



CATE Storage and Backup Engineering
Citigroup Architecture & Technology Engineering

EMC VNX for SAN

CSBE-11059-SG

Solutions Guide

Version 2.0

©Citigroup Inc. 2015, All Rights Reserved.

This document contains material internal to Citigroup. The contents are protected by specific employee and contractor agreements regarding confidentiality. Reproduction for distribution outside of Citigroup is prohibited.

1 Version Control

Version	Date	Author	Reviewed	Comments
0.1	09-Sep-2011	Quang Le	Tim Armstrong	Peer Review
0.2	16-Sep-2011	Quang Le	Global Stakeholders	Release for Comments
1.0	03-Oct-2011	Quang Le	-	Release for Production
1.1	20-Jan-2012	Tim Armstrong Sean Tai	Gary Johnson Stephen Hopkins	Peer Review
1.2	19-Apr-2012	Tim Armstrong	Global Stakeholders	See Document Version Control Release for Testing
1.3	21-May-2012	Tim Armstrong	-	Release for Production
1.4	18-Apr-2013	Tim Armstrong	David Lee	VNX Block OE 32 Release Peer Review
1.5	30-Apr-2013	Tim Armstrong	Global Stakeholders	See Document Version Control Release for Testing
1.6	31-May-2013	Tim Armstrong	-	See Document Version Control Release for Production
1.7	29-Nov-2013	Tim Armstrong	Global Stakeholders	See Document Version Control Release for Testing
1.8	16-Dec-2013	Tim Armstrong	-	See Document Version Control Release for Production
1.9	19-Sep-2014	Tim Armstrong	Global Stakeholders	See Document Version Control Release for Testing Release for Production
2.0	30-June-2015	Tim Armstrong Gary Johnson	Global Stakeholders	See Document Version Control Release for Testing Release for Production

2 Legal Notice

Information contained herein is for Internal Use and may be used only for business purposes authorized by Citigroup.

CITI, CITIBANK, CITICORP, CITIGROUP and the Arc Device are trademarks and service marks of Citicorp and are used and registered throughout the world. Citibank, Citicorp, Citigroup and their subsidiaries also claim rights in certain other trademarks and service marks contained in these pages.

3 Contact

Please address any comments or questions regarding the information contained in this document via email to [*CATE Storage and Backup Engineering](#).

4 Latest Versions

It is recommended that the latest version of the document be used from the [CATE Storage and Backup Engineering website](#). Newer versions may include certification on additional platforms, or comments on issues discovered after deployment.

5 Regional Reviewers

The purpose of this section is to list all the stakeholders who reviewed this document during the Request for Comment process.

Region	Current Version Reviewed by	Previous Version Reviewed
Asia Pacific	No feedback	
EMEA	No feedback	
Japan	No feedback	
Latin America / Mexico	No feedback	
North America	No feedback	

Table 1 - Regional stakeholder review

6 Table of Contents

1	Version Control	2
2	Legal Notice	2
3	Contact.....	2
4	Latest Versions	3
5	Regional Reviewers	3
6	Table of Contents.....	4
7	Introduction	8
8	Assumptions and References	9
8.1	Assumptions.....	9
8.2	How to use this document.....	9
8.3	References	9
9	General Design	10
9.1	Overview	10
9.2	Citi Support and Management	10
9.2.1	Array Management	10
9.2.2	Enterprise Monitoring and Alerting	10
9.2.3	Email Alerting.....	10
9.2.4	Audit Logs	11
9.2.5	Reporting, Performance, Change Management and Capacity Planning.....	11
9.3	Remote Vendor Support	11
9.3.1	Email Alerting.....	11
9.3.2	EMC Secure Remote Support	11
9.4	Host Connectivity	11
9.5	Software and Features	12
9.6	Hardware	14
9.6.1	VNX 2 Hardware	14
9.6.2	VNX Hardware	17
10	Storage Best Practices	20
10.1	Virtual Provisioning and Traditional (Classic) Storage Provisioning	20
10.2	Drive Selection	20
10.3	SAS Back-end	21
10.3.1	Enclosure ID's	21
10.4	Vault Drives and DPE / DAE OS	22
10.4.1	Vault Drive Replacement.....	22
10.5	Placement of Drives	22
10.5.1	General Disk Placement Guidelines.....	22
10.5.2	Placement of Flash drives for FAST Cache	22
10.6	Drive Sparing.....	23
10.6.1	Drive Sparing - Multicore RAID (VNX OE for Block 33)	23
10.6.2	Drive Sparing - FLARE RAID (VNX OE for Block 32 and earlier)	25
10.7	SP Cache	26
10.7.1	Multicore Cache (VNX OE for Block 33)	26
10.7.2	FLARE cache (VNX OE for Block 32 and earlier)	29
10.8	FAST Cache.....	31
10.8.1	FAST Cache Overview	31
10.8.2	Multicore FAST Cache (VNX OE for Block 33)	31
10.8.3	Supported FAST Cache Disk Configurations	31
10.8.4	FAST Cache Guidelines	32
10.8.5	FAST Cache Failure Handling.....	32
10.8.6	FAST Cache Monitoring in Unisphere.....	33
10.8.7	FAST Cache Configuration.....	33
11	Storage Provisioning Best Practices.....	34
11.1	Storage Provisioning Options.....	34
11.2	Virtual Provisioning.....	34

11.2.1	Pool LUNs (Virtual Provisioning)	34
11.2.2	Thick LUNs (Virtual Provisioning).....	35
11.2.3	Thin LUNs (Virtual Provisioning)	35
11.2.4	Storage Pools (Virtual Provisioning).....	35
11.2.5	Storage Pools with VNX File	36
11.2.6	Protection Levels / RAID Types (Virtual Provisioning)	36
11.2.7	Storage Pool Creation (Virtual Provisioning)	37
11.2.8	Storage Pool Expansion (Virtual Provisioning).....	38
11.2.9	Limits with Virtual Provisioning	38
11.2.10	LUN Creation (Virtual Provisioning)	39
11.2.11	LUN Expansion and Shrinking (Virtual Provisioning)	39
11.2.12	Monitoring, Adding and Deleting Pool Capacity	40
11.3	Traditional/Classic Storage Provisioning	41
11.3.1	Traditional/Classic RAID Groups Overview.....	41
11.3.2	RAID Level (Traditional/Classic RAID Groups)	41
11.3.3	Sequential Optimization.....	42
11.3.4	Traditional / Classic LUNs	42
11.3.5	Creating a Classic LUN	43
11.3.6	LUN HLA (or HLU).....	44
11.3.7	MetaLUNs.....	44
11.3.8	LUN Migration (Classic LUNs)	44
12	Data Protection Best Practices	45
12.1	Local Protection.....	45
12.2	Remote Protection.....	45
12.3	Backup Standards	45
13	Security Best Practices	46
13.1	Security Standards	46
13.2	Audit Logging	47
13.2.1	SEMS Audit Log Reconciliation.....	47
13.3	Data Security.....	48
13.3.1	Disk Sanitization	48
13.3.2	Disposal of failed disks	48
13.4	CITMP Information Security Policy Convergence	48
13.4.1	Information Ownership	48
13.5	Role Based Access Controls (RBAC)	49
13.5.1	Roles Defined in UniSphere	49
13.5.2	Roles defined for use in Citi.....	50
13.6	Functional IDs	50
13.7	Storage Domains.....	51
13.7.1	Setting Up the Storage Domain.....	51
13.7.2	Adding Storage Arrays to an Existing Domain	52
13.7.3	Setting Up Multi-Domain Management	52
13.7.4	Adding a Domain User	54
13.7.5	Global/Local User Accounts	57
13.8	Communication Interfaces	57
13.9	Unisphere Management Server Configuration.....	58
13.9.1	Management Server Listening Port.....	58
13.9.2	Configuring the Management Access Control List (ACL).....	58
13.9.3	SSL/TLS Certificate Renewal	59
14	Alerting and Monitoring	60
14.1	SNMP and Email Alerting.....	60
14.1.1	Adding an Array to the Monitoring System.....	60
14.1.2	SNMP alerting with Tivoli NetCool	60
14.1.3	Creating an Alerting Template.....	60
14.1.4	SNMP MIB	63
14.2	Call Home and Email Alerting	64
14.3	Remote Support (EMC ESRS).....	64
14.4	Event Monitor	64

14.4.1	Configuring Central Monitoring Server	64
14.4.2	Configuring and Applying Event Monitoring Template	66
14.5	Performance Monitoring using Unisphere Analyzer	68
14.5.1	Unisphere Analyzer Settings	69
15	Host Connectivity and Configuration	70
15.1	Symmetric Active/Active - Multicore RAID (VNX OE for Block 33)	70
15.2	Host file system alignment	70
15.3	Unisphere Host Agent	70
15.4	SAN Fabric and Zoning Guidelines	71
15.5	Host Registration	72
15.5.1	Initiator Settings	72
15.5.2	ALUA	72
15.5.3	Manually Registering a New Host	73
15.5.4	Replacing an HBA on a Host	73
15.5.5	Removing an HBA	73
15.5.6	Verify Host Registration	74
15.5.7	Changing the Failover Mode for a Host	74
15.6	Allocating Storage to Hosts	74
16	Local Protection	75
16.1	SnapView Configurations	75
16.2	Snapshots	76
16.2.1	Reserved LUN Pool	76
16.2.2	Snapshot Sessions	77
16.2.3	Snapshot Configuration	77
16.3	Clones	78
16.3.1	Clone Guidelines	78
16.3.2	Clone States	78
16.3.3	Unisphere Clone Wizard	79
17	Off-Array Management Server	80
17.1	Server Requirements	80
17.2	Off-Array Management Station Software Installation	81
17.2.1	Unisphere Client Installation Procedure	81
17.2.2	Unisphere Host Agent Installation Procedure	82
17.2.3	Navisphere CLI Installation Procedure	84
17.2.4	Upgrade Procedure	85
17.3	Modifying agent.config File	86
17.4	Navisphere CLI Configuration	87
17.4.1	Creating the User's Security File	87
18	Setting Up the Storage System	88
18.1	Physical Setup	88
18.2	Upgrading Array Microcode	88
18.3	Array Software	88
18.4	DPE/DAE - OS Configuration / Initialization	89
18.5	Storage System Properties	90
18.5.1	General Settings	90
18.5.2	SP Cache	91
18.5.3	FAST Cache	91
18.5.4	Software	91
18.5.5	Time Settings	91
18.6	Service Processor (SP) Properties	92
18.6.1	SP Network Settings	92
18.6.2	ALPA	92
18.6.3	Agent	92
18.6.4	Fibre Speed	92
19	Appendix A: iSCSI for VNX	93
19.1	Restrictions for iSCSI	93
19.2	VNX Array Configurations for iSCSI	93
19.2.1	100TB Useable Configuration	93

19.3	Standard Array Setup	94
19.4	Pool Configuration for iSCSI	94
19.5	How to setup iSCSI on the VNX.....	94
19.5.1	ISCSI Port Configuration	94
19.5.2	iSCSI CHAP Security	98
19.5.3	iSCSI ISNS Service	98
19.6	How to register iSCSI Hosts.....	99
19.7	How to allocate LUNs to iSCSI Hosts	101
19.8	Host Configurations for iSCSI (SAN Stacks)	103
19.9	Host Data Migration.....	104
19.10	SRM Reporting.....	104
20	Document Version Control.....	105

7 Introduction

This document should serve as guide when configuring VNX series storage arrays in Citi's storage environment for block-based storage utilizing Storage Area Networks (SAN). The information compiled in this document has been tested and certified by CATE Storage and Backup Engineering.

This is a living document and will be updated accordingly as new products and code versions are certified in the Citi environment. Its main purpose is to provide the reader with a comprehensive view of the VNX storage arrays and their usage at Citi. However, the document does not cover the internals of the VNX.

This document is intended for use by all offices within Citi who intend to purchase and implement a VNX Mid-tier Storage Infrastructure for SAN (block storage) purposes.

For a detailed list of changes and enhancements in this document, please refer to the [Document Version Control](#) section.

8 Assumptions and References

8.1 Assumptions

This document assumes that the reader is familiar with the basics of VNX block storage, including: hardware platforms, VNX OE for Block, RAID groups, LUNs and SnapView snapshots/clones operations. Although experience of a previous version of VNX OE for Block or FLARE is not a requirement, familiarization with the product is strongly recommended before continuing to read this document.

8.2 How to use this document

This document is not intended to be read cover to cover, it is purely a reference document designed to augment the product documentation provided by EMC; provide focused and detailed insights of storage management and to consolidate important reference materials required by Citi's Information Technology Management Policies (CITMP).

8.3 References


Item	Location / Value
CATE Storage and Backup Engineering site	https://catecollaboration.citigroup.net/domains/platstor/gdse/products/default.aspx
CATE Storage VNX SAN standards	https://catecollaboration.citigroup.net/domains/platstor/gdse/products/techstdshmpgs/clariion.aspx
CATE Storage VNX NAS standards	https://catecollaboration.citigroup.net/domains/platstor/gdse/products/techstdshmpgs/celeerra.aspx
CATE Storage ECC standards	https://catecollaboration.citigroup.net/domains/platstor/gdse/products/techstdshmpgs/ecc.aspx
CATE Storage ESRS standards	https://catecollaboration.citigroup.net/domains/platstor/gdse/products/techstdshmpgs/esrs.aspx
CATE Storage SANscreen standards	https://catecollaboration.citigroup.net/domains/platstor/gdse/products/techstdshmpgs/sanscreen.aspx
CATE Storage SAN Stacks standards	https://catecollaboration.citigroup.net/domains/platstor/gdse/products/techstdshmpgs/SANStacks.aspx
CATE Storage Cisco MDS standards	https://catecollaboration.citigroup.net/domains/platstor/gdse/products/techstdshmpgs/ciscomds.aspx
EMC Navisphere Command Line Interface (CLI) Reference	https://support.emc.com/
Introduction to the New EMC VNX Series VNX5400, VNX5600, VNX5800, VNX7600, and VNX8000 – A Detailed Review – White Paper	https://support.emc.com/
VNX MCx Multicore Everything – White Paper	https://support.emc.com/
EMC VNX Best Practices for Performance and Availability	https://support.emc.com/
EMC VNX FAST Cache – A Detailed Review	https://support.emc.com/
EMC VNX Virtual Provisioning – White Paper	https://support.emc.com/
EMC Unisphere Online Help	Available with the Unisphere Client software.
Citi DNS and NTP configuration standards	http://ipns.ny.ssmb.com/documents/Product-Development/CA-NTP/GNAST/NTP_design_RFP.doc

9 General Design

This section explains the general design for the implementation of an EMC VNX for block-based storage, utilizing a storage area network (SAN) for host and storage connectivity.

9.1 Overview

The diagram below depicts an overview of the infrastructure implemented for a standard Citi deployment, highlighting data and management communication.



9.2 Citi Support and Management

9.2.1 Array Management

The VNX storage array is managed with the Unisphere application, which can be accessed as a standalone client installed on an Off-Array Management Server (recommended) or by typing the array's IP address in a browser. Unisphere is capable of managing multiple arrays as well as multiple storage domains. Please refer to the [Off-Array Management Server](#) section of this guide for further details.

9.2.2 Enterprise Monitoring and Alerting

SNMP alerting should be configured by Citi Storage team to send appropriate SNMP traps to the Enterprise Monitoring System. The monitoring infrastructure can then filter the alerts to generate notifications or problem tickets as appropriate. Please refer to the [Alerting and Monitoring](#) section of this guide for further details.

9.2.3 Email Alerting

Email alerting should be configured by Citi Storage team to send notifications to the appropriate Citi Storage Support team. Please refer to the [Alerting and Monitoring](#) section of this guide for further details.

9.2.4 Audit Logs

Off-array management servers should be selected as central event monitoring servers to capture VNX audit logs in the Windows Application log. Audit log reconciliation for the EMC VNX environment is to be completed through Citi's Secure Event Monitoring System (SEMS). Please refer to the [Audit Logging](#) section of this guide for further details

9.2.5 Reporting, Performance, Change Management and Capacity Planning

Unisphere Analyzer gathers storage system performance statistics from which trend analysis, reporting and system capacity management reports can be generated. These may be leveraged to improve performance and assist in the hardware capacity planning process. Please refer to the [Performance Monitoring using Unisphere Analyzer](#) section for configuration information.

Where available, it is recommended to use EMC SRM or ECC with VNX for reporting purposes and also for historical data collection as part of performance and capacity planning. Please refer to the [EMC SRM site](#) or [ECC Solutions Guide](#) for more information.

SANscreen is a change-monitoring, analysis and validation tool that correlates various changes in the open-systems SAN environment and indicates violations and vulnerabilities based on pre-determined and authorized parameters. Please refer to the current [SANscreen Solutions Guide](#) for more information.

9.3 Remote Vendor Support

9.3.1 Email Alerting

Email home should be configured by EMC certified engineer to use Citi's SMTP server to send the alerts to EMC support. Please refer to the Call Home and Email Alerting section for details.

9.3.2 EMC Secure Remote Support

ESRS will also be used to provide EMC support with a secure method of accessing the VNX storage array for remote troubleshooting and problem resolution. The [ESRS Solutions Guide](#) should be referenced for detailed configuration instructions.

9.4 Host Connectivity

It is recommended that hosts be configured with redundant Host Bus Adapters connected via redundant Fibre Channel SAN Switches, with paths to both Service Processors of the Storage Array. Please refer to the [CSBE SAN Stack](#) Solutions Guide for the specific host platform for the certified hardware/software configurations for connectivity to the VNX.

9.5 Software and Features

The table below lists the VNX Software components and their current status in Citi.

Feature	Item	Function	Citi Status
VNX OE	VNX Block OE	VNX Block Operating Environment	Required
	Unisphere Block	Unified VNX block management	Required
VNX OE Capacity	Performance Tier / TB	Licensed Per TB of raw SAS & SSD	Required ¹
	Capacity Tier / TB	Licensed Per TB of raw NL-SAS	Required ¹
FAST Suite	Unisphere Analyzer	Trend analysis and reporting	Recommended
	FAST Cache	Cache extension using SSD's	Optional
	FAST-VP	Automated Storage Tiering	Not Supported ²
	Unisphere QOS Manager	Quality of Service management	Not Supported
Local Protection	Snapview	Block storage Snapshots + Clones (Copy On First Write)	Optional
	VNX Snapshots	Redirect On Write Snapshots	Not Supported
	Recoverpoint/SE CDP		Not Supported
Remote Protection	Mirrorview , SANcopy Recoverpoint/SE CRR	Block storage replication	Not Supported
Application Protection	Application Protection Suite	Integrated Application Protection	Restricted ³
	Admsnap Utility	Replication Manager snap admin	Restricted ³
Virtual Provisioning	Virtual Provisioning Storage Pools	Thick and Thin Pool LUNs	Recommended ⁴
	LUN Compression		Not Supported
	De-Duplication		Not Supported
Reporting	ECC / ProSphere	Centralized reporting	Recommended
Monitoring	SNMP Alerting	Centralized Enterprise Monitoring	Required
	ESRS IP Client for VNX	Monitors VNX and legacy systems	Not Supported
Connectivity	Powerpath	Host Multi-Pathing	Refer to platform SAN Stack
	FC Connectivity	Supported connectivity protocol	Recommended
	FCoE Connectivity	Unsupported connectivity protocols	Not Supported
	iSCSI Connectivity	Restricted connectivity protocol	Restricted ⁵
	Windows ODX Copy	Windows Offload Data Transfer	Not Supported
Management	Unisphere	Web-based interface used to manage storage systems	Required
	Unisphere Service Manager	Java-application used to maintain system hardware and software	Not Supported

Note: It is recommended that consideration be given to the intended use of the system to determine whether purchasing a license bundle is appropriate / cost effective.

¹ Not required by vendor for VNX 5300 model

² FAST VP be re-assessed at later date following the slice size reduction from 1GB to 256MB

³ Application Protection is Restricted to [Exchange 2010](#) environments using [Replication Manager](#).

⁴ The use of Virtual Provisioning with VNX OE 31 is Restricted to [Exchange 2010](#) environments.

⁵ The use of iSCSI is restricted to W2K12,ESX5.5,and RHEL hosts with dual 10Gb connectivity to the Citi Storage IP Network.

For the currently certified and code versions, please reference the [CSBE VNX Block and CLARiiON Supported Configuration Matrix](#).

