



DXi-SERIES CORE CONCEPTS EXPLAINED

FOR DXi4800, DXi9000, AND DXi V-SERIES SYSTEMS

EMAIL US AT LEARNING.ADMIN@QUANTUM.COM
LAST UPDATED: JULY 2020

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.

CONTENTS

Overview	3
Edge to Core Solutions	4
DXi4800	4
DXi9000	4
DXi V-Series	4
Writing and Accessing Data	5
Presentation Type	5
Virtual Tape Library (VTL)	5
Network Attached Storage (NAS)	5
Application-Specific Storage	6
OpenStorage (OST)	6
The Concept of Deduplication	6
Deduplication Terminology	6
Data Flow	7
How Data Is Stored to Disk	7
Duplicate Blocks	7
Replicating Data	8
Replication Requirements	9
Using Replication with VTL, NAS, and Application Specific Presentations	10
Replication Types	11
Continuous/Namespace Replication	11
File/Cartridge Based Replication	13
Replicating Data with the OST Presentation	14
Optimized Duplication	14
Automated Image Replication (AIR)	16
Concurrent Optimized Duplication	16
Moving Data to Tape	17
Backup Application Specific Path to Tape	17
OST Direct To Tape	18
Managing Disk Space	19
Space Reclamation	19
What Is It?	19
When Does It Occur?	20
Low Capacity Management	21
What Is Low Capacity Management?	21
What Is the Impact?	21
Encrypting and Securing Data	22
Data-in-Flight Encryption	23
What Is It?	23
Data-at-Rest Encryption	24
Running Application Environments	25
What is Veeam?	25
The Company	25
The Software	25
Example Veeam Integration	25
What Is the Dynamic Application Environment (DAE)?	26
Example DAE Integration	26
Finding Additional Information	27

OVERVIEW

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

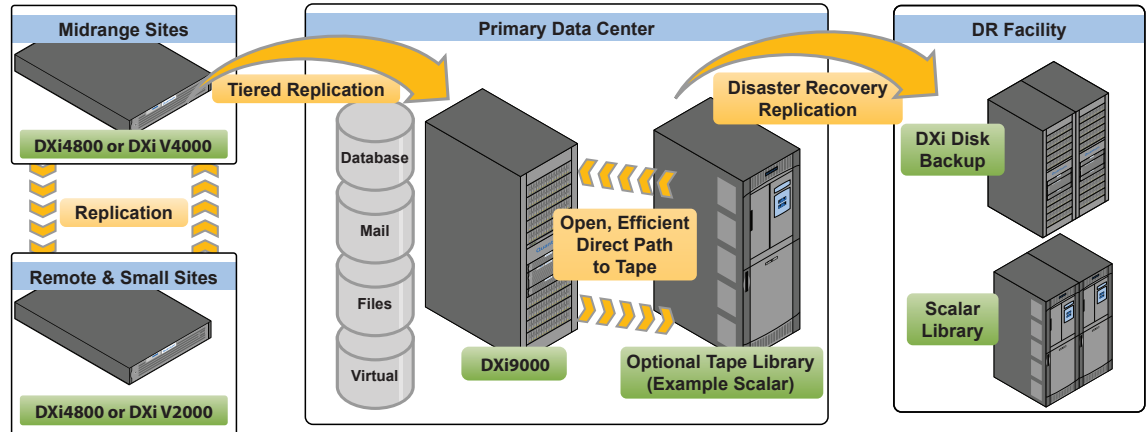
- Quantum's edge-to-core solutions
- Writing and accessing data
- Replicating data
- Moving data to tape
- Managing disk space
- Encrypting and securing data
- Running application environments

The core concepts information in this document applies to DXi4800, DXi9000 Series, and DXi V-Series products. For detailed, step-by-step configuration procedures for the core concepts described in this document, refer to the DXi documentation for your product on the Quantum Documentation Portal at www.quantum.com/documentation.



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.



DXi4800

The DXi4800 serves as the core of data protection solutions for small- to medium-sized businesses and remote site protection where customers want to provide disaster recovery (DR) by replicating data to a larger data center. The DXi4800 uses variable-length deduplication to maximize data reduction. It provides high performance and capacity-on-demand scalability and can run backup applications as a VM.

DXi9000


DXi9000 systems, which include DXi9000 Standard Density (SD) and DXi9000 High Density (HD) systems, serve as the core of multi-site enterprise data protection solutions. They provide industry-best density and deliver ultra-fast performance by leveraging SSDs for metadata operations. DXi9000 systems use variable-length deduplication to maximize data reduction and offers unique capacity-on-demand scalability. They can also run backup applications as a VM.

DXi V-Series

DXi V-Series systems (DXi V1000, DXi V2000, and DXi V4000) are designed for small businesses or remote site protection. DXi V-Series is a deduplication solution in virtual appliance model that combines the power of variable-length deduplication with the simplicity and flexibility of a virtual machine. DXi V-Series provides affordable backup and DR protection for any physical or virtual data. In addition, for DR protection, DXi V-Series 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.

