

Note: Due to a lack of significant product updates, the Dell EMC Unity Product Analysis is no longer receiving updates as of September, 2023.

Dell EMC Unity XT - SAN Storage and NAS System

Dell EMC announced the midrange unified storage system called Unity XT in May of 2016. The Dell EMC Unity XT system is derived from the Dell EMC VNX/VNXe series of storage systems with a new operating environment and new hardware configuration. Unity XT systems have both all flash and hybrid models and are scale-up systems with dual active controllers. The unified systems support block and file with the file software originally released in the VNXe3200.

There are 8 models, 380/380F, 480/480F, 680/680F, 880/880F.

Many of the software functions of the Dell EMC Unity XT systems were extracted from VNX software. The Unity XT systems use a SuSE Linux kernel and support Docker containers. The filesystem software supports multiple NAS Servers, which were previously called virtual data movers. Unisphere is the element manager as a web GUI, implemented with HTML5.

The dual controllers are packaged in a 2U enclosure along with either 25 2.5" devices or 12 3.5" devices that are SAS attached. Additional devices can be added with Device Array Enclosures of either 15 3.5" devices in a 3U enclosure or 80 2.5" devices in a 3U enclosure. Software features are included in the base price and include snapshots, synchronous and asynchronous remote replication, VVOLs and VASA support. Telemetry data is collected and analyzed with a SaaS product called CloudIQ. Health checks are done and proactive maintenance is scheduled based on the analytics.

The SSDs in the Dell EMC Unity XT have a lifetime warranty as long as a maintenance contract is in place. In addition, the maintenance price is guaranteed not to be increased from the initial purchase.

The Unity XT systems provide a wide set of connectivity options for block access and for file access as a NAS system.

Evaluator Group Coverage

Related coverage on the Evaluator Series Research website includes:

SAN and NAS matrices

SAN and NAS Evaluation Guides

Dell EMC Unity XT Product Brief

Competitive Product Analysis and Product Briefs



Highlights

- Dell EMC Unified Midrange System derived from VNX software
- Block - Storage Processors software – MCx environment
- File – Updated 64 bit OS
- Container based – Linux OS
- All flash or hybrid models and VSA version
- Support for RAID 0/1, 5, 6
- HTML5 embedded Unisphere management
- Encryption in SAS controller hardware with internal key management
- Block storage support over iSCSI and Fibre Channel
- File support over CIFS, NFS, FTP/SFTP
- VMware vSphere API's for Array Integration (VAAI) for hardware array support of advanced features
- VASA 2.0 and VVOL support
- Dell EMC Support Remote Services – monitoring and CloudIQ link
- Software functions included in system
- Snapshots, remote replication, file and data migration
- Storage Tiering – Virtual Pools (FAST VP)
- FAST Cache which uses SSD drives as cache extension
- Internal SSD M.2 for OS and cache backup on power downs
- Battery backup for graceful shutdown
- 2U Enclosure with 2 controllers – Intel Haswell processor
- 25 SSDs in controller enclosure
- 12 3.5" in 3U DAE or 80 2.5" in 3U DAE added for scale up
- Lifetime flash guarantee while under support maintenance
- Online data-in-place controller upgrades
- Guarantee maintenance price – no increases
- QoS – Host I/O limits on LUNs, filesystems, snapshots
- Inline compression and deduplication in software for all flash storage pools
- Policy-base file tiering to clouds

Contents

Dell EMC Unity XT – SAN Storage and NAS System	1
Highlights	2
Overview	5
Dell EMC Unity XT Model Comparison	5
Product Architecture – Unity XT	6
Disk Processor Enclosure (DPE)	8
Device Array Enclosures (DAE)	10
Flash Drives	10
Power Supplies	10
M.2 SSD	10
Battery Backup Unit	11
Unity XT Configuration	11
Software Architecture	12
Unity XT Operating Environment	12
VMware VAAI Support	13
• Full copy	13
• Block zeroing	13
• Hardware-assisted locking	13
• Dead Space Reclamation	14
• Out of Space Condition Support	14
VMware VASA Support	14
VMware VVOL Support	14
Storage Pools (Virtual Pools)	14
Virtual Provisioning (Thin Provisioning)	15
Inline Data Reduction	15
RAID Protection	15
Advanced Features	16
Included Software	16
Unisphere	16
Unisphere Central	17
Data Migration	18
Fully Automated Storage Tiering – Virtual Pools (FAST VP) for Hybrid Unity XT Systems	18
FAST Cache	19
Encryption – Data @ Rest Encryption	19
RecoverPoint Basic and VMs	19
Snapshot (redirect on write)	19
Native Synchronous Replication – Block Only	20
Native Asynchronous Replication – Block (LUN) and Filesystem replication	21
CloudIQ	22
Unity XT NAS	22
NAS Protocols	23



Data Protection	23
NDMP	23
MetroSync File Replication	24
VAAI Support for NAS	24
Quotas	24
Anti-Virus Scanning	25
Virtual Storage Appliance (VSA)	27
Reliability, Availability, Serviceability Features	28
Redundant write cache	28
PowerStore Metro Node	28
Performance	29
Evaluator Group Comments	30

Overview

The Dell EMC Unity XT series are multi-platform modular storage systems providing unified block and file access. There are a number of different models and configurations can vary based on the interfaces, capacity, and in the case of the Unity XT series, selections for block and/or file support.



Figure 1: Dell EMC Unity XT Flash Systems (source: Dell EMC)

The Unity XT system is a traditional two-controller system with an operating environment that supports both block and file access. The storage processors run a SuSE Linux embedded operating system and the Unity XT environment in a Linux container. The Dell EMC System Remote Support software executes in a Docker container.

Dell EMC Unity XT Model Comparison

Table 1: Dell EMC Unity XT X80(F)

Model	Dell EMC Unity 380/380F	Dell EMC Unity 480/480F	Dell EMC Unity 680/680F	Dell EMC Unity 880/880F
Processor	2 x Intel CPUs, 12c/1.7GHz	2 x dual-socket Intel CPUs, 32c/1.8GHz	2 x dual-socket Intel CPUs, 48c/2.1GHz	2 x dual-socket Intel CPUs 64c/2.1GHz
Memory	128 GB	192 GB	384 GB	768 GB
Max Drives	500	750	1000	1500
Max Capacity (RAW)	2.4 PB	4.0 PB	8.0 PB	16.0 PB

Table 1: Dell EMC Unity XT Model Comparison

Evaluator Group Comment: *The Unity XT systems can scale up by adding more capacity with additional enclosures and devices. It does not have the ability to scale out where additional controllers can be added with capacity to maintain the performance to capacity ratio. This means that to achieve more performance when capacity is added, the controllers must be replaced in a move to the next larger model. It is up to the customer to order the correct product that has the capacity and performance to meet current and future needs.*

Product Architecture - Unity XT

The architecture of the Unity XT system from a hardware standpoint is a traditional dual controller storage system with devices attached using SAS. These are scale-up systems with the ability to add more capacity. Increased performance is achieved by replacing the controller with the next larger model, ending with the largest model.

The system diagram that follows includes the different software elements and the operational functions for controller-to-controller communications.

