

# VMAX3 AND VMAX ALL FLASH WITH CLOUDARRAY

## HYPERMAX OS Integration with CloudArray

### ABSTRACT

With organizations around the world facing compliance regulations, an increase in data, and a decrease in IT spending, there is growing need for cost effective, online storage that can be used to scale out current infrastructures in order to meet business requirements. Dell EMC CloudArray is the gateway that enables users to take advantage of cloud storage with the benefits of near limitless capacity, reduced management, and lower cost. HYPERMAX OS provides support for CloudArray to be federated behind a VMAX3 or VMAX All Flash storage array, giving users more available storage while taking advantage of the arrays high performance and availability for critical applications. This paper explains how HYPERMAX OS integrates with CloudArray.

September, 2016.

The information in this publication is provided "as is." EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

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

EMC<sup>2</sup>, EMC, the EMC logo, FAST, Unisphere, VMAX, VMAX All Flash and VMAX3 are registered trademarks or trademarks of EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners. © Copyright 2016 EMC Corporation. All rights reserved. Published in the USA. 09/16, technical note, H14247.3

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

Dell, EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries.

## TABLE OF CONTENTS

<b>EXECUTIVE SUMMARY .....</b>	<b>4</b>
AUDIENCE .....	4
<b>CLOUDARRAY OVERVIEW .....</b>	<b>4</b>
CloudArray Appliance.....	4
CloudArray Features .....	5
<b>VMAX3 AND VMAX ALL FLASH OVERVIEW .....</b>	<b>5</b>
HYPERMAX OS .....	5
DX Emulations .....	5
eDisks .....	6
External Disk Group.....	6
Virtual RAID Group.....	6
Thin Devices (TDEV) .....	6
Data Devices (TDAT) .....	6
<b>CONFIGURING THE VMAX3/VMAX ALL FLASH AND SAN FOR INTEGRATION WITH CLOUDARRAY .....</b>	<b>6</b>
Configuring DX Directors.....	6
Zoning .....	7
Single Fabric with two external Fibre Channel Controller ports.....	8
Dual Fabric with two external Fibre Channel Controller ports.....	8
<b>EXTERNAL PROVISIONING .....</b>	<b>9</b>
<b>ENSURING DATA INTEGRITY .....</b>	<b>10</b>
<b>VMAX3 AND VMAX ALL FLASH CONFIGURATION CONSIDERATIONS.....</b>	<b>10</b>
<b>CLOUDARRAY CONFIGURATION CONSIDERATIONS.....</b>	<b>10</b>
<b>GENERAL CONSIDERATIONS AND LIMITATIONS .....</b>	<b>11</b>
<b>FACTORY RESTRICTIONS .....</b>	<b>11</b>
<b>MANAGING EXISTING DATA .....</b>	<b>11</b>
<b>SOFTWARE AND HYPERMAX OS REQUIREMENTS .....</b>	<b>11</b>
<b>CONCLUSION .....</b>	<b>11</b>

## EXECUTIVE SUMMARY

With the exponential growth of data being seen in enterprise environments, cloud storage is rapidly becoming a viable method for online, cost effective storage. Organizations are looking to reduce the high costs of IT spending and decrease footprint without sacrificing capability, scalability, availability and replication feature sets. HYPERMAX OS solves these problems by allowing CloudArray to be federated behind VMAX3 and VMAX All Flash storage arrays. This allows users to leverage the high performance and availability of primary storage while taking advantage of the benefits of cloud storage.

HYPERMAX OS integrates with CloudArray allowing tiering to the cloud or object storage in the form of a public or private cloud. Cloud enabled VMAX3 and VMAX All Flash storage lowers the total cost of ownership by moving infrequently used data to lower cost cloud storage. CloudArray addresses all the requirements for connecting to cloud based storage. Designed to combine the resource efficiency of the cloud with traditional, on premise storage, CloudArray enables organizations to scale their enterprise storage capacity without increasing physical footprint.

