



Storage and Backup Engineering
Citi Architecture & Technology Engineering (CATE)

EMC VNX NAS

CSBE-11066-SG

Solutions Guide

Version 1.7

Version Control

Version	Date	Author	Reviewed	Comments
0.1	09/16/2011	Gary Johnson	Tim Armstrong Jackson Tan Quang Le Sean Tai	Peer Review
0.2	09/26/2011	Gary Johnson		RFC
1.0	10/24/2011	Gary Johnson	Citi Storage Community	RFP
1.1	05/25/2012	Gary Johnson	Tim Armstrong Sean Tai Stephen Hopkins	Unified Release
1.2	06/21/2012	Gary Johnson	Citi Storage Community	RFP
1.3	04/26/2013	Gary Johnson	David Lee	Storage Pools
1.4	11/27/2013	Gary Johnson	David Lee	VNX NG
1.4.1	12/18/2013	Gary Johnson David Lee	Citi Storage Community	Additional VNX NG File Filtering
1.5	09/15/2014	Gary Johnson David Lee	Citi Storage Community	
1.5.1	02/27/2015	Tan Jackson	Tim Armstrong	Updated for CAVA 6.5.0 RFT Updated VNX to AV Scanner Ratio Added dedicated FLR infrastructure requirement
1.6	03/25/2015	Tan Jackson	n/a	Release for Production
1.7	08/28/2015	David Lee	Citi Storage Community	Release for Production

Legal Notice

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

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, Citi and their subsidiaries also claim rights in certain other trademarks and service marks contained in these pages.

Contact

Please address any comments or questions regarding the information contained in this document via email to the “CATE Storage and Backup Engineering” distribution list or alternatively to gdse@citi.com.

Latest Versions

It is recommended that the latest version of the document be used from the CATE Storage Engineering website

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

Newer versions may include certification on additional platforms, or comments on issues discovered after deployment.

Table of Contents

1	Introduction	9
1.1	Architecture and Performance	9
1.1.1	Components.....	9
1.1.2	FAST Cache Architecture	12
1.1.3	Model Details and Views	12
1.2	Currently Supported VNX and VNX NG Models	23
1.3	Current Software Packages	24
2	Network and System Configuration	25
2.1	Accessing the Array via Unisphere or SSH	25
2.1.1	Accessing the Array from a Browser	25
2.1.2	Using the Unisphere Client	27
2.1.3	Logging Into Unisphere.....	28
2.1.4	Logging Into SSH.....	28
2.1.5	Navigation	28
2.2	Devices and Interfaces.....	28
2.3	Gigabit Ethernet (GigE) NIC Configurations and Naming.....	28
2.4	Aggregated (Trunked) Network Interfaces	29
2.5	Failsafe Network Planning	30
2.6	Control Station Naming Standards.....	30
2.6.1	Primary (Production) Site.....	30
2.6.2	Backup / Disaster Recovery Site	30
2.7	VDM Naming Standards	31
2.8	CIFS Server Naming Standards.....	31
2.8.1	Primary (Production) Site.....	31
2.8.2	Backup / Disaster Recovery Site	31
2.9	Domain Name and Host Resolution (DNS).....	31
2.10	Time Synchronization and Time Zones	32
2.10.1	NTP.....	32
2.10.2	Time Zones.....	32
2.11	VNX Standby Operations (Local Failover)	34
2.11.1	Failover Concepts.....	34
2.11.2	Data Mover Failover	34
2.11.3	How Standbys Work	35
2.11.4	Choose Standby Data Movers	35
2.11.5	Failover Policies	35
2.12	Control Station Failover Considerations (Dual Control Stations)	35
2.13	Forcing Failovers and Takeovers on Dual Control Stations	36
2.13.1	Takeover.....	36
2.13.2	Failover	36
2.13.3	Status.....	37
2.14	VNX Database.....	37
2.15	Using Robocopy with VNX (FAT Timestamp).....	37
2.16	Quotas.....	37
3	Storage Pools and Volume Management	38
3.1	Storage Concepts: VNX vs. Filer	38
3.2	AVM (Automated Volume Management)	38
3.3	File Storage Pools (Legacy Storage Pools for File)	39
3.4	Block Storage Pools	39
3.4.1	Creating Block Storage Pools.....	39
3.4.2	Expanding Block Storage Pools	43
3.4.3	Deleting Block Storage Pools	43
3.5	Supported RAID Configurations, Spares, and Vault	44
3.5.1	Spare Drives	45
3.5.2	Vault Drives and VNX Private Space.....	46
3.5.3	Supported Drive Types	46
3.6	Storage Pool Selection.....	46

