication and storage in the **Blockpool**.

In this example, data stays within the DXi, because the NBU Media Server has direct access to a NAS share. Some DAE installations may not have direct access to the DXi software, and would communicate only with the LAN ecosystem.



## FINDING ADDITIONAL INFORMATION

Product documentation, how-to videos, and other resources are available on the Quantum Support Web site at [www.quantum.com/serviceandsupport](http://www.quantum.com/serviceandsupport).

Training courses for your DXi system are available on the Quantum StorageCare Learning page at [www.quantum.com/serviceandsupport/storagecarelearning/index.aspx](http://www.quantum.com/serviceandsupport/storagecarelearning/index.aspx).

# Quantum®

For education and training information, visit [quantum.com](http://quantum.com) or email [Learning.Admin@quantum.com](mailto:Learning.Admin@quantum.com)

### About Quantum

Quantum (NYSE: QTM) a global leader in storage, delivers highly reliable backup, recovery and archive solutions that meet demanding requirements for data integrity and availability with superior price/performance and comprehensive service and support. Quantum is the world's largest supplier of tape drives, and its DLT®, LTO Ultrium, DAT/DDS and Travan-based technologies set the standards for tape backup, recovery and archive of business-critical data for the small business to mid-range enterprise. Quantum offers the broadest portfolio of tape autoloaders and libraries and is one of the pioneers in the disk-based backup market, providing solutions that emulate a tape library but are optimized for backup and recovery.