Whether used for archive/backup, file, or secondary data, CloudArray ensures instant accessibility allowing users to take advantage of the cloud's greater scalability, lower costs, and reduced maintenance.

## AUDIENCE

This white paper is intended for:

- Customers, including IT planners, storage architects, and administrators involved in evaluating, acquiring, managing, operating, or designing and implementing HYPERMAX OS and CloudArray solution.
- Dell EMC staff and partners, for guidance and development of proposals.

## CLOUDARRAY OVERVIEW

CloudArray is a storage software technology that easily integrates cloud-based storage into traditional enterprise IT environments. CloudArray is designed to combine the resource efficiency of the cloud with traditional, on-site storage, allowing organizations to scale their infrastructure and plan for future data growth. CloudArray makes cloud object storage look, act, and feel like local disk, seamlessly integrating with existing applications, giving a virtually unlimited tier of storage in one easy package.

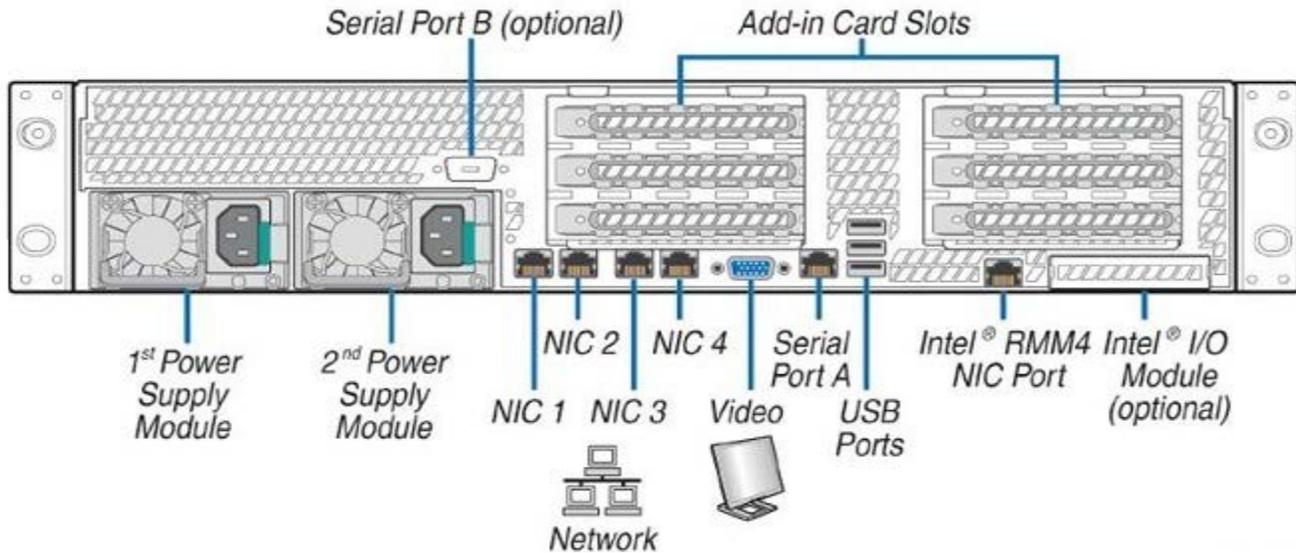
### CloudArray Appliance

Integration of HYPERMAX OS and CloudArray requires a hardware CloudArray appliance. The appliance supplies the physical connection capability from the storage array to cloud storage by presenting the appliance as an external device with Fibre Channel controller cards that allows the VMAX3 and VMAX All Flash to connect with the appliance.

The hardware CloudArray appliance is a 2U server that consists of:

- 20TB or 40TB usable local cache depending on licensing. 12x4TB drives in a RAID6 configuration
- 192GB RAM
- 2x2 port 8Gb Fibre Channel cards which will be configured in add-in slots on the physical appliance

Figure 1 shows a rear view of the CloudArray physical appliance.