3.7	VNX LUNs on Gateway Models (VG2, VG8)	46
3.8	Zoning and Registering with VNX Block Storage for Gateway Models	47
3.8.1	Zoning on Gateway Models	47
3.8.2	Registering the VNX within Unisphere on Gateway Models.....	47
3.9	FAST Cache Configuration	48
3.10	Adding more storage to VNX	49
3.11	File systems	49
3.12	File System Size Guidelines	49
3.13	File Systems and Automatically Extending File Systems	50
3.14	CIFS	51
3.14.1	Share Naming Recommendation	51
3.14.2	CIFS File Permissions	51
3.14.3	Using NFS and CIFS File Naming Conventions.....	51
3.14.4	Resolving File Naming Conflicts.....	51
3.14.5	Creating CIFS Servers	52
3.14.6	Using SMB Protocols (SMB 2.0, SMB 2.1, SMB 3.0)	53
3.14.7	Renaming a CIFS Server (Changing Compname).....	54
3.14.8	File Extension Filtering	55
3.15	NFS	59
3.15.1	NFS Overview	59
3.15.2	LDAP name service	60
3.15.3	Configuring LDAP	61
3.15.4	Debugging/Testing LDAP	65
3.15.5	NFSv3 and NFSv4 configuration	65
3.15.6	Access-security policy types.....	72
3.15.7	Configuration to support concurrent Multi-protocol access to data	75
3.15.8	Use of Local Files	84
3.15.9	Multi-protocol permission and access checking	85
3.15.10	Assign an interface to a VDM	90
4	Alerting	91
4.1	System Alerting via SNMP and Email	91
4.1.1	Alerting with SNMP Traps	91
4.1.2	System Alerting with Email	93
4.2	File System Capacity Alerting	103
4.2.1	Example SNMP Traps	104
4.2.2	Example Email Alert.....	104
4.2.3	System Warning Alert	104
4.3	Storage Pool Utilization Alerting.....	104
4.3.1	System Warning Alert	105
4.3.2	Email Examples for Utilization	105
4.3.3	Trap Examples for Utilization	105
4.4	SavVol Utilization Alerting	106
4.5	CAVA and DEDUPE Alerting (Event Notifications).....	107
4.5.1	Email Examples for CAVA and Dedupe	108
4.5.2	SNMP Trap Examples for CAVA and DEDUPE	109
4.6	DataMover CPU/Memory Alerting	110
5	Remote Support via ESRS	111
5.1	EMC Secure Remote Support Gateway (ESRS)	111
5.1.1	EMC VNX Remote Support Access.....	111
5.1.2	Gateway Server Agent.....	111
5.1.3	Gateway to EMC communication	111
5.1.4	Device Setting	112
5.1.5	Access Rights	112
5.2	Sending Alerts to EMC	112
6	Audit and System Logging	113
6.1	Audit Logging via SYSLOG	113
6.1.1	Message Format	113
6.1.2	Audit Scripting Setup	114

6.1.3	Syslog Server Setup	115
6.1.4	Audit Log Scripting Options	116
6.2	System Log Files	116
6.2.1	System Log Files Overview	116
6.2.2	Interpreting System Log Files	116
7	Reporting	117
7.1	Reporting via ECC	117
7.2	Reporting using Unisphere	118
7.3	EERS Entitlement Reporting	118
7.4	Performance Reporting	118
7.4.1	Summary Views	119
7.4.2	Table Views	119
8	Virtual Data Mover (VDM)	131
8.1	Terminology	131
8.2	Restrictions	131
8.3	Virtual Data Mover Concepts	132
8.4	Creating VDMs	133
8.5	Replicating and Failing Over VDMs	134
8.6	Using VDMs with 3DNS	136
8.6.1	Setting the Aliases (SPNs) for 3DNS	136
8.6.2	Disable Name Checking on Datamover for 3DNS	137
8.6.3	3DNS Port/IP Monitoring	138
8.6.4	Verify the User Connections (CIFS / SMB Connections)	138
8.6.5	Verifying the 3DNS Failover	139
8.6.6	Troubleshooting	139
8.6.7	Preventing Client Redirection during DM Reboots or Panics	139
9	File-Level Retention (FLR / WORM)	141
9.1	Compliance version of File Level Retention (FLR-C)	142
9.2	Enterprise version of File Level Retention (FLR-E)	142
9.3	Configuration	143
9.4	Replication	143
9.5	FLR Activity log	143
9.6	Anti-Virus	145
9.7	Deduplication	145
10	NDMP Backups	146
10.1	NDMP Overview	146
10.2	Authorizing Access to the NDMP Host	146
10.2.1	Authorization with NIS or LDAP	146
10.2.2	Authorization without NIS or LDAP	146
10.3	Backup of Checkpoints	147
10.4	Backup configurations	147
10.4.1	Remote NDMP to Tape	147
10.4.2	Remote NDMP to Tape	148
10.5	NDMP Operating Notes and Restrictions	149
10.6	NDMP Volume Backup (NVB)	149
10.7	NDMP Datastreams	150
10.8	Average File Size	151
10.9	Exclusions Lists for Backups	151
10.9.1	Directory Exclusions	152
10.9.2	File Exclusions for Backups	152
10.9.3	Valid path rules for excluding directories from backup	152
11	SnapSure Checkpoints (Snapshots)	154
11.1	Planning Considerations for SnapSure	154
11.2	Use of NL-SAS	154
11.3	Checkpoint Operations	154
11.4	Disk Space for Creating and Extending SavVols (Automatic)	154
11.5	SavVol Size Customization (Manual)	155
11.18	Recommended SavVol Sizes	156