9.6 Hardware

9.6.1 VNX 2 Hardware

9.6.1.1 VNX 2 Models

	VNX 5200	VNX 5400	VNX 5600	VNX 5800	VNX 7600	VNX 8000
Max # Drives ⁵	125	250	500	750	1000	1000 ⁶
Max FAST Cache	600 GB	1T B	2 TB	3 TB	4.2 TB	4.2 TB
Processors Per Array	2 x 4-core 1.2 GHz	2 x 4-core 1.8 GHz	2 x 4-core 2.4 GHz	2 x 6-core 2.0 GHz	2 x 8-core 2.2 GHz	4 x 8-core 2.7 GHz
Physical Memory per SP	16 GB	16 GB	24 GB	32 GB	64 GB	128 GB
Max Write Cache per SP ⁷	4.5 GB - 10.5 GB	4.5 GB - 10.5 GB	9.7 GB - 16.9 GB	13.7 GB - 22.8 GB	31.9GB - 36.6 GB	50 GB
Max # Flex I/O Modules per SP	3	4	5	5	5	11
Embedded 6Gbps SAS Back-End Ports per SP	2 x 4-Lane	2 x 4-Lane	2 x 4-Lane	2 x 4-Lane	2 x 4-Lane	-
Base 6Gbps SAS Back-End Buses per Array	2 x 4-Lane	2 x 4-Lane	2 x 4-Lane	2 x 4-Lane	2 x 4-Lane	8 x 4-Lane
Max 6Gbps SAS Back-End Buses per Array	2 x 4-Lane	2 x 4-Lane	6 x 4-Lane or 2 x 4-Lane + 2 x 8-Lane	6 x 4-Lane or 2 x 4-Lane + 2 x 8-Lane	6 x 4-Lane or 2 x 4-Lane + 2 x 8-Lane	16 x 4-Lane or 8 x 8-Lane
Max Total Ports per Array	28	36	44	44	44	88
Max 2/4/8Gbps FC Front-End Ports per SP	16	16	20	20	20	36
Array Enclosure	3U DPE 25 x 2.5" drives	3U DPE 25 x 2.5" drives	3U DPE 25 x 2.5" drives	3U DPE 25 x 2.5" drives	3U DPE 25 x 2.5" drives	4U SPE (no drives)
Standby Power System (SPS)	Battery Onboard	Battery Onboard	Battery Onboard	Battery Onboard	Battery Onboard	2 x 2U 2.2KW
Max # SAN Hosts	1024	1024	1024	2048	4096	8192
Max # Pools	15	15	20	40	40	60
Max # LUNs (Pool)	1000	1000	1000	2000	3000	4000
Max # LUNs (RAID group)	2048	2048	2048	4096	4096	8192

9.6.1.2 VNX 2 Disk Array Enclosures

	25 x 2.5" Drive DAE	15 x 3.5" Drive DAE
Drive Types Supported	2.5" SSD / Flash 2.5" 15K Rotating 2.5" 10K Rotating 2.5" Near-line Rotating	2.5" SSD / Flash ⁸ 2.5" 15K Rotating ⁸ 3.5" 15K Rotating 2.5" 10K Rotating ⁸ 3.5" Near-line Rotating
Enclosure Height (Rack Units)	2U	3U
Controller Interface	6 Gb SAS	6 Gb SAS

Note: The 60 x 3.5" Drive DAE is not currently certified for use in Citi.

⁵ The maximum number of drives is dependent on the Disk Array Enclosure used

⁶ Support for 1500 drives in the VNX8000 is expected in Q1 2014.

⁷ The cache available decreases with the number of services in use (Thin, FAST cache, FAST VP, Compression, De-duplication, ...)

⁸ The use of a 2.5" drive in a 3.5" DAE requires the use of a 3.5" carrier

9.6.1.3 VNX 2 Disk Drive Types

The tables below lists all disk drives certified for use with the VNX 2 storage arrays for installation into Disk Processor Enclosure (DPE) or a Disk Array Enclosure (DAE).

Notes:

- To reduce datacenter footprint, it is recommended to use 2.5" drives in environments not requiring larger 2TB/3TB NL-SAS 3.5" drives
- SSDs are only certified for FAST Cache at the current time (using 100/200GB SLC SSDs).

Disk Drives for 2.5" DPE / DAE

Nominal Capacity	100 GB SSD ⁹	200 GB SSD ⁹	300 GB 15K	600 GB 10K	900 GB 10K	1TB 7.2K
Supported in 25 Drive DAE	✓	✓	✓	✓	✓	✓
Drive Form Factor	2.5"	2.5"	2.5"	2.5"	2.5"	2.5"
Rotational Speed	Solid State	Solid State	10K RPM	10K RPM	10K RPM	7.2K RPM
Interface	6 Gb SAS	6 Gb SAS	6 Gb SAS	6 Gb SAS	6 Gb SAS	6 Gb SAS
Usable Capacity(GB) ¹⁰	91.71	183.43	268.39	536.79	820.60	917.16
Nominal Power (W)	100 GB SSD	200 GB SSD	300 GB 15K	600 GB 10K	900 GB 10K	1TB 7.2K
Operating Mode	4.97	4.97	9.07	5.6	5.6	7.44
Idle Mode	1.36	1.36	5.25	3.1	3.1	4.84

Note: 300 GB 10K 2.5" SAS drives are no longer sold with new VNX series systems

Disk Drives for 3.5" DAE

Nominal Capacity	100 GB SSD ⁹	200 GB SSD ⁹	300 GB 15K	600 GB 15K	600 GB 10K	900 GB 10K	2 TB 7.2K	3 TB 7.2K
Supported in 15 Drive DAE	✓	✓	✓	✓	✓	✓	✓	✓
Drive Form Factor	3.5"	3.5"	2.5"	3.5"	2.5"	2.5"	3.5"	3.5"
Rotational Speed	Solid State	Solid State	15K RPM	15K RPM	10K RPM	10K RPM	7.2K RPM	7.2K RPM
Interface	6 Gb SAS	6 Gb SAS	6 Gb SAS	6 Gb SAS	6 Gb SAS	6 Gb SAS	6 Gb SAS	6 Gb SAS
Usable Capacity(GB) ¹⁰	91.71	183.43	268.39	536.79	536.79	820.60	1834.34	2751.51
Nominal Power (W)	100 GB SSD	200 GB SSD	300 GB 15K	600 GB 15K	600 GB 10K	900 GB 10K	2 TB 7.2K	3 TB 7.2K
Operating Mode	4.97	4.97	9.07	16.35	5.6	5.6	12.2	12.2
Idle Mode	1.36	1.36	5.25	11.68	3.1	3.1	8.0	8.0

Note: 1 TB 7.2K 3.5" NL-SAS drives are no longer sold with new VNX series systems

⁹ 100GB/200GB SSDs are available in SLC or eMLC technology. 400GB SSD is eMLC technology. eMLC can only be used for FAST VP use cases (not currently supported in Citi)

¹⁰ Vendor provided values from EMC VNX5400, VNX5600, VNX5800, VNX7600 and VNX8000 Series Storage Systems Disk and OE Matrix

9.6.1.4 VNX 2 Storage Processor I/O Slot Availability

Unified Configuration

	VNX 5200	VNX 5400	VNX 5600	VNX 5800	VNX 7600	VNX 8000					
SAS back-end buses	1-2	1-2	1-6 ¹¹	1-6 ¹¹	1-6 ¹¹	1-8 ¹²	9-12		13-16		
Number of X-blades	1-2 ¹³	1-2 ¹³	1-2 ¹³	1-3 ¹³	2-4 ¹³	2-4	5-8	2-4	5-8	2-4	5-8
# Slots for Host I/O Connectivity	2 ¹⁴	3 ¹⁵	3	3	3	8	7	7	6	6	5

Block Configuration

	VNX 5200	VNX 5400	VNX 5600	VNX 5800	VNX 7600	VNX 8000					
SAS back-end buses	1-2	1-2	1-6 ¹¹	1-6 ¹¹	1-6 ¹¹	1-8 ¹²	9-12		13-16		
# Slots for Host I/O Connectivity	3 ¹⁴	4 ¹⁵	4	4	4	9	8		7		

¹¹ VNX5600/5800/7600 reserve Slot 0 for 4-port 6Gb SAS I/O module

¹² VNX8000 is shipped with a minimum of two 4-port 6Gb SAS I/O modules

¹³ VNX5200/5400/5600/5800/7600 reserve Slot 4 for 4-port 6Gb FC I/O module

¹⁴ VNX5200 supports a maximum of 3 I/O modules per SP and no I/O module in slots 0 and 1

¹⁵ VNX5400 supports a maximum of 4 I/O modules per SP and no I/O module in Slot 0

9.6.2 VNX Hardware

9.6.2.1 VNX Hardware Models

The table below lists the VNX models and configurations that are supported by CSBE

	VNX 5100	VNX 5300	VNX 5500	VNX 5700	VNX 7500
Front-End FC Ports Onboard	4 @ 2/4/8Gbps per SP	4 @ 2/4/8Gbps per SP ¹⁶	4 @ 2/4/8Gbps per SP ¹⁶	-	-
Front-End FC Ports I/O Module(s)	-	4 @ 2/4/8Gbps per SP	4 @ 2/4/8Gbps per SP	4/8/12 @ 2/4/8Gbps per SP	4/8/12/16 @ 2/4/8Gbps per SP
Back-End SAS Disk Ports	2 @ 6Gbps per SP	2 @ 6Gbps per SP	2/6 @ 6Gbps per SP	4 @ 6Gbps per SP	4/8 @ 6Gbps per SP
Total Memory	8 GB	16 GB	24 GB	36 GB	48/96 GB
Max Read/Write Cache per SP	801 MB (571 MB ¹⁷)	3997 MB (3129 MB ¹⁷)	6988 MB (5324 MB ¹⁷)	10906 MB (8446 MB ¹⁷)	14250 MB (10830 MB ¹⁷)
Processors	1 x 2-core 2.0 GHz per SP	1 x 4-core 1.66 GHz per SP	1 x 4-core 2.13 GHz per SP	1 x 4-core 2.4GHz per SP	1 x 6-core 2.8 GHz per SP
Array Enclosure	3U DPE 15 x 3.5" or 25 x 2.5" drives	3U DPE 15 x 3.5" or 25 x 2.5" drives	3U DPE 15 x 3.5" or 25 x 2.5" drives	2U SPE (no drives)	2U SPE (no drives)
Max # Drives ¹⁸	75 (DAE15) 75 (DAE25)	120 (DAE15) 125 (DAE25)	240 (DAE15) 250 (DAE25)	495 (DAE15) 500 (DAE25)	990 (DAE15) 1000 (DAE25)
Max # LUNs	512	2048	4096	4096	8192
Max # LUNs per RAID group	256	256	256	256	256
Max # LUNs per Storage Pool	256	512	512	1024	1024
Max # Hosts	512	2048	4096	4096	8192
Max # Block I/O Modules	0 per SP	2 per SP	2 per SP	5 per SP	5 per SP
Standby Power System (SPS)	1U 1.2KW	1U 1.2KW	1U 1.2KW	1U 1.2KW	1U 1.2KW

9.6.2.2 VNX Disk Array Enclosures

	15 x 3.5" Drive DAE	25 x 2.5" Drive DAE
Drive Types Supported	3.5" 15K Rotating 2.5" 10K Rotating ¹⁹ 3.5" Near-line Rotating 3.5" SSD / Flash	2.5" 10K Rotating 2.5" SSD / Flash
Enclosure Height (Rack Units)	3U	2U
Controller Interface	6 Gb SAS	6 Gb SAS

Note: The 60 x 3.5" Drive DAE is not currently certified for use in Citi.

¹⁶ Two FC ports are reserved for Data Mover connectivity (not used for host connectivity)

¹⁷ The lower value indicates the usable cache when maximum FAST cache is installed, additional reduction using FASTVP/Thin/Compression

¹⁸ The maximum number of drives is dependent on the Disk Array Enclosure used

¹⁹ The use of a 2.5" drive in a 3.5" DAE requires the use of a 3.5" carrier

9.6.2.3 VNX Disk Drive Types

The tables below lists all disk drives certified for use with the VNX storage array for installation into Disk Processor Enclosure (DPE) or a Disk Array Enclosure (DAE).

Note: SSDs are only certified for FAST Cache at the current time.

Disk Drives for 3.5" DPE / DAE

Nominal Capacity	100 GB SSD	200 GB SSD	300 GB 15K	600 GB 15K	300 GB 10K	600 GB 10K	1 TB 7.2K	2 TB 7.2K	3 TB 7.2K
Supported in 15 Drive DAE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Drive Form Factor	3.5"	3.5"	3.5"	3.5"	2.5"	2.5"	3.5"	3.5"	3.5"
Rotational Speed	Solid State	Solid State	15K RPM	15K RPM	10K RPM	10K RPM	7.2K RPM	7.2K RPM	7.2K RPM
Interface	6 Gb SAS	6 Gb SAS	6 Gb SAS	6 Gb SAS	6 Gb SAS	6 Gb SAS	6 Gb SAS	6 Gb SAS	6 Gb SAS
Formatted Capacity(GB) ²⁰	93.1	186.31	272.59	545.19	272.59	545.19	926.04	1852.09	2794.50
Nominal Power (W)	100 GB SSD	200 GB SSD	300 GB 15K	600 GB 15K	300 GB 10K	600 GB 10K	1 TB 7.2K	2 TB 7.2K	3 TB 7.2K
Operating Mode	4.97	4.97	12.92	16.35	6.15	5.6	12.2	12.2	12.2
Idle Mode	1.36	1.36	8.74	11.68	3.5	3.1	8.0	8.0	8.0

Disk Drives for 2.5" DPE / DAE

Nominal Capacity	100 GB SSD	200 GB SSD	300 GB 10K	600 GB 10K
Supported in 25 Drive DAE	✓	✓	✓	✓
Drive Form Factor	2.5"	2.5"	2.5"	2.5"
Rotational Speed	Solid State	Solid State	10K RPM	10K RPM
Interface	6 Gb SAS	6 Gb SAS	6 Gb SAS	6 Gb SAS
Formatted Capacity(GB) ²⁰	93.1	186.31	272.59	545.19
Nominal Power (W)	100 GB SSD	200 GB SSD	300 GB 10K	600 GB 111111
Operating Mode	4.97	4.97	6.15	5.6
Idle Mode	1.36	1.36	3.5	3.1

²⁰ Vendor provided value calculated using 520 bytes/sector and 1 MB = 1,048,576 bytes

9.6.2.4 VNX Storage Processor I/O Slot Availability

Unified Configuration

	VNX5300	VNX5500		VNX5700	VNX7500			
SAS back-end buses	1-2	1-2	3-6	1-4	1-4	5-8		
Number of X-blades	1-2	1-3	1-3	2-4	2-4	5-8	2-4	5-8
# Slots for Host I/O Connectivity	2	2	1	3	3	2	2	1

Block Configuration

	VNX5300		VNX5500				VNX5700		VNX7500			
SAS back-end buses	1-2		1-2	3-6		1-4		1-4	5-8			
Slots reserved for unified upgrade?	N	Y	N	Y	N	Y	N	Y	N	Y	N	Y
# Slots for Host I/O Connectivity	2		2	1		4	3	4	3	3	3	2

File Configuration

	VNX5300	VNX5500		VNX5700	VNX7500			
SAS back-end buses	1-2	1-2	3-6	1-4	1-4	5-8		
Number of X-blades	1-2	1-3	1-3	2-4	2-4	5-8	2-4	5-8
# Slots for Host I/O Connectivity	1	1	1	1	1	1	1	1

Note: With a file only no host connectivity is allowed. The above table shows 1 configurable slot available for RecoverPoint FC (not certified in Citi).

10 Storage Best Practices

10.1 Virtual Provisioning and Traditional (Classic) Storage Provisioning

Physical disk drives can be grouped into either Storage Pools using Virtual Provisioning or Traditional RAID Groups. Storage Pools can be used for thick or thin pool LUN allocation. Traditional RAID Groups are used for Traditional/Classic LUN allocations.

Virtual Provisioning is required to enable a number of features on the VNX block array

- Thick and Thin Pool LUNs
- LUN compression
- LUN De-duplication (VNX OE for Block 33)
- FAST-VP (Automatic Storage Tiering)

Further analysis is planned to assess the performance overhead of LUN compression and De-Duplication as well as the cost/benefit of storage tiering.

Given the distinct methods of provisioning, a separate section has been created. Please refer to the [Storage Provisioning Best Practices](#) section.

10.2 Drive Selection

There are three disk types that can be installed in the VNX storage array:

- Solid State Drives (SSD)
- Serial Attached SCSI (SAS)
- Near-line SAS (NL-SAS).

NL-SAS drives are the lowest performing drive of the three drive types; thus, they should only be used for applications that are not performance sensitive. For example use NL-SAS for less active data, well-behaved streaming data, aging data, archive purposes and backups.

SSDs have the highest performance of the three drive types but are significantly more expensive than SAS drives. Typical use of SSDs is for FAST Cache and to provide a higher performing tier if the FAST feature is used with Virtual Provisioning. SSDs are only certified for use as FAST Cache at the current time. Note: With VNX 2 arrays SAS Flash / FAST Cache SSD (SLC) are required for FAST Cache.

Typically, drives with lower capacity and higher RPM should be used for greater performance. To save on costs, high capacity drives should be used if high performance is not required. Note that a failure of a high capacity drive will have a higher impact on applications residing on the storage system since more data is affected and the recovery process (rebuild) for the failed drive takes longer.

For a list of all disk drives certified for use in Citi with the VNX storage arrays please refer to Disk Drives Type section appropriate for the VNX model.

10.3 SAS Back-end

The VNX range introduces the use of a 6 Gb/s Serial Attached SCSI (SAS) backend, replacing the previously used Fibre Channel (FC) backend in the CLARiiON range.

The SAS backend has different operational characteristics to the FC protocol used with the previous generation of CLARiiON systems. Each SAS port can address 4 storage devices at the same time on a port, as compared to a single device on an FC bus.

To some degree, balancing across ports is “less critical” on the VNX because of the use of SAS as well as the increased I/O bandwidth available on the VNX. However there is a general performance advantage to evenly distributing the I/O to the storage devices evenly across all of the available back-end ports. It should be noted that it is not necessarily the number of drives or capacity, but it is more that the I/O that should be balanced across the backend ports.

When selecting the drives to use in a RAID group, drives can be selected either from a DAE or DAEs on the same bus (horizontal positioning), or from multiple DAEs on different buses (vertical positioning). There are now almost no performance or availability advantages to using vertical positioning. Therefore:

- Use the default horizontal positioning method of drive selection when creating RAID groups.
- If a single DAE does not have enough available drives to create the RAID group, selecting drives from a second DAE is acceptable.

It is recommended to utilize all available back-end ports where possible. When all backend ports are in use there is no performance advantage of using any one port over another. However where not all ports are used on VNX models with greater than two backend SAS ports per SP, there is a performance advantage to using specific ports due to the use of a shared/common 8 lane SAS controller on ports 0/1 and on ports 2/3.

- If using only 2 backend ports on a VNX 5700 you should use ports 0 and 2 (or alternatively 1 and 3) for best performance
- If using 4 or fewer ports on a VNX 7500 you should alternate ports on the backend I/O module for highest performance.

For details on the number of backend ports per model, please refer to the specifications table in the Hardware section appropriate for the VNX model.

10.3.1 Enclosure ID's

Each enclosure has an address identifier known as an Enclosure ID which is algorithmically auto-configured by the system. Newly added enclosures will receive the lowest available Enclosure ID for that back-end port. The Enclosure ID may not be in the same order as they are cabled or racked. Cabling order only affects the Enclosure ID setting the first time the system is booted. Adding another enclosure to a backend port does not result in the renumbering of the enclosures on that port. Enclosure IDs cannot be changed by the user.

Note: It is not recommended to relocate enclosures that have been previously been provisioned with storage.

10.4 Vault Drives and DPE / DAE OS

The VNX vault is located on the first four drives (0-3) on bus 0, enclosure 0 which is either the DPE or the first DAE (DAE OS with VNX5700/7500/8000). These disks are protected with parity protection and hold the reserved area for write cache de-staging. In addition, these disks are used for system information, such as storing the operating system, Persistent Storage Manager (PSM), File Control LUNs, and operating environment database.

The standard configuration of the vault DPE or DAE (Bus 0 Enclosure 0) should be as follows:

- Vault drives should not be used for host data and cannot be part of a storage pool.
- The vault drives should be configured as a 3+1 RAID 5 RAID Group.
- For the remaining slots in the enclosure, please refer to the below guidelines on placement of drives.

10.4.1 Vault Drive Replacement

- Failed drives should be replaced by the vendor.
- Before and after replacing failed drives, ensure that a full set of spcollects are generated.
- **Note:** To avoid the risk of a DU/DL, ensure that 30 seconds elapses before inserting any replacement drive.

10.5 Placement of Drives

10.5.1 General Disk Placement Guidelines

When initially placing drives in the array:

- Place highest performing drives in lowest numbered enclosures on each bus.
- Spread each drive type across all available buses.
- Avoid placing drives in the DPE or DAE-OS enclosure (0_0) that will be mirrored with drives in another enclosure. For example, avoid mirroring a disk in 0_0 with a disk in 1_0.
- Avoid placing drives in the DPE or DAE-OS enclosure (0_0) if they will be in a parity RAID Group where one disk is placed outside of that enclosure.

10.5.2 Placement of Flash drives for FAST Cache

10.5.2.1 VNX 2 FLASH Drive Placement (VNX OE for Block 33)

When Flash drives will be included for use as FAST Cache

- Spread Flash drives evenly across all available buses.
- Add no more than 8 FAST Cache flash drives per bus, including any unused flash drives for use as hot-spares
- Use DAE 0 on each bus when possible for flash drives (has the lowest latency)

10.5.2.2 VNX FLASH Drive Placement (VNX OE for Block 32 and earlier)

When Flash drives will be included for use as FAST Cache

- For up to 8 Flash drives
 - Place all Flash drives in enclosure 0_0
- For > 8 Flash drives
 - Spread Flash drives evenly across all available buses.
 - Avoid placing drives in the DPE or DAE-OS enclosure (0_0) that will be mirrored with drives in another enclosure.
 - Use DAE 0 on each bus when possible for flash drives (has the lowest latency)

10.6 Drive Sparing

10.6.1 Drive Sparing - Multicore RAID (VNX OE for Block 33)

The way in which Hot Spares is implemented on the VNX 2 (VNX OE for Block 33) is different to the earlier VNX arrays (VNX OE for Block 32 and 31). These differences as well as guidelines for each are documented below.