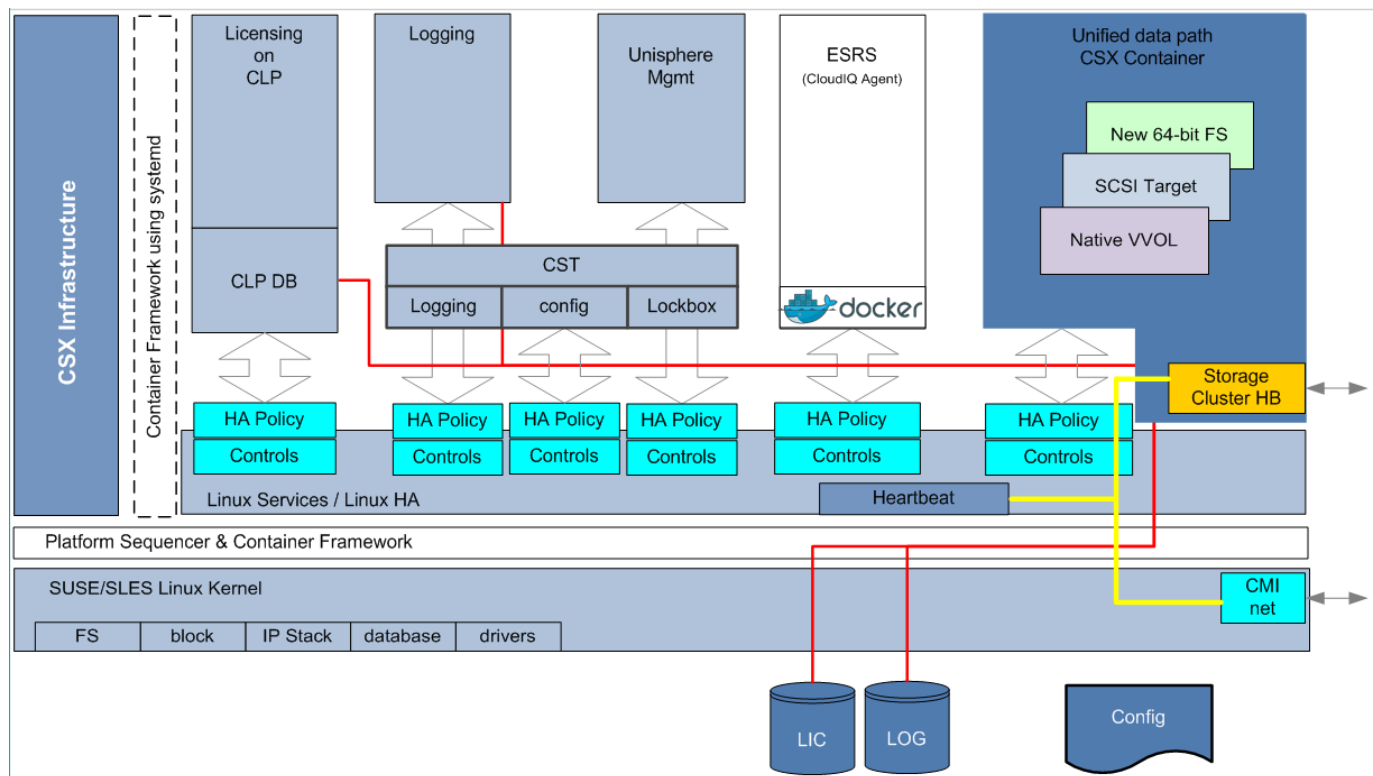


Figure 2: Dell EMC Unity XT Architecture (source: Dell EMC)

- Unity XT Hardware - 2U Disk Processor Enclosure (DPE) with two controllers and 25 SSDs for the all-flash system and for the hybrid system, 25 2.5" devices or 15 3.5" devices.
- Optional Disk Array Enclosures (DAE) with 80 2.5" devices in a 3U enclosure or 15 3.5" devices in a 3U enclosure.

The Unity XT systems all consist of the following modular components:

All components can be housed in the Dell EMC rack or any standard 19-inch rack



Some of the elements are hot pluggable as customer replaceable units and some are not. DAEs and the DPE enclosures are field replaceable units.

Hot Pluggable CRUs	Not Hot Pluggable CRUs
Devices	Battery Backup Unit
Power Cables	Fans
Power Supplies	Memory DIMMs
SAS Cables	M.2 SSD for OS and cache storage on power loss
SFP interface plugs	I/O Modules
Storage Processors (controllers)	

Table 2: Pluggable Elements for Unity XT

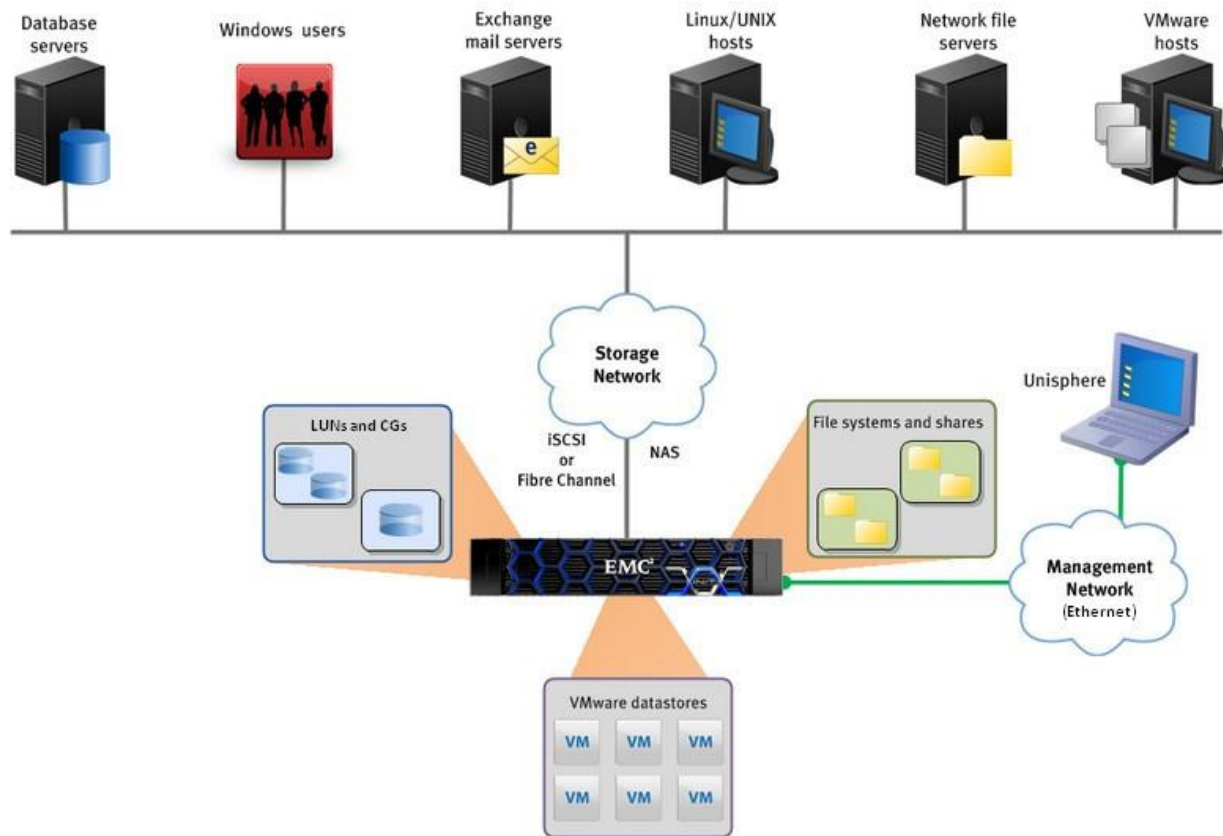


Figure 3: Unity XT Overview Source: Dell EMC

Disk Processor Enclosure (DPE)

The DPE components include Storage Processors (SP), which are the central processing units (CPU) for the system, power supplies, fan modules, I/O modules, memory DIMMs, an M.2 SSD, a battery backup unit and device connections. There are two storage processors, SP A and SP B, in each model and the SP is considered the control center of the system.