11.19	Performance Considerations	156
11.20	Guidelines	157
11.21	Recommended Schedule of Checkpoints.....	157
11.22	Restrictions	157
11.22.1	SnapSure and VDMs.....	157
11.22.2	SnapSure and Replication.....	157
11.22.3	Checkpoints and Backup Restores	158
11.22.4	User Access for checkpoints (User Managed Restores)	158
12	VNX IP Replication	160
12.1	How VNX Replication sessions work	160
12.2	Space requirements	160
12.3	Required Communication Ports	160
12.4	Replicator/SnapSure	160
12.5	Controlling replication sessions with replication policies.....	160
12.6	Update policy.....	161
12.6.1	Max_time_out_of_sync.....	161
12.6.2	Manual refresh.....	161
12.7	Bandwidth schedule	161
12.8	General Restrictions	162
12.8.1	Filesystem Restrictions.....	162
12.8.2	VDM Replication Restrictions	162
12.9	Read-Only COB Optional Replication	163
12.9.1	Read-Only COB Impact to VDMs.....	163
12.9.2	Configuring the Read-Only Replication (VDM replication independent)	163
12.9.3	Configuring the Read-Only or Read-Write COB Test (with VDM replication) ...	165
13	VNX Anti-Virus Agent (CAVA)	169
13.1	Integration with Citi AV	169
13.2	CAVA System Requirements	169
13.3	Supported 3rd Party AV Engines	169
13.4	VNX to AV scanner ratio	170
13.5	CAVA File Exclusions	170
13.6	CAVA Restrictions.....	170
13.7	CAVA Setup Procedures.....	170
13.8	CAVA Domain Functional ID	170
13.9	Create CIFS Server for CAVA.....	171
13.10	Create Local User Group on CAVA CIFS Servers	171
13.11	Assign EMC Virus-Checking Right to Local Group	171
13.12	Assign Local Admin Rights to CAVA Domain User on CAVA Server	172
13.13	CAVA Setup on a scanning server	172
13.14	CAVA Server Network Configuration.....	172
13.15	ePO Policy on CAVA Server.....	173
13.16	CAVA Configuration file	173
13.16.1	RPCTimeOut and RPCRetryTimeout.....	174
13.16.2	Shutdown Parameter.....	174
13.16.3	Parameter File Location	174
13.16.4	Unique File for each Datamover.....	174
13.17	Create Local User on VNX	175
13.18	Starting and Stopping CAVA	175
13.19	Updating Configuration File	175
13.20	Deleting CAVA	176
13.21	Checking the status of CAVA	176
13.22	Alerting on CAVA events	177
13.23	Confirmation of Virus Scanning Protection	177
13.23.1	Test Virus Scanning	177
13.23.2	Log Files	178
13.24	CAVA Sizing Tool	178
13.24.1	Enable the CAVA Sizing Tool Registry Option.....	178
13.24.2	Interactive CAVA Sizing Tool	179