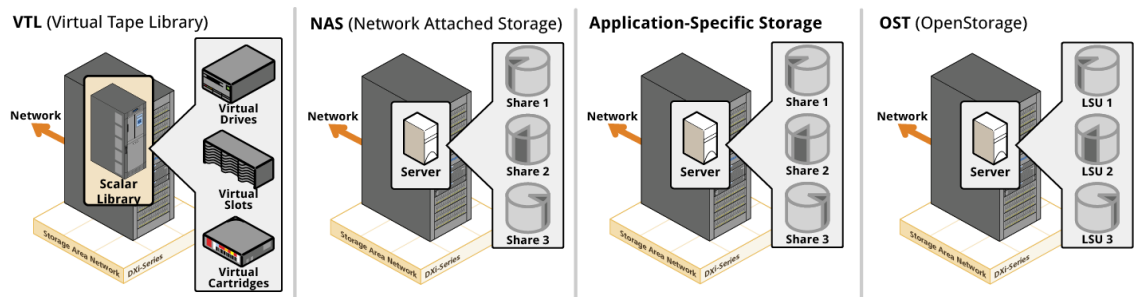


Presentation Type

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

- Virtual Tape Library (VTL)
- Network Attached Storage (NAS)
- Application-Specific Storage
- Open Storage (OST)

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



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.

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. You must select either the **Common Internet File System (CIFS)/Server Message Block (SMB)** protocol or **Network File System (NFS)** protocol when you configure a NAS share on the DXi system. This allows the DXi to be used as a NAS appliance for backup.

Application-Specific Storage

The Application Specific presentation allows the DXi to use application-specific shares to be used by Oracle™ Recovery Manager (RMAN) or Accent File System (AccentFS).

Oracle RMAN allows Oracle servers to integrate with DXi systems. Once installed and configured, an Oracle server can manage backups through the DXi system and take advantage of the system's capabilities such as data deduplication and replication. Enabling Oracle RMAN requires a Quantum plug-in to be installed on the Oracle server. Shares on the DXi that will be used for Oracle backup must be configured as Application Specific 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.

AccentFS includes your servers in the deduplication process, to minimize bandwidth and send only unique data over the network. Since the Accent file system appears as a native file system on the client, or backup host, any program that can write to a file system can use AccentFS.

AccentFS is intended to be used by backup applications not supported directly by Quantum application-specific plug-ins. The AccentFS Plug-in must be installed on the media server and then the DXi can be configured with an Application Specific share.

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 (7.6 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 back up data to the DXi using OST, 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.

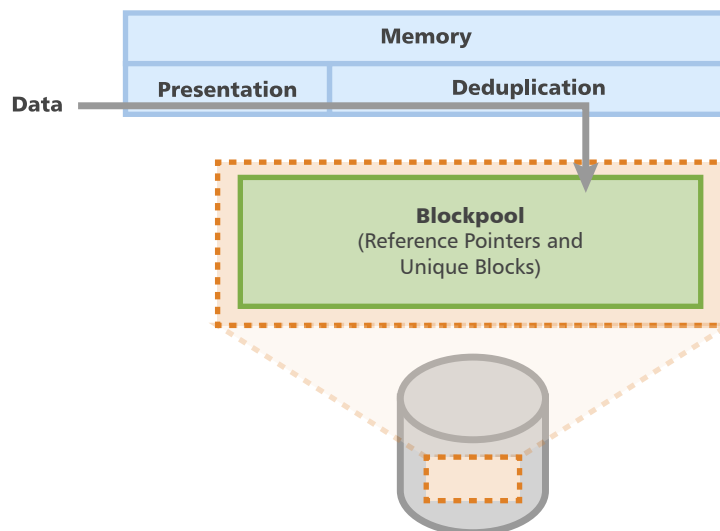
The Concept of Deduplication

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 VTL, NAS, Application Specific, 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 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.

REPLICATING DATA


Data replication is the process of copying deduplicated data from one DXi system to another for purposes of DR. The DXi copies only the blocks of deduplicated data that are different between source and target, which results in a very efficient use of network bandwidth.



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. 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 30 different source DXi systems for the DXi4800 and up to 50 systems for the DXi9000.