Each SP consists of one CPU module with the model-determined CPU size and system memory, the I/O module carrier with the appropriate number of base I/O module slots. The SP senses the speed of the incoming host I/O and sets the speed of its front-end ports to the lowest speed it senses.

There are two different enclosures used for the DPE. One has 80 2.5" devices in 3U and the other has 12 3.5" devices in 2U. All devices are connected with SAS. There will be a future option for 8 internal NVMe drives with PCIe connection.




Figure 4: DPE with 80 Devices – top view



Figure 5: DPE Enclosure with 12 Devices

The rear of the DPE contains connections to hosts and additional device enclosures and has the hot plug service access.

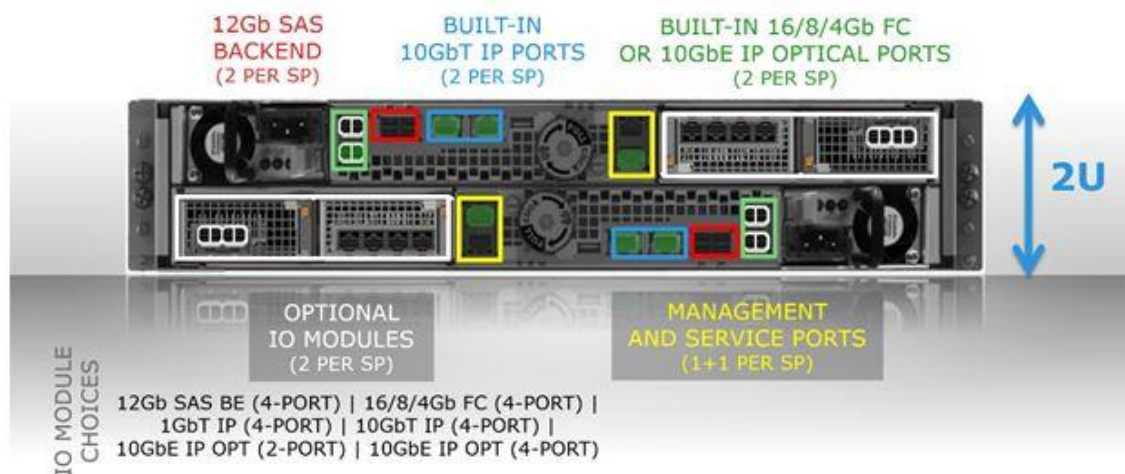


Figure 6: DPE Enclosure Rear View (source: Dell EMC)

The management module has two GbE Ethernet LAN ports, one for management and one for service. The management port is used for managing the storage processor with the Unisphere software, while the service port is for the use of Dell EMC Service if access is needed to the SP.

Device Array Enclosures (DAE)

Additional device array enclosures may be attached to the DPE by use of 12 Gb/s SAS over four lanes as the interconnect interface. A device enclosure connects to another DAE or a DPE and is managed by storage system software. The minimum number of drives for a system is four, used for storing system configuration information. The DAE has two versions where one contains up to 15 3.5" devices in a 3U form factor and contains up to 80 2.5" devices in 3U. Any unoccupied device slot has a filler module to maintain airflow.

Flash Drives

Dell EMC utilizes solid-state disk (SSD) drives utilizing flash storage technology. Flash drives will allow users to define a new higher speed access, persistent tier of storage. The use of this type of storage is suited for applications with high transaction rates and those requiring the fastest possible retrieval and storage of data. Flash drives are used both for FAST Cache and as storage devices.

Power Supplies

There is a power supply for each SP controller with the ability to power the entire DPE in the case of a failure of one of the power supplies.

M.2 SSD

The M.2 SSD internal to each storage processor is used to store the operating system software and cache data in the case of a power loss.



Figure 7: M.2 SSD Card (source: Dell EMC)

Battery Backup Unit

The single battery backup unit provides power to the DPE in the case of power loss to be able to write cache data to the M.2 SSD cards to avoid potential data loss.

Unity XT Configuration

Model	VSA	380 / F	480 / F	680 / F	880 / F
LUNs	64	1,000	1,500	2,000	6,000
LUN Size	16 TB	256 TB	256 TB	256 TB	256 TB
VVOLs	1,750	9,000	9,000	13,500	30,000
Snapshots	128	1,000	8,000	8,000	16,000
Replication Sessions	16	1,000	1,000	1,500	2,000

Table 3: Unit Configuration Limits

Software Architecture

Unity XT Operating Environment

The Unity XT Operating Environment uses an embedded SuSE Linux with containers for Unity XT software and other functionality. Currently, the additional functionality is Dell EMC System Remote Support Service (ESRS) running in a Docker container. The Unity XT software uses elements from Dell EMC VNXe including the 64-bit file system software, software functions extracted from Dell EMC VNX, and newly developed software such as the VVOL implementation. The following diagram illustrates these elements.

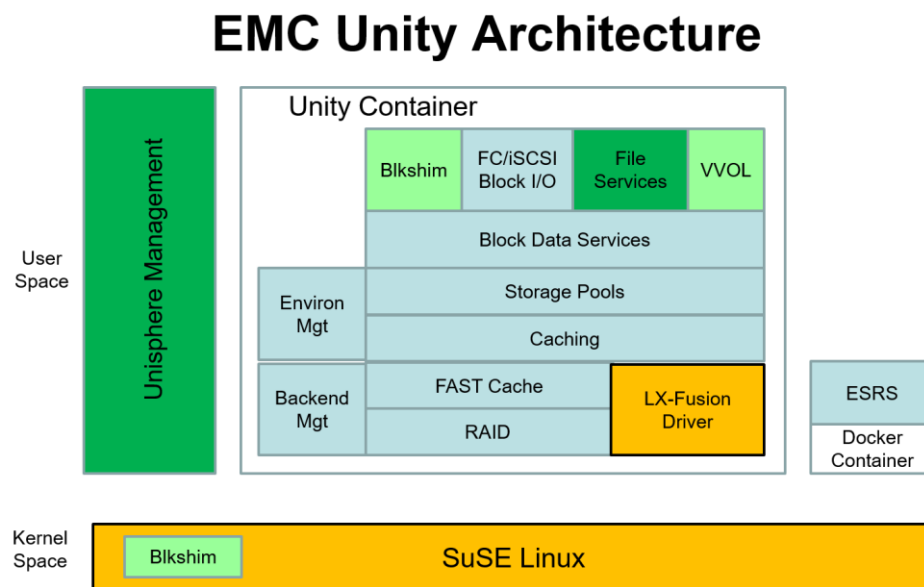


Figure 8: Dell EMC Unity XT Software Architecture

The Unity XT software optimizes usage of multiple processor cores and multi-threading for a performance increase over earlier generation. The Unity XT systems supports non-disruptive upgrades for both firmware and hardware.

- 64 bit based operating environment for enhanced memory and I/O address space support
- Support for multi-core Intel processors
- Flash SSD drive support
- The ability to intermix iSCSI and FC networks attachments simultaneously using UltraFlex I/O modules
- Virtual provisioning support
- Support for RecoverPoint with a write splitter

The Unity XT includes all software function in the base price.