10.6.1.1 Hot Spares - Multicore RAID (VNX OE for Block 33)

Sparing is the process of rebuilding the data on a failed drive onto a system-selected compatible drive. Any unbound drive can be considered for sparing.

When a drive fails, a suitable permanent replacement drive is selected from the unbound drives. The selected drive (in a different location²¹) then becomes part of the RAID group (substituting for the failed drive). When the failed drive is removed from the array and replaced, the new drive simply becomes unbound.

The following guidelines should be followed:

- A minimum of 2 drives should be unbound (considered for sparing) in any given array.
- One drive should be left unbound for every 60 disks provisioned
 - In remote sites where there may be longer delays in replacing failed disks, consideration should be given to leaving 1 unbound for every 30 disks.
- Unbound drives (spares) should be evenly distributed across available buses.
- A drive of each drive type (SSD, SAS, or NL-SAS) in the array should be left unbound.
 - SAS Flash / FAST Cache SSD (SLC) must spare for SAS Flash / FAST Cache SSD (SLC)
 - SAS Flash VP/ FAST VP SSD (eMLC) must spare for SAS Flash VP / FAST VP SSD (eMLC)
 - SAS must spare for SAS (regardless of rotational speed)
 - NL-SAS must spare for NL-SAS
- The unbound drive capacity should be equal to or larger than the provisioned drives
 - Always leave a drive with the largest capacity of each drive type unbound since large drives can spare for smaller drives.
- If maintaining performance is essential during a drive failure, ensure that a drive is left unbound with the same type, speed, and capacity as the drives in the RAID Groups requiring the sustained performance.

²¹ If for some reason it is required that the physical position remain fixed for all drives in a RAID group it is possible to copy the data from the selected spare back to the new disk manually using the CLI (with the *copytodisk* option within *naviseccli*).

10.6.1.2 Hot Spare Policy - Multicore RAID (VNX OE for Block 33)

Multicore RAID introduces the ability to provide a warning to an administrator about the state of usable spare drives in the system through the use of a *Hot Spare Policy*. A hot spare policy monitors a number of unbound drives (potential spare drives) of a certain type in the array and posts a warning when the number is low. There are 3 Hot Spare Policies:-

Policy	Description
Recommended	1 hot spare per 30 drives of every type (this is the EMC standard recommendation)
Custom	Allow an administrator to specify the ratio (required to configure as per Citi standards)
No Hot Spares	Turn off policy monitoring (Note: Sparing will continue function as long as a suitable unbound drive is found)

Policies are configured by selecting **Hot Spare Policy** under **Hardware** on the **System** tab in Unisphere. The recommended configuration should be based on

Disk Type	Hot Spare Policy	Keep Unused per Policy
SAS	Custom	1 per 60
NL-SAS	Custom	1 per 60
SAS Flash	Recommended	1 per 30

Note: Due to a limitation in the current Unisphere GUI code, it is not possible to set the Custom policy to 1 per 60 in the GUI when you have less than 60 drives of that type installed. However, this Custom ratio can be set using the CLI:

First list the policies (*hotsparepolicy -list -allentries*) to obtain a list of the appropriate ID's
Then set the ratio to 1 to 60 (*hotsparepolicy -set <id> -keep1unusedper 60*) for each ID
Once set, the ratio can be verified using either the Unisphere GUI or Navisphere Secure CLI.

10.6.1.3 RAID 6 Parallel Rebuild - Multicore RAID (VNX OE for Block 33)

Multicore RAID supports the rebuild of 2 drives in parallel in RAID 6 configurations. This reduces the time taken to reconstruct and restore the degraded RAID group.

10.6.1.4 Drive Mobility - Multicore RAID (VNX OE for Block 33)

Drives Mobility (also known as Portable Drives) drives to be physically relocated to any slot within the same array (as long as they fit in the available space).

During an online move the RAID data continues to be available after the drive is removed. The moment a drive becomes unavailable (failed / removed / firmware error / ...) a 5 minute timer begins. If not reinserted within 5 minutes, drive sparing is invoked.

Note: Removing any active drive will degrade the RAID group in which it belongs. It is recommended to avoid moving drives where possible, however if it is deemed necessary to move a drive then it is recommended that this should be performed as an offline activity to reduce the risk of Data Unavailability.

10.6.2 Drive Sparing - FLARE RAID (VNX OE for Block 32 and earlier)

Hot spares are drives that temporarily take the place of failed drives until the failed drives are replaced. In order for a hot spare to be invoked, it must be of equal or higher capacity to the total capacity of the LUNs configured on the failed drive. Once the hot spare is invoked, the data on the failed drive must be fully rebuilt on the hot spare before equalizing (copying back) to the replacement drive. Once the failed drive is replaced, the hot spare is made available again to replace other failed drives.

The following guidelines should be followed:

- A minimum of 2 hot spares should be allocated in any given array.
- One hot spare should be provided for every 60 disks
 - In remote sites where there may be longer delays in replacing failed disks, consideration should be given to using one spare for every 30 disks.
- A hot spare of each drive type (SSD, SAS, or NL-SAS) in the array should be created. Always choose a drive with the largest capacity of each drive type for a hot spare since large drives can spare out smaller drives.
- If maintaining performance is essential during a drive failure, ensure that a hot spare is configured with the same type, speed, and capacity as the drives in the RAID Groups requiring the sustained performance.
- Hot Spare RAID Group ID should start from the highest ID in descending order (e.g., 245, 244 ...)
- Hot Spare LUN ID should follow the same concept for Hot Spare RAID Group in which the LUN IDs should start with the last number and so on (e.g. 8121, 8120, ...)
- Vault drives (Disks 000-003) cannot be used as hot spares.

10.7 SP Cache

SP cache operates on the VNX 2 (VNX OE for Block 33) differently to the earlier VNX arrays (VNX OE for Block 32 and 31):-

- VNX OE for Block 33 introduces Multicore Cache with a shared “adaptive” cache.
- VNX OE for Block 32 and earlier split cache into Read cache, Write cache and peer SP mirror cache. This is also referred to as FLARE cache.

10.7.1 Multicore Cache (VNX OE for Block 33)

Multicore cache is shared for both reads and writes.

When a host requests a read, it will be satisfied from Read or Write cache if the data is present there, otherwise it will be loaded from disk into Read cache and returned to the host.

When a host requests a write, the incoming IO is marked as “dirty” when written to the SP cache and is mirrored to the peer SP cache. Both SPs track which pages they own and which are owned by the peer SP. Each SP is responsible for copying/flushing the dirty pages it owns to the back-end, which clears the dirty page flag. If one of the SPs becomes unavailable, the surviving SP will take ownership of all of the dirty pages.

Multicore cache keeps statistics on every page and uses an aging algorithm to determine which pages should remain in cache.

For a more detailed description of Multicore Cache, it is recommended to read through the VNX MCx *Multicore Everything White Paper* available from EMC.

10.7.1.1 SP Mirrored Cache

The SP write cache is mirrored between the two SPs via the Common Management Interface (CMI) between the two SPs. Backup batteries are used to protect the dirty pages in write cache. In the event of a power fault, the SP’s and the DPE (where the system drives are located) is kept online to allow the dirty pages in cache to be de-staged to the vault.

When a hardware fault occurs, the cache state will be set to one of three states:

State	Behavior
Enabled	• Cache remains fully functional
Disabled	• Dirty pages flushed to disk • All data writes go directly to disk (Write-Through mode) • Reads continue to be cached
Dumping	• Enters into a quiescing 5 second interval for hosts to complete inflight IOs (if possible these are mirrored to peer SP cache) • No new IO accepted to cache • All background processes (drive rebuilds, background verifies, ...) stopped • All dirty pages flushed to the vault Note: A full system reboot is required to recover the array following a vault dump.

Note: In the event of a temporary power loss of less than 15 seconds, the array will operate using the built in batteries. If power has not resumed once the 15 second timeout is reached, the cache will be set to the *Dumping* state.

Please refer to the [Power Faults Table](#) for a full list of hardware faults and how the resultant cache state.

10.7.1.2 Read/Write Cache Configuration

Multicore Cache read cache is always enabled. Write cache has two modes of operation:

- Write cache enabled
- Write-Through mode (Write cache disabled)

It is recommended that Write cache be enabled.

Write Cache is configured by selecting **Manage Cache** under **System Management** on the **System** tab in Unisphere

10.7.1.3 Page Size

Multicore Cache uses two (fixed) page sizes

- 8KB Physical Page
- 64KB Logical Page (containing 1 to 8 physical pages)

Cache Page Size is no longer a configurable parameter.

10.7.1.4 Watermarks

Multicore Cache uses Dynamic Watermarks.

High and Low Watermarks are no longer configurable parameters.

10.7.1.5 Power Faults Table

All power faults in the below table are for a period longer than 15 seconds. Any faults less than 15 seconds are accommodated by on the battery backup within the array.

Use Case	Description	Multicore Cache State
Single AC power failure, peer SP still running	AC power is removed from one SPS in a system with multiple SPS (1 or 2 per SP)	Enabled
Single SPS power failure, peer SP not running	AC power is removed from one SPS in a system with multiple SPS (1 or 2 per SP) and the peer SP is not present	Vault Dump
AC power failure to all SPS	AC power has been removed from all SPS	Vault Dump
Single SP over temperature	This SP or the power supply for this SP is reporting an over temperature condition	Enabled
Dual SP over temperature	An over temperature condition is reported by all SP and power supplies	Vault Dump
DAE 0 (0_0) over temperature	An over temperature condition is reported by one or more power supplies in the first drive enclosure	Vault Dump
Single SPE/DPE power supply fault	A single power supply is reported as faulted in the SPE	Enabled
Single 60-drive DAE power supply fault ²²	A single power supply is reported as faulted in the first DAE and the first enclosure is a 60 drive enclosure	Vault dump
Multiple power supply faults	Multiple power supplies reported as faulted	Enabled
Single fan fault	A single fan has faulted	Enabled
Multiple fan faults	Multiple fans reported as faulted	Enabled
Degraded vault	A single system drive is offline or has been removed	Enabled
Faulted vault	Multiple system drives are offline or have been removed	Disabled
NDU without power supply firmware upgrade	Array software being changed	Enabled
NDU with power supply firmware upgrade	Array software and power supply firmware is being changed	Disabled
Single SP fault	A single SP in a dual SP system is removed or rebooted or panics or has a hardware fault	Enabled
Single SPS communication fault	An SP is unable to communicate with a single SPS. Other SPS OK.	Enabled
Multiple SPS communication faults	Neither SP is able to communicate with the SPS(s) connected to that SP.	Disabled
Single SPS cabling fault	An SPS is incorrectly cabled	Enabled
Multiple SPS cabling faults	Multiple SPS are incorrectly cabled on both SP	Disabled

²² The 60-drive DAE is not currently certified in Citi, however has been included for completeness

10.7.2 FLARE cache (VNX OE for Block 32 and earlier)

VNX OE for Block 32 and earlier was based on the FLARE code used in the earlier generation CLARiiON arrays. FLARE divides its cache into Read cache, Write cache and peer SP mirror cache.

The main function of read cache is prefetching. When a host requests a read, it will be satisfied from Read or Write cache if the data is present there, otherwise it will be loaded from disk into Read cache and returned to the host.

When a host requests a write, the incoming IO is marked as “dirty” when written to the SP Write cache and is mirrored to the peer SP cache. A Write Hit can occur if a dirty page is overwritten whilst in Write cache. Writes are destaged/flushed to disk using a Least Recently Used algorithm. Once written to disk these oldest, unused pages are discarded from the (mirrored) write cache on both SP’s.

10.7.2.1 Read/Write Cache Configuration

Read cache should initially be set using the following guidelines

- Install add-on software enablers before setting read and write cache
- Minimum 200 MB for the block only VNX5100 or 256 MB for File enabled systems
- Maximum 1024 MB
- Allocate 10% of available cache to read (set per SP) and the rest to write (set per array).

Following the above guidelines will lead to the following initial read cache values

Per SP Memory	VNX 5100	VNX 5300	VNX 5500	VNX 5700	VNX 7500
Initial Read Cache	200 MB	400 MB	700 MB	1024 MB	1024 MB

Write cache should be set to the maximum level allowed after the read cache is set.

Array features, such as FAST Cache, will reduce memory available for Read and Write cache.

If there is memory remaining in the **Free Memory** pool after the read and write cache are set, the remainder of the available cache should be distributed evenly between the SPs for read cache so that all the cache is used on the storage system.

The read/write cache mix can be adjusted to better fit the workload of the applications accessing the array. For example: if the storage system is servicing more reads than expected, more of the cache can be allocated for reads to improve the performance. Make sure that all of the cache is allocated after every adjustment to the cache settings

Note: Changes to cache should only be made during planned maintenance windows since SP cache is disabled during the process.

10.7.2.2 SP Mirrored Cache

The SP write cache is mirrored between the two SPs via the Common Management Interface (CMI) between the two SPs.

Cached data is protected as long as the surviving SP has enough time to detect the failure on the other SP. In the event of an emergency, such as a power outage or temperature alarm on both SPs, the Standby Power Supply (SPS) will kick in, and the cache will be de-staged to the VAULT drives. The SPS is designed to keep the array up long enough for this de-staging to complete. This operation can take anywhere from several seconds to several minutes, depending on the model, cache capacity, and type of VAULT drives. Once power is restored, any writes that were de-staged are reconciled and persisted to the target backend disks to ensure no data is lost.

VNX features Greater Write Cache Availability (GWCA), where the write cache remains enabled through a wide range of failures, including a single SP failure, single VAULT drive failure, and single SPS or power supply failure. In cases of a single SP failure, the cache is no longer mirrored across the SPs. To reduce the risk of losing write cache data if an SP reboots, the software saves the write cache data at the time of the reboot for retrieval on the next boot. Write cache is disabled only in circumstances where there is a power loss, when only a single fan remains, or both SPs become overheated.

10.7.2.3 Page Size

For typical deployments of shared VNX arrays, the cache Page Size should be left at the default value which is 8k. This setting can be adjusted for arrays deployed specifically for an application (e.g. Exchange) to better match the applications I/O profile.

For unified (both file and block) or for file only deployments, the default of 8k is recommended

10.7.2.4 High/Low Watermark

The watermark settings in the **SP Cache** tab of the **Storage System Properties** window in Unisphere are used to control the flushing behavior of write cache.

The two watermark settings are:

- **High watermark:** Percentage of write cache utilization where the data in write cache starts to destage (i.e. flush) to disk.
- **Low watermark:** Percentage of write cache utilization that the destaging of data in write cache stops.

It is suggested to start with

- High watermark of 80%
- Low watermark of 60%

This is the default level and is expected to be suitable for the general workloads. The recommended values do not change when VNX for File is involved.

The need may arise to vary the values.

- If frequent forced flushing occurs, reduce watermark values.
- Maintain a difference between low and high of about 20 percent.
- AVOID drastic changes to these values unless advised by CSBE and EMC Support.

10.8 FAST Cache

10.8.1 FAST Cache Overview

FAST Cache utilizes SSDs to serve as secondary cache and is separate from the storage systems primary cache (DRAM in the SP). The FAST Cache algorithm identifies most frequently used data (in 64KB chunks) and promotes them from the hard drives to FAST Cache where they will get the added performance of SSDs. If access to the data chunks decrease, the FAST Cache algorithm will demote the data back to the hard drives.

FAST Cache is only beneficial for workloads that access the same blocks on a frequent basis. If the workload is highly sequential, FAST Cache will not improve performance since data blocks are not accessed repeatedly for sequential workloads. Preferred applications for FAST cache:

- Random small block workloads with high locality
- High frequency of access to the same data
- Systems where current performance is limited by HDD capability, not SP capability.

Avoid enabling FAST Cache for LUNs that are not expected to benefit, such as when:

- The primary workload is sequential.
- The primary workload is large-block I/O

Note: Avoid enabling FAST Cache for small block sequential workloads such as Database logs, Circular logs, LUNs in SnapView Reserved LUN Pool, and VNX File SavVol LUNs.

For more detail on FAST Cache, refer to the EMC white paper [VNX FAST Cache – A Detailed Review](#).

10.8.2 Multicore FAST Cache (VNX OE for Block 33)

With changes introduced in VNX OE for Block 33,

- FAST Cache Warmup allows all new I/O to be promoted until 80% full
 - After 80% full, standard threshold of statistics is reinstated
 - Happens only one time, upon creation of FAST Cache
- FAST Cache will coalesce data when destaging I/O's larger than 64KB and will perform 4MB write to data drive where possible.

10.8.3 Supported FAST Cache Disk Configurations

The following table shows the supported FAST Cache configurations for each VNX model:

Model	Max Capacity 100 GB SSDs	Max Capacity 200 GB SSDs	Max Number 100 GB SSDs	Max Number 200 GB SSDs
VNX5200	300 GB	600 GB	6	6
VNX5400	500 GB	1000 GB	10	10
VNX5600	1000 GB	2000 GB	20	20
VNX5800	1500 GB	3000 GB	30	30
VNX7600	2100 GB	4200 GB	42	42
VNX8000	2100 GB	4200 GB	42	42
VNX5100	100 GB	-	2	-
VNX5300	500 GB	400 GB	10	8
VNX5500	1000 GB	1000 GB	20	10
VNX5700	1500 GB	1400 GB	30	14
VNX7500	2100 GB	2000 GB	42	20

10.8.4 FAST Cache Guidelines

The following guidelines should be followed when configuring FAST Cache on the system:

- An additional SSD should be installed to spare for the SSDs being used for FAST Cache
- Configure with an even number of drives (SSDs in FAST Cache are mirrored).
- Do not mix drive sizes / types
- Due to its high cost, FAST Cache should only be configured for storage systems that are experiencing low performance where it is determined that the bottleneck is determined to be drive-related.
- CPU utilization should be taken into consideration as increasing IOPS will also increase CPU utilization.
 - Systems should be sized such that maximum sustained utilization is 70%.
 - On existing systems, a thorough review should be conducted before enabling FAST cache if utilization is above 70%.
- For data residing on existing EMC VNX/CX4 arrays, EMC Pre-Sales have tools that can help determine active data set size. The potential improvement and expected impact from introducing FAST cache can be estimated by generating a NAR file and having this analyzed by EMC.
- When using Virtual Provisioning, FAST cache is enabled at the Pool level and is therefore enabled for all LUNs in that pool. Consideration should be given to creating separate pool(s) where not wishing to enable FAST cache for data LUNs.
- When using Traditional Storage Provisioning, FAST Cache is enabled at the LUN level. Any new RAID Group LUNs created after FAST Cache has been enabled on the storage system will have FAST Cache enabled by default. For any existing Traditional/Classic LUNs that you wish to use FAST Cache with, you must enable this on a per LUN basis.
- When adding FAST cache to a running system, it is recommended to enable it on a few LUNs at a time and wait until these LUNs have equalized. Check CPU utilization prior to enabling on more LUNs.
- Use separate storage pool or traditional RAID groups for SnapView Snapshots or VNX File SavVol storage. FAST Cache should not be enabled for Snapview Reserved LUN Pool or SavVol storage.

10.8.5 FAST Cache Failure Handling

If an SSD drive fails in the FAST Cache RAID group, VNX global hot sparing provides an online rebuild of the RAID group. Where it is identified that a drive is nearing failure, proactive hot sparing will activate to preemptively copy the contents of the drive before it fails.

If an SSD hot spare is not available, a single pair of FAST Cache drives will be in a degraded mode at which time the cache-page cleaning algorithm will increase the rate at which FAST Cache pages are copied out to HDDs. Only read operations will be allowed from this degraded RAID group. Write operations destined for the degraded RAID group will be serviced by the HDDs. Any other RAID groups in FAST Cache will continue to operate as read/write.

10.8.6 FAST Cache Monitoring in Unisphere

Unisphere Analyzer gathers statistics that can be used to monitor FAST Cache Performance, which are viewable with Analyzer's Advanced mode enabled.

The following FAST Cache statistics are available at the storage processor level:

- FAST Cache Dirty Pages (%)
- FAST Cache MBs Flushed (MB/s)

The following FAST Cache statistics are available for RAID group LUNs and storage pools:


- FAST Cache Read Hits/s
- FAST Cache Read Misses/s
- FAST Cache Read Hit Ratio
- FAST Cache Write Hits/s
- FAST Cache Write Misses/s
- FAST Cache Write Hit Ratio

10.8.7 FAST Cache Configuration

The FAST Cache enabler must be installed on the storage system in order to use FAST Cache.

Note: FAST Cache should only be enabled during planned maintenance windows since SP cache is disabled during the process.

FAST Cache is configured by selecting **Manage Cache** under **System Management** in Unisphere. Below is a screenshot of the configuration window for configuring FAST Cache:



Note: FAST cache requires additional memory to be dedicated for SP usage. The amount of SP memory required to run FAST Cache varies by model.

After FAST Cache is initialized and shows as enabled on the storage system, it can be enabled for existing storage pools and RAID Group LUNs.

11 Storage Provisioning Best Practices

11.1 Storage Provisioning Options

It is important to understand the application requirements to select the approach to meet these requirements, however should conditions change it is possible to migrate data between thick LUNs, thin LUNs and traditional/Classic LUNs without application downtime.