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 VTL, NAS, and Application Specific Presentations

To be replicated, VTL partitions, NAS shares, and Application Specific 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 VTL partition, NAS share, or Application Specific share 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 4800 web interface. The top navigation bar includes 'Quantum. DXi: 4800' and 'Normal'. 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, UTILITIES, WIZARDS, MAP, Advanced Reporting, and Best Practices Guide. The main content area is divided into two panels. The left panel is 'Share Settings' for 'share_1', showing options like 'Enforce read-only access', 'Hide from network browsing', 'Deduplication' (Enabled), and 'Protocol' (CIFS/SMB). The right panel is 'Replication' and is highlighted with a red box. It has a 'Send' section with a checked checkbox 'Enable replication to the replication target system' and an unchecked checkbox 'Enable Directory/File Based replication to target'. Below this is a 'Sync ID' field with 'share_1'. A table titled 'Replication Target systems' is shown with the following data:

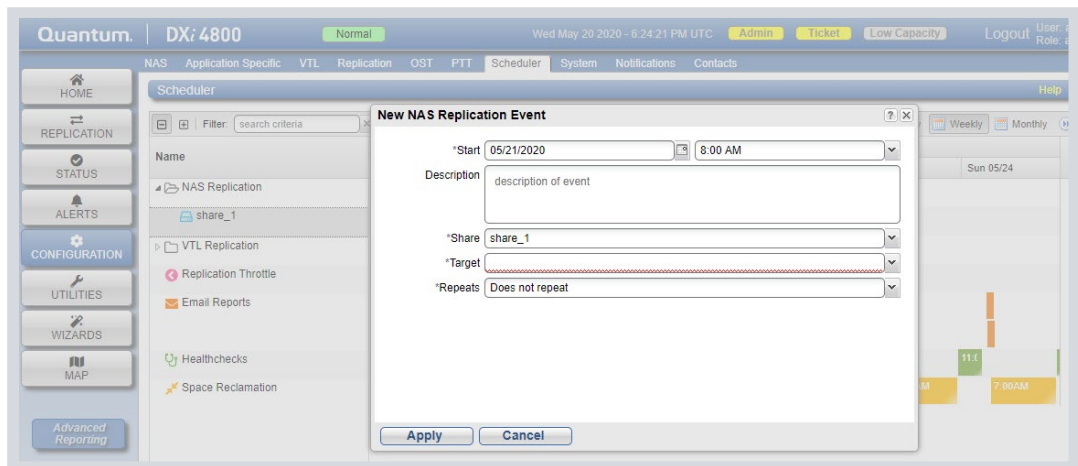
Replication Target systems	Status	Encryption
<input type="checkbox"/> 10.20.162.129	Ready	None
<input checked="" type="checkbox"/> 10.20.162.238	Ready	AES 256-bit

Below the table is a 'Receive' section with an unchecked checkbox 'Enable Directory/File Based replication to this system' and a 'Sync ID' field with 'share_1'. At the bottom of the page are 'Apply', 'Reset', and 'Cancel' buttons.

How Do You Schedule Namespace Replication?

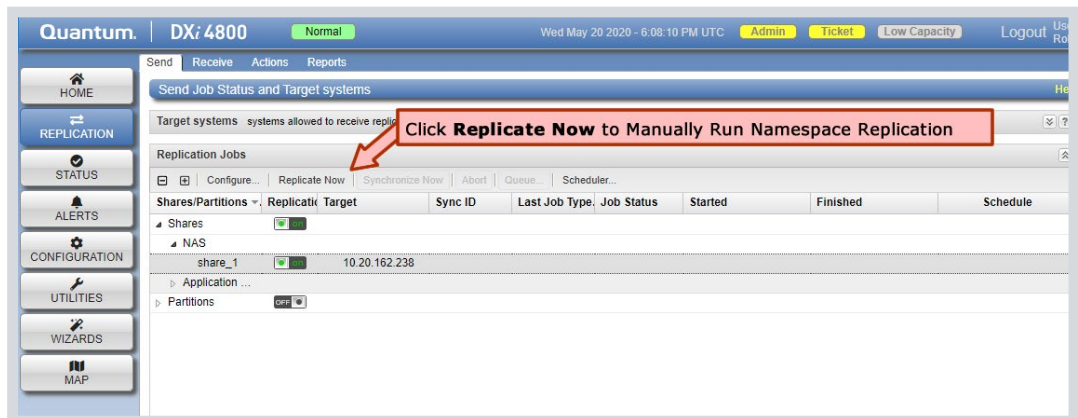
To schedule Namespace Replication, go to the **Scheduler** page (**Configuration > Scheduler**), expand either **NAS Replication** or **VTL Replication**, and select your share or partition.

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/Namespave Replication, *File/Cartridge Based Replication* sends data from a NAS share, Application Specific 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/SMB 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?