Included Software Area	Function
Management	<ul style="list-style-type: none">• Unisphere element manager• Unisphere Central – Consolidated management across multiple systems• Proactive Assist – configure remote support, online service• Quality of Service• Storage Analytics Adapter
Protocols	<ul style="list-style-type: none">• File – NFS, CIFS/SMB, FTP/SFTP• Block – Fibre Channel, iSCSI• VVOL
Local Protection	<ul style="list-style-type: none">• Encryption – Controller-based• Snapshots• Anti-virus linkage• Quotas
Remote Protection	<ul style="list-style-type: none">• Native Asynchronous – block and file• Native Synchronous – block and file• RecoverPoint Basic• RecoverPoint for VMs• Cloud Data Protection

Table 4: Included Software for Unity XT

VMware VAAI Support

The Unity XT system supports the VMware vSphere APIs for Array Integration (VAAI), a set of mechanisms that allow processing for certain data-related services—copying data when creating a new VM, for example—to be offloaded from the ESX host to a storage array. The intent of these APIs is to streamline the functioning of the ESX server and speed-up delivery of storage-supported services.

- **Full copy** — Enables the storage system to make full copies of data within the storage system without having the ESX host read and write the data. A single SCSI Extended Copy command is used to copy a set of contiguous blocks.
- **Block zeroing** — Enables storage systems to zero out a large number of blocks to speed provisioning of virtual machines. The SCSI Write Same command is used to copy a zero block for the number of blocks required
- **Hardware-assisted locking** — provides an alternative means to protect the metadata for VMFS cluster file systems, thereby improving the scalability of large ESX host farms sharing a datastore.



Beginning with vSphere 5.0, enhancements for thin provisioning were added that are supported by Unity XT.

- **Dead Space Reclamation** — Reclaims blocks from VMFS deleted files. The SCSI Unmap command is used to free physical space.
- **Out of Space Condition Support** — A VM is paused if disk space is exhausted as reported by the storage system. This allows administrators to mitigate the situation rather than causing VM to fail.

VMware VASA Support

The Unity XT systems support the VMware vSphere APIs for Storage Awareness (VASA) version 2.0 with an internal provider. VASA can display the features of the physical storage devices, allowing vSphere administrators insight into storage capabilities. The Unity XT system has a plugin called a vendor providers, which is the integration 'glue' that sits between vCenter and the storage array. Vendor providers retrieve the storage capabilities from the array and pass these onto vCenter, which in turn can display these capabilities in the user interface. VASA may also provide information about storage health status, configuration information and capacity.

VMware VVOL Support

Unity XT support VVOLs either through block mode VMDK or attached through the NFS interface. Volumes for VMware virtual machines can be defined within VMware without the need for storage administrator to configure LUNs. Instead, the storage administrator creates a Protocol Endpoint which is part of the physical storage fabric and treated like a LUN and is discovered during a scan by ESXi. VASA APIs will bind VVOLs to the Protocol Endpoint (PE). Within the Protocol Endpoint, a Storage Container is created which is a logical grouping of the VVOLs that will be created in VMware by the VMware administrators. A VASA provider for the storage system discovers the Storage Containers and reports to vCenter. There may be multiple Storage Containers per storage system.

Evaluator Group Comment: With the relationship (ownership) between Dell EMC and VMware, it should be expected that the Unity XT system will always be a leading system for integrating capabilities needed for a VMware environment.

Storage Pools (Virtual Pools)

Unity XT uses storage pools to abstract the capacity of disks or solid state drives. Capacity is allocated based on chunks, which represent fixed capacity sizes from the drives that comprise the storage pool. An administrator creates storage pools and Unisphere will create RAID groups across the devices used in the pools. LUNs or filesystems are created from the chunks in a pool. Pools can be created across



multiple media types but Dell EMC recommends creating pools from the same media type because of different capacity sizes.

Space is assigned to LUNs and filesystems in 256MB for thin provisioned LUNs. I/O is distributed across the LUN, which is called a Dynamic LUN. The same RAID protection must be used across the entire pool. LUNs may be reduced in size (shrunk) through Windows Server Manager and requires a VDS provider. LUNs and filesystems may be expanded if capacity is available.

Adding additional drives to a storage pool will rebalance data across all drives in the pool automatically. The addition of drives must be in the increment of the minimum number of drives for the RAID type of the pool. The system will create additional RAID group(s) in the pool when drives are added.

Virtual Provisioning (Thin Provisioning)

Thin provisioning allocates capacity as needed by adding chunks from the storage pool. The minimum capacity consumed for thin LUNs is 1.75 GB – .25 GB of usable capacity and 1.5 GB of overhead for each thinly provisioned LUN. Capacity is added in 256 MB increments. A chunk is 8K in the Unity XT thin provisioned volumes. Additionally, monitoring and threshold enforcement to prevent abuse of “runaway” consumption of storage are included.

Space reclamation returns freed capacity back to the storage pool. Reclamation can be done by shrinking a LUN or filesystem or with software support using an API or a SCSI UNMAP command. Currently Symantec Storage Foundation thin replication API is support for VxVM and VxFS and Windows is supported with add-in software.

Inline Data Reduction

Inline compression and deduplication for Unity XT works for all flash and hybrid storage pools and is selectable at the volume level. Software is used to check for duplicate data and found using a hash of the data, it is not stored and the block table is updated. Data that is not duplicate, is compressed using software algorithm. Highly active data is written directly without the additional processing overhead for compression and deduplication. Data stored that is not compressed may be compressed and deduplicated as a background task. The system automatically selects data that is not highly active to be compressed and deduplicated within the storage group. Compression and deduplication can work independently of all other software features. A pretest of data is performed and if data is not compressible, it is left uncompressed. A compression estimator tool is available from Dell EMC. Writes of zero blocks is also detected and does not result in a write action, only a table update.

RAID Protection

Multiple RAID protection types are supported. RAID groups are limited to 16 disks. The following table explains the supported configurations.



RAID Type	Default Configuration	Supported Configurations
RAID 1/0	4+4	1+1*, 2+2, 3+3, 4+4
RAID 5	4+1	4+1, 8+1, 12+1
RAID 6	6+2	4+2, 6+2, 8+2, 10+2, 12+2, 14+2

Table 5: Unity XT RAID Protection (source: Dell EMC)

Advanced Features

The advanced features of Unity XT provide the distinguishing elements of the product. Dell EMC has used many of the software features previously available in Unity XT and are now included the base system. Many of the feature names have been changed.

Included Software

Unisphere

The element manager for Unity XT is Unisphere which is an HTML5 implementation, differing from the Unisphere used for VNX. Unisphere is targeted for the IT generalist with a simple control interface, dashboards for management and status, and wizards to accomplish common tasks. From the highly customizable user interface, simple, intuitive management for the systems can be performed.