The following is a table of general use cases for each option²³

Provisioning / LUN Type	General Use Case
Virtual Provisioning Thick LUNs	- For applications that require good performance. A thick LUN's performance is comparable to the performance of a Traditional/Classic LUN and is better than the performance of a thin LUN. - For applications with random workloads
Virtual Provisioning Thin LUNs	- For applications with moderate performance requirements - When storage efficiency requirements outweigh performance requirements. - For applications where space consumption is difficult to forecast - Do not use with VNX OE for File; use only thick LUNs when using pool LUNs with VNX OE for File
Traditional RAID Group Traditional/Classic LUNs	- For applications that require extreme performance (for example when milliseconds of performance are critical) - For most predictable performance - For precise data placement on physical drives and logical data objects - For physical separation of data (can also be achieved by using separate disk pools)

11.2 Virtual Provisioning

Virtual Provisioning provides pool-based storage provisioning. Pools are a collection of physical disks from which Thick and Thin Pool LUNs can be created.

For detailed description of Virtual Provisioning including a list of terminology, please refer to the [EMC VNX Virtual Provisioning White Paper](#).

11.2.1 Pool LUNs (Virtual Provisioning)

When using Virtual Provisioning a pool LUN is similar to a traditional/Classic LUN in many ways. Many of the same Unisphere operations and Secure CLI commands can be used. Most functions work the same way including underlying data integrity features, LUN migration and local protection. Pool LUNs are available as choices in Unisphere Tasks and Unisphere Analyzer. Features such as hot sparing and proactive sparing also work with pool LUNs.

Pool LUNs are comprised of a collection of slices and have the option to be thin or thick. A slice is a unit of capacity which is allocated from the private RAID Groups to the pool LUN when it needs additional storage. Starting with VNX Operating Environment for Block Release 33, the slice size has been reduced to 256 MB, in earlier releases (VNX OE 31 and 32) the slice size is 1 GB.

²³ Based on extract from *EMC VNX Virtual Provisioning Applied Technology White Paper*

11.2.2 Thick LUNs (Virtual Provisioning)

When using Virtual Provisioning a thick LUN has higher and more predictable performance in comparison to a thin LUN. Thick LUNs reserve capacity and will be fully allocated (all slices) from a pool at time of creation. The writes to the slices are organized according to the Logical Block Address range and therefore the considered to be deterministic (i.e. the slice where data is stored is determined by the LBA).

11.2.3 Thin LUNs (Virtual Provisioning)

Thin LUNs can present more storage to an application than is physically available.

- Data and LUN metadata is written to thin LUNs in 8 KB chunks. Thin LUNs consume storage on an as-needed basis from the underlying pool.
- The *user capacity* (storage perceived by the application) can be larger than the *consumed capacity* on the storage subsystem
- Thin LUNs generally do not perform as well as Thick LUNs because of the overhead of mapping the blocks of data to the slice in which they are contained.
- Can be sized to accommodate data growth with regard for currently available assets

11.2.4 Storage Pools (Virtual Provisioning)

The following guidelines should be followed when implementing Virtual Provisioning

- Create homogeneous, single tier pools
 - Consistent performance within a pool, using the same disk (size and speed) and RAID protection level.
 - Heterogeneous pools (different disk types) are not currently supported in Citi.
 - Pools can be as large as the maximum number of drives (except vault drives and hot spare) allowed per system type (refer to VNX Hardware Models section for max drive numbers per VNX model)
- Create separate pools
 - For different disk types
 - To separate workloads with different I/O profiles.
 - Predominantly sequential workloads should be placed in dedicated pools where possible.
 - For Reserved LUN pool for snapshots
 - Where using FAST Cache and wish to enable for some LUNs and not others
 - Whilst technically both Thick and Thin LUNs can exist in the same pool it is recommended that they are separated into different pools where possible to simplify management of oversubscription.
 - To minimize failure domains. A fault domain refers to data availability. A Virtual Provisioning pool is made up of one or more private RAID groups. A pool's fault domain is a single pool private RAID group. That is, the availability of a pool is the availability of a single private RAID group.
 - The risk of the loss of a single private RAID group in a pool is lower when using RAID 6. RAID 5 is not a recommended protection level when implementing pools in Citi.
 - Consideration should be given to creating multiple smaller pools. Using 2 separate pools on different physical disk shelves would reduce the risk of data loss in certain applications (for example separating DB data from DB logs).

11.2.5 Storage Pools with VNX File

A Mapped Pool is a Storage Pool for File that is created during the storage discovery process for use on VNX for File. It is a one-to-one mapping with a VNX Storage Pool for Block

Where block storage pools are used with VNX File

- Separate pools should be used for file and block data.
- Avoid enabling FAST cache for VNX SavVol. Where possible use a separate pool for SavVol (checkpoint) storage.
- Only use Thick Pool LUNs (do not use Thin, compressed or de-duplicated LUNs)
- Balance LUN ownership between SPs
- Pro-actively create LUNs and assign them to VNX OE for File, so that VNX OE for File has available space for file system creation and extension, snapshots, and so on
- Allow LUN creation process (thick LUN slice allocation) to complete before adding LUNs to the File storage group
 - Creation status can be checked using `naviseccli lun -list -opDetails`

Please refer to the [EMC VNX NAS Solutions Guide](#) for further information, including LUN sizes, LUN quantities and volume management configuration.

11.2.6 Protection Levels / RAID Types (Virtual Provisioning)

The below table provides information on general use case and guidelines for protection levels when implementing Virtual Provisioning:

RAID Type	Drive Count	General Use Case and Guidelines
RAID 6	14+2 6+2	RAID 6 should be the preferred RAID type (especially for large capacity NL-SAS drives) as it is more resilient than RAID 5 • Use larger R6 RAID Group (14+2) where possible. • Larger R6 RAID Groups are typically better for random workloads and lowers the overall cost of storage by increasing the data-to-parity drive ratio.
RAID 5	4+1 8+1	RAID 5 is not recommended for use in Citi with Virtual Provisioning. • RAID 5 has a higher risk of data loss in the event of a double disk failure than RAID 6. The failure of a private RAID group within a storage pool would result in the entire Storage Pool becoming invalid.
RAID 1/0	4+4	Due to the high cost of RAID 1/0 and the lower level of robustness, it should only be used for applications with high performance requirements that cannot be appropriately met using RAID 6. Generally RAID 1/0 works best for heavy transactional workloads with high random write workloads.

Note: When adding disks to a pool, it is recommended to only use multiples of the desired drive count to ensure all private RAID groups are created the same size for consistent performance. For example if using RAID 6 14+2, then create pools that are multiples of 16 disks.

11.2.7 Storage Pool Creation (Virtual Provisioning)

When creating a pool, the following Pool Attributes are required

- Pool Name
- Disks (number and type)
- Protection Level

Drive Type	VNX Classification
Flash / SSD	Extreme Performance
SAS	Performance
NL-SAS	Capacity

Larger pools must be created using multiple operations. Pools can be created using the maximum allowed drive increment and then expanded to reach the desired number of drives in a pool. Please refer to the Max Drive Increments per model shown in the table in the [Limits with Virtual Provisioning](#) section.

11.2.7.1 Virtual Provisioning Free Space Recommendation

Whilst a storage pool can operate with 100% consumed space, it is recommended to leave

- 5% unallocated for VNX OE for Block 32 and above
- 10% unallocated for VNX OE for Block 31

11.2.7.2 Virtual Provisioning Overhead

When provisioning LUNs from a block storage pool there is a Metadata overhead that needs to be taken into consideration. There is both a fixed and a variable component for the overhead.

The fixed overhead is

- 2GB for Thick LUNs
- 2GB for Thin LUNs > 900GB
- 3GB for Thin LUNs < 900GB

The variable overhead is

- Approximated to be 3% of the LUN's user capacity (this is considered to be a conservative approximation for thin LUNs 250 GB and larger, while smaller LUNs may have slightly higher variable overhead).

This formula only approximates the storage consumed from the pool if you write to the entire LUN. Below are some examples for different sized User LUNs (size visible to the host), showing the consumed space for entire Thick and Thin LUNs.

User LUN (GiB)	Thick (GiB)	Thin (GiB)
50	54.11	55.11
100	105.21	106.21
250	259.51	260.51
500	516.01	517.01
750	772.51	773.51
1024	1054.05	1054.05
1536	1579.08	1579.08
2048	2106.11	2106.11
2400	2466.81	2466.81
3072	3157.15	3157.15
3584	3684.18	3684.18

11.2.8 Storage Pool Expansion (Virtual Provisioning)

Storage pools can be expanded non-disruptively by adding drives to the pool. The expansion of a storage pool will automatically invoke a one-time rebalancing/relocation of the slices within the pool to reduce the risk of inconsistent performance within the single tier storage pool.

It is recommended that pool expansions be scheduled during a period of low disk activity due to the additional disk activity during rebalancing/relocation.

This feature does not require the FAST VP license to be installed.

Note: When adding disks to a pool, it is recommended to only use multiples of the desired drive count to ensure all private RAID groups are created the same size for consistent performance. For example if using RAID 6 14+2, then expand pools in multiples of 16 disks.

11.2.9 Limits with Virtual Provisioning

The table below lists the limits with VNX Operating Environment for Block 33 for each of the VNX models.

System	Max Pools	Max Drive Increments	Max Disks per pool	Max LUNs per pool	Max LUNs per Array	Max User visible LUNs per Array ²⁴	Max RAID Groups
VNX 5200	15	80	121	1000	2048	3816	62
VNX 5400	15	80	246	1000	2048	3816	62
VNX 5600	20	120	496	1000	2048	3816	125
VNX 5800	40	120	746	2000	4096	6864	250
VNX 7600	40	120	996	3000	4096	8376	500
VNX 8000	60	180	996	4000	8192	14496	1000

The table below lists the limits with VNX Operating Environment for Block 32 for each of the VNX models.

System	Max Pools	Max Drive Increments	Max Disks per pool	Max LUNs per pool	Max LUNs per Array
VNX 5100	10	20	71	512	512
VNX 5300	20	40	121	512	512
VNX 5500	40	80	246	1024	1024
VNX 5700	40	120	496	2048	2048
VNX 7500	60	180	996	2048	2048

²⁴ This includes FLUs + pool LUNs + Snapshot LUNs + MirrorView/A LUNs

11.2.10 LUN Creation (Virtual Provisioning)

When creating a pool LUN, the following inputs are required

- Pool selection
- Amount of host visible user capacity
- Pool LUN name (optional)
- Thin, or not (Thick)

Note: With VNX Operating Environment from Block 33, by default Thin is selected when creating a Pool LUN. Please be sure to un-tick this when creating Thick LUNs in a pool.

The following restrictions apply to pool LUNs (enforced through Unisphere)

- Maximum size 16TB
- Minimum size 1MB (Note that the *consumed capacity* for the smallest LUN is 3GB due to space being allocated for pool LUN metadata)
- Cannot be a component of a metaLUN
- Thin LUNs cannot be used in the reserved LUN pool (Thick LUNs can be used) for SnapView snapshots
- Default cache options are not tunable. FAST cache is enabled at a pool level and therefore applies to all LUNs in the pool.

The following are recommended best practices for managing pools

- It is recommended not to change the default owner of the pool LUN once it has been provisioned as it adversely affects performance
 - This will change the underlying private structures, which provide storage for the pool LUN which are controlled by the original SP.
 - If it is desired to change the ownership of a LUN after it is created, this should be achieved by using LUN migration to migrate the LUN to a new LUN with the desired SP owner (this may then require a LUN trespass to the new owner).
- It is recommended to allow the system to balance pool LUNs on the SPs (this is the default)
- Create at least 2 LUNs per Storage Pool to allow ownership by both SPA and SPB.
- You should take into consideration the metadata overhead when provisioning pool LUNs. Refer to the [Virtual Provisioning Overhead](#) section.

11.2.11 LUN Expansion and Shrinking (Virtual Provisioning)

Whilst the VNX has the capability to expand and shrink LUNs (depending on host OS), this feature requires further testing to be completed and best practices developed in conjunction with the various platform teams.

Note: LUNs that are assigned to VNX for File cannot be expanded. If additional capacity is required on VNX for File, new LUNs should be created and assigned for AVM to manage.

11.2.12 Monitoring, Adding and Deleting Pool Capacity

Terminology

Term	Definition
Total Capacity	Total physical capacity available to all LUNs in the pool
Total Allocation	Total physical capacity assigned to all LUNs in the pool
Subscribed Capacity	Total host reported capacity supported by the pool
Over-Subscribed Capacity	Amount of user capacity for LUNs that exceeds the physical capacity in the pool. Thin LUNs with larger user capacity than the physical capacity can result in an oversubscribed pool.

Monitoring Storage Pools

The Storage Pool Properties shows parameters such as Free Capacity, Percent Full, Total Allocation, Total Subscription, Percent Subscribed and Oversubscribed Capacity.

Pools are monitored using the **% Full Threshold Alert** setting and storage capacity reports. This alert is only active if there are one or more thin LUNs in a pool, since thin LUNs are the only way to oversubscribe a pool (the alert is not active if the pool only contains thick LUNs). You can specify the value for **% Full Threshold** (Total Allocation / Total Capacity) when a pool is created.

Alerts are visible using the **Alerts** tab in Unisphere. Alerts should generate an SNMP trap to be sent to the Enterprise Monitoring System. Below is the information on the thresholds and their settings.

Threshold Type	Threshold Range	Default	Alert Severity	Alert Severity
User settable	1% - 84%	70%	Warning	None
Built-in	N/A	85%	Critical	Clears user settable alert

Allowing total allocation to exceed 90% of total capacity puts you at risk of running out of space and impacting all applications using LUNs in the pool.

As over-subscribed pools can run out of space, it is important to have considered trigger points (thresholds) at which time additional storage will be added to the pool to increase the capacity available to the LUNs in that pool. Appropriate trigger points are yet to be agreed to within Citi and will be added to the guide at a later date. It is not recommended to use over-subscribed pools for critical applications.

11.3 Traditional/Classic Storage Provisioning

11.3.1 Traditional/Classic RAID Groups Overview

Traditional block storage provisioning should be used for systems running VNX OE for Block 31 or where it is determined that the performance requirements cannot be met using Virtual Provisioning.

A Traditional/Classic RAID Group is a logical group of physical disks that will be used to bind one or more LUNs. Each LUN that is bound in a RAID Group will be distributed equally between the physical drives in that group. The RAID Group's RAID type is determined by the RAID type of the first LUN that is bound to it or specified at the time the RAID Group was created.

11.3.2 RAID Level (Traditional/Classic RAID Groups)

The below table provides information on general use case and guidelines for implementing Traditional RAID Groups:

RAID Type	General Use Case and Guidelines
All	- RAID groups can have a maximum of 16 drives. - All drives within a RAID Group should be of the same size, speed, and type. - RAID groups in the vault enclosure should not span enclosures. This is to prevent inconsistent data in the case where power to the array is lost and the backup battery power is used to keep the vault enclosure alive while the cache is being destaged to the vault drives.
RAID 6	RAID 6 should be the preferred RAID type (especially for large capacity NL-SAS drives) as it is more resilient than RAID 5 and has comparable performance - Use an even number of disks with a minimum of 6 and a maximum of 16 total disks. - R6 2+2 RAID is not recommended (cost benefit of R6 RAID Groups will not be realized). - Create larger R6 RAID Groups (10+2, 12+2, and 14+2) where possible. - The added cost of storing the additional parity in RAID 6 as compared to RAID 5 can be offset by creating larger RAID 6 RAID Groups. e.g. Instead of R5 (4+1) use R6 (8+2). - Larger R6 Raid Groups are typically better for random workloads and lowers the overall cost of storage by increasing the data-to-parity drive ratio. Note: Once RAID 6 RAID Groups are created, they cannot be expanded. If additional space is required in a RAID 6 RAID group, it is possible to free up capacity by using LUN Migration to migrate or more LUNs to a different RAID Group.
RAID 5	RAID 5 has a higher risk of data loss in the event of a double disk failure than RAID 6. - RAID 5 should not be used for high capacity NL-SAS drives as the risk of a double-disk failure is too high due to the longer rebuild times. - RAID 5 should only be implemented as an exception following receipt of appropriate approval. - For random large-block random read activity with RAID 5 (4+1), select the option to use large element size of 1024 blocks (512KB). The use of large element size should be avoided for all other workloads.
RAID 1/0	Due to the high cost of RAID 1/0 (twice the amount of useable space), it should only be used for applications with high performance requirements that cannot be met using RAID 6. - RAID 1/0 is also the only protection type for FAST Cache. - RAID 1/0 RAID Groups should not comprise of more than ten disks. - M1 and M2 of a RAID 1/0 RAID Group should reside on different DAE's and on different buses if possible.

11.3.3 Sequential Optimization

In environments with predominantly large-block sequential workloads, optimization may occur when the array can perform full-stripe operations. Whilst it may be possible for this to occur with non-preferred drive counts, the following configurations should be considered for large-block sequential workloads:

RAID Type	Preferred drive count for Large-block Sequential Workloads
RAID 6	RAID 6 has a preference of 8+2
RAID 5	RAID 5 has a preference of 4+1 or 8+1
RAID 1/0	RAID 1/0 has a preference of 4+4

11.3.4 Traditional / Classic LUNs

This section will cover the configuration guidelines for LUNs created in Traditional RAID Groups.


- EMC has adopted the terminology “Classic LUN” to be *logical unit of storage created on a user-defined RAID group. The amount of physical space allocated is the same as the user capacity seen by the host server. Classic LUNs cannot be created on a pool; they are always created on a RAID group.*
- A LUN has a RAID type and other properties that define it.
- A LUN is carved out of a single RAID Group. The RAID type of the first LUN that is bound to a specific RAID Group will define the RAID type of the RAID Group. Thus, all LUNs in a given RAID Group are of the same RAID type. Only one SP has ownership of the LUN at a given time, but this ownership can be transferred to the alternate SP either manually or automatically in the case of an SP or host path failure. This transfer of ownership is called a LUN trespass.

11.3.4.1 LUN distribution

LUNs for a given host should be distributed between different RAID Groups (if the RAID Groups are of equal configuration), where possible. This will ensure the workload for the host will be spread across more physical drives.

11.3.4.2 SP Ownership

When binding LUNs, the default owner (or SP) is specified. To evenly distribute the load, the auto option should be selected when choosing the Default Owner of the LUN.



11.3.4.3 LUN Sizing Guidelines


The LUN size can be as large as the amount of free space available in the RAID Group. Host OS and application limits need to be taken into consideration when determining LUN size.

11.3.4.4 LUN Element Size

The element size is the number of disk sectors (512 bytes) that the storage system can read or write to a single disk without requiring access to another disk. The default is 128 sectors (or 64 KB) and cannot be changed.

11.3.5 Creating a Classic LUN

Creating LUNs is accomplished through a process called binding. The RAID type of the first LUN that is bound to a RAID Group will determine the RAID type of the RAID Group. When creating the LUN, the following parameters should be specified:



1. Specify the **RAID Type** (6, 5, or 1/0) of the LUN.
2. Select the **RAID Group**.
3. Specify the **LUN Properties**:
 - a. User Capacity
 - b. LUN ID
 - c. Number of LUNs to create
4. Specify the following in the Advanced Tab
 - a. Unless you want to disable any Caching, leave all the Cache options selected. Deselecting any of the Cache options will degrade performance for the LUN.
 - b. **Enable Auto Assign** should remain unchecked.
 - c. **No Initial Verify** should remain unchecked.
 - d. Choose **Auto** to have the array assign the **Default Owner** of the LUN. This will provide the most even distribution of the LUNs between the two SPs

11.3.6 LUN HLA (or HLU)

On UNIX hosts, the “d” portion of a “ctd” address is defined in the Storage Group properties. This can be changed when adding a LUN to a Storage Group via the CLI (-hlu parameter) or when selecting LUNs to add to the Storage Group.

On Wintel hosts, this HLU value and the Windows physical disk number do not always match and are negotiated as the LUNs are added to the host.

The HLU value for LUNs assigned to a host can be determined by issuing the following CLI command

```
# naviseccli -h <SP IP address> storagegroup -list -gname <Storage Group Name>
```

11.3.7 MetaLUNs

A metaLUN is a logical volume comprised of multiple LUNs, each from a different RAID Group. They can either be striped, concatenated, or hybrid (concatenating two striped metaLUNs). All LUNs comprising the metaLUN must be on the same SP, so all I/O going to the metaLUN will be going through one SP.

If possible, use a host-based volume manager to group multiple LUNs since the LUNs comprising the volume can be distributed evenly between the SPs of the VNX. If a volume manager is not available on the host, a striped metaLUN should be used. Striped metaLUNs are preferred over concatenated since striped metaLUNs will provide better distribution of the I/O across the spindles than concatenated metaLUNs. Use either a host volume manager or metaLUNs to group LUNs, not both; using both can cause cross-plaiding where stripe members can reside on the same spindle which will negatively impact performance. The LUN Migration feature should also be the preferred method for LUN expansion over LUN expansion by forming a metaLUN.

11.3.8 LUN Migration (Classic LUNs)

The LUN Migration feature allows for the migration of data from one LUN or metaLUN to another. LUNs can only be migrated from LUNs of equal or greater size. This may be useful in the following cases:

- Migrating the LUN to another disk type (i.e. SATA disk to FC disk)
- Migrating the LUN to another RAID type (i.e. RAID 1/0 TO RAID 5)
- Migrate to a LUN of larger size (this method is preferable to expansion by the use of meta-LUNs).
- Performance related issues.