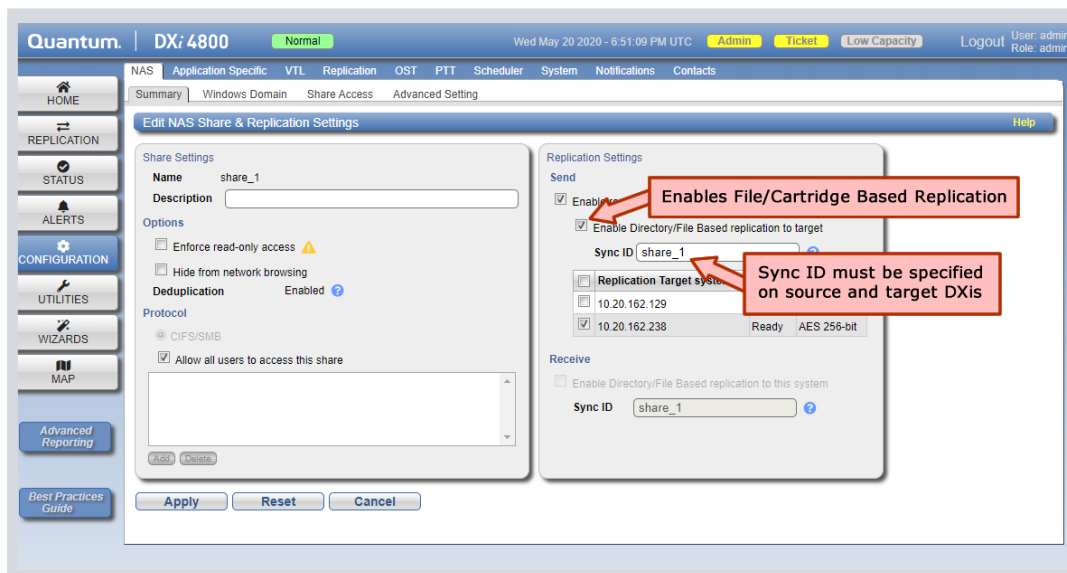
You can manually synchronize data between the source DXi system and target DXi system, if needed. For example, you should manually synchronize data when File/Cartridge Based Replication is first enabled for a share or partition. In addition, if a File/Cartridge Based Replication job fails or is disabled 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, Application Specific 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.

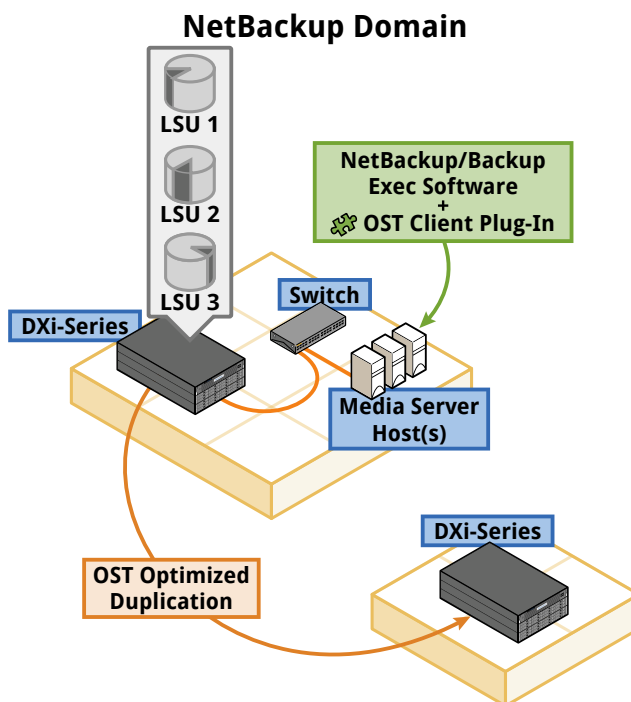
Optimized Duplication

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 and Backup Exec—it does not work with other backup applications. See your DXi Release Notes to identify the supported versions of NetBackup and Backup Exec.

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.

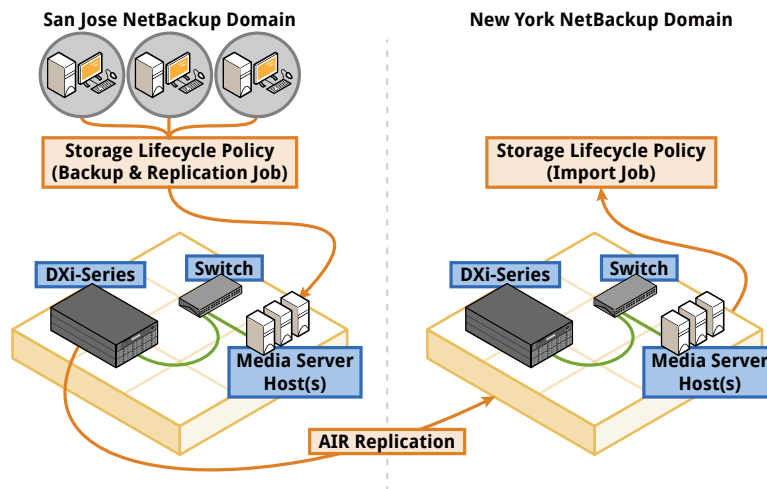
Accessing Replicated Data

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)

Specific versions of Veritas NetBackup enable you to configure an LSU for Automatic Image Replication (AIR). When AIR 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 replication, the backup images that are replicated, and the destination of the replication (all configured master server domains or specific master server domains) are determined by the storage lifecycle policies (SLPs) that you configure in NetBackup.