**Figure 1. CloudArray Physical Appliance Rear View**

**CloudArray Features**

CloudArray is more than just a gateway to the cloud. CloudArray provides a complete storage solution that encompasses a range of features and allows VMAX3 and VMAX All Flash users to quickly begin connecting existing applications to cloud storage.

CloudArray uses local caching to reduce bandwidth and latency. HYPERMAX OS connects to the 40TB CloudArray disk-based cache through a fibre channel connection. This cache can be used for multiple data volumes according to use cases and stores a portion of users' most recently-accessed data locally to maintain performance while simultaneously replicating to the cloud.

In addition to the local caching to reduce bandwidth and latency, CloudArray also offers 256 AES in flight and at rest encryption, data compression and bandwidth scheduling to optimize network utilization.

**VMAX3 AND VMAX ALL FLASH OVERVIEW**

HYPERMAX OS allows an external disk array to provide physical storage for VMAX3 and VMAX All Flash volumes. This implementation required the development of new entities within the array that would allow it to attach to external array storage ports and configure external volumes to be used as physical storage.

**HYPERMAX OS**

Connectivity is implemented entirely in HYPERMAX OS and does not require any additional storage array hardware. Connectivity with an external array is established through the same fibre optic I/O modules currently used for configuring FAs for host connectivity or RFs for SRDF connectivity. Instead of running FA or RF emulation, however, the processors run a different type of emulation.

**DX Emulations**

An emulation, referred to as DX (DS eXternal), has been developed that adapts the traditional DS emulation model to act on external volumes as though they were physical drives. The fact that a DX is using external logical units, instead of a DS using internal physical disks, is transparent to other director emulations and to the HYPERMAX OS infrastructure in general. With respect to most non-drive-specific HYPERMAX OS functions, a DX behaves the same as a DS, which is the internal VMAX3 and VMAX All Flash Disk Adapter for connectivity to physical drives.

---

**Note: A DS is equivalent to a DA, or Disk Adapter, in previous generation VMAX arrays.**

---

## **eDisks**

An eDisk is a logical representation of an external volume when it is added into the configuration. The terms eDisk and external spindle both refer to this external volume once it has been placed in an external disk group and a virtual RAID group.

## **External Disk Group**

External disk groups are virtual disk groups that are created by the user to contain eDisks. Exclusive disk group numbers for external disk groups start at 512. External volumes and internal physical spindles cannot be mixed in a disk group.

## **Virtual RAID Group**

An unprotected virtual RAID group gets created for each eDisk that gets added to the system. The RAID group is virtual because eDisks are not protected locally by the array.

## **Thin Devices (TDEV)**

Thin devices are the host addressable Virtual Provisioning component. They are created with a size but no assigned RAID protection. They inherit the RAID protection of the Data devices contained in the pool where they are bound. In VMAX3 and VMAX All Flash, all host-addressable devices are thin devices.

## **Data Devices (TDAT)**

Data devices are a type of internal device that are dedicated to providing the storage for thin devices in a VMAX3 and VMAX All Flash array. They are configured in HYPERMAX OS as part of adding internal or external storage to the array and are configured automatically when an eDisk is virtualized into an external connectivity environment. There is a 1:1 relationship between a Data device and an eDisk within the configuration.

# **CONFIGURING THE VMAX3/VMAX ALL FLASH AND SAN FOR INTEGRATION WITH CLOUDARRAY**

## **Configuring DX Directors**

DX emulations are configured in dual initiator (DI) pairs like traditional DSs. They are fully redundant like DSs and, when necessary, a failing director fails over to the other, fully functioning director in the DI pair. Dell EMC requires a minimum of four paths to external devices. This means that at least four ports belonging to a single DX dual-initiator pair must be configured.

DX DI pairs are always configured on different directors on the same engine. For example, the engine shown in Figure 2 has two directors, each containing four I/O Modules with four ports each. Because it is valid to add DX emulation to both director 1 and director 2 and use any two ports on each of those directors, it is possible to create a valid DX configuration with two ports that physically reside on the same I/O Module. For example, using Director 1, ports 4 and 7, which are on the same I/O Module along with Director 2, ports 25 and 31, which are not, will pass pathing compliance check. This configuration is allowed even though an I/O Module replacement could affect both ports 4 and 7 requiring that all paths to external devices through Director 1 fail over to Director 2.