Note: Host I/O performance will be negatively impacted during the migration process.

12 Data Protection Best Practices

12.1 Local Protection

The VNX Service Processors are clustered so as to protect from the failure of a single Service Processor (SP). In the case of an SP failure, the LUNs owned by the failed SP will be taken over by the surviving SP (termed LUN trespassing). Each Service Processor also has built in redundancy with dual power supplies, mirrored write cache as well as multiple paths connecting to the disk subsystem.

For cluster support at the host level, please refer to the [CSBE SAN Stack](#) Solutions Guide for the specific host platform for the certified hardware/software configurations.

For details on Local Data Protection using Snapshot and Clone technologies, please refer to the Local Protection section.

12.2 Remote Protection

Replication at the storage array level is not certified for use in Citi for EMC VNX.

System/data owners should ensure that any requirement for replication for COB purposes is addressed at the host or application level. For details on certified replication approaches please consult with the application and platform owners.

Further information is also available in the [Open Systems Data Replication Playbook](#)

12.3 Backup Standards

All data is accessed via the hosts to which LUNs are allocated. For a complete list of backup options, please refer to the [CSBE Backup Hardware and Software Products](#) site.

13 Security Best Practices

13.1 Security Standards

Citi Information Security Standards (“CISS”) are available from [Information Security Services \(“ISS”\) website](#).

Citi’s SAN, NAS and Backup infrastructure stores and processes information across the full range of information security classifications.

For data classified as Public or Internal, no specific provision is required for data encryption either in-flight or at-rest.

For Confidential and Confidential PII support, all new applications are required to encrypt confidential PII data in transit (refer to Encryption Requirements Section of [CISS](#)). For Restricted or Authentication data, encryption of data is required both in-flight and at-rest.

The recommended approach for Restricted, Authentication and new applications storing Confidential or Confidential PII data would be to encrypt the data within the application prior to transmission over the network. This approach has the benefit of also meeting the requirement in all these cases for persistent encryption when non-Citi managed storage infrastructure is used. At the time of writing this document there are no tested or approved solutions at the storage infrastructure layer to encrypt data in-flight or to encrypt the data at rest, meaning that where required, this functionality needs to be applied by the application.

This classification, by definition, infers a “Medium Risk” status on all Citi’s SAN, NAS and Backup infrastructure.

13.2 Audit Logging

Off-array management servers can be setup as central event monitoring servers to capture VNX audit logs in the Windows Application log. Audit log reconciliation for the EMC VNX environment is done through Citi's Secure Event Monitoring System (SEMS).

13.2.1 SEMS Audit Log Reconciliation

Citi's Security Event Monitoring System (SEMS) is used to centrally collect all security events. Custom reports can then be generated from the events collected by SEMS and sent to the log review teams for reconciliation. The following steps need to be followed to integrate the EMC VNX Block systems with SEMS:

1. Central event monitoring must be configured to send all EMC VNX event logs to the off-array management server's Windows Application Event Log (refer to the [Event Monitor](#) section for setting up the Event Monitor).
2. The SEMS IDVA team should be contacted to initiate the request for integrating the VNX Block with SEMS for log reconciliation. The EMC VNX management server should be specified as the source device in the request. Use the following distribution: ***CTI Global SOC SDM**
3. The SEMS IDVA team will submit a request for the system administrator of the EMC VNX management server to create a local functional ID for the SEMS system.
 - a. The functional ID should be named **arcsight_app**
 - b. The functional ID should be given read-only access to the Windows Application Event Log.
4. The IDVA team will then configure Arcsight to import EMC VNX audit events and generate daily reports to the regional log review teams.

The IDVA team will filter the report to only include audit events. Please refer to the [VNX Block and CLARiiON SEMS Events Include List](#).

13.3 Data Security

The recycling or reusing of LUNs from one host/application to another is prohibited. The recycling of LUNs compromises the data written by the initial host/application.

It is recommended that all LUNs must be deleted/unbound from the RAID group when no longer required by the host to which it was initially presented.

It is also recommended that the host SA should follow guidelines consistent with the data erasure of a physical drive prior to deleting the LUN.

13.3.1 Disk Sanitization

Disk sanitization is the process of physically obliterating data on disks in a manner that precludes recovery of that data by any known recovery methods.

Disk sanitization is a requirement mandated by CITMP IS Policy.

As there is currently no feature available within the FLARE Operating Environment that enables overwriting the contents of a disk, it is recommended to follow an approved process in accordance Global Data Center Secure Media Disposal policies to

- degauss the disks
- use an approved 3rd party to securely erase or dispose of the disks (compliance with Global Data Center Secure Media Disposal policies should be verified and approved by an appropriate Information Security Officer)

NOTE: Current CITMP limitations restricting movement of electronic media warrant that all disk sanitization be completed without the transportation of disks.

13.3.2 Disposal of failed disks

All failed disks within every CTI data center is subject to stringent sanitization or degaussing procedures. As such, all failed media should be processed in accordance Global Data Center Secure Media Disposal policies.

13.4 CITMP Information Security Policy Convergence

13.4.1 Information Ownership

The head of Distributed Engineering has information ownership of Storage products and all related engineering documents. Storage products are available in the Citi Technology Catalogue. All security and technology best practices for Storage products are the responsibility of CATE Storage and Backup Engineering and can be referenced from the [CSBE website](#).

13.5 Role Based Access Controls (RBAC)

13.5.1 Roles Defined in UniSphere

The following roles are defined for users and groups in UniSphere²⁵

VNX Role	Privileges
Operator	Read-only privileges for storage and domain tasks. No privileges for security tasks.
Network Administrator	Full privileges to configure DNS, IP settings, and SNMP. Read-only privileges for storage and domain tasks. No privileges for security tasks.
NAS Administrator	Full privileges for File tasks. Read-only privileges for Block tasks. No privileges for security tasks.
SAN Administrator	Full privileges for Block tasks. Read-only privileges for File tasks. No privileges for security tasks.
Storage Administrator	Full privileges for File and Block tasks. No privileges for security tasks.
Security Administrator	Full privileges for security tasks including domains. No privileges for File and Block tasks.
Administrator	Full privileges for File, Block and security tasks.
Local Data Protection	Privileges only to do Snapview (Snapshots and Clones) and SnapSure (Checkpoints) tasks. Data recovery operations like rollback a snapshot or reverse synchronize a clone are not allowed. No privileges to create new storage objects. No privileges for security tasks
Data Protection	All privileges included in Local Data Protection plus MirrorView and SAN Copy tasks. Data recovery operations like promoting a secondary and fracturing a mirror are not allowed. No privileges to create new storage objects. No privileges for security tasks.
Data Recovery	All privileges included in the Local Data Protection and Data Protection roles plus privileges for data recovery tasks. No privilege to create new storage objects. No privileges for security tasks.

²⁵ The default roles and descriptions on VNX can be obtained from Unisphere Help

13.5.2 Roles defined for use in Citi

For everyday administration, Citi administrators should perform one of the following 3 roles:

- Security Administrator
- Storage Administrator
- Storage Monitor

This section breaks down the responsibilities and corresponding UniSphere role for each role.

Citi Role	Description	UniSphere Role
Security Administrator	The role of the security administrator is to • Manage storage domains • Create / manage user accounts (including role association) • Manage audit / event logs	Security Administrator
Storage Administrator	The role of the Storage Administrator is to • Make configuration changes to the storage system. UniSphere will allow these two roles to perform configuration changes to the storage system but does not allow this user to manage user accounts nor logs.	SAN Administrator (block only) Storage Administrator (block + file)
Storage Monitor	The storage monitor will only be given read-only access to the storage system for troubleshooting and reporting purposes	Operator
System Administrator	A system account should be maintained to be used in the unexpected event of being unable to manage the system using the above roles. The system account should be used to create the initial Security Administrator accounts and then should be immediately handed over to an Information Security Officer/Administrator to ensure that these accounts are not used for system administration. Should there be a need to use the system account; the requester should use an auditable and traceable Change Management system to request the use of the account.	Administrator

13.6 Functional IDs

The following accounts (functional IDs) will need to be created for applications that need to access the VNX storage array.

Application	Account	Interactive / Non-Interactive	Unisphere Role	Password Change Frequency
ESRS	emc	Interactive	Administrator	In accordance to CISS for privileged accounts
ECC	ecc	Non-Interactive	Operator	Not required
SANscreen	sanscreen	Non-Interactive	Operator	Not required
Arcsight	arcsight_app	Non-Interactive	Operator	Not required
Replication Manager ²⁶	rmagent	Non-Interactive	SAN Administrator	Not required

²⁶ The use of Application Protection is Restricted to the Exchange 2010 environments. For further details please refer to the [Exchange 2010 Storage Design Document](#)

13.7 Storage Domains

Storage Domains need to be created for central management of security settings. Storage Domains allow for the creation of local and global user accounts for storage system management. Unisphere Manager V1.1.25 and above allows for managing multiple VNXs in a single domain. Multi-domain management can be used to manage VNX domains and legacy (CLARiiON and Celerra) domains in a single Unisphere Manager session.


If there are multiple domains in the environment, multi-domain management should be configured so that the off-array management server can be leveraged to manage multiple domains in the environment. The management server should always be running an equal or later version of the management software than the arrays being managed.

13.7.1 Setting Up the Storage Domain


The procedure to setup a single VNX storage domain is discussed below.

Only the Security Administrator or Administrator users can configure domains. EMC should configure the network settings of the VNX to allow remote connections for management at array initialization.

1. Launch Unisphere by typing the SPA IP address of the EMC VNX in the URL of Internet Explorer.
2. Click Yes at the following screen:



3. Enter the initial Unisphere Account information. This is the initial Global Administrator user for the domain. One Global Administrator account must exist in the domain but this user account should be used in a limited basis. The Security Administrator user should be created immediately after and should be used to create the necessary user accounts needed to manage the storage arrays.



- The array should show up as the master at domain:

Node Information	System Name	Provider Information
10.40.42.208	VNX5500-1868	AnalyzerProvider (7.31.0.4.16), i
10.40.42.207	VNX5500-1868	AnalyzerProvider (7.31.0.4.16), i


The domain is now setup and additional storage arrays can be added to the domain.

13.7.2 Adding Storage Arrays to an Existing Domain

Note: If the array being added to the domain exists in another domain, all of its domain information from the prior domain will be lost.

Storage arrays can be added to the existing domain as follows:

- From Unisphere, select “All Systems” from the drop down box at the top
- Select the “Domain” tab
- Select the domain to add new array
- Select “Add/Remove Systems” from the navigation list on the right



Note: It may take a few minutes for the added array to show up in the domain.


13.7.3 Setting Up Multi-Domain Management

If there are multiple domains in the environment, multi-domain management should be set up on the off-array management server to manage these domains. In order to manage arrays in other domains, the user account used to log on to the off-array’s domain must also exist in the domain it wants to manage. Only the Security Administrator is allowed to configure domains.


The following are steps to configure multi-domain management:

- Use Unisphere to connect to the storage array running the latest code. This will allow the configuration of multi-domain management using the Security Administrator account of the domain.


- Under “Domains” tab, click on “Manage Multi-Domain Configuration”.




- In “Multi-Domain Management”, enter the IP address of the master node and name of the domain being added. Click the right-pointing arrow.



The Domain should appear in the “Selected Domains” box. Click “OK”.




- The newly added domain should show up in the list of domains in the main Unisphere window after a while. If it does not, click on the green “Refresh” icon on the right.



5. You should be able to see all storage arrays that are managed in the multiple domains listed in the “All Systems” drop down box. If not, you might need to do a manual login to the domain with storage arrays not showing.

Click on the domain and click on “Login” as shown in the previous screenshot. Login to the domain with an account that has “Security Administrator” or “Administrator” privilege in the domain.



The arrays can now be managed using multi-domain management when connected to the array where multi-domain management was configured.

13.7.4 Adding a Domain User

13.7.4.1 User Accounts

VNX should be setup to authenticate users with Citi's Active Directory (AD) for the following reasons:

- Central management of user accounts.
- Enforces user account and password policies as defined in AD
- Assign Unisphere roles to AD groups instead of having to map roles to each individual user.
- Provides both segregation of duties and makes it easier for Storage Operations to outsource the user management functions including entitlement reviews to ISA.


13.7.4.2 Unisphere LDAP User Authentication

Unisphere has the ability to integrate with AD to authenticate users. To setup Unisphere AD user authentication, the ISA will need to do the following:


1. Unisphere will need a functional ID and an associated password to be created in AD which it will use to bind to the directory. For further information on AD LDAP, please refer to the [CATE Microsoft Active Directory site](#).
2. AD Groups should be created for each of the Unisphere roles needed to manage a block device: Operator, SAN Administrator, Storage Administrator and Administrator.
3. Once the groups have been created in the AD directory, users should be assigned membership to the AD group with their corresponding role.

4. This is a one-time configuration setting. Once this setting is made, all user management tasks will be done on AD, not Unisphere. To configure AD user authentication in Unisphere, do the following:


- a. Log into Unisphere using the default global Administrator account.
- b. Under the **Domains** tab, select **Manage LDAP Domain for Block**:



- c. Enter the AD connection settings and certificate. All of this information should be provided by CATE Directory Engineering through the request process and certificate added for the SSL connection to the AD server:



- d. In the **Role Mapping** tab, map the Unisphere roles to their corresponding LDAP group. For example:



- e. In the **Advanced** tab, put the following entries:

10.40.38.46 - Modify LDAP Connection Settings

Server | Role Mapping | **Advanced**

Advanced Fields

User ID Attribute: sAMAccountName

User Name Attribute: cn

Group Name Attribute: cn

Group Member Attribute: member

User Object Class (optional):

Group Object Class: group

Additional Functions

Duplicate the current service settings [Duplicate...]

- f. Click **OK** and a success message will appear. The AD server should now show up as the primary server. Hit the synchronize button to synchronize the user database with the AD server
- g. Modify the **Synchronization Rate**. This specifies how often Unisphere will synchronize the user information with the AD server. This example specifies to synchronize every hour but this may have to be increased if it generates too much load on the replicated AD infrastructure.

Configure LDAP for Block Storage Systems

LDAP Services

Service Connection(s):	Ranking:
10.40.38.46	primary
10.40.38.79	alternate

Add a new LDAP service connection [Add...]

Modify selected connection setting [Modify...]

Remove selected connection setting(s) [Remove...]

Synchronize accounts with current service [Synchronize...]

Synchronization Rate


1440 [v] ☒ minute(s) ☐ hour(s) ☐ day(s)

☐ Bypass Certificate Verification for LDAP

[OK] [Apply] [Cancel] [Help]

- h. Repeat these steps to setup the connection to the secondary LDAP server.

Once the AD settings have been configured and successful connection to the replicated SSO LDAP infrastructure has been established, users can log in with their AD login. For example:



Note that the **Use LDAP** checkbox is selected

13.7.5 Global/Local User Accounts

A global user account gives the user the same access privileges to all storage systems in the domain.

A local user account gives the user access privileges to a specified subset of storage systems in the domain.

Once AD authentication has been configured, only the default Administrator global account (to be used on a limited basis in cases where all other access is not possible) and the ISA Security Administrator global accounts should reside in the domain.

Note: All other users should be managed in AD.

13.8 Communication Interfaces

The following table lists information on the interfaces used to communicate with the EMC VNX storage arrays.

Component	Interface Type	Communication Ports
Unisphere	IP	80/443 or 2162/2163 ²⁷
Navisphere CLI	IP	6389, 80/443 or 2162/2163 ²⁷
Unisphere Host Agent	IP	6389
ESRS	IP	5414 and 7338
Host connectivity to LUNs (Powerpath as the multipathing software)	FC	N/A
Service Laptop direct connection (Null modem or crossover cat5)	Point-to-Point	N/A

²⁷ 2162/2163 are alternate port pairs that may be used to hide the storage system from attacks that target the default HTTP and SSL/TLS ports

13.9 Unisphere Management Server Configuration

The Unisphere Management Server is an application that runs on the storage arrays. The Unisphere Management UI allows the user to connect to the Management Server software to manage the storage arrays.

13.9.1 Management Server Listening Port

The Management Server listens on TCP ports 80/443. This is the default setting.

13.9.2 Configuring the Management Access Control List (ACL)

IP ACLs can be configured by a user with Administrator or Security Administrator privileges. This allows the user to specify the IP address of servers that are allowed to connect to the VNX to administer storage arrays. Only the IP address of servers/workstations requiring management access to the VNX should be listed in the ACL. To configure this:

1. Open the setup page of the VNX in IE. <https://<IP Address of SP>/setup>
2. Click the **Set Administration Access Restrictions** button
3. Enter the IP addresses of all off-array management stations, SANscreen servers, and ECC Agent Servers that are managing and monitoring the storage arrays and check the **Enable** radio button. Click **Apply Settings**.

Warning, values entered for Filters should be static IP addresses (not dynamic addresses obtained using DHCP). In addition, temporary IPv6 addresses are not recommended due to their limited lifetime. Please use global unicast addresses which will also suffice for their link-local counterpart.

Filters that apply to the connected storage system only

Filters that apply to all storage systems in the domain

10.40.43.29			

IP Filtering State: ☒ Enable ☐ Disable

Apply Settings

13.9.3 SSL/TLS Certificate Renewal

The Security Administrator user should renew the SSL/TLS certificate used for Unisphere SSL/TLS communication between the off-array management station and the storage arrays. This is a manual process that needs to be done at least once every two years per Citi's security requirements. An operations level 3 process should be documented to ensure that the certificate gets renewed at least once every two years. To do this:

1. Log into the setup page of the Management Server on the off-array management station as Security Administrator.
2. Click on the **Manage SSL/TLS Certificate** button.
3. Click on the **Generate a Self-Signed Certificate** button.
4. A confirmation message will appear:

Self Signed Certificate

A new self-signed certificate has been successfully generated.

Back

14 Alerting and Monitoring

14.1 SNMP and Email Alerting

The VNX Block SNMP alerts should be configured to send to an Enterprise Monitoring trap host. The received traps are then filtered and (as appropriate) will create a problem ticket (VT) or generate an SMS / Paging alert.

14.1.1 Adding an Array to the Monitoring System

Tivoli monitoring and alerting service is provided by Citi's ESM team. The ESM team can be contacted at ***ESM Distributed Event Delivery** for instructions on requesting the systems to be monitored. Once the request has been submitted, the ESM team should provide the Tivoli server that will be used to send the VNX SNMP traps to.

It should also be requested to the monitoring team to have the storage arrays setup in the ping scanning list. i.e. if a ping of the array does not respond then this should be configured to generate an alert.

14.1.2 SNMP alerting with Tivoli NetCool

When there is an alert event, the array will send an SNMP alert to the trap host configured in the system. In order for the trap host to create a problem ticket (VT) or send out SMS/ Paging alert, it is important to check that a valid entry exists for each of the SPs in [Aperture](#).

For more information on NetCool Impact Architecture, please refer to the [CATE Netcool site](#).


Notes

- It is only possible to configure a single SNMP trap host Server IP Address (DNS names are not supported).
- Resiliency at the Tivoli NetCool (trap host) level is not available at this time. i.e. if the configured SNMP trap host server is down, all SNMP alerts sent by the storage system will not be received by the Tivoli NetCool environment²⁸

14.1.3 Creating an Alerting Template

The following steps should be taken to configure SNMP alerting on the VNX:

1. Create a template for SNMP alerting.
 - a. From Unisphere, navigate from **System -> Monitoring and Alerting** and select **Notifications for Block**



²⁸ When configuring a new array, it is suggested to check with ESM team if a solution with failover/load-balancing is available for the trap host (not available at time of documentation)

- b. On Configure tab, click on Create Template, select General Events, check all items under Event Security and Event Category, enter a template name

- c. Select the **SNMP** tab on the Template properties window, enter SNMP Management Host (Tivoli NetCool) IP address and check with ESM team on the appropriate Community string to be used, click Apply and OK

- d. On Unisphere, **Distributed Monitors** tab, select one of the SP listed and click on **Use Template**. Then Select the name of the template created from the drop down list and click **OK**.

- e. Under Templates in Use, the template name should appear

Event Notification for Block Storage Systems

Configure | Centralized Monitors | Distributed Monitors | Notification Templates

Distributed Monitors

Filter for

Name	Monitored System
VNX7600-0470-A	E13A1-EMCVNX2-01
VNX7600-0470-B	E13A1-EMCVNX2-01

1 Selected | View Events | Send Test Event | Use Template

Templates In Use

Name
SNMP_All_Events

0 Selected | Stop Using

- f. Select the other SP and verify the same template was assigned automatically

Event Notification for Block Storage Systems

Configure | Centralized Monitors | Distributed Monitors | Notification Templates

Distributed Monitors

Filter for

Name	Monitored System
VNX7600-0470-A	E13A1-EMCVNX2-01
VNX7600-0470-B	E13A1-EMCVNX2-01

1 Selected | View Events | Send Test Event | Use Template

Templates In Use

Name
SNMP_All_Events

0 Selected | Stop Using

- g. Click **Send Test Event**, an alert should be received on Tivoli NetCool console

14.1.4 SNMP MIB

The below table describes the SNMP traps that are generated by the storage array.