It is important to remember that with AIR, 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 AIR, 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 AIR 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.

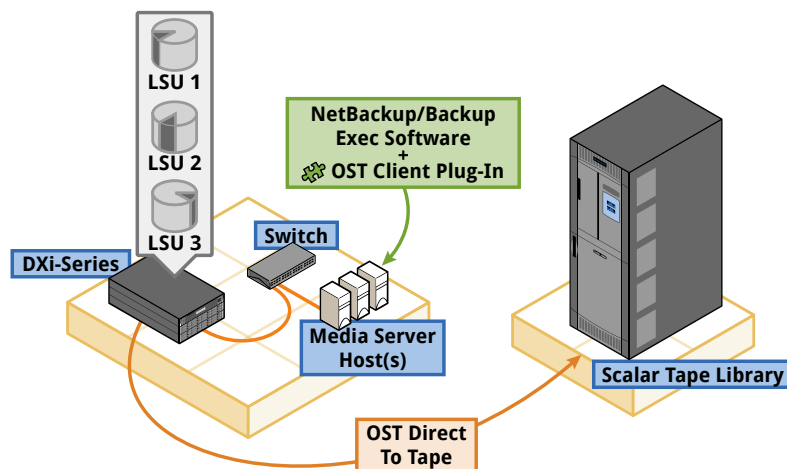
OST Direct To Tape

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:

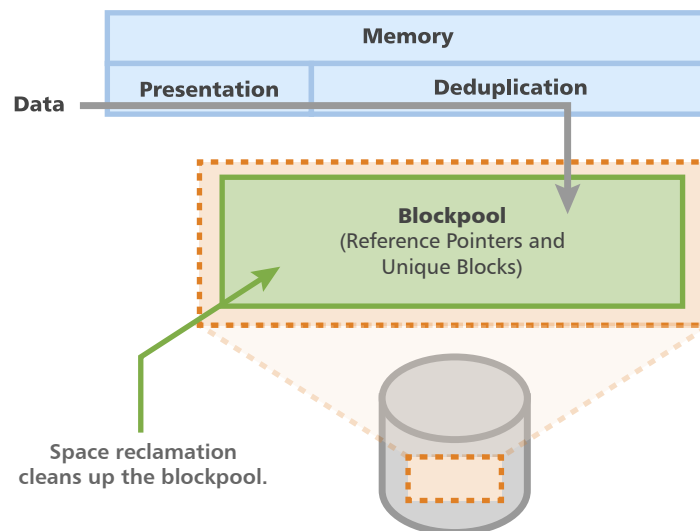
- Space reclamation
- Low capacity management

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 five stages:

- **Stage 1: Reclaiming 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 5.
- **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.**
- **Stage 4: Liberating Candidates.** The DXi detects the unreferenced blocks and marks them as unreferenced with a zero-reference count.
- **Stage 5: Reclaiming Disk Space (Compaction).** Unique blocks with zero reference pointers are

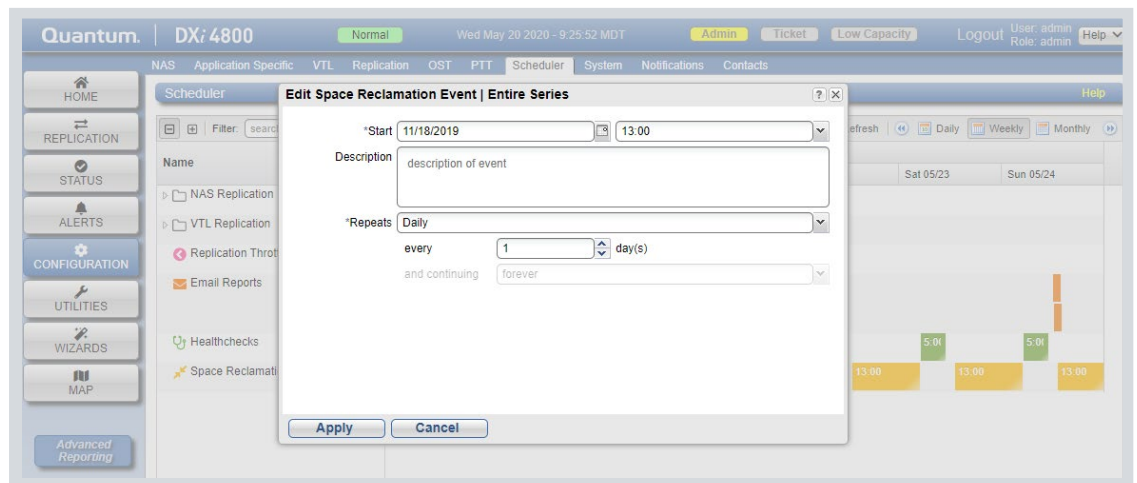
removed from the blockpool and the remaining data is compacted, thereby creating space for new unique data.