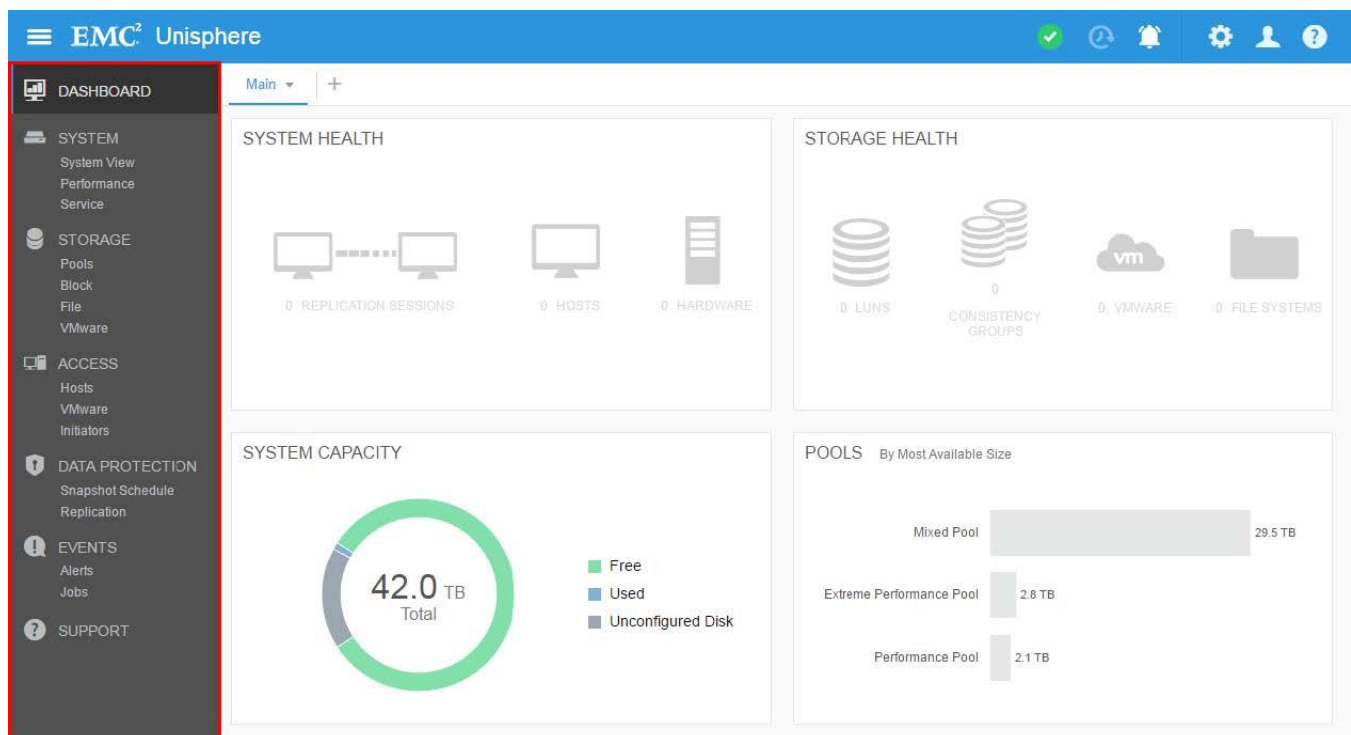


Figure 9: Unisphere Management

Hardware representations are available in Unisphere, providing a very simple element management experience.

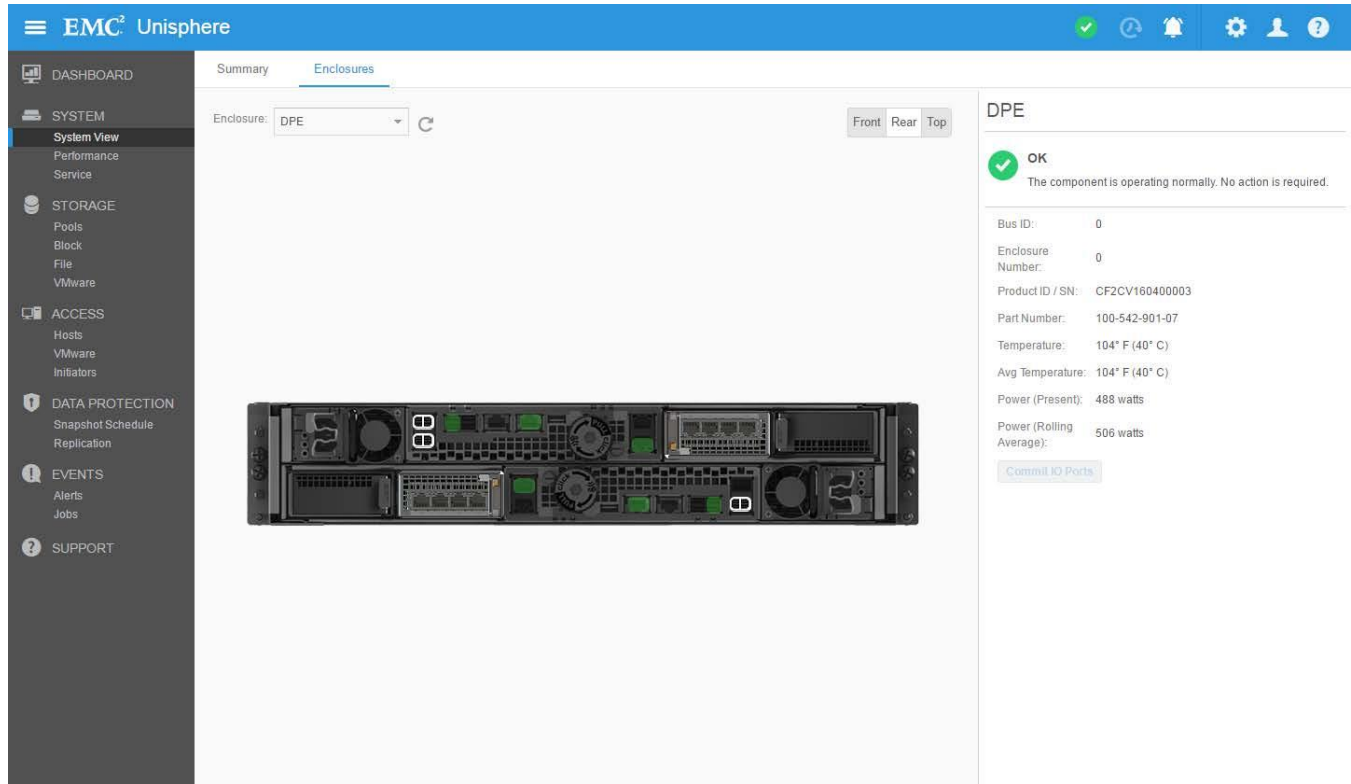


Figure 10: Unisphere Example

Unisphere is integrated with VMware vSphere to be aware of the virtual machines and provide storage system management. Additional awareness for management of applications controlling storage is built into Unisphere.

Evaluator Group Comment: *This new management tool answers many of the criticisms of customers where each different storage system had its own management software and “style.” This is an improvement over the earlier Unisphere. Overall, this helps make storage easier to manage. The features look very advanced and should help new customers that are not already used to existing tools. Existing customers should find the new tools to be much easier to use.*

Unisphere Central

Unisphere Central is a software application that remotely monitors status and activity of Unity XT, VNX, VNXe, and Clariion systems. Central can aggregate status, alerts, host details, performance, capacity, and storage usage information from systems, allowing administrators a view across the entire midrange Dell EMC environment. Unisphere Central is delivered as an OVF file that can be installed in an ESXi virtual machine.



Data Migration

The VNX SAN Copy feature is now Unity Data Migration. Data Migration is a storage system-based LUN replication that moves LUNs/Volumes between Unity XT and legacy Dell EMC VNX systems. At some point, additional storage system sources will be supported.