OID/ Trap#	Trap Name	Description	Variables	VNX Default Severity	NetCool Default Severity
2	EventMonitor Trap	An EventMonitorTrap is generated in response to a user-specified event. Details are found in Variables data.	hostName, deviceID, eventID, eventText, storageSystem	Severity unable to be determined	Minor
3	EventMonitor TrapInfo	An Informational EventMonitorTrap is generated in response to a user-specified event. Details are found in Variables data	hostName, deviceID, eventID, eventText, storageSystem	Informational	Warning
4	EventMonitor TrapWarn	A Warning EventMonitorTrap is generated in response to a user-specified event. Details are found in Variables data	hostName, deviceID, eventID, eventText, storageSystem	Warning	Warning
5	EventMonitor TrapError	An Error EventMonitorTrap is generated in response to a user-specified event. Details are found in Variables data	hostName, deviceID, eventID, eventText, storageSystem	Error	Minor
6	EventMonitor TrapFault	A Fault EventMonitorTrap is generated in response to a user-specified event. Details are found in Variables data	hostName, deviceID, eventID, eventText, storageSystem	Critical Error	Critical

Whilst the VNX default severity can be determined from the trap, it should be noted that the Variables data contains the relevant information on the event that trigger the trap:

Variable Name	Description
hostName	Array Host Name
deviceID	Array Device ID
eventID	Event ID
eventText	Event description
storageSystem	Storage System Name

The various events have been reviewed within Citi to determine where there should be changes to the default severity. This is then used by the ESM team to create a rule file that maps an eventID to the Citi severity level. If a trap is received that is not defined in the rule file, the above NetCool Default Severity level is used.

In order for the ESM team to complete the SNMP alerting setup, they also require the MIB file for the array. Both the MIB file and the list of events with reviewed severity levels are available on the [CATE VNX Block and CLARiiON site](#).

14.2 Call Home and Email Alerting

Email home should be configured on the storage array by an EMC support engineer for remote alerting. A template should be setup to also send an email to Citi's support team responsible for monitoring the VNX environment each time the VNXs email home to EMC. Citi's SMTP server should be specified to send out the email home alerts in the email home template.

Notes

- The address in the From field must be a valid external email address.
- It is only possible to configure a single SMTP Server IP Address (DNS names are not supported). When selecting a mail relay host, ensure that a Virtual IP (VIP) address is used for the SMTP Server to provide resiliency.
- Do not use an SMTP server that is using storage on the array you are configuring.

14.3 Remote Support (EMC ESRS)

Once the email alerts are received by EMC support, ESRS will provide the vendor with a secure connection to the storage arrays for troubleshooting and problem resolution.

Please reference the [ESRS Solutions Guide](#) for more details. EMC should only be using Unisphere and CLI to troubleshoot issues remotely since user actions can be logged. EMC remote access should not be allowed via ESRS (this is configured in ESRS).

14.4 Event Monitor

For Audit Log Reconciliation, it is required that central event monitoring be configured for all arrays. The off-array management server should be used as a central event monitoring server. This will allow events on all the storage arrays to be sent to the management server's Windows Application Event log. These are then forwarded on to a central system for reconciliation as part of the Audit Log Reconciliation process covered in the [Audit Logging](#) section.

Note that the Unisphere Host Agent is required for setting up a central event monitoring server. Refer to [Unisphere Host Agent](#) section for more details.


Each storage array should be configured to send events to only one central monitoring server. Multiple central monitoring servers may be configured in an environment as long as they are monitoring a separate set of arrays. As such, it is imperative to take into consideration of redundancy setup. This could be in the form of having different central monitoring servers configured for Production and COB sites.

The steps to setup central event monitoring are discussed below. There are two parts to the process:


14.4.1 Configuring Central Monitoring Server

1. In **System** tab, click on **Monitoring and Alerts**
2. Click on **Notifications for Block**
3. In **Event Notification for Block Storage Systems**, click on **Centralized Monitors** tab. Click on **Configure** button.
4. In **Configure Centralized Monitor**, right click on **Portals** and click **Add Portal...**


5. In **Add Portal**, select the array from the **Available Systems** list and add into the **Selected Systems** list. Click **OK**
6. In **Configured Centralized Monitor**, the array should now appear in the **Portals** list
7. Right click on the array and click **Add Centralized Event Monitor**



8. In **Add Event Monitor Agent**, under **Select Systems**, enter the IP Address of the central monitoring server. Click on the right pointing arrow. The central monitoring server should appear in the **Selected Systems** list. Click **OK**




9. In **Configure Centralized Monitor**, the central monitoring server should now appear under the array in the **Portals** list.




Click **OK**

10. In **Event Notification for Block Storage Systems**, under **Centralized Monitors** tab, the IP Address of the central monitoring server should appear in the **Centralized Monitors** list



- a) If you are adding the host for the first time, **Select** the host, click on **Monitor System**. At the pop up, **Select** the storage system from the drop down list and click **OK**




11. Make sure that there is an entry in the central monitoring server's agent.config for each array being monitored to allow for the communication (refer to the [Modifying agent.config File](#) section).


Example: user system@10.40.42.207

14.4.2 Configuring and Applying Event Monitoring Template

1. In **Event Notification for Block Storage Systems**, click on **Configure** tab and click **Create Template**
2. In **Template Properties**, under **General** tab, choose **General** for Events and check all boxes for Event Severity and Event Category. Also check the box for **Log to System Log**. Click **OK**




3. In **Event Notification for Block Storage Systems**, under **Centralized Monitors** tab, right click the IP Address of the central monitoring server configured in earlier and select **Select Global Template**




Note: If you are setting this up for the first time, you must complete the following before you are allowed to select a global template


- a) Right click on the monitoring server's IP address and select Monitor System



- b) Select the system to monitor and click OK




4. In **Select Template**, look for the template created earlier. Click **OK**



5. In **Event Notification for Block Storage Systems**, under **Centralized Monitors** tab, click on the IP Address of the central monitoring server

The event monitoring template should appear in **Details** under **Global Templates**



6. Centralized Monitoring is now setup, click on **Send Test Event** to ensure that setup is working as designed




14.5 Performance Monitoring using Unisphere Analyzer

Unisphere Analyzer is an add-on software feature that gathers performance statistics at the LUN level. The analyzer data is stored in reserved space on the storage system. The reserved space for performance data cannot be increased nor decreased. Once the space is taken up, Analyzer will overwrite the performance data in a circular manner, meaning the oldest data will be deleted and replaced with the newer data.

The Unisphere Analyzer enabler must be installed on the storage system in order to use this feature. To verify that the enabler has been installed, check the software list in the **Storage System Properties** window in Unisphere Manager: **Unisphere Analyzer** should be listed and should show as **Active**.

14.5.1 Unisphere Analyzer Settings

Analyzer provides real-time performance statistics as well as archived performance statistics. The sampling intervals for real-time statistics and archived statistics are specified separately. The duration of the sampling period for the archived data is dependent on the archive interval. The longer the interval, the longer the sampling period.



The sampling intervals for Unisphere Analyzer can be set by opening up the **Performance Data Logging** window in Unisphere Manager. The default values for the sampling intervals are:

Real Time Interval: 60 Seconds
Archive Interval: 120

The duration of the archive sampling period will change if the archive interval is changed and applied. The smaller the interval, the shorter the sampling period.

Statistics logging needs to be enabled on the storage system in order for Unisphere Analyzer to gather performance data. This is done by checking the **Statistics Logging** checkbox under the **General** tab of the **Storage System Properties** window in Unisphere Manager. There is a performance impact when data logging is turned on. If the SP's are not heavily utilized then leave the statistics logging running. If however, you notice that the CPU utilization is abnormally high then you may want to have the stats logging enabled for a duration that gives you enough data points for troubleshooting and reporting purposes.

Analyzer's Advanced mode provides a larger range of statistics. Advanced mode can be enabled by using the following steps:

1. In Unisphere, click the System tab.
2. Click Monitoring and Alerts.
3. Click Statistics for Block.
4. Click Customize Charts.
5. Click the General tab.
6. Select the Advanced checkbox.
7. Click OK to apply the settings.

15 Host Connectivity and Configuration

A server connected to an EMC VNX storage array is called a host. Currently, CSBE has only certified host connectivity to the EMC VNX storage arrays via a fibre channel SAN.

Once fibre channel connectivity is established between the host and the storage system, the host must register its initiators (HBAs) with the storage system (see [Host Registration](#) for instructions).

Please refer to the [CSBE SAN Stack](#) Solutions Guide for the specific host platform for the certified hardware/software configurations for connectivity to the VNX. After the host successfully registers with the storage system, storage can be provisioned to the host. It is recommended to upgrade an existing host to the latest SAN stack where possible before continuing.

15.1 Symmetric Active/Active - Multicore RAID (VNX OE for Block 33)

Symmetric Active/Active allows hosts to access a Classic LUN simultaneously through both SPs. Since all paths are active, this removes the need to “trespass” LUNs between storage processors during certain path failures or in the event of an SP failure.

- Stripe locking service used for parallel access
- Write I/O operation acquires a lock on LBA address on both SPs (sent over CMI)
- Symmetric Active/Active is not yet supported for pool LUNs nor with Classic LUNs that use data services (such as Snapview or Clones).

For more information on Symmetric, please refer to *VNX MCx Multicore Everything – White Paper*.

Note: PowerPath 5.7 or higher is required to support ALUA state transition on some path(s) from Symmetric Active/Optimized to Active/Non-optimized

15.2 Host file system alignment

File system alignment is covered in detail in the [CSBE SAN Stack](#) Solutions Guide.

In general:

- Windows Server 2008, and later, automatically aligns, as do more recent Linux environments.
- VNX OE for File automatically aligns when it creates file systems.
- When provisioning LUNs for older Windows and Linux operating systems that use a 63-block header, the host file system needs to be aligned manually.

Alignment practices:

- Use host-based methods to align the file system.
- EMC recommends aligning the file system with a 1 MB offset.

15.3 Unisphere Host Agent

To eliminate the administrative overhead involved in managing host-side software, the Unisphere Host Agent is not to be installed on hosts connecting to the VNX.

The exception to this is the off-array management server which also functions as a central event monitoring server on which the installation of the host agent is required.

A packaged version is available in TPM. Please refer to the [CSBE VNX Block and CLARiON Supported Configuration Matrix](#) for the recommended version.

15.4 SAN Fabric and Zoning Guidelines

The following guidelines should be adhered when connecting a host to a VNX array.

- The VNX storage systems and hosts should be connected to two fabrics for redundancy.
- For high availability, hosts should have at least 2 front end port connections to each storage processor from 2 separate HBAs
- Balance front-end host port connections across UltraFlex I/O modules. The front end ports selected on each storage processor should be on separate IO modules where available.

Note: With Multicore RAID on VNX 2 models, host port connections effect starting CPU core assignment.

- First use the even numbered SP ports on each Ultra-Flex I/O module before using the odd numbered ports.
- For VNX8000, balance host port assignments across UltraFlex slots 0-5 and 6-10
- WWPN zoning is the **ONLY** supported zoning configuration, Port zoning is not recommended in Citi.
- A host should have at least 2 connections to the SAN, one connection to each fabric.
- One zone should be created for each host initiator (i.e. port).
- Each host initiator should be zoned to a port on SPA and a port on SPB.
- Avoid zoning every host port to every SP port

The following is an example of the 2 zones created for one host:

- Zone 1 (in Fabric A): HBA0 -> SPA0, SPB4
- Zone 2 (in Fabric B): HBA1 -> SPA4, SPB0

The zones should be created so that host traffic will be evenly distributed between all front-end ports on the storage array. For example, if connecting hosts to a VNX5500 with 8 front-end ports per SP connected to the SAN:

- Host #1 would be zoned to SP ports 0 and 4 on both SPs
- Host #2 would be zoned to SP ports 1 and 5 on both SPs
- Host #3 would be zoned to SP ports 2 and 6 on both SPs
- Host #4 would be zoned to SP ports 3 and 7 on both SPs

15.5 Host Registration

Application hosts will need to have their HBAs registered with the VNX storage array before any VNX storage can be allocated to them.

Host registration should be done manually via Unisphere Manager. Careful planning and adherence to operational procedures will be essential to ensure that proper connectivity is established between the application hosts and the VNX storage array for new host connectivity, HBA removal/addition and HBA replacement.

The exceptions to manual registration are

- VMware ESX hosts for which auto-registration is recommended (no agent required)
- Off-array management / central event monitoring server

15.5.1 Initiator Settings

Settings for the **Failover Mode** and **Array Commpath** must be setup correctly for host level multi-pathing (e.g. Powerpath). This setting must be applied for each host.

- The **Initiator Type** should always be set to **CLARiiON Open**.
- The **Array Commpath** should be enabled to allow hosts that do not have allocated LUNs to communicate with the array.
- The **Failover mode** option affects behavior required by different path management software packages on the host (0, 1, 2, 3, or 4).

The following table lists the failover modes for each Citi certified host platform.

Host OS	Failover Mode
ESXi 5.x / 4.x	4
Windows 2012 / 2008 / 2003	4
Windows 2000	1
RedHat 6.x / 5.x	4
AIX 7.x / 6.x / 5.x / VIO 2.x	4
Solaris 10 / 8	4
HP-UX 11.x	1

These settings are configured when registering the application hosts with the storage system. Please refer to the SAN stack for the platform for the appropriate version of PowerPath (as required).

15.5.2 ALUA

Failover mode 4 is specified for the Asymmetric Logical Unit Access (ALUA) feature. Not all host platforms support the new ALUA functionality. With the exception of ESXi and Windows 2012, ALUA functionality requires PowerPath to be installed and at a recent version. If the host platform has a PowerPath version prior to the minimum required for ALUA, it is recommended to upgrade to the current SAN stack for the platform (until which time failover mode 4 should not be used).

15.5.3 Manually Registering a New Host

Each HBA connected to the VNX storage array will need to register its information with the storage array. The hostname, host IP address, and failover settings are also specified during registration.

1. Ensure that the host's HBAs have a connection to the fabric and that they are zoned to both service processors on the array (see [SAN Fabric and Zoning Guidelines](#)).
2. In the **Connectivity Status** window in Unisphere, select one of the WWNs of the host that's being registered and click on the **Create** button
3. Set the **Initiator Type** to **CLARiiON Open** for host platforms other than HP-UX. HP-UX hosts should set the **Initiator Type** to **HP No Auto Trespass**.
4. Set the **Failover Mode** according to the table in the previous section.
5. Leave the default advanced settings
6. Set **Array CommPath** to **Enabled**.
7. The **Unit Serial Number** should always be set to **Array**.
8. If the WWNs being registered belong to a host that is already registered, select the **Existing Host** select box; otherwise, select the **New Host**.
9. If this is an existing host, select the hostname in the **Browse Hosts** window; otherwise, enter the **Hostname** and **IP Address** of the new host.
10. Register the rest of the host's HBAs under the hostname specified in the previous steps.

15.5.4 Replacing an HBA on a Host

Follow the steps below when replacing a host HBA:

1. Deregister the replaced HBA by doing the following:
 - a. Log into Unisphere Manager.
 - b. Go into Engineering mode by pressing Ctrl-Shift-F12 and entering the password.
 - c. In the Connectivity Status window, select the row with the initiator record of the replaced HBA (The icon for the initiator should be faded to show that it is no longer logged on) and hit the Deregister.
2. Register the replacement HBA by following the steps in the previous section to.

15.5.5 Removing an HBA


Follow the steps below when removing a host HBA:

1. Remove the HBA entry from the zone created for the host's VNX connectivity.
2. Deregister WWN of the removed HBA in the Connectivity Status window.

15.5.6 Verify Host Registration

To verify that a host is properly registered do the following:


1. The following screen shows a host named “gdse-ecc04” that is properly registered:



Note that the **Logged In** and **Registered** columns both have **Yes** for a host that is properly registered.

15.5.7 Changing the Failover Mode for a Host

In the case where the failover mode for a host was not set correctly, it can be modified by using the **Failover Setup Wizard** in Unisphere.



15.6 Allocating Storage to Hosts

A Storage Group specifies which LUNs hosts have access to. Hosts must be registered with the array and LUNs must be bound before they're able to get added to storage groups.

The following steps should be followed to allocate storage to a host:

1. Create a Storage Group - Storage Group names should be meaningful and should contain the name of the host or application that is accessing the LUNs contained in the Storage Group at a minimum.
2. Add LUNs to the Storage Group
3. Add the host to the Storage Group
4. The LUNs are now available for the host to discover.

16 Local Protection

SnapView is a storage-system-based software application that allows you to create a copy of a LUN by using either clones or snapshots. The implementation of SnapView requires at least two hosts: one production and one secondary host. Clones and snapshots can only be mounted by the secondary host.

16.1 SnapView Configurations

SnapView Supported Configurations VNX 2 (VNX OE for Block 33)

	VNX 5200	VNX 5400	VNX 5600	VNX 5800	VNX 7600	VNX 8000
Snapshots ²⁹ per storage system	256	256	512	512	1024	2048
Snapshots ²⁹ per source LUN	8	8	8	8	8	8
Snapview sessions per source LUN ²⁹	8	8	8	8	8	8
SnapView Snapshot source LUNs	256	256	256	256	512	512
SnapView Snapshot sessions per array	512	512	512	512	1024	2048
SnapView Snapshot source LUNs is a Consistent Start	32	32	32	32	64	64
Reserved LUNs ³⁰ per storage system	128	128	256	256	512	512
Source LUNs with Rollback active	300	300	300	300	300	300
Clones per storage system	256	256	1024	2048	2048	2048
Clones per source LUN	8	8	8	8	8	8
Clone source LUNs	128	128	512	1024	1024	1024
Clones in a Consistent Fracture	32	32	32	32	64	64
Clone private LUNs ^{31, 32} per storage system (required)	2	2	2	2	2	2

SnapView Supported Configurations VNX (VNX OE for Block 32 and earlier)

	VNX 5100	VNX 5300	VNX 5500	VNX 5700	VNX 7500
Snapshots ²⁹ per storage system	256	512	512	1024	2048
Snapshots ²⁹ per source LUN	8	8	8	8	8
Snapview sessions per source LUN ²⁹	8	8	8	8	8
Reserved LUNs ³⁰ per storage system	128	256	256	512	512
Clones per storage system	256	256	1024	2048	2048
Clones per source LUN	8	8	8	8	8
Clones per consistent fracture	32	32	32	64	64
Clone groups per storage system	128	128	512	1024	1024
Clone private LUNs ^{31, 32} per storage system (required)	2	2	2	2	2

²⁹ The limits for snapshots and sessions include SnapView snapshots or SnapView sessions as well as reserved snapshots or reserved sessions used in other applications SAN Copy incremental sessions, MirrorView /Asynchronous applications)

³⁰ A thin LUN cannot be used in the reserved LUN pool

³¹ Each clone private LUN must be at least 1GB

³² A thin or thick LUN may not be used for a clone private LUN

16.2 Snapshots

Snapshots allow access to point-in-time data of a source LUN. The cache for the snapshots is configured in an area called the Reserved LUN Pool. Only changed data is copied from the source LUN to the Reserved LUN Pool on a Copy-On-First-Write (COFW) basis. With COFW, data is copied the first time a data chunk is changed. Any subsequent changes to that same data chunk will not be copied.

16.2.1 Reserved LUN Pool

The Reserved LUN Pool (RLP) is used by SnapView to store the original data chunks from a source LUN that have changed since a snapshot session was activated. Each RLP can have session(s) from only one source LUN at a time.

The Reserved LUN Pool size should not be less than 20% of the total size of the disk space of the source LUNs that are being used for snapshots. The RLP is shared between the SPs. EMC's recommended LUN size for the LUNs in the Reserved LUN Pool is:

$$(\text{Estimated Reserved LUN Pool Size}) / (2 * \# \text{ of source LUNs on the SP})$$

The Snapshot Wizard in Unisphere simplifies the process of configuring the Reserve LUN Pool by calculating the space needed from the following data provided by the user:

- Estimated number of snapshots per source LUN.
- Estimated percentage of data change per source LUN.

Though using the wizard simplifies the process, it does not give the user the ability to specify the RAID Group for the RLP. The inability to specify the RAID Group for the RLP may result in snapshot space using up disk space that was intended to be used for primary storage. For this reason, the use of the Snapshot Configuration Wizard is not recommended.

For SnapView:

- Use SAS drives for reserve LUN pool (RLP) configurations, with write cache enabled LUNs.
- Match the secondary side RLP configuration to the primary side.
- Avoid configuring RLP on the same drives as the primary and secondary LUNs to avoid drive contention.
- Do not enable FAST Cache on RLP LUNs. RLP exhibits multiple small-block sequential streams that are not suited for FAST Cache.
- Do not enable FAST Cache on clone private LUNs (CPL). Optimal performance for the CPL I/O pattern is achieved by SP Cache algorithms.

16.2.2 Snapshot Sessions

A snapshot session preserves the data of the source LUN at the point-in-time that the session was started. After the session is started, the original data from modified data chunks gets copied to the Reserved LUN Pool.

Every snapshot session is started in persistent mode. This mode cannot be disabled on VNX storage arrays. Persistent mode allows a session to withstand the following failures and LUN trespasses:

- SP Reboot or failure
- Array reboot or power failure
- Host I/O trespass to the peer SP
- Allows for rollback to the point-in-time the session was started

When the consistent snapshot option is selected, writes going to the specified LUNs in the SnapView session are buffered until the process of starting the session is completed. This will capture the state of all the LUNs in the SnapView session at the exact point-in-time the session was initiated. It is still recommended that the application is quiesced on the host-side before starting a SnapView session to provide a restartable snapshot of the application.