A better choice for the two ports from Director 1 would be ports 4 and 9. If that was the configuration, any of the four front-end (FE) fibre channel I/O Modules could be replaced without requiring a fail over to that director's DI partner. Recovering from a DI failover requires manual customer service intervention. It is not disruptive if a DX fails over and it is not disruptive to recover from this state, but it should be avoided, if possible.

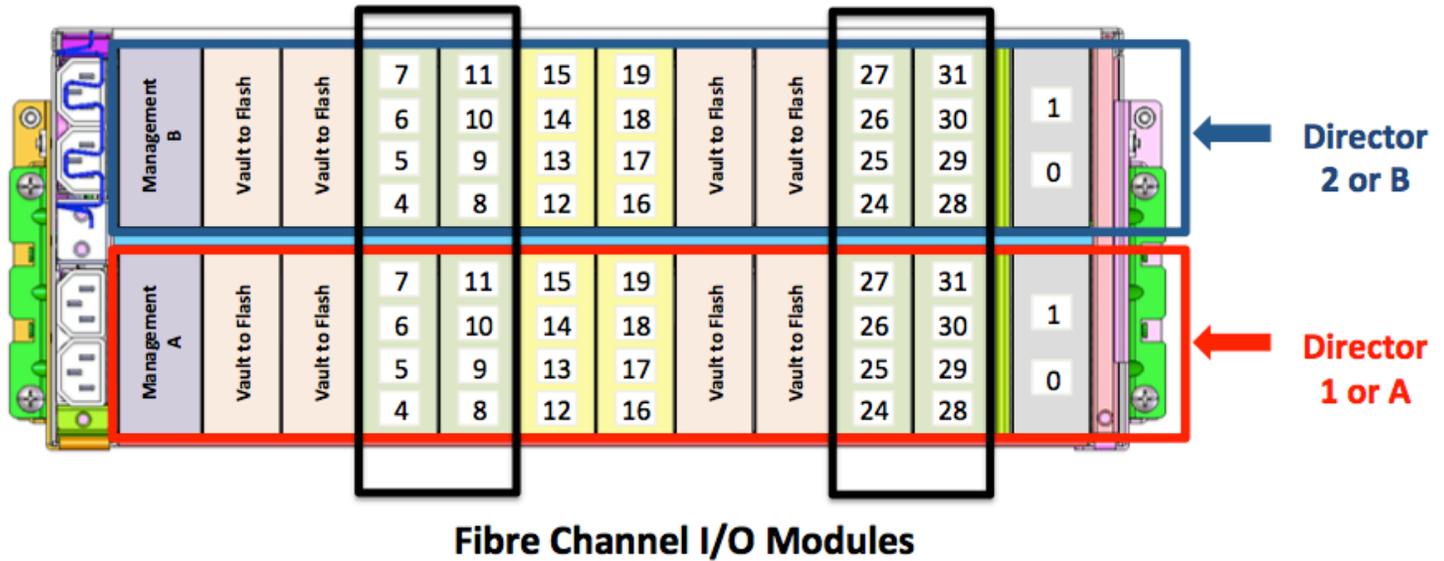


Figure 2. Single Engine with 8 Fibre Channel I/O Modules

**Note:** DX directors are not user-configurable. Dell EMC Customer Service must create them.

If converting FA ports to DX ports, any previously assigned devices must be unmapped and unmasked and the FA ports must be removed from any port groups.

### Zoning

When configuring zoning, using a single-target per single-initiator (1:1) zoning scheme is preferred. If the FC switch zone count limitation has been reached, it is possible to use single-target per multiple-initiator (1:Many) zoning.