Fully Automated Storage Tiering - Virtual Pools (FAST VP) for Hybrid Unity XT Systems

The FAST VP function is an embedded software feature that provides tiering within the hybrid Unity XT systems. The tiering is between the different classes of devices that may be configured SSDs, SAS disks, and NL-SAS disks. The tiering is using policies to control placement of data to provide the tradeoff between the cost of the different classes of storage devices and the different performance characteristics. Data can be migrated between tiers in one gigabyte increments based on policies. This sub-LUN tiering across defined virtual pools provides the ability to move highly active data into the highest performing drives (usually SSDs) and less active data to high capacity drives. The default is automatic tiering with manual controls if desired.

Evaluator Group Comment: This is another refinement on the internal tiering that is more sophisticated in understanding data usage at the block level rather than a gross LUN level. This is being done with no hints or interaction from the operating system or applications. Certainly that would be a better architecture but that is not possible in an open systems environment with a general purpose storage system.

In operation, FAST moves a slice, which is 256MB of data (of 8KB chunks), based on the activity for data in the slice. The activity is an I/O count of reads or writes to data in the slice and is done continuously as a background collection activity on LUNs in a storage pool. The I/O count is maintained over time and the recent activity has a higher weighting factor than earlier activity. The activity is counted over a 24-hour period but activity that occurs 24 hours earlier has less than half weight. The collected data is analyzed on an hourly basis to create an ordered list of active slices within a pool.

Relocation of tiers can be either automatic through the scheduler or initiated manually. 1GB slices are promoted to higher tiers based on the order in the list. Slices are only demoted to a lower tier if there is less than 10% of the tier space for promoting other, higher priority slices.

FAST Policy - The Flash 1st policy starts with putting data in the highest tier, assumed to be solid-state disks and then move to lower tiers over time based on decreasing activity for data. Another policy change allows slices to be migrated to different tiers to auto load-balance across tiers. The FAST policies are:

Auto-tier – Default setting when LUNs are created in a virtual pool. Slices will be automatically located based on activity levels.

Highest available tier – Set for LUNs that require the highest performance, FAST will prioritize movement of slices for those LUNs over any other settings.

Lowest available tier – Set for LUNs that do not require high performance. FAST will leave slices for those LUNs on the lowest tier regardless of activity level.



FAST Cache

Fast Cache is an additional feature added to allow SSDs to be used as an extension to the system cache. This feature can provide a very large read/write cache in the Unix to optimize specific workloads. Fast Cache uses 64KB slices of LUNs for the granularity of caching. FAST Cache and FAST VP can be used together to handle bursty data activity in cache and the more frequently referenced data over time to be moved into higher tiers.

The Unity XT has a warm-up feature for cache where a single hit for access to data will put data into cache until cache is 80% full. After 80% full, it takes three hits on data before it is put into cache.

Encryption - Data @ Rest Encryption

Encryption of data at the array level is supported using function in the SAS adapters. All keys are managed internal to the Unity XT system. AES-256 encryption is used and the implementation has been submitted for FIPS-140-2 level 1 certification.

RecoverPoint Basic and VMs

RecoverPoint Basic is the product designed for continuous data protection and for replication of LUNs and VMware VMDK datastores. The RecoverPoint family is based on the replication software developed by Kashya Corporation, which Dell EMC acquired in 2006. The base RecoverPoint license can be upgraded at an unspecified cost to the user to support a heterogeneous any-to-any replication environment (RecoverPoint). This software supports asynchronous replication with consistency groups.

RecoverPoint Basic supports synchronous local copy (within the same system) and CRR is asynchronous remote copy (to another Unity XT either in the same location or at a separate location over distance). A RecoverPoint Basic system consists of the following components: at least one RecoverPoint appliance (RPA); RecoverPoint software installed on the RPA; RecoverPoint enabler for the array-based write splitter; RecoverPoint host agent software; and RecoverPoint storage replication adapter for VMware Site Recovery Manager is necessary.

RecoverPoint Basic is a scaled down version of RecoverPoint, designed specifically for systems, which support write-splitting technologies. RecoverPoint Basic supports a maximum of one array for continuous data protection and two systems for continuous remote replication.

RecoverPoint Basic replicates data by intercepting the application writes through the use of write-splitting modules. The Unity XT array-based splitter runs in each storage processor of the array and will "split" or mirror all writes to Unity XT sending one copy to the original target and the other copy to the RecoverPoint appliance.

Snapshot (redirect on write)

The redirect on write solution is a space efficient copy that only requires space for the blocks that are modified. Unity XT snapshots are used for both LUNs and filesystems – a single implementation for the unified environment. Snapshot works with the storage pooling implementation and allocates 8K blocks as minimum increment for snapshots. A Snap Mount Point is created as the LUN or filesystems reference for the snapshot. This allows usage of the LUN or filesystem without requiring another SCSI scan.

All snapshots can be configured with automatic deletion policies to contain the proliferation of snapshots and unneeded consumption of storage. Snapshots can be controlled by Unisphere, CLIs, or through the RESTful API.

Read/write access to the snapshot is allowed. The user is able to do a rollback to restore a snapshot copy to the original volume. Cascading snapshots or a snap of a snap can be performed. Consistency groups are supported for coordinated snapshots.

Native Synchronous Replication

Synchronous replication for Unity XT is called Native Synchronous Replication. Previously, the VNX had an optional feature named MirrorView/S used for synchronous replication. The Unity XT replication feature is native to the data services and works for LUNs only (block mode).

Synchronous replication involves replicating write data from one storage system to another before signaling complete to the host that originated the write operation. The following diagram depicts the synchronous operation. With Native Synchronous Replication, write ordering is preserved for the volumes within the single originating storage system. A maximum of 100Km distance or less than 10ms response time is supported for the replication but the actual distance depends on the amount of delay for the synchronous operation that can be tolerated for the application.

Native Synchronous Replication can be initiated and controlled by Unisphere, a CLI, or RESTful APIs. The first Fibre Channel port configured on the system is used for synchronous replication. VLAN tagging is not supported.

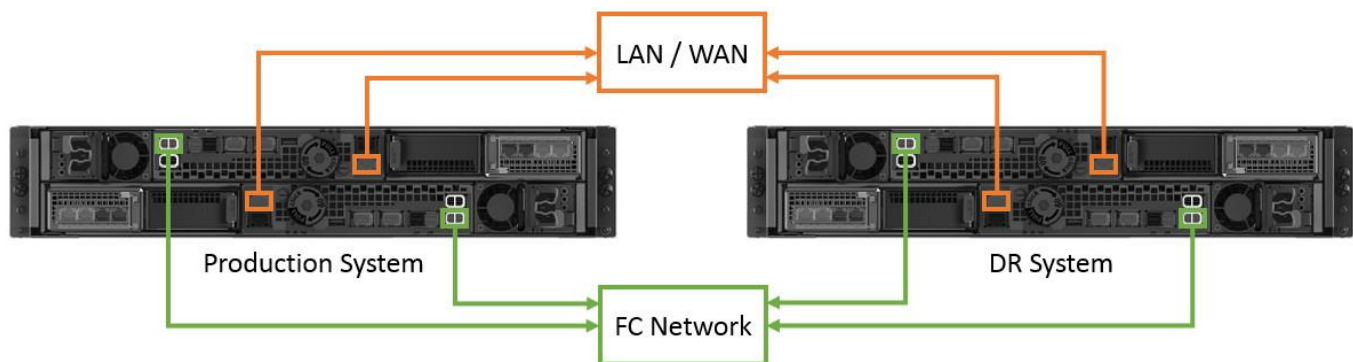


Figure 11: Synchronous Replication Connections

Consistency Groups that contain one or more LUNs are supported. The source and destination system replication resources must be of the same type. For Unity XT, these resources are: LUNs, Consistency Groups, and VMware VMFS Datastores. Both uni-directional and bi-directional replication are supported.