Session name should follow the standard **servername_session#**

16.2.3 Snapshot Configuration

- Up to 8 snapshots can be created per source LUN and only one snapshot can be active per session of a source LUN.
- Snapshot names should follow the standard **servername_snap_lun#**.
- Snapshots can only be mounted on a host other than the host that owns the source LUN.

The following steps should be followed to configure snapshots:

1. Create the Reserved LUN Pool – Should equal about 20% (assuming 20% data change) of the total capacity of all source LUNs. The size of the LUNs in the reserved LUN pool can be calculated using the following formula:

$$(\text{Estimated Reserved LUN Pool Size}) / (2 * \# \text{ source LUNs on SP})$$
2. Start a SnapView session - The **Consistent** option can be chosen if the session contains more than one LUN to ensure consistency between all the LUNs.
3. Create a snapshot for the LUN – This will be the LUN that will be allocated to the Storage Group of the secondary host. The Storage Group can be specified at snapshot creation or at a later time.
4. Activate the snapshot on a session - The snapshot can be activated on a single session. Multiple sessions can be started on the source LUN but only one snapshot can be activated on that session.
5. Assign the snapshot to a Storage Group – The snapshot can now be assigned to the Storage Group of the secondary host which will make it accessible to the secondary host.

Space requirements will double per server in clone configuration.

Snapshot cache size will vary based on use and may cause snap sessions to end due to under-sizing of the Reserved LUN Pool.

16.3 Clones

A clone is a complete copy of a source LUN. A source LUN would need to be part of a Clone Group in order to create a clone. The software assigns a clone ID to each clone added to the clone group. The ID remains until the clone is removed from its group.

16.3.1 Clone Guidelines

- A **Clone Private LUN (CPL)** must be allocated to each SP before clones can be created. Each CPL should be 1GB in size. The CPL is used to record information about changed data on the source LUN and the clone LUN once the clone is fractured. The information contained in the CPL will then be used at the time of synchronization or reverse synchronization. One CPL should be assigned to each SP.
- **Protected Restore** feature should always be enabled before reverse synching a clone. This will ensure that writes to the source LUN will not be written to the clone LUN during a reverse synchronization. Be sure to enable **Protected Restore in Clone Feature Properties** for the storage array.
- The recommended storage group naming for clone LUNs is: **servername_clone**³³
- A clone LUN created on lower performing drives than the source LUN's drives will affect performance on the source LUN. Therefore, the clone LUN should be created on a drive of similar type and capacity.

16.3.2 Clone States


There are 6 possible clone states see table below:

Clone State	Description	Cause of State	Permitted Operations
Reverse Synchronizing	A clone is in the process of copying its data to its source LUN	• Manually administered with Navisphere or CLI	• Fracture
Reverse Out-of-Sync	A failure during a reverse synchronization process which makes the source unusable	• A reverse synchronization operation failed to complete successfully.	• Reverse synchronize • Remove • Fracture
Synchronized	A clone is a byte-for-byte copy of its source.	• After synchronization • A clone is added to Clone Group without setting "Initial Sync Required"	• Fracture • Remove
Synchronizing	A clone is in the process of copying data from its source LUN	• Automatically started when you add a clone to a clone Group with the Initial Sync Required properly selected.	• Fracture
Out-of-Sync	A clone was in the process of synchronizing but failed. Therefore clone is unusable.	• A synchronization operation failed to completed successfully	• Synchronize • Remove • Fracture
Consistent	A clone is in a synchronized state that receives host I/O from the source.	• A clone is fractured while in a Synchronized state. • A clone is unfractured and has yet to transition reached a Synchronized state.	• Fracture • Remove


³³ The LUN name is used for display via the GUI and cannot be used as a CLI argument.

16.3.3 Unisphere Clone Wizard

Unisphere has a Clone Wizard which simplifies the process of creating clones.



If CPLs have not been created before using the wizard, the CPLs will be bound when the first clone is created using the wizard. The following screen provides the actions that are taken when creating clones using the Clone Wizard.



After the clones have been created using the wizard, they will be synchronized. Once synchronized, the clones can be fractured and be made available to a secondary server.

17 Off-Array Management Server

It is recommended to manage and monitor the storage arrays through the management software installed on the Off-Array Management Server.

Unisphere Manager is the main software used to manage the EMC VNX storage array. It connects to the EMC VNX array using the IP address of either SP.

The Unisphere GUI is accessible via:

- Local UI on the storage array, using Internet browsers (e.g. Microsoft Internet Explorer)
- Stand-alone Unisphere Client package (discussed below)

The advantages of using the stand-alone Unisphere Client compared to the local UI on storage array includes:

- Less management overhead incurred on the storage array
- No reboot required for Unisphere Client updates (storage UI can only be updated through a VNX OE upgrade which requires reboot)
- Faster startup time

17.1 Server Requirements

The off-array management server's requirements are as follows:

	Virtual (Recommended)	Physical (Exception)
Operating System	Windows 2008 R2 ³⁴	Windows 2008 R2
Hosting Server CPU	1 x vCPU (2800 MHz)	2.8 GHz Dual Core
Hosting Server Memory Size	2 GB	2 GB

Both physical and virtual instances of the management server are supported. Where it is determined that a physical installation is required, please refer to the [Wintel Engineering site](#) for the recommended enterprise-class rack mounted Intel server.

TCP/IP network connectivity is required to all VNX Storage Arrays and to each system or host managed by the management server.

Terminal Services should be configured on the off-array management server for Storage Administrators to access the management software.

³⁴ Note: Windows 2008 with ESX 4.0 is not supported.

17.2 Off-Array Management Station Software Installation

There are 3 main management software packages that can be obtained from TPM for installation on the management server.

Software	TPM package
Unisphere Client	CITI_WIN_<OS ver>_UnisphereClient_<ver>_<TPM ver>
Unisphere Host Agent	CITI_WIN_<OS ver>_UnisphereHostAgent_<ver>_<TPM ver>
Navisphere CLI	CITI_WIN_<OS ver>_NavisphereCLI_<ver>_<TPM ver>

Please refer to the [CSBE VNX Block and CLARiiON Supported Configuration Matrix](#) for the recommended versions.

If the management server is also used for managing existing CLARiiON environments, then the software for CLARiiON (Navisphere Host Agent and Navisphere Client GUI) can remain on the server for management purposes.

Note: Always upgrade the software on the off-array management server first before upgrading the software on the storage array. New versions of management software are typically released to support new array microcode.

Application installations are via TPM. Please contact your local Wintel SA team for installation. The installation procedures have been provided to the Wintel SA team but are also provided below for reference.

17.2.1 Unisphere Client Installation Procedure

Prerequisites

1. JRE 1.7.0_update 60 or higher
2. Internet Explorer 8.0

Installation Steps

1. Deploy the TPM package for Unisphere Client
CITI_WIN_<OS ver>_UnisphereClient_<ver>_<TPM ver>

Please refer to the [CSBE CLARiiON Block and CLARiiON Supported Configuration Matrix](#) for the recommended version.

2. At the **Introduction** page, click **Next**
3. At the **Choose Install Folder** page, use the default path and Click **Next**
4. At the **Pre-Installation Summary** page, click **Install** to start the installation process
5. Wait for the installation to complete.

17.2.2 Unisphere Host Agent Installation Procedure

The Unisphere Host Agent is required on all off-array management servers to facilitate central event monitoring.

Prerequisites


If this is an upgrade, make a backup of the agent.config and navimon.cfg files located in the installation directory. (Default directory: C:\Program Files\EMC\HostAgent)

Installation Steps

1. Deploy the TPM package for Unisphere Host Agent
CITI_WIN_<OS ver>_UnisphereHostAgent_<ver>_<TPM ver>

Please refer to the [CSBE VNX Block and CLARiiON Supported Configuration Matrix](#) for the recommended version.

2. At the “Introduction” page, click “Next”
3. At the “Choose Install Folder” page, please enter the folder path to install, or use the default path. Click “Next”



4. At the “Unisphere Host Agent Configuration” page,
 - If this is an upgrade, copy the agent.config backup file into the installation directory and choose “Preserve Unisphere Host Agent settings”
 - If this is a fresh installation, choose “Create a new agent.config file and add entries”.

You can add the list of IP addresses now or manually update the configuration file later

Click “Next”

EMC Unisphere Host Agent 1.3.3.1.0072

Unisphere Host Agent Configuration (Step 3 of 5)

☐ Preserve Unisphere Host Agent settings.
☐ Add entries to an existing agent.config file.
☒ Create a new agent.config file and add entries.

List of attached system IP with privileged user
 Enter an IP address below for each storage processor (SP) authorized to run privileged commands remotely (e.g. updating agent config) on the local agent.

(C) 2014 EMC Corporation All Rights Reserved. Back Next Cancel

5. At the “Pre-Installation Summary” page, click “Install” to start the installation process

EMC Unisphere Host Agent 1.3.3.1.0072

Pre-Installation Summary (Step 4 of 5)

Review the following settings before proceeding.

Product Install Path:
C:\Program Files (x86)\EMC\Unisphere Host Agent

Product Install Version:
1.3.3.1.0072

Preserve Settings:
FALSE

Edit Existing Settings:
FALSE

Overwrite Settings:
TRUE

(C) 2014 EMC Corporation All Rights Reserved. Back Install Cancel

6. Wait for the installation to complete

17.2.3 Navisphere CLI Installation Procedure

Installation Steps


1. Deploy the TPM package for Navisphere CLI
CITI_WIN_<OS ver>_NavisphereCLI_<ver>_<TPM ver>

Please refer to the [CSBE VNX Block and CLARiiON Supported Configuration Matrix](#) for the recommended version.

2. At the “Introduction” page, click “Next”
3. At the “Choose Install Folder” page, use the default path.

Check the box for “Include Navisphere CLI in system environment path”. Click “Next”


4. At the “Navisphere CLI Configuration” page, choose “Do not create security credentials file”
Click “Next”



5. At the “Pre-Installation Summary” page, click “Install” to start the installation process
6. Wait for the installation to complete

17.2.4 Upgrade Procedure

1. Copy the agent.config and navimon.cfg files to a safe location as a backup copy.
(Default directory: C:\Program Files\EMC\HostAgent\)
2. Export all monitoring templates to a backup directory on the management server:
System > Monitoring and Alerts > Notifications for Block



3. Uninstall all Navisphere software that is installed on the local server.
4. Go through the installation of the new software by following the steps in the previous section.
5. Stop the Navisphere Agent service.
6. Copy the backup files into the installed directory.
7. Restart the Navisphere Agent service in Computer Management.
8. Import the templates that were backed up in Step 2

17.3 Modifying agent.config File

The agent.conf file specifies which users and systems can communicate with the Unisphere Host Agent.

It is located in the *<install directory>\EMC\HostAgent* directory.

An entry in the format **user system@<sp IP address>** should be entered in the agent.config file for every SP that the management server manages.

The following is an example of what the agent.config file may look like:

```
clarDescr
clarContact
poll 60
baud 9600
eventlog 2048
user system@10.40.42.207
user system@10.40.42.208
user system@10.40.53.52
user system@10.40.53.53
user system@10.40.43.115
user system@10.40.43.116
user system@10.40.53.27
user system@10.40.53.28
```

The “poll” interval should be set to 60 seconds to improve the Unisphere GUI refresh rate.

Note: You might see “Agent denied request - caller not privileged” error message when you try to send a test event for Centralised Monitoring. Make sure that the IP addresses entered is correct, “user” should be in lower case and always finish the last entry with an “enter” before saving the file.

17.4 Navisphere CLI Configuration

This section covers the configuration of Navisphere CLI (or naviseccli) to improve the ease of use of the CLI.

17.4.1 Creating the User's Security File

The security file is an encrypted file that resides on the system from which the naviseccli command is to be run. This file is created on the local server by using the **naviseccli –addusersecurity** command. It contains the user's NT login credentials as well as the user's associated Unisphere account information.

Ensure that the user has an existing Unisphere account before creating the security file. It is recommended that every Unisphere user (who will use naviseccli) create a security file. Creating the security file allows the user to issue naviseccli commands without having to enter their user credentials along with the command every time they issue the command. The user will be able to issue naviseccli commands that correspond to their Unisphere role (i.e. Security Administrator, Administrator, SAN Administrator, or Operator).

Note: Each user can only use the security file that they created; they cannot use a security file created under another user's profile. The command will fail if the user attempts to use a security file that was created under another user's profile.

Note: DO NOT use the **–password** parameter in the naviseccli command because the password will be displayed on the screen in clear text. Instead, enter the command with only the **–user** parameter and the user will be prompted for a password which is masked when typed in at the prompt. This also applies to scripting naviseccli commands so that the password will not be compromised. If naviseccli commands are to be scripted, ensure that the script will run under the correct user and that the security file is used to execute the command.

To create a security file, do the following:

1. Log into the management server using your SOE ID.
2. Issue the following naviseccli command on the command prompt (DO NOT type in the **–password** parameter with this command) with the username and scope (0=global, 1=local, 2=ldap):

```
naviseccli –addusersecurity –scope <0|1|2> -user <user name>
```

Note: Most users will be using the scope of 2 since Unisphere should be configured to authenticate users with Citi's Active Directory.

3. Enter the password at the password prompt.
4. The file named **SecuredCLISecurityFile.txt** is created in the user's home directory.

Note: If the security file was created with an LDAP user, any changes to that user account will require a new security file to be created. To create a new security file, simply repeat the steps above.

18 Setting Up the Storage System

This section contains the configuration guidelines for setting up the supported models of the VNX storage system.

18.1 Physical Setup

Item	Recommendation
Rack mount	Determined by Hardware Group
Power	Determined by Hardware Group
Cooling	Determined by Hardware Group
Enclosure Addresses	Reference the Hardware Reference Manual for the array model. This address is typically set by the EMC certified engineer at the time of installation.
Cabling	Reference the Hardware Reference Manual for the array model. This is typically done by the EMC certified engineer at the time of installation.
Switch Connections	Determined by design on growing basis. Reference the CSBE Cisco MDS Fibre Channel Switch guidelines for more information.
North America	No feedback

18.2 Upgrading Array Microcode

When upgrading VNX OE for Block code on the storage arrays, ensure that the management software running on the off-array management station is running the same or higher version of software required for the version of VNX OE for Block that is being upgraded to.

Off-array management servers cannot manage arrays running higher levels of code. Always upgrade the software on the off-array management server first before upgrading the software on the storage array (See [Off-Array Management Station Software Installation](#) section for instructions on upgrading software on the off-array management station).

Note: Before beginning any upgrade, ensure that the system has been reviewed by EMC to identify any known issues that may occur during the upgrade process.

As an example, see Primus emc308955

ETA emc308955: VNX Series: SMLink_check tool should be run to check condition of Virtual Provision (VP) Pools on VNX OE 05.31 prior to upgrade to Release 05.32.

The SMLink_check tool is designed to be run by customers to verify if the array in question may be affected by Thin or Thick LUNS going offline after upgrading from VNX OE Release 05.31 to VNX OE Release 05.32. The tool basically checks if pool LUNS were created before Release 31 P502 (05.31.5.000.502) and if the FAST VP feature is installed on the array. The tool indicates whether further investigation is required prior to the NDU

18.3 Array Software

Please refer to the [Software and Features](#) section for recommendations relating to Software features.

18.4 DPE/DAE - OS Configuration / Initialization

Item	Recommendation
Initialization	EMC initializes the array.
Vault / System Space	- The first four drives (Disk IDs 000-003) on the storage system are reserved to store the configuration database for the storage system as well as vault space for data that is flushed from write cache. - NL-SAS drives are not recommended to be used for the vault drives. - Whilst larger SAS drives are supported for vault drives, consideration should be given to using smaller drives to reduce cost. - Although not all the space on the vault drives is reserved for the storage system, it is not recommended to bind LUNs on the remaining storage space on the vault drives.
Physical Disks	Maximum number of disk drives per DAE is 15 or 25 (depending on DAE model). See Disk Drive Types and Drive Selection sections for supported disk drives and recommended configurations.
Settings	See Storage System Properties section
Compatibility	The VNX supports the following two different DAE models. 2U DAE that can house 25 x 2.5" disk drives 3U DAE that can house 15 x 3.5" disk drives. See VNX Disk Array Enclosures section.

18.5 Storage System Properties

18.5.1 General Settings

The following screen shows the properties window of the VNX storage array in Unisphere in which the storage system settings can be modified.

SP	IP Addresses	DNS Hostname
SP B	10.40.42.208	EMC_NAS_RESERVE06.nam.nsroot.net
SP A	10.40.42.207	EMC_NAS_RESERVE05.nam.nsroot.net

18.5.1.1 Array Name

A meaningful name should be given to the storage system and its Service Processors. For example VNX5500-1868-A.

This ensures that all log entries identify the array rather than just “SPA” or “SPB”.

Please contact the global SAN Support team to confirm current global naming standard.

Note: Changing the SP name will require a reboot of the SP, so should be planned accordingly.

18.5.1.2 Performance Statistics Logging

Statistics logging needs to be enabled on the storage system in order for Unisphere Analyzer to gather performance data.

For further information refer to Section [Performance Monitoring using Unisphere Analyzer](#).

18.5.2 SP Cache

Please refer to section [SP Cache](#) for recommendations relating to the configuration of SP cache.

18.5.3 FAST Cache

Please refer to section [FAST Cache](#) for recommendations relating to the use of FAST cache.

18.5.4 Software

Please refer to section [Software and Features](#) for recommendations relating to Software features.


18.5.5 Time Settings

The storage system's time is always set to UTC at installation. When the time is displayed in Unisphere from a client's browser, the system's Java Runtime Environment is used to gather the client's time zone information and the converted time is displayed. The VNX can be configured to synchronize its time clock with a Network Time Protocol (NTP) server. NTP should be enabled for all VNXs.

18.5.5.1 Configuring NTP

NTP settings for the domain members can be configured by a user with the role of Security Administrator or Administrator.

1. Select **All Systems** in the top left drop-down menu in Unisphere and select the **Domains** tab.
2. Select **Configure NTP** in the right pane and **Enable NTP**, set the **Time Sync Interval**, and enter the IP addresses of the NTP servers:



3. Check the SP's event log to verify that the NTP time synchronization was successful.

18.6 Service Processor (SP) Properties

18.6.1 SP Network Settings

SP Network Name

A meaningful name should be given to the storage system and its Service Processors.

Set the network name to each SP to a meaningful name that will help identify the array it belongs to. This ensures that all log entries identify the array rather than just "SPA" or "SPB". Changing the name of an SP will require it to be rebooted.

Please contact the global SAN Support team to confirm current global naming standard.

Management Port Settings

Set the Requested Settings to Auto

SNMP Settings

Enable Processing of SNMP MIB read requests. The SNMP Community string set appropriately (should not be left blank). Please refer to section [Alerting and Monitoring](#) for details on configuring SNMP monitoring.

18.6.2 ALPA

The ALPA address only applies in an arbitrated loop configuration. It does not serve a purpose if the storage array is attached to a switched fabric. Since CSBE supports switched configurations only, ALPA settings do not need to be modified.

18.6.3 Agent

Set the Contact Person and Description as appropriate to the regional Operational standards. Privileged Users should not be configured.

18.6.4 Fibre Speed

The SP port speeds can be set to auto-negotiate. The switch settings should be set for auto-negotiation.

19 Appendix A: iSCSI for VNX

This section describes the use of iSCSI as alternative host connectivity under certain restrictions. It is noted that this connectivity can be used only on the newer Arista Networks and only certain datacenters that have Storage IP Fabrics implemented. Therefore, the majority of hosts cannot take advantage of this feature.

19.1 Restrictions for iSCSI

- Only larger VNX2 models will be allowed due to increased back-end efficiency and additional CPU/memory. See [VNX Array Configurations for iSCSI](#) for more information.
- The VNX must have at least 2x 10GigE (Fiber) per SP directly connected to Storage Server Leaf switches (Arista).
- Only the array configurations listed in this section are allowed.
- Only (Windows 2012, RHEL 6.6+, ESX 5.5) hosts with at least 2x 10GigE with connectivity to Storage Client Leaf switches (Arista) can use iSCSI. This is initially limited to T++ pods.
- Network ACLs must be in place to limit access to the iSCSI ports to a known set of hosts.
- All LUNs are traditional thick devices.
- The Pool for iSCSI must be comprised of R10 SAS 10K (or higher grade of disk).
- The Pool for iSCSI must be dedicated.
- VNX arrays with VMware clients will be dedicated for VMWare

Additional restrictions and limitations may be found in the [iSCSI Storage Design Document](#).

Any exceptions to the above restrictions must be approved by CATE Storage and Backup Engineering.

19.2 VNX Array Configurations for iSCSI

The arrays used for iSCSI are purpose built and should be dedicate for this role. However, traditional (thick LUN) block access from the same pool is allowed with CSBE approval. The following are the only supported configurations. The higher model number is required to accommodate the workload (increased memory and CPU).