The zoning examples provided below allow the servicing of the components within the environment without incurring data unavailability and are required in order for the configuration to be supported. The potential service activities include:

- Cable changes and individual FC port servicing
- VMAX3 and VMAX All Flash director and I/O module replacement
- CloudArray firmware upgrade
- SAN fabric servicing

Proper zoning also ensures that a failing switch or storage controller will not cause a DX to fail over, which requires manual recovery.

Figures 3 through 5 below show examples of how to meet these requirements in single and dual-fabric environments that represent common configurations. These zoning requirements are in addition to existing connectivity requirements consisting of two, physically independent SCSI I-T nexus per DX for each external volume to be configured through a DX pair. These base connectivity requirements are checked by HYPERMAX OS during the eDisk configuration process. If the zoning does not pass HYPERMAX OS compliance checks, attempts to configure eDisks will fail.

### Single Fabric with two external Fibre Channel Controller ports

Single fabric connectivity is supported, though it does not provide the redundancy of a dual fabric configuration. As an example, take a HYPERMAX OS with CloudArray environment configured with a single-engine array with DX emulation running on directors 1 and 2 with a CloudArray, which will contain two Fibre Channel controllers.

Figure 3 shows an example of single-fabric zoning with this configuration.

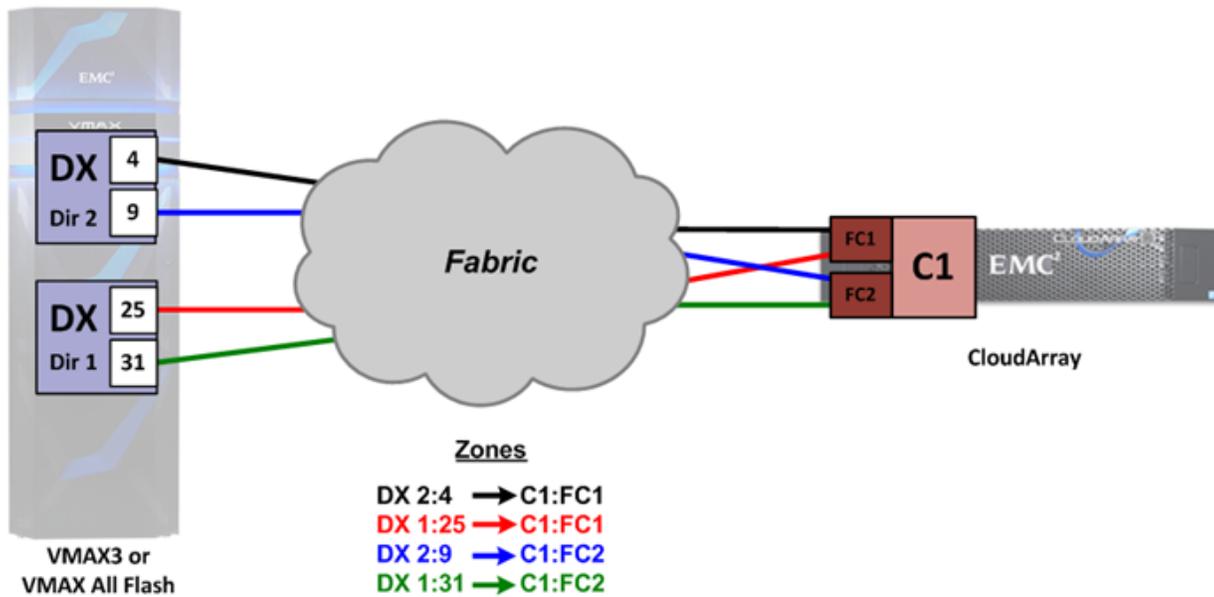


Figure 3. Single Fabric Zoning with CloudArray

### Dual Fabric with two external Fibre Channel Controller ports

Though single-fabric connectivity is supported, best practice for redundancy is achieved by using dual fabrics. One DX initiator port from each DX director pair must connect to one fabric while the other DX initiator port connects to the other fabric.

As an example, take an environment configured with a single-engine VMAX3 or VMAX All Flash array with DX emulation running on Directors 1 and 2 with a single CloudArray, which contains two Fibre Channel controllers.