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

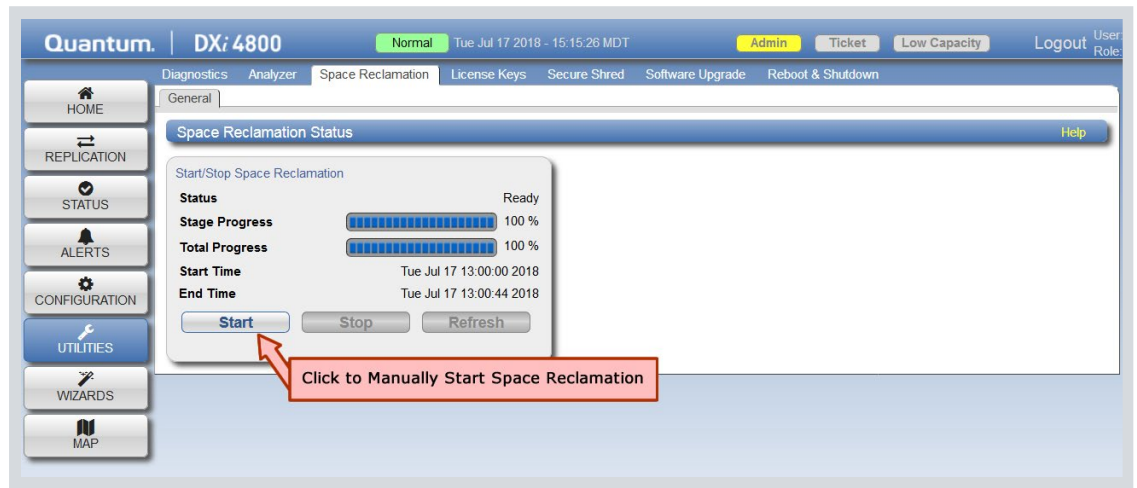
- **New / Normal Mode:** To increase performance, when space reclamation is initiated manually on the **Space Reclamation** page or as a scheduled event, only Stage 2, Stage 3, and Stage 4 are run. Stage 1 and 5 are not required in **Normal** mode because the DXi can automatically compact reclaimable 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 stages of space reclamation are run.

When Does It Occur?

Quantum recommends to run space reclamation on a weekly basis to prevent the system from running out of space. 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 when your DXi is not busy, for example, 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.



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. For example, on a DXi4800 system, as disk capacity is used and free disk space approaches low levels, the following will occur:

- The **Low Capacity** condition indicator is lit 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:
 - **850 GB + 100 GB x (total usable capacity in TB / 10 TB)** - This is referred to as **Low Space**. Backup ingest and replication occur as normal. Space reclamation is automatically started.
 - **Free Space is less than 250 GB** - This is referred to as **Critical Reserve Space**. Backup ingest and replication are stopped. Space reclamation is automatically started. **Note:** VTL cartridge metadata files may still be updated.
 - **Free Space is less than 10 GB** - This is referred to as **No Space**. Backup ingest and replication are stopped. Space reclamation is automatically started. **Note:** VTL cartridge metadata files may still be updated.