The Operating Environment 5.2 update added additional support for file synchronous replication.

Native Asynchronous Replication - Block (LUN) and Filesystem replication

Asynchronous replication for Unity XT is called Native asynchronous Replication. Previously, the VNX had an optional feature named MirrorView/A used for asynchronous replication. The Unity XT replication feature is native to the data services and works for LUNs (block mode) or filesystems (file mode). Asynchronous replication is done over Ethernet IP connections only. Any Ethernet port may be used for asynchronous replication.




Figure 12: Asynchronous Replication Connectivity

Unity XT snapshots are used for asynchronous replication providing a periodic remote copy implementation. The first replication will be for a full image and subsequent copies will be for changed data only. Dell EMC has provided the following diagram to explain operation.




Figure 13: Asynchronous Replication Operation (Source: Dell EMC)

Consistency Groups that contain one or more LUNs are supported. The source and destination system replication resources must be of the same type. For Unity XT asynchronous replication, these resources are LUNs, Consistency Groups, VMware VMFS Datastores, Filesystems, NAS servers, and VMware NFS Datastores.



Both uni-directional and bi-directional replication are supported as well as one-to-many and many-to-one. With asynchronous periodic replication, updates at the target site occur at a set interval based on the defined Recovery Point Objective.

CloudIQ

Dell EMC uses telemetry data from storage systems including PowerMax for monitoring and storage analytics. Functions include proactive support with health checks and reporting such as capacity consumption and performance. Telemetry data is collected and analyzed with a SaaS product called CloudIQ.

Unity XT NAS

The NAS implementation for Unity XT is a 64-bit filesystem that operates in the Unity XT environment, executing on the storage processor. Virtual file servers are known as NAS Servers and may operate on either storage processor. A NAS Server must be configured before creating a filesystem and is logically isolated from other NAS Servers. Each NAS Server has its own unique information such as name services, file sharing protocol, active directory domain settings, Unix directory services, user mappings, data protection settings and up to 10 interfaces.

NAS Servers use LUNs configured in the Unity XT system as filesystem storage, leveraging storage pooling, FAST VP, and FAST cache functions.

Model	VSA	380 / F	480 / F	680 / F	880 / F
Filesystems	32	500	750	1,000	1,500
Filesystem Size	50 TB	256 TB	256 TB	256 TB	256 TB
NAS Servers	16	90	90	128	128
Subdirectories per directory	100M	100M	100M	100M	100M
Files per filesystem	32B	32B	32B	32B	32B
File names per directory	10M	10M	10M	10M	10M
ACL IDs	4M	4M	4M	4M	4M

Table 6: Unity XT NAS Specifications



NAS Protocols

Unity XT NAS supports a number of protocols including the following: NFS v3, & v4/V4.1 CIFS/SMB1, 2, 3, FTP, SFTP.

Data Protection

Unity XT offers methods to protect user data including the following:

- NDMP Backup – two-way NDMP is supported
- MetroSync file replication

NDMP

The Network Data Management Protocol (NDMP) was developed primarily for backing up NAS systems from remote management applications. NDMP not only can retain permissions for both NFS and CIFS clients, but it also provides the ability to share backup infrastructure, including the backup applications along with SAN tape drives and libraries.

Unity XT NAS requires a read-only filesystem for an NDMP backup, which may be accomplished most easily by utilizing snapshot to create a point in time consistent copy that is used to backup. A snapshot checkpoint is automatically created when an NDMP backup is initiated.

Unity XT NAS backups, including NDMP backups, operate as a traditional type, known as PAX, or a new volume based backup method known as Volume Backup (VBB). The PAX based method creates backups at the file level, and recursively descends through each directory, backing up each file in that directory.

NDMP Volume Backup reads a set of disk data blocks in an efficient manner compared to the method used for traditional, file-based backups. NDMP Volume Backup works only with Dell EMC-qualified vendor backup software, as listed in the Dell EMC NAS Interoperability Matrix. Unity XT NAS NDMP Volume Backup can be used when performing a full backup or an incremental backup.

Types of NDMP backup

- Three way backup – the backup is created on tape drives that are SAN attached to Unity XT
- Local Backup – the backup is created on tape drives attached to the Unity XT system
- Remote Backup – the backup is created on remote Unity XT systems tape drives that are attached to the remote system

The following type of restore operations are supported with NDMP volume based backups (VBB):

- Full destructive restore
- File Level restore

Types of PAX Restore:

- Normal
- DAR - direct access restore
- DDAR – directory DAR, and improved form of DAR that enhances performance



The Directory Direct Access Restore (DDAR) optimizes recovery by allowing the NDMP client to directly access backed up data anywhere on the tape set when used with a backup vendor that supports DDAR. DDAR is not compatible with volume backups.

Each Data Mover supports four concurrent NDMP sessions at one time. For example, a system can run three backup sessions and one restore session, simultaneously.

NDMP restrictions:

- NDMP based backups and restores of iSCSI LUNs is not supported
- NDMP does not follow symbolic links within filesystems, therefore NDMP backups do not include the target of the link

MetroSync File Replication

Ported from the earlier VNX2 software, Metrosync uses MirrorView/S to create a replica copy for a Virtual Data Mover (VDM). The replication of a VDM includes file systems, checkpoints, exports, and other control information. Failovers can be supported at the file system granularity and a third DR site can be configured when used in combination with MirrorView/A. The replication works with a second (or third) Unity XT system.

VAAI Support for NAS

The Unity XT system support VMware vSphere APIs for Array Integration (VAAI) for NAS introduced starting with vSphere 5.0. The features are exploited by VMware to accelerate performance and provisioning time.

- Full File Clone – enables the system to make copies of virtual disks (VMDK files) in the storage system
- Extended Statistics – enables vSphere to query space utilization. Information returned includes the size of the file and the space consumed by the file.
- Space Reservation – enables thick virtual disk files to be provisioned with the lazy-zeroed option.

Quotas

Both soft (advisory) and hard (restrictive) quotas may be set for the amount of storage space consumed, or the number of files utilized. These quota restrictions may be applied on a per user, a per group basis or in combinations.

The Unity XT NAS can track quotas using either a block or a file quota policy. The default policy is set to track the number of blocks utilized. Any file smaller than a filesystem block size of 8 KB is counted as 8KB towards the quota since the file consumes one 8 KB block on disk.

If quota policy is set to file-size, disk usage is calculated in logical file size, which means a 1 KB file counts as 1 KB in the quota. When using the file-size quota policy, Unity XT NAS quota usage is unaffected by file migration and recall operations because these operations do not impact the logical size of files.



Anti-Virus Scanning

Unity XT NAS supports the use of external virus (AV) scanning engines through a mechanism known as AVA (or an Anti-Virus Agent). AVA does not actually perform anti-virus scanning; instead, it acts as an interface and method for allowing an AV agent to scan a file. AVA is an Dell EMC proprietary interface and method of virus scanning. Another developing industry standard is the ICAP interface (Internet Content Adaptation Protocol).

The AV software on an AVA server in conjunction with virus scanning software such as Symantec Norton, McAfee or others is configured to scan network drives on access.