Figure 4 shows an example of dual-fabric zoning with this configuration.

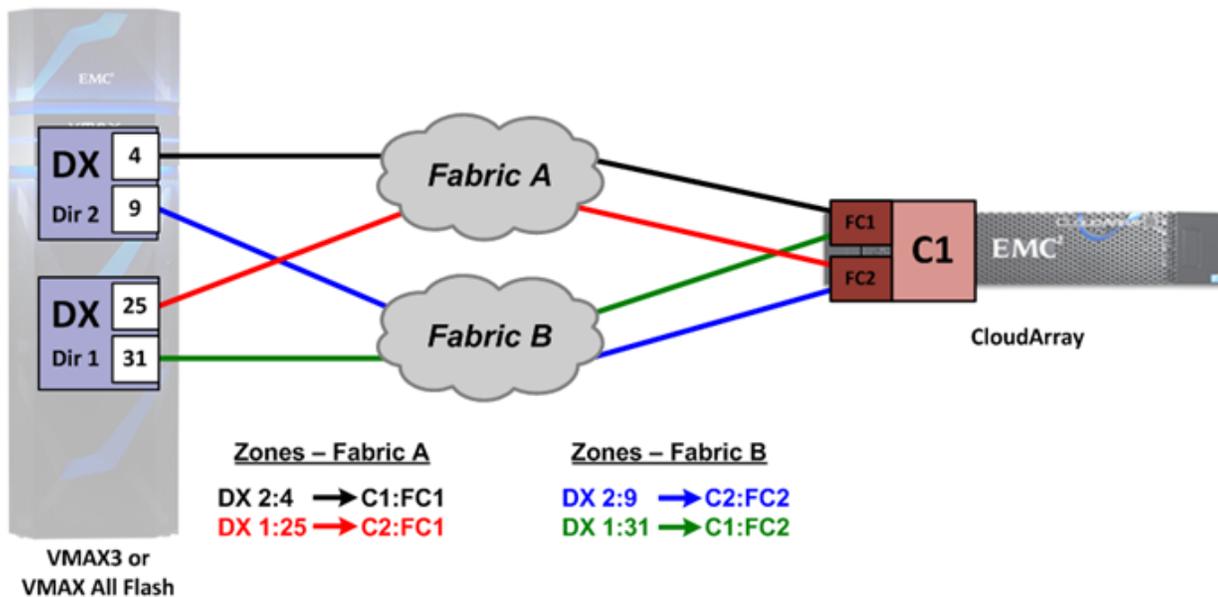


Figure 4. Dual Fabric Zoning with CloudArray

Note: Figures 3 and 4 show the logical, not physical connections from the switch to the CloudArray. In both diagrams, from the storage array, there are four physical connections to the fabric, one from each DX port to the switch(es). However, there are only two physical connections, one for each external storage port, from the fabric(s) to each CloudArray Fibre Channel controller.

## EXTERNAL PROVISIONING

External Provisioning is the mode of operation that HYPERMAX OS uses to virtualize volumes on the CloudArray array as eDisks on the storage array. It allows the user to access these logical units that exist on CloudArray as raw capacity for new devices. When configuring an external volume, HYPERMAX OS creates an eDisk and adds it to the specified external disk group. There is a 1:1 relationship between the virtualized volume on the CloudArray and the eDisk that gets created on the VMAX3 or VMAX All Flash.

External disk groups are separate from disk groups containing internal physical disks, and start at disk group number 512. Because RAID protection is provided by the external array, each eDisk is added to its own unprotected virtual RAID group. Thin volumes can then be created to present to hosts using capacity from the external disk group. External provisioning should only be used with external volumes that contain no data or unwanted data. These devices are VTOCed as part of the eDisk configuration process, and any data residing on the volume prior to adding it as an eDisk will be destroyed.

Figure 5 shows HYPERMAX OS with CloudArray external provisioning.