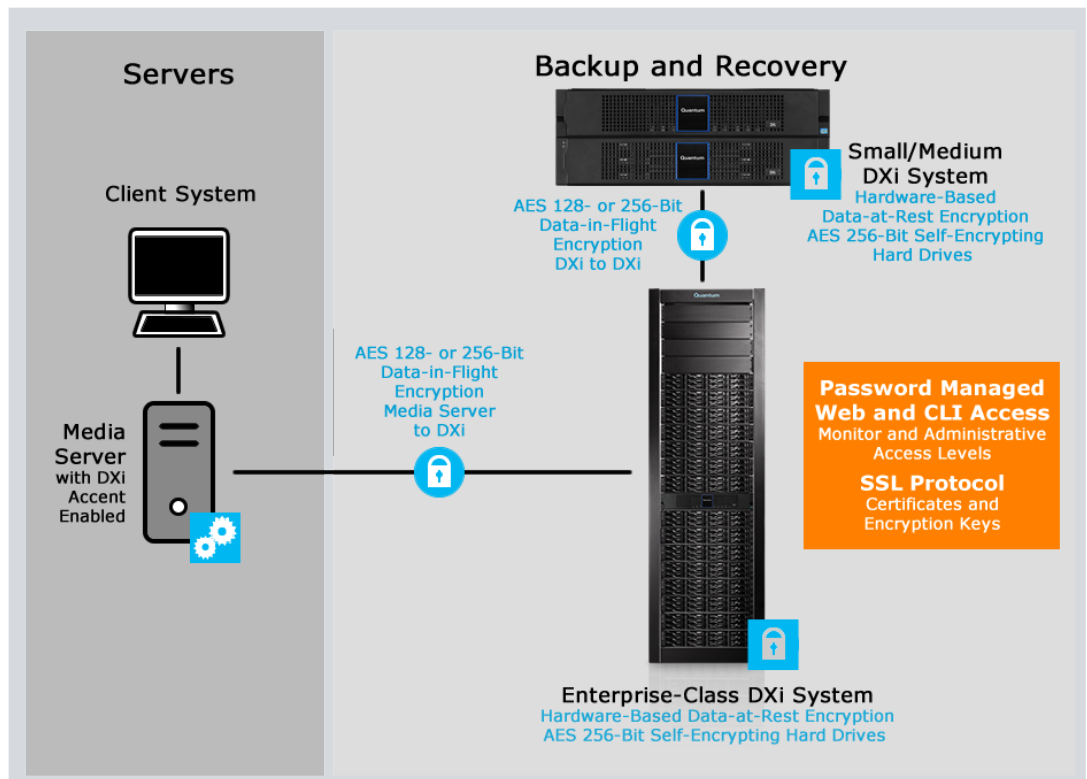


ENCRYPTING AND SECURING DATA

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:

- **Data-in-Flight Encryption:** 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:** Self Encrypting Drive (SED) technology to secure all customer data stored on the DXi.



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

- **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

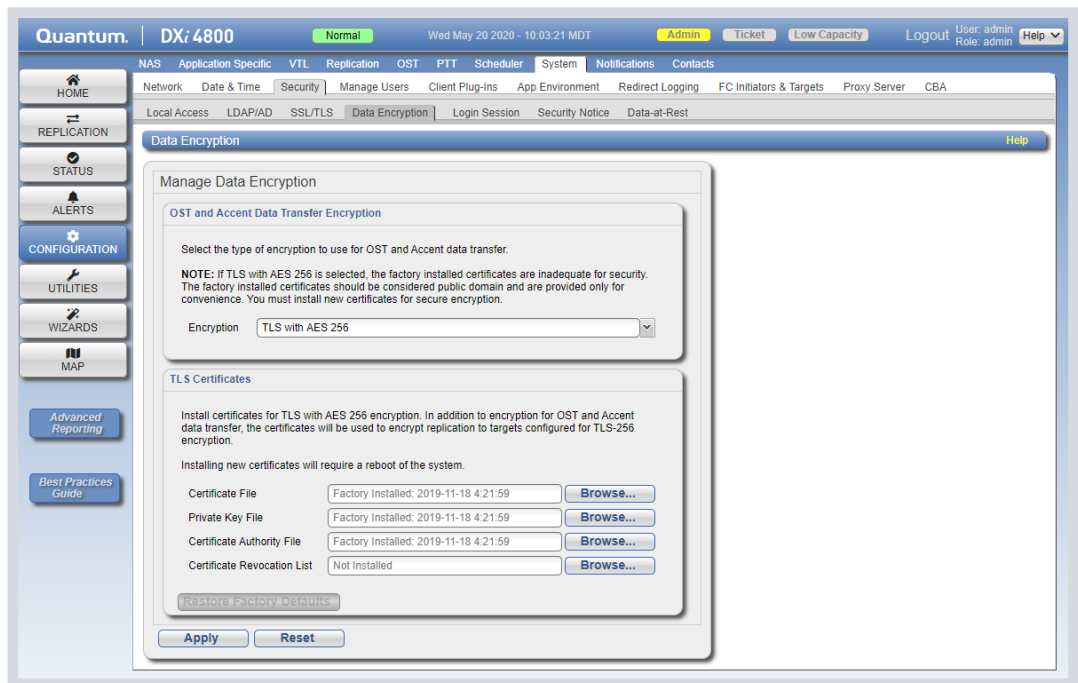
What Is It?

Replication and Failback encryption have always been available on DXi systems. For DXi4800 and DXi9000, Data-in-Flight Encryption is 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. Transport Layer Security (TLS) with AES 256 requires you to supply your own digital certificates for use with DXi Accent encryption.

After the Data-in-Flight Encryption license is installed, you can configure the OST DXi encryption settings from the **Data Encryption** page (**Configuration > Security > Data Encryption**).

OST and Accent Encryption Settings Example



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 DXi4800 and DXi9000 systems. This encryption is not available in DXi V-Series systems or in certain 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.



Self-Encrypting Drives

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

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.