There are two primary methods for utilizing AV scanning, either real-time (as files are accessed) or scheduled scanning. These methods may be used individually or together.

The order of operations is:

1. File open via CIFS to NAS
2. NAS locks the file and ships it to the AVA server(s) via the agent service (AVA) on a remote server
3. Real-time AV software scans the file
4. Once determined clean, AVA agent tells NAS to release the lock

AVA is not a feature that is installed on the operating environment, it is an agent installed on a separate server, together with anti-virus software on a standalone server. AVA by itself does not perform virus scanning.

Policy-Based File Migration

Unity XT Policy-Based File Migration (formerly File Mover) enables administrators to place file based data according to policies. Administrators may create rules, known as policies that determine whether files should be migrated from primary storage to slower, less-expensive storage devices and back again. Commonly used migration policies are those that migrate older, less frequently used files or files of a particular type.

File mover reads file attributes and will move files by matching the rules or policies that have been established. File mover is able to utilize NFS, CIFS or HTTP to scan files. File Mover will not operate on a read-only filesystem, nor does it support the use of the NFS v4 protocol.

File Mover has added support to migrate files to clouds. Cloud targets include VirtuStream cloud Amazon S3 and Microsoft Azure.

When the policy and migration software finds a match, it migrates or moves the file from Unity XT. It does so by reading the file and then writing it to the secondary file server using NFS or CIFS. Next, the policy and migration software issues a command via the File Mover API to convert the file to a stub file. This releases the storage associated with the file and completes the migration operation. When a file is migrated, its data is moved to a file on secondary storage, and the original file is replaced with a stub file. The stub file contains all the metadata associated with the original file including permissions, timestamps, size and other attributes. The stub file also contains additional data associated with the File Mover functionality, including the location of the file content on secondary storage, the identity of the policy and migration software that migrated the file's content, and an information field in which the policy and migration software can store any additional information. A stub file consumes one inode



(one 8KB data block), plus the size of any alternate data streams associated with the file, in the primary file system on the Unity XT system.

When a write operation occurs, Unity XT uses the information in the stub file to locate the file on secondary storage. It then compares the modification time (mtime) and size it has stored to the current mtime and size of the file on secondary storage. This comparison ensures the stub file and migrated file remain synchronous. When these attributes match, file data is returned to the client to satisfy the read request. If the attributes differ, an error message appears.



Virtual Storage Appliance (VSA)

The software for Unity XT is also available as a virtual storage appliance, targeted for usage in remote replication from remote offices and department and for test/dev evaluations. The Unity VSA software runs in a VMware ESXi virtual machine. Most of the Unity XT features are available except for Fibre Channel interfaces. The follow chart gives more details about VSA requirements and support. Unity VSA may be freely downloaded and only requires licensing when put into production.

- Community Edition – Free
 - 4TB limit, single node, community support, test/dev usage
- Professional Edition – subscription 10, 25, 50 TB
 - Single node, Dell EMC support, production usage
- Requirements
 - 2 vCPUs (2+GHz)
 - 12 GB DRAM
 - vNICs – 4 @ 1 or 10 GigE
 - Protocols – iSCSI, NFS, SMB



Reliability, Availability, Serviceability Features

The Unity XT storage systems maintain data integrity and provide data availability via the following mechanisms:

- Hot swappable storage processors
- Mirrored write cache
- RAID protection levels 0/1, 5, 6
- Battery backup and cache vault to M.2 SSD in case of power failure
- Redundant power supplies
- Redundant cooling modules
- Redundant data paths
- Online upgrades
- Support for data-in-place upgrades

Redundant write cache

The Unity XT system has dual write cache that enhances data integrity and provides for higher performance for mirroring and data replication. The Unity XT cache also has battery backup and is destaged to the internal M.2 SSD to provide persistence.

PowerStore Metro Node

PowerStore systems can achieve business continuity through active/active stretched clusters with the use of an additional hardware appliance called Dell EMC PowerStore Metro Node. PowerStore metro node is a 1U appliance with support for 32Gb/s Fibre Channel. Two 1U metro nodes are required for each PowerStore cluster to achieve business continuity.



Performance

Dell EMC does not support running standard benchmarks of their systems and does not report standard benchmark data. Dell EMC does provide marketing information regarding performance however. Dell EMC has stated that the Unity XT 880F obtained approximately 900,000 IOPs at 8K datalength with 80% read and 20% write.

For the all flash models 380F, 480F, 580F and 680F system, the 8KB 100% read IOPs performance is listed at 260K, 610K, 790K and 880K respectively.

Evaluator Group Comment: The performance is really an unknown because industry standard benchmarks such as SPC (Storage Performance Council) and IOMARK have not been run and reported. This leaves customers with taking a chance on the vendor numbers or running the system in their own environment to determine if the system will meet requirements. Industry standard benchmark information would be very useful for customers.

Evaluator Group Comments

Dell EMC has answered critics of the implementation of “Unified” in the VNX system with a true unified block and file system appropriately called Unity. Unity differs from the earlier VNX significantly in that it has a new operating environment with an embedded Linux and containers for execution of the Unity XT software and other capabilities such as ESRS. Evaluator Group expects additional containers to execute storage tiering to cloud functions and data protection software.

Using elements of VNX and VNXe software in the new environment contributes to the stability of the system and, obviously, to the timely delivery of the system. Advanced data services come from the earlier implementation. Additionally, the 64-bit file system in use had been previously introduced on the VNXe3200.

Strengths:

- *A mature set of functional storage software providing low-risk for customers*
- *Dell EMC support and service*
- *Optimized Operating Environment to take advantage of processor cores and multi-threading and running containers for isolation of different software function.*
- *Focus on Flash as the primary offering. Hybrid offering as still available but will have no advantages as flash prices continue to decline and data reduction becomes available*
- *Virtual (thin) provisioning: A provisioning storage pool can be configured and through policy specification, applications may be enabled to use this pool for their storage requirements. The Virtual Provisioning support provides excellent reporting and management for the storage pool.*
- *Integration with VMware*
- *VAAI support for improved virtual machine provisioning*
- *Fully Automated Storage Tiering – Virtual Pool (FAST-VP) support for automated tiering within the hybrid Unity XT systems across storage devices of different type*
- *Unisphere implemented with HTML5 as a single management interface for storage. Advanced wizards to simplify deployment and provisioning*
- *Necessary advanced features of snapshot and remote replication*

Perceived Challenges:

- *The Unity XT systems are scale-up only. To increase performance in lockstep with capacity increases, the controllers must be replaced with the next bigger model up to the largest model. A scale-out implementation would have more options and allow greater granularity.*
- *There currently is no SAS host connection with the systems. We would think this would be very important for some customers.*
- *The largest model does not exactly match the largest VNX. This may not be a big issue but does change the thinking.*



More detailed information is available at <http://evaluatorgroup.com>

Copyright 2023 Evaluator Group, Inc. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose without the express written consent of Evaluator Group Inc. The information contained in this document is subject to change without notice. Evaluator Group assumes no responsibility for errors or omissions. Evaluator Group makes no expressed or implied warranties in this document relating to the use or operation of the products described herein. In no event shall Evaluator Group be liable for any indirect, special, consequential or incidental damages arising out of or associated with any aspect of this publication, even if advised of the possibility of such damages. The Evaluator Series is a trademark of Evaluator Group, Inc. All other trademarks are the property of their respective companies.