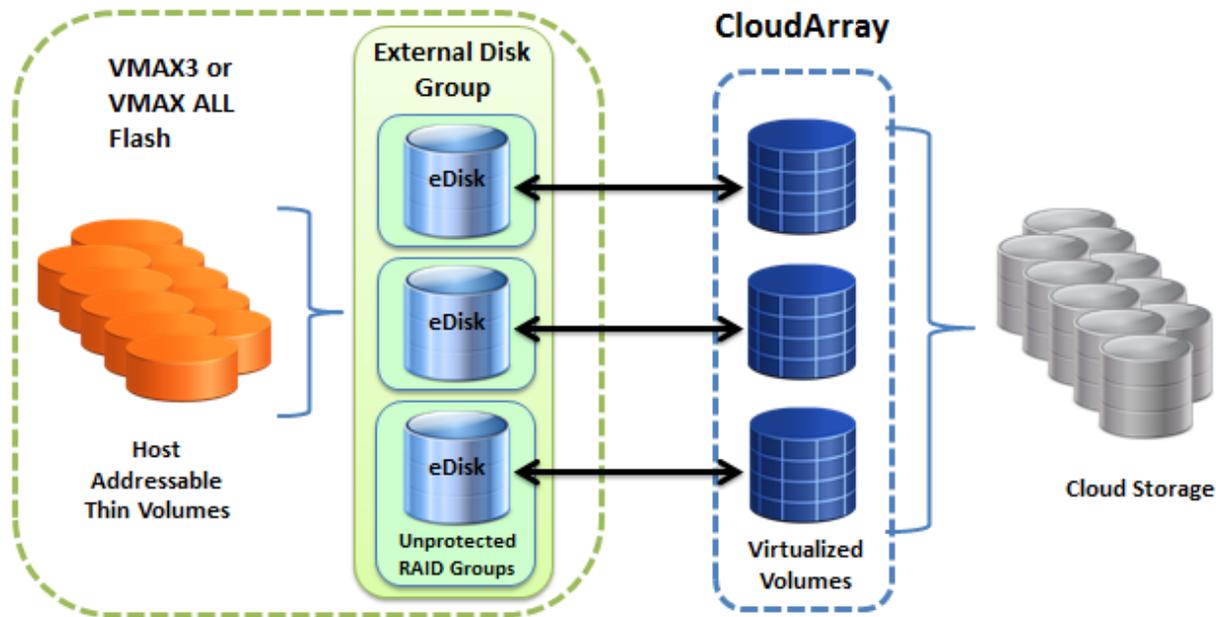


Figure 5. HYPERMAX OS with CloudArray External Provisioning

## ENSURING DATA INTEGRITY

VMAX3 and VMAX All Flash use a basic CRC mechanism to detect any data corruption that may occur. When data is read into cache from local disk, it is protected by a standard block-level CRC, which is based on the industry standard T10 Data Integrity Field (DIF) block.

CloudArray's encryption is also a critical component in ensuring data integrity. CloudArray segments its cache into cache pages and, as part of the encryption process, generates and assigns a unique hash to each cache page. The hash remains with the cache page until that page is retrieved for access by a requesting initiator. When the page is decrypted, the hash must match the value generated by the decryption algorithm. If the hash does not match, then the page is declared corrupt. This process helps to prevent any data corruption from propagating to an end user.

## VMAX3 AND VMAX ALL FLASH CONFIGURATION CONSIDERATIONS

- CloudArray capacity must be configured into its own SRP. Multiple CloudArray appliances can have their capacity virtualized in the same SRP, but if an SRP has any CloudArray capacity in it, it must be the only type of storage in this SRP. This requirement is mandatory, but is not enforced by HYPERMAX OS or management software.
- An SRP configured for HYPERMAX OS with CloudArray only supports an Optimized service level. Attempting to bind a storage group with a different SLO is rejected.
- No local or remote replication (including SRDF, TimeFinder, and SnapVX) is allowed using VMAX3 or VMAX All Flash volumes with capacity provisioned from a CloudArray SRP. This restriction is mandatory, but is not enforced by HYPERMAX OS management software.
- The cumulative front-end throughput limit of all storage groups provisioned to any given CloudArray appliance is 400MB/s.