19.2.1 100TB Useable Configuration

VNX2 7600 Block Version with Optical 10GbE iSCSI (Hardware Components)

Qty	EMC Part Number	Description
1	VNXB76DP25F	VNX7600 DPE 25X2.5" DRIVE SLOTS-FLD IN
17	VNXB6GSDAE25F	VNXB 25X2.5 6G SAS EXP DAE-FIELD INST
1	VNXB6GSDAE25W	VNXB 25X2.5 6G SAS PRI DAE-FIELD INST
408	V4-2S10-600	VNX 600GB 10K SAS 25X2.5 DPE/DAE
1	V-V4-260010	VNX 600GB 10K VAULT 25X2.5 DPE/DAE
43	FLV42S6F-200	VNX 200GB FAST CACHE 25X2.5 DPE/DAE
20	C13-PWR-12	2 C13 CORDS NEMA 5-15 125V 10A - NON DPE

1	C13-250V-ULCSA	250V PWR CRD C13 TO 6-15P UL/CSA
1	VNX76-KIT	VNX7600 Documentation Kit=IC
2	VSPBMXGI2A	VNXB 2 PT OP 10G ISCSI IO MOD PR-SFP

Note: Other configurations may be added after a pilot has been implemented. Changes to the above configuration may also occur.

Note: Any deviations from the above configurations must be approved by CATE Storage and Backup Engineering.

19.3 Standard Array Setup

There will be the normal 100Mb or 1GigE management interfaces to the SPs that are independent of the new iSCSI ports. The array will be configured as a normal VNX block array for DNS, NTP, and AD access. Audit logging, SRM, and any other security aspects will be configured as normal. The differences in setup are listed in the following sections.

19.4 Pool Configuration for iSCSI

The storage pool to be used for iSCSI will need to be configured as RAID 10. This is due to the higher performance criteria (currently Silver Tier 450 IOPS/uTB OLTP2HW) set by the application classes that will use the storage.

1. Configure the storage pool as RAID 10.
2. Create only traditional thick LUNs in the pool to be used by the iSCSI hosts.

19.5 How to setup iSCSI on the VNX

19.5.1 iSCSI Port Configuration

1. Go to Unisphere → Settings → Network → Settings for Block

The screenshot shows the EMC Unisphere web interface, specifically the 'Settings for Block' page under 'Network'. A table lists various ports and their configurations:

Physical Location	SP-Port	Type	Speed	IP Addresses	IQN/WWN
Slot A0, Port 2	A-Bus 2	SAS	6Gbps	N/A	N/A
Slot A0, Port 1	A-Bus 1	SAS	6Gbps	N/A	N/A
Slot A0, Port 0	A-Bus 0	SAS	6Gbps	N/A	N/A
Slot B4, Port 3	B-7	Fibre	8Gbps	N/A	50:06:01:60:C7:20:33:3E:50:06:01:6F:4...
Slot B4, Port 2	B-6	Fibre	8Gbps	N/A	50:06:01:60:C7:20:33:3E:50:06:01:6E:4...
Slot B4, Port 1	B-5	Fibre	8Gbps	N/A	50:06:01:60:C7:20:33:3E:50:06:01:6D:4...
Slot B4, Port 0	B-4	Fibre	8Gbps	N/A	50:06:01:60:C7:20:33:3E:50:06:01:6C:4...
Slot B1, Port 3	B-3	Fibre	8Gbps	N/A	50:06:01:60:C7:20:33:3E:50:06:01:6B:4...
Slot B1, Port 2	B-2	Fibre	8Gbps	N/A	50:06:01:60:C7:20:33:3E:50:06:01:6A:4...
Slot B1, Port 1	B-1	Fibre	8Gbps	N/A	50:06:01:60:C7:20:33:3E:50:06:01:69:4...
Slot B1, Port 0	B-0 (MirrorView)	Fibre	8Gbps	N/A	50:06:01:60:C7:20:33:3E:50:06:01:68:4...
Slot A4, Port 3	A-7	Fibre	8Gbps	N/A	50:06:01:60:C7:20:33:3E:50:06:01:67:4...
Slot A4, Port 2	A-6	Fibre	8Gbps	N/A	50:06:01:60:C7:20:33:3E:50:06:01:66:4...
Slot A4, Port 1	A-5	Fibre	8Gbps	N/A	50:06:01:60:C7:20:33:3E:50:06:01:65:4...
Slot A4, Port 0	A-4	Fibre	8Gbps	N/A	50:06:01:60:C7:20:33:3E:50:06:01:64:4...
Slot A1, Port 3	A-3	Fibre	8Gbps	N/A	50:06:01:60:C7:20:33:3E:50:06:01:63:4...
Slot A1, Port 2	A-2	Fibre	8Gbps	N/A	50:06:01:60:C7:20:33:3E:50:06:01:62:4...
Slot A1, Port 1	A-1	Fibre	8Gbps	N/A	50:06:01:60:C7:20:33:3E:50:06:01:61:4...
Slot A1, Port 0	A-0 (MirrorView)	Fibre	8Gbps	N/A	50:06:01:60:C7:20:33:3E:50:06:01:60:4...
Slot B2, Port 0	B-8 (MirrorView)	iSCSI	10Gbps	10.40.89.135	iqn.1992-04.com.emc:cx.apm001216007...
Slot A2, Port 0	A-8 (MirrorView)	iSCSI	10Gbps	10.40.89.136	iqn.1992-04.com.emc:cx.apm001216007...
Slot B2, Port 1	B-9	iSCSI	N/A	N/A	iqn.1992-04.com.emc:cx.apm001216007...
Slot A2, Port 1	A-9	iSCSI	N/A	N/A	iqn.1992-04.com.emc:cx.apm001216007...

The interface also shows a sidebar with 'Wizards' (Edit Data Move, Interface Wizard, Device Wizard, Route Wizard), 'Security Settings' (Manage Idle Ti, Manage LDAP), 'Network Settings' (Edit Network S, Edit Network S, Control Station, Configure DNS, Ping - SPA, Ping - SPB, Ping - Data Mo, Trace Route, Trace Route), 'More Settings' (Manage License), and 'iSCSI Management' (CHAP Management, iSNS). The bottom status bar indicates 724 Alerts, 717 Critical, 3 Certificates, and 2 Unaccepted. The user is identified as g05529.

You should see the 4x iSCSI ports listed. Un-configured ports will not have a speed or IP assigned. Select one of the un-configured ports and then select the *Properties* button.

2. On the properties screen the speed and MTU will be greyed out until you create a virtual port.

c5-vnx5700-cs0 - SP A: Port 1 [A-9] - iSCSI Port Properties

Port Identification

IQN: iqn.1992-04.com.emc:cx.apm00121600765.a9

Alias: 0765.a9

Physical Port Properties

Mac Address: 00-60-16-46-B1-F1

Actual Speed (Mbps): N/A

Selected Speed (Mbps): 10000

MTU (bytes): 1260

Virtual Port Properties

Virtual Port	VLAN ID	IP Addresses
--------------	---------	--------------

Properties Add Delete

Connector

EMC Part Number: 019-078-041

EMC Serial Number: 00000000000

Vendor Part Number: AFBR-703ASDZ-E2

Vendor Serial Number: AGL1343D1000810

OK Apply Cancel Help

3. Select the Add button under Virtual Port Properties to add an IP to the iSCSI port.

IQN: iqn.1992-04.com.emc:cx.apm00121600765.a9

Virtual Port

IPv4 Configuration

IP Address: . . .

Gateway: . . .

Subnet Mask: . . .

VLAN Configuration

☐ Enable VLAN Tagging


VLAN ID: (1-4094)

Authentication

☐ Require initiator authentication

OK Cancel Help

4. Fill in the IP Address, Gateway, and subnet mask per assigned network IP. VLAN tagging is not used on the Storage IP Fabric. Keep the *Require initiator authentication* unchecked (this is for CHAP and is not used at Citi).
5. Select OK to make the changes effective. You will also have to acknowledge that you are working with Network staff. This change may take several minutes to complete.



The previous screen will not reflect the new information and will need to be closed and reopened to see the new data.

6. The MTU value will initially be set to 1260 or 1500. This should be changed to 8000. Hosts may be configured for an MTU of 8800, but it will be negotiated down to 8000 as configured here.

Port Identification

IQN: iqn.1992-04.com.emc:cx.apm00121600765.a9

Alias: 0765.a9

Physical Port Properties

Mac Address: 00-60-16-46-B1-F1

Actual Speed (Mbps): 10000

Selected Speed (Mbps): 10000

MTU (bytes): 8000

Virtual Port Properties

Virtual Port	VLAN ID	IP Addresses
Virtual Port 0		10.40.89.253

Properties Add Delete

Connector

EMC Part Number: 019-078-041

EMC Serial Number: 00000000000

Vendor Part Number: AFBR-703ASDZ-E2

Vendor Serial Number: AGL1343D1000810

OK Apply Cancel Help

7. Select Apply to make the changes effective. This change may take several minutes to complete after you confirm you are working a network resource.

Make note of the IQN number for the port and the IP. This will need to be relayed to any hosts.

8. Repeat these steps for any additional ports (2 per SP, 4x total). The first port for each SP should be on the same subnet (Storage Server subnet A) and the second port for each SP should be on a different subnet (Storage Server subnet B).

19.5.2 iSCSI CHAP Security

CHAP – Challenge Handshake Authentication Protocol. CHAP security consists of a user name and password (secret). You can configure and enable CHAP security for initiators and for targets. The CHAP protocol requires initiator authentication. Target authentication is optional.

CHAP is not in use at Citi. All access security is handled by the IQN identifier and network ACLs on the Storage Leaf switches for iSCSI ports.

19.5.3 iSCSI ISNS Service

ISNS Internet Storage Naming Service. This service prevents misconfiguration of IQNs with typos.


ISNS is not in use at Citi. All access security is handled by the network ACLs on the Storage Leaf switches for iSCSI ports.

19.6 How to register iSCSI Hosts


Similar to fiber Channel hosts, iSCSI hosts are organized into host storage groups. These host storage groups are registered to an iSCSI port and a host IQN. This will be a manual registration process.

Create a Storage Group


1. Go to Unisphere → Hosts → Storage Groups




2. Select the **Create** button.



then select **OK**.




Select **Yes** on confirmation window.




Select No. The host will need to be registered from a different screen.

3. Go to Hosts → Initiators




Select *Create* to manually register a host.



- WWN/IQN: Use the iqn from the host. Most hosts will only have a single iSCSI initiator.
- SP – port: Select the first iSCSI port from SPA
- Initiator Type: CLARiiON/VNX

- Failover Mode: Active-Active mode (ALUA) – failover mode 4 (Older hosts may not be compatible with this mode)
- New Host button selected, fill in host name and IP in open fields below button. You will need to fill out both hostname and IP on the first entry to get the OK button to be selectable. The IP address is not important, as long as there is an entry here (for example, most hosts will have multiple IPs due to having multiple NICs).



Select Yes on the confirmation screen.

4. Repeat the entries above by selecting Create again, but use the remaining iSCSI ports each time (you will have to select new host each time with same info if host does not show up in existing list):
 - 1st iSCSI port SPA
 - 1st iSCSI port SPB
 - Optional:
 - 2nd iSCSI port SPA
 - 2nd iSCSI port SPB

There should be a total of 2 (or 4) entries for each host IQN. Two ports should be sufficient performance for most applications and is the default standard.


If using 2 ports per host, then some manual balancing of the ports should be followed. The second set of iSCSI ports should be used on the next host (1st host use 1st ports, 2nd host use 2nd ports, 3rd host use 1st ports, 4th host use second ports, ...)

If using 4 ports per host, then balancing between the ports is not needed.


19.7 How to allocate LUNs to iSCSI Hosts

LUNs are allocated to host storage groups identically to how LUNs are allocated to fiber channel host groups. *You may not be able to assign a host to the Storage Group until you have allocated LUNs to the Storage Group.*


1. Go to Unisphere → Hosts → Storage Groups. Select the storage group created in the previous section. Then select the Connect LUNs button.



- Expand the LUN list under SPA. Select (highlight by clicking and holding CTRL key) the LUNs you want to allocate to the host/server/cluster. Expand the LUN list on SPB. Add the LUNs you want to allocate to the host (with CTRL key).



- Select the Add button for each LUN or group of LUNs you need to add to the Storage Group. Select the OK or Apply button. Confirm the change (Yes).
- Select the Storage Group again and then select the Connect Hosts button. The host you created in the previous section (by manually registering the host) should now be on the available hosts list. Move it to the right side (Hosts to be Connected) and then select OK or Apply.



Confirm the change (Yes).

5. Hosts should now be able to connect to the iSCSI ports and see the allocated LUNs.

19.8 Host Configurations for iSCSI (SAN Stacks)

The following physical Host OS configurations are in scope for support with VNX iSCSI storage:

- Windows 2012 Server R2
- RHEL 6 (update 6 and higher)
- ESX 5.5 (update 2 and higher)

The following OS level clustering solutions are currently out of scope for SAN attachment to VNX iSCSI:

- All CATE certified versions of Veritas for Red Hat Linux
- All CATE certified versions of Microsoft Failover Cluster

The current versions of the SAN Connectivity stacks for each OS version in scope requires configuration and software changes to accommodate reliable iSCSI connectivity. A SAN stack for RHEL 6 (excluding Veritas support) will be released first followed by ESX 5.5 and Windows 2012 R2. Documentation will be released and updated on the following CSBE SharePoint site:

<https://catecollaboration.citigroup.net/domains/platstor/gdse/products/services/default.aspx>

As this is a new technology with in Citi, you should expect frequent updates to the host SAN stack documentation on that site. Please refer directly to the documents on the website and do NOT save your own local copy.

19.9 Host Data Migration

There may be a need for migrating data to/from iSCSI LUNs depending on the nature of the applications and any history. The following table summarizes the options in this area:

Product	Migration Method	Migrating Group	Features/Functions	Considerations
EMC Open Replicator	Array-based solution (SAN)	Storage Administrator	Requires both source and target array to be on the same network (FC)	Depending on location in the DC, the target VNX would need to be connected to a FC SAN to allow for migration to take place over FC. Host connected to VNX over iSCSI would be able to access data over iSCSI connectivity
Hitachi UVM	Array-based solution (SAN)	Storage Administrator	Requires both source and target array to be on the same network (FC). Also requires HDS hardware.	Although this approach is similar to open replicator, it requires the introduction of an additional component
Volume Manager	Host level (IP)	Systems Administrator	Provides flexibility in LUN sizes during migration	Existing production servers are on the Cisco Network switches. Routing between Cisco and Arista network would be required to allow migration to occur since iSCSI servers will be on the T++ pods.
Robocopy	Host level (IP)	Systems Administrator	Provides flexibility in LUN sizes during migration	Existing production servers are on the Cisco Network switches. Routing between Cisco and Arista network would be required to allow migration to occur since iSCSI servers will be on the T++ pods.

19.10 SRM Reporting

In SRM, LUNs will be reported in the same way that they are reported for other block-based LUNs. There is no information (configuration or performance) specific to the iSCSI protocol being used. Arrays will be dedicated to iSCSI. Naming the LUNs as iSCSI LUNs would be helpful in a multi-protocol environment.

20 Document Version Control

This section contains a detailed list of the updates made during each document release. Section 1 is purely an overview.

Version	Changes
0.1	Initial draft – general CATE peer review
1.0	Release For Production – no changes from global stakeholder review
1.1	- Reformatted to new template - Created new sections headers for grouping of sub sections - Added Section Assumptions and References - Added Section General Design - Added Section Software and Features - Updated Section Virtual Provisioning and Traditional (Classic) Storage Provisioning - Added Section Continuity Of Business - Added Section Data Security - Updated Section Role Based Access Controls (RBAC) - Updated Section Storage Domains - Updated Section Communication Interfaces - Updated Section SNMP and Email Alerting - Updated Section Event Monitor - Updated Section Host Connectivity and Configuration - Updated Section Off-Array Management Server
1.2	- Updates following peer review - Added SnapView Configurations
1.3	Release For Production – no changes from global stakeholder review
1.4	- Updated General Design to include SANscreen - Updated VNX Disk Drive Types to include support for 1TB and 3TB drives - Added section VNX Storage Processor I/O Slot Availability - Updated Software and Features to include supported connectivity protocols - Updated Software and Features and Virtual Provisioning to include restricted use of Virtual Provisioning and Application Protection for EXCH2010 - Added section SAS Back-end - Added section Placement of Drives - Updated FAST Cache, in particular added File Recommendations - Updated SEMS Audit Log Reconciliation reference to SEMS Events List - Updated Functional IDs to add rmagent for EXCH 2010 - Updated SNMP and Email Alerting to add note on use of VIP for SMTP host - Updated Performance Monitoring using Unisphere Analyzer - Updated Event Monitor - Updated Host Connectivity and Configuration - Updated Off-Array Management Station Software Installation Included screenshots of installation from new EMC binaries Removed version column in Unisphere Software Table - Updated Unisphere Client Installation Procedure - Updated Array Name and SP Network Name - Updated SP Cache including read/write cache as well as high and low watermarks recommendations
1.5	- Changes following peer review including - Updated Drive Sparing to accommodate for smaller and remote sites - Updated Host Registration to indicate that ESX hosts should auto register

	- Updated Initiator Settings to include ESX - Updated Upgrading Array Microcode to recommend EMC review prior to upgrade - Updated SNMP MIB to show default NetCool severity
1.6	Updated Software and Features to include restricted use of admsnap utility
1.7	- Added section Hardware VNX 2 Hardware - Updated Software and Features to include De-Duplication and Windows ODX - Updated Drive Sparing guidelines to cover VNX OE for Block 33 including - Permanent Drive Sparing Hot Spare Policy - Multicore RAID (VNX OE for Block 33) RAID 6 Parallel Rebuild - Multicore RAID (VNX OE for Block 33) Drive Mobility - Multicore RAID (VNX OE for Block 33) - Updated SP Cache guidelines to cover Multicore Cache - Updated FAST Cache guidelines to cover Multicore FAST Cache - Updated Virtual Provisioning guidelines to cover VNX OE for Block 33 including - Change in slice size for VNX OE for Block 33 Limits with Virtual Provisioning using VNX OE for Block 33 with each model LUN Creation (Virtual Provisioning) Note related to Thick LUN creation - Added section Symmetric Active/Active - Multicore RAID (VNX OE for Block 33) - Updated Host Registration to include Windows 2012 - Updated SnapView Configurations to cover VNX OE for Block 33 - Updated Reporting, Performance, Change Management and Capacity Planning - Updated Security Standards - Updated Server Requirements to Windows 2008 R2 as recommended standard. - Updated Off-Array Management Station Software Installation to include OS version in package names and Java version - Updated DPE/DAE - OS Configuration / Initialization to remove requirement for 15K 300GB SAS drives for vault disks
1.8	- Updated SnapView Configurations for VNX OE for Block 32.207 - Updated Placement of Flash drives for FAST Cache
1.9	- Updated Unisphere Client Installation Procedure Java prerequisite - Updated Unisphere Client Installation Procedure TPM package naming format - Release For Production – no changes from global stakeholder review
2.0	- Fix references in 1.7 above - Update references - Support for 4TB drives on VNX2 - Add in all items that are covered in GNAST doc (see ScaleIO from Scott Torg) - Update SRM site link - Test Unisphere with IE 11 and Chrome - Update comments on storage-efficiency following release notes testing for VNX2 - Update AD LDAP configuration to show FQDN - Support for additional drives - Add in HPNA details. - Add in Aperture details - Add section for Aperture to indicate dependency for Inventory, SNMP and configuration validation - LUN Migration currently only in Classic LUNs - Confirm Max LUN size of 16TB - Add multipathing (PowerPath / Native MPIO/ Native ESX) column to ALUA table - SAS back-end recommendations for VNX NG - Review Reserved LUN Pool guideline - Update SANscreen version required for VNX Next Gen - Add in Drive Zeroing, LUN Zeroing, Rebuild Logging, Write Journaling? - Write Journal protects data due to no hot spare available or whilst RAID group is in a degraded state. Write Journaling comes with some performance overhead: every write is now doubled. - Add SMTP and SNMP port information to ports table - Updated Added Section 11.2.7.1 Virtual Provisioning Free Space Recommendation - Add reference to validation checklist

	• Add comment that Deduplication has been introduced and will be assessed at later date: ◦ Deduplication of data at 8K block granularity ◦ Settable per pool LUN ◦ Freed up blocks are available for re-use within storage pool ◦ Deduplication occurs out of band • Block OE 32.201 is only compatible with Powerpath version 5.7 or later • Potentially include drive space map to show LUNs, LUN metadata, RAID group metadata, Virtual Drive metadata (see slide 37 from EMC presentation to AP) • Sniff is now drive based • FAST VP has changed and will be assessed at later date: ◦ slice size reduction from 1GB to 256MB ◦ Makes more precise relocation decisions ◦ Better space efficiency ◦ Less data movement impact ◦ Potentially takes less time to complete the relocation ◦ Lower copy overhead • Add section on Failure Domains • Update 8.3 References for more recent documents • Change version of Unisphere host to Windows 2008 R2. (confirm support 4.1 and 5.0 with 2008 SP2 and R2) • Include storage size for off array and update VMware specs • Update call home to mail via ESRS gateway • Update functional ID's to include region <i>all accounts globally have an entry in the root namespace (nsroot.net), although the accounts themselves are created in the regional domain, there is a requirement to keep every account name unique across the regional domains</i> <i>The accounts must have a unique suffix indicating the domain</i> • SAN copy • iSCSI for VNX (Appendix A)
Future	• 
	•