13.24.3	Scripting CAVA Sizing Tool	180
14	DeDuplication (Dedupe) and Compression	182
14.1	Configuration	182
14.2	Compression	183
14.3	Enabling Deduplication	183
14.4	Processing Time.....	184
14.5	Performance.....	184
14.5.1	Read Access	184
14.5.2	Write Access.....	184
14.6	File System Selection.....	184
14.7	Additional Configuration and Deduplication Parameters	185
14.8	File System Space Usage	187
14.8.1	Auto-extend file systems	187
14.8.2	Non auto-extend file systems	187
14.8.3	Working space from File System.....	187
14.9	SavVol Space Usage	187
14.10	Methodology for Initial Deduplication (Due to Space Limitations)	188
14.10.1	File System Workarounds	188
14.10.2	SavVol Workarounds.....	189
14.11	Controlling When Deduplication is Run	189
14.12	Deduplication and Backups	190
14.12.1	Network Share Based Backups.....	190
14.12.2	NDMP Based Backups	190
14.12.3	NDMP Volume Backups (NVB or VBB).....	190
14.13	Checkpoints and Deduplication	190
14.14	Replication and Deduplication	190
14.15	File Level Retention (FLR) and Deduplication	191
14.16	Alerting on DEDUPE events	191
15	Upgrading and Downgrading the VNX.....	192
15.1	Upgrades.....	192
15.2	Downgrades	192
15.3	Upgrade Procedures	192
16	Security Best Practices	193
16.1	Functional IDs	193
16.2	Root Access for Storage Administrator and EMC support accounts (Sudo)	193
16.3	Control Station Security and Configuration	196
16.3.1	Control Station Setup of DNS and NTP	196
16.3.2	Access Control List (ACL) via IPTables	197
16.3.3	Control Station Authentication & Authorization via Active Directory	202
16.3.4	Default Password Expiration	207
16.3.5	Idle Session Timeout	207
16.4	Role-Based Access Control	207
16.4.1	Groups and Roles.....	207
16.5	VNX Audit Log Files	209
16.6	Terminal Server (Uplogix)	209
16.7	Standard Build Security Checklist.....	210
16.8	Required Scripting.....	210
16.9	Disabling the FTP service on the Data Movers.....	210
16.10	Apache Web Server Security.....	210
16.11	Altering the Local User Path (ECC SSH account)	211
16.12	Security Standards.....	211
16.13	Information Classification Certification.....	211
17	Unisphere Navigation.....	212
17.1	Unisphere Navigation	212
17.1.1	HOME	212
17.1.2	Systems List	212
17.1.3	Dashboard	213
17.1.4	SYSTEM LIST	213

17.1.5	SYSTEM	213
17.1.6	STORAGE	213
17.1.7	HOSTS	213
17.1.8	SETTINGS.....	214
17.1.9	SUPPORT	214
17.2	Customizing the Dashboard.....	214
18	Unisphere Storage Domains.....	216
18.1.1	DART and Unisphere	216
18.1.2	Legacy and Multi Domains	216
18.1.3	Domain Attributes	216
18.1.4	Storage Domain Design	216
18.1.5	Reporting	216
18.1.6	Storage Domain Masters.....	217
18.1.7	Local vs. Global	217
19	Unified Storage Considerations	218
19.1	Unified Storage Support Ownership.....	218
19.2	Unified Storage Notifications	218
19.3	Unified Storage Audit Logging	218
19.4	Unified Storage Activities and Roles	218
19.5	Unified Storage NAS / SAN Team Coordination.....	219
19.6	Unified Storage Security	220
19.7	Unified Storage Drive Configuration.....	220
20	APPENDICES	221
20.1	APPENDIX A: <i>SNMP MIB</i>	221
20.2	APPENDIX B: <i>Operational Commands</i>	225
20.2.1	Check Versions of Software	225
20.2.2	Check System Status	225
20.2.3	Collect Support Materials and SP Collect	230
20.2.4	Backing Up the VNX Database	232
20.2.5	Restart Control Stations	232
20.2.6	Restart Datamovers.....	232
20.2.7	Detailed Storage Configuration	233
20.3	APPENDIX C: <i>Standard Build Security Checklist</i>	235
20.4	APPENDIX D: <i>Additional VNX Scripting</i>	237
20.4.1	Check Configuration (Recommended)	237
20.4.2	Collect Logs (Required).....	239
20.5	APPENDIX E: <i>Sample Configuration of VNX Models</i>	242
20.5.1	VNX 5300 Sample Configuration	242
20.5.2	VNX 5500 Sample Configuration	242
20.5.3	VNX 5700 Sample Configuration	243
20.5.4	VNX 7500 Sample Configuration	244
20.5.5	VNX 5200 Sample Configuration	245
20.5.6	VNX 5400 Sample Configuration	245
20.5.7	VNX 5600 Sample Configuration	246
20.5.8	VNX 5800 Sample Configuration	246
20.5.9	VNX 7600 Sample Configuration	247
20.5.10	VNX 8000 Sample Configuration	247
20.6	APPENDIX F: <i>Software Enablers of VNX Models</i>	249
20.7	APPENDIX G: <i>CAVA 4.6.8 CAVA Support and Dual CPU Recommended Ratio</i>	250
20.7.1	CAVA System Requirements	250
20.7.2	VNX to AV scanner ratio.....	250
21	Revision History	251

1 Introduction

The operating system used by EMC VNX NAS is DART 7. This guide is a general release solution implementation for all Citi VNX NAS systems. This guide describes the technical aspects for EMC VNX's that run DART 7 software, now called OE for File. The document is intended for Systems and Storage Administrators responsible for the management of Citi VNX NAS. The goal of this release is to provide supported and certified software configurations for use on all Citi EMC VNX NAS platforms.

To access the solutions guide for Block (or SAN) access see [EMC VNX for SAN Solutions Guide](#).

There is a section in this guide that describes the Unified configuration aspects. The VNX Unified array is a VNX that handles both Block and File. For more information about the Unified configuration see [Unified Storage Considerations](#).