Data-at-Rest Configuration Settings Example

Quantum | DXi 4800 | Normal | Wed May 20 2020 - 10:11:21 MDT | Admin | Ticket | Low Capacity | Logout | User: admin | Role: admin | Help

NAS | Application Specific | VTL | Replication | OST | PTT | Scheduler | System | Notifications | Contacts

Network | Date & Time | Security | Manage Users | Client Plug-Ins | App Environment | Redirect Logging | FC Initiators & Targets | Proxy Server | CBA

Local Access | LDAP/AD | SSL/TLS | Data Encryption | Login Session | Security Notice | Data-at-Rest

Data-at-Rest

Current Data-at-Rest Encryption **Disabled** [Check Status](#)

Enable Data-at-Rest Encryption

*Passphrase

*Confirm Passphrase

Download/Send Data-at-Rest Recovery File

Download Email

*Email Recipient

*Password

*Confirm Password

RUNNING APPLICATION ENVIRONMENTS

DXi software enables you to license and install one of two optional virtual machine features. These features expand the capability of the DXi4800 and DXi9000 systems by enabling them to run either the Veeam Agent or the Dynamic Application Environment (DAE). DXi V-Series systems do not support application environments.

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

What is Veeam?

The Company

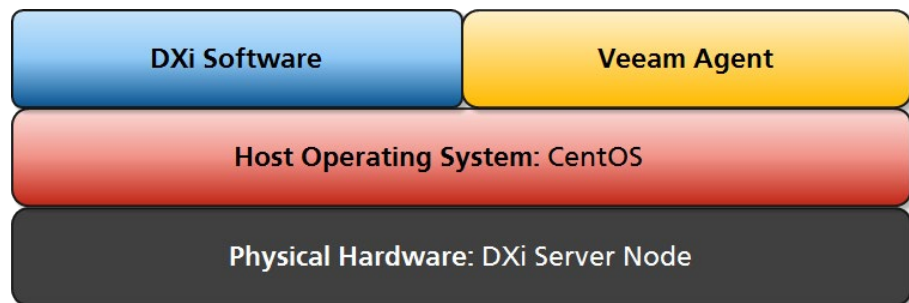
Veeam Software develops backup, disaster recovery, and virtualization management software for VMware and Hyper-V virtual environments.

The Software

Veeam Backup and Replication is proprietary software that performs backup, replication, and disaster recovery for 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 software helps Veeam Backup and Replication move VMs from the ecosystem to the backup share in the DXi.




Example Veeam Integration

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**.

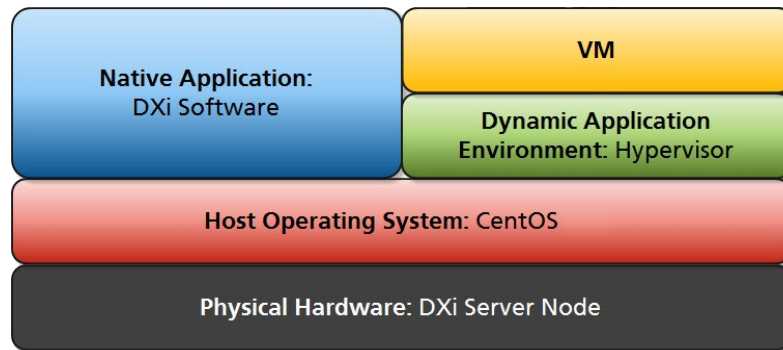


What Is the Dynamic Application Environment (DAE)?

The DAE is a hypervisor that allows you to run applications in a VM environment on DXi4800 and DXi9000 Series systems.

When DAE is licensed and installed on a DXi system, 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.



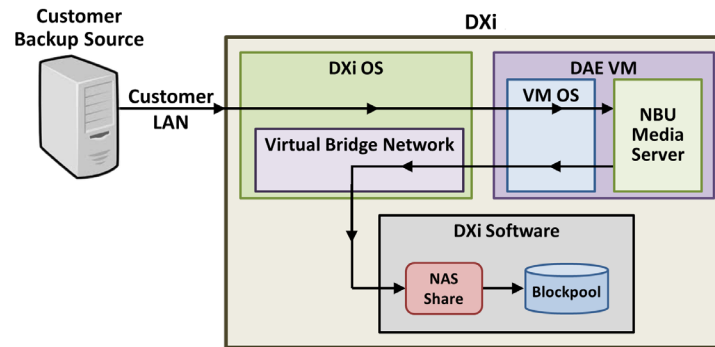
Example DAE Integration

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

DXi product documentation and other resources are available from the Quantum Documentation Portal at www.quantum.com/documentation.

Quantum®

ABOUT QUANTUM

Quantum technology and services help customers capture, create and share digital content – and preserve and protect it for decades at the lowest cost. Quantum’s platforms provide the fastest performance for high-resolution video, images, and industrial IoT, with solutions built for every stage of the data lifecycle, from high-performance ingest to real-time collaboration and analysis and low-cost archiving. Every day the world’s leading entertainment companies, sports franchises, research scientists, government agencies, enterprises, and cloud providers are making the world happier, safer, and smarter on Quantum. See how at www.quantum.com.

www.quantum.com • 800-677-6268