## CLOUDARRAY CONFIGURATION CONSIDERATIONS

- The CloudArray appliance supporting fibre channel connectivity and qualified for HYPERMAX OS integration comes in two licenses. One has 20TB of CloudArray cache and one has 40TB. The maximum qualified capacity of these appliances is 120TB and 240TB, respectively.
- The CloudArray appliance used for HYPERMAX OS with CloudArray must be dedicated to it.
- 5 caches of 4TB each should be configured for the 20TB license and 10 caches of 4TB should be configured for the 40TB license.
- The maximum cache to cloud capacity ratio that is qualified is 6:1.
- Each CloudArray volume should be 4TiB and should only be expanded by a multiple of that value.

- There are a minimum number of CloudArray volumes that are required for HYPERMAX OS with CloudArray. This minimum depends on the DX configuration in the VMAX3 or VMAX All Flash array. There must be two CloudArray volumes for each engine in the system. For example, this means that on a single VMAX3 or VMAX All Flash engine, the minimum will be 2 volumes, on an eight engine system, it will be 16 volumes (a minimum of 64TiB – 16\*4TiB virtualized).
- Maximum capacity is reached with 30 volumes virtualized for the 20TB license and 60 volumes for the 40TB license.
- Once the minimum number of volume has been virtualized, any numbers of 4TiB volumes can be added up to the maximum allowed by the license.
- Volumes should be allocated to the cache in a round robin fashion: The first 5 volumes (10 volumes for the 40TB) end up with 1:1 cache ratio. The next 5 (20 for the 40TB license) will get 2:1, and so on, until a ratio of 6:1 is reached, matching the 120TiB or 240TiB maximum licensed capacity.

## GENERAL CONSIDERATIONS AND LIMITATIONS

- HYPERMAX OS with CloudArray supports fabric connectivity only. Direct attach is not supported.
- A maximum of 1024 external volumes per DX emulation (2048 per system) can be virtualized as eDisks.
- The maximum number of logical paths to each external SCSI Logical Unit is 4.
- The maximum external logical unit size is 64TiB.
- A supported cloud provider must be used in a CloudArray configuration. For more information on compatible providers, go to <http://www.emc.com/collateral/technical-documentation/h14387-emc-cloudarray-support-matrix.pdf>.

## FACTORY RESTRICTIONS

No pre-configuration is done in the factory. Some elements of a HYPERMAX OS and CloudArray SRP configuration require customer service engagement, but can be done online at any time after the deployment of the system providing that the VMAX3 and VMAX All Flash along with its cache have been correctly sized.

## MANAGING EXISTING DATA

VMAX3 and VMAX All Flash users can move existing data applications to the cloud once HYPERMAX OS and CloudArray are configured. Since CloudArray is bound to its own SRP, users can move an application storage group in and out of this SRP allowing legacy applications to be migrated to the cloud. All applications remain available during and after the move of a storage group.

## SOFTWARE AND HYPERMAX OS REQUIREMENTS

VMAX3, VMAX All Flash, and CloudArray require the following minimum software versions:

- HYPERMAX OS 5977.691.684
- CloudArray 6.0
- Solutions Enabler 8.1.0
- Unisphere 8.1.0

## CONCLUSION

HYPERMAX OS with CloudArray offers easily integrated cloud-based storage into traditional enterprise IT environments. CloudArray utilizes a hardware appliance that supplies physical connection capabilities from VMAX3 and VMAX All Flash to cloud storage using Fibre Channel controller cards. The CloudArray appliance leverages 40TB local disk-based cache that can be used as multiple data volumes, storing the user's most recently accessed data locally while replicating to the cloud. HYPERMAX OS with CloudArray combines the efficiency of cloud storage along with onsite primary storage, allowing organizations to scale to near limitless storage capacity, reduced management, and lower cost.