1.1 Architecture and Performance

The VNX is designed for file only, block only, and unified (block and file) implementations. The EMC VNX series arrays try to achieve scalability and efficiency in a simple intuitive user interface for the mid-tier storage customer. The VNX series utilize Intel's Xeon 5600 processor's higher clock speeds and more cores to handle the demands of SAN and NAS storage. Increased memory capacity, better use of FLASH drives, a SAS 6Gb/s (4x lanes) back end, PCIe Gen 2 I/O, and denser disk enclosures help provide improved scalability and efficiency. It is estimated that the VNX series performs two to three times faster than previous-generation CLARiiON and Celerra systems.

When in degraded mode (single SP), all VNX models perform better compared to its (non-degraded) predecessor on IOPS intensive workloads. Bandwidth-intensive workloads will suffer some performance degradation, but to lesser of a degree than the previous generation in degraded mode.

The backend for the VNX is a 4x channel 6GB/s SAS data bus.

When using FLASH, the VNX can achieve more IOPS than previous generations making FLASH Cache much more efficient for leveling peak I/O requests.

1.1.1 Components

1.1.1.1 Storage Processors

The storage processors (SPs) support block data via FC, iSCSI, and FCoE protocols. The SPs provide access for all external hosts and for the block storage behind the file die of the VNX array.

Note: iSCSI is a restricted protocol and must be approved by CSBE

1.1.1.2 Disk Processor Enclosure (DPE)

The disk processor enclosure houses each storage processor and the first tray of disks. It is used by the 5200, 5300, 5500, 5400, 5600, 5800, and 7600. This component has dual power supplies.

1.1.1.3 Storage Processor Enclosure (SPE)

The storage processor enclosure is used by some of the higher-end models VNX 5700, VNX 7500, and VNX 8000 houses each storage processor (no drives). This component has dual power supplies.

1.1.1.4 Data Movers (DM)

The data movers (or blades or X-blades) access data from the SPs and provide host access via CIFS, NFS, and other protocols. Since the x-blades do not manage the back-end storage, they are more scalable and are able to provide N+1 redundancy.

**Note that while the VNX/VNX2 supports iSCSI, HTTP and FTP protocols, these protocols are not certified for use within Citi on the NAS platform.*

1.1.1.5 Data Mover Enclosure (DME)

The datamover enclosure houses the data movers and is similar in form to the SPE. This component has dual power supplies.

1.1.1.6 Stand-by Power Supply (SPS)

The stand-by power supplies provide enough power to each SP to ensure that any data in flight is de-staged to the vault area in the event of a power failure and thereby preventing any lost writes. Upon restart of the array, the pending writes are reconciled and persisted.

1.1.1.7 Control Stations

The control stations provide management functions to the file-side components of the array. The control station is responsible for data mover failover. Control stations can be configured in a dual-mode to ensure additional redundancy.


1.1.1.8 Disk Array Enclosures (DAE)

The disk array enclosures house the drives used in the array. There are 15-drive (3.5" NL-SAS drive or 2.5" SAS or 3.5" SSD) and 25-drive (2.5" SAS) variants. There is also a 60-drive variant that is not supported by CATE.

25-drive DAE (2U)


The 25-disk DAE can hold up to 25 2.5" SAS, NL-SAS, or SSDs with 6Gb/s connectivity.

Front View (25 Disk DAE):



Item	Description
1	2.5" 6Gb/s SAS, NL-SAS, or SSD drive
2	DAE Fault status LED (amber)
3	DAE Power status LED (green)
4	Disk Drive fault LED (amber)
5	Disk Drive status / activity LED (blue)

Rear View (25-Disk DAE):




Item	Description
1	LLC B Power Supply status LED (green)
2	LLC B Power Supply fault LED (amber)
3	LLC B AC Power Supply Plug
4	LCC B SAS Connector Link LED (blue)
5	LCC B SAS Connector Input
6	LCC B SAS Connector Output
7	LCC B Bus ID
8	LCC B Power and Fault LEDs
9	DAE Enclosure ID or Address
10	LCCB Management (RJ-12) Connector to SPS
11	LCC A Power Supply Latch Handle
12	LCC A Right Latch Handle

15-disk DAE (3U)


The 15-disk DAE can hold up to 15x 3.5" (or 2.5") SAS, NL-SAS, or SSDs with 6Gb/s connectivity.

Front View (15 Disk DAE):



Item	Description
1	2.5" 6Gb/s SAS, NL-SAS, or SSD drive
2	DAE Fault status LED (amber)
3	DAE Power status LED (green 2Gb / blue 6Gb)
4	Disk Drive fault LED (amber)
5	Disk Drive status / activity LED (green)

Rear View (15-Disk DAE):



Item	Description
1	LLC B AC Power Supply Plug
2	LLC B Power Supply Fan fault (amber)
3	LLC B Power Supply status (green)
4	LCC B Power Supply fault (amber)
5	LCC B Management (RJ-12) Connector to SPS
6	LCC B SAS Connector Link LED
7	LCC B SAS Connector Output
8	LCC B SAS Connector Input
9	LCC B Bus ID
10	LCC B Bus fault LED (amber)
11	LCC B Bus status LED (green)
12	DAE Enclosure ID or Address

1.1.1.9 Power Supplies

The power supplies on the array are redundant and provide both power and cooling. They are also 7-10% more efficient than their predecessors and provide for adaptive cooling (variable fan speeds).

1.1.2 FAST Cache Architecture

The VNX series arrays use FAST Cache to extend the arrays read-write cache and ensure that unpredictable I/O spike are provided services at the highest tier speed. FAST Cache works at a granular level by copying 64KB chunks onto Flash drives reserved for FAST Cache depending on the I/O characteristics.



Repeated (3x) access of the same 64KB chunk of data causes the policy engine to promote that data to FAST Cache. It has a very quick response time in the order of milliseconds. FAST Cache reacts to I/O spikes and maintains access levels by acting as an extension to the existing cache. FAST Cache currently supports SLC SSD drives, so larger 400GB+ MLC SSD drives cannot be used. Under DART 8, if the utilization is < 80% then promotions occur for all requests until 80% is reached. Checkpoints are excluded from FC.

1.1.3 Model Details and Views

Each array model has specific limits on the number of drives supported and datamovers. Most models have a DPE and DAEs for drives, whereas some of the the larger models (5700,7500, and 8000) use DAEs for all drives. The VNX5100 does not support file services so it is not described here.



1.1.3.1 VNX5300

The VNX5300 is a dual datamover array with up to 125 drives. There are 3 configurable I/O slots per data mover. Each Datamover has 6GB system memory. There is a DPE for the first 15 drives (3.5" slots). The following is an example configuration with dual SPS; 15 disk drive DPE with two storage processors; two control stations; one datamover enclosure with two datamovers.





1.1.3.2 VNX5500

The VNX5500 is a triple datamover array with up to 250 drives (240x 3.5" or 250x 2.5"). There are 4 configurable I/O slots per data mover. Each Datamover has 12GB system memory. There is a DPE for the first 15 drives (3.5" slots). The following is an example configuration with dual SPS; 15 disk drive DPE with two storage processors; two control stations; two datamover enclosures with three datamovers.





1.1.3.3 VNX5700

The VNX5700 is a quad datamover array with up to 500 drives (495x 3.5" or 500x 2.5"). There are 4 configurable I/O slots per data mover. Each Datamover has 12GB system memory. This array uses DAEs for all of its drives (has an SPE instead of DPE). The following is an example configuration with dual SPS; SPE with two storage processors; two control stations; two datamover enclosures with four datamovers.



1.1.3.1 VNX7500


The VNX7500 is an 8 datamover array with up to 1000 drives (990x 3.5" or 1000x 2.5"). There are 5 configurable I/O slots per data mover. Each Datamover has 24GB system memory. This array uses DAEs for all of its drives (has an SPE instead of DPE). The following is an example configuration with dual SPS; SPE with two storage processors; two control stations; four datamover enclosures with eight datamovers.




1.1.3.1 VNX 5400

The VNX 5400 is a dual datamover array with up to 250 drives. There are 3 configurable I/O slots per data mover. Each Datamover (up to 2) has 6GB system memory. There is a Disk Processor Enclosure (DPE) for the first 25 drives (2.5" slots). The following is an example configuration with dual SPS; 25 disk drive DPE with two storage processors; two control stations; one datamover enclosure with two datamovers.

Front View (5400):




Rear View (5400):




1.1.3.2 VNX 5600

The VNX 5600 is a dual datamover (3 allowed by RPQ for Citigroup) array with up to 500 drives. There are 3 configurable I/O slots per data mover. Each Datamover (up to 3 by RPQ) has 12GB system memory. There is a Disk Processor Enclosure (DPE) for the first 25 drives (2.5" slots). The following is an example configuration with dual SPS; 25 disk drive DPE with two storage processors; two control stations; one datamover enclosure with two datamovers.

Front View (5600):



Rear View (5600):




1.1.3.1 VNX 5800


The VNX 5800 is a triple datamover array with up to 750 drives. There are 4 configurable I/O slots per data mover. Each Datamover has 12GB system memory. There is a Disk Processor Enclosure (DPE) for the first 25 drives (2.5" slots). The following is an example configuration with

dual SPS; 25 disk drive DPE with two storage processors; two control stations; two datamover enclosures with one and two datamovers.

Front View (5800):




Rear View (5800):




1.1.3.1 VNX 7600

The VNX 7600 is a quad datamover array with up to 1000 drives. There are 4 configurable I/O slots per data mover. Each Datamover has 24GB system memory. There is a Disk Processor Enclosure (DPE) for the first 25 drives (2.5" slots). The following is an example configuration with dual SPS; 25 disk drive DPE with two storage processors; two control stations; two datamover enclosures with two datamovers.

Front View (7600):




Rear View (7600):




1.1.3.1 VNX 8000

The VNX8000 is an 8 datamover array with up to 1000 drives (expected to increase in future). There are 5 configurable I/O slots per data mover. Each Datamover has 24GB system memory. This array uses DAEs for all of its drives (has an SPE instead of DPE). The following is an example configuration with dual SPS; SPE with two storage processors; two control stations; two datamover enclosures with eight datamovers.

Front View (8000):



Rear View (8000):



1.2 Currently Supported VNX and VNX NG Models

The following chart lists the models supported by CATE Storage Engineering. Any legacy or current hardware platforms should be upgraded to the current recommended DART level. “X-Blade” and “Data Mover” are synonymous terms. The distinction between the VNX and VNX NG

models is minimal, though the NG series has improved back-end and the microcode was rewritten to take advantage of multi-core technology.

Array Model	Control Stations	Storage	Data Movers	CPU/Mem per DM	Flex IO ports per DM
VNX					
5300	1-2*	Internal (up to 125 drives, 2x SAS Bus)	1-2**	Xeon 5600 4C 2.8GHz / 6GB	3
5500	1-2*	Internal (up to 250 drives, 2x/6x SAS Bus)	1-3**	Xeon 5600 4C 2.8GHz / 12GB	4
5700	1-2*	Internal (up to 500 drives, 4x SAS Bus)	2-4**	Xeon 5600 4C 2.8GHz / 12GB	4
7500	1-2*	Internal (up to 1000 drives, 4x/8x SAS Bus)	2-8**	Xeon 5600 4C 2.8GHz / 24GB	5
VG2	1-2*	External (VNX, CLARiiON)	1-2**	Xeon 5600 4C 2.4GHz / 6GB	3
VG8	1-2*	External (VNX, CLARiiON)	2-8**	Xeon 5600 4C 2.4GHz / 6GB	5
VNX NG					
5200	1-2*	Internal (up to 125 drives, 2x SAS Bus)	1-2**	Xeon 5600 4C 2.4GHz / 6GB	3
5400	1-2*	Internal (up to 250 drives, 2x SAS Bus)	1-2**	Xeon 5600 4C 2.4GHz / 6GB	3
5600	1-2*	Internal (up to 500 drives, 2x/6x SAS Bus)	1-3***	Xeon 5600 4C 2.4GHz / 12GB	3
5800	1-2*	Internal (up to 750 drives, 2x/6x SAS Bus)	1-3**	Xeon 5600 4C 2.4GHz / 12GB	4
7600	1-2*	Internal (up to 1000 drives, 2x/6x SAS Bus)	2-4**	Xeon 5600 4C 2.4GHz / 24GB	4
8000	1-2*	Internal (up to 1000 drives, 2x/16x SAS Bus)	2-8**	Xeon 5600 4C 2.4GHz / 24GB	5

*per CSBE standards, 2x control stations are required.

**per CSBE standards, 2x DMs are the minimum required.

***up to 3x DMs are allowed on the 5600 by RPQ for Citigroup.

The Flex IO ports per Datamover (X-Blade) can be configured as:

- Four ports 10/100/1000 BaseT (RJ-45 Copper)
- Two ports 10/100/1000 BaseT and 2 ports 1Gb Ethernet Optical
- Two ports 10Gb Ethernet Optical
- Four ports 2/4/8Gb/s Fiber Channel

The Citi standard for VNX X-Blades will be:

- 1x Four ports 8Gb/s Fiber Channel
- 2x Two ports 10Gb Ethernet Optical
- Or
- 1x Four ports 8Gb/s Fiber Channel
- 2x Four ports 10/100/1000 BaseT

1.3 Current Software Packages

The following are the current software packages needed for optional or required functionality:

Software	Version	Platform	Package / SPAC
EMC VNX CAVA Agent	6.5.0	Win 2008/ 2012 x64	CITI_WIN_X64_CAVA_650_00
Tectia SSH Client	6.3.8	Win 2008/2012 x64	CITI_WIN_TECTIASSHCLIENT_638_00
Active Perl	5.1.42	Win 2008/2012 x64	CITI_WIN_ActivePerl_5142_00
McAfee Virus Scan	8.8i	Win 2008/2012 x64	CITI_WIN_VirusScanENT_8801_02

2 Network and System Configuration

2.1 Accessing the Array via Unisphere or SSH


2.1.1 Accessing the Array from a Browser

In order to access the management screen on a VNX, open an IE (or FireFox) browser to the IP address of the primary control station.


Note: If the array has already been configured for security, any new client would need to be added to the ACL to allow for HTTPS or SSL connections. For more information see [Access Control List \(ACL\) via IPTables](#).

It is recommended that you access the array via designated management servers and not directly from workstations to avoid frequent updates to the ACL lists and to minimize consistency issues due to local workstation security settings.


Using a standard Citi build server as the client will require that you add the IP address (or a pattern of IP addresses) to the trusted sites list for intranet. This will enable JavaScript and other security settings needed by moving the site security to Medium.



To add the IP (or pattern of IPs) to the trusted sites list within IE, go to Tools → Options → Security.



Select Local Intranet and then select the Sites button.




You can add a https entry for the IP address of the primary control station. The following is a pattern example for all systems in a particular subnet:

://10.40.42.

Proxy Server Settings

In order to bypass the proxy server when accessing the arrays, the default proxy settings should be changed. To modify the proxy settings within IE, go to Tools → Options → Connections. Select the LAN Settings button.




Automatically detect settings should be unchecked.

- Use automatic configuration script should be unchecked.
- Use a proxy server for your LAN should be checked.
- Address: webproxy.wlb.nsroot.net (or a similar domain specific web proxy server)
- Port: 8080
- Bypass proxy server for local addresses should be checked.

To create a pattern for local addresses, select the Advanced button

Under Exceptions, you can put in a simple patterns such as 10.* for all 10.x.x.x network addresses.



2.1.2 Using the Unisphere Client

The use of the client is very similar to the above procedure. The Unisphere Client is a pre-loaded copy of the JAVA Unisphere client that is specific to certain releases of DART. The interface is nearly identical to running a browser pointing to the array IP address. You will need to open IE and configure the trusted sites. You will probably not see the JavaScript error as above. Run the client. You may have to acknowledge ActiveX scripting several times to see the first screen. It will ask for the IP address of the array.

Once configured, the Unisphere Client may improve stability and may avoid having to modify the proxy settings. The Unisphere Client has a limited range of compatibility, but generally a newer client can access an older DART/FLARE or OE File/ OE Block version. Version incompatibility

will prevent you from logging onto that array. You can go back to the direct browser access or use an SSH client (such as Tectia SSH).

2.1.3 Logging Into Unisphere

The login screen asks for username and password. If using a local account, do not check the LDAP box. Global scope should be used (unless it is a local account to only a single array).

If using an AD account, do not refer to the domain and only use the username. Check the LDAP box for these accounts only. The scope will not be selectable.

2.1.4 Logging Into SSH

Using the Tectia SSH client is no different than other SSH sessions. You will want to provide the IP address of the primary control station. If using an AD account, you will need to refer to the AD domain (i.e. domain\user or user@domain) when logging on.

2.1.5 Navigation

For more details on the menu structure for Unisphere, see [Unisphere Navigation](#).

2.2 Devices and Interfaces

The VNX Network Server (integrated and gateway) NAS systems support several different types of NICs. Each Data Mover (X-blade) model supports multiple NICs.

Flex IO Network Module Choices	1Gb Configuration (2 slots per DM)	10Gb Configuration (2 slots per DM)
Four ports 10/100/1000 BaseT (RJ-45 Copper) Two ports 10/100/1000 BaseT and 2 ports 1Gb Ethernet Optical Two ports 10Gb Ethernet Optical	1. Four ports 10/100/1000 BaseT (RJ-45 Copper) 2. Four ports 10/100/1000 BaseT (RJ-45 Copper)	1. Two ports 10Gb Ethernet Optical 2. Two ports 10Gb Ethernet Optical

For each physical port on a NIC, the Datamover software defines a network device. The device name consists of a device prefix, representing the type of NIC, and a number that is sequentially appended, starting with 0 (zero), for each port on the NIC. For example, the first Ethernet port is cge0. For a device with four Ethernet ports, the remaining ports will be named cge1, cge2, and cge3, respectively.

2.3 Gigabit Ethernet (GigE) NIC Configurations and Naming

On each datamover, there will a total of 8x GigE ports or 4x 10GigE ports. Use of these ports should be spread across the cards to account for card or port failure. The GigE (or 10GigE) ports should be configured (trunked and) in an Active-Standby configuration for redundancy from both the datamover and network switch perspective.

A minimum of two connections from each datamover **MUST** be connected to different network switches for redundancy in case of switch failure.

All data movers **MUST** be on the same network segment (subnet) to support a data mover fail-over scenario, the standby data mover has to takeover the name and IP address of the failed data mover.

When configuring the virtual interfaces, single-mode trunks/fail safe networks should be used (active-passive).

All trunked and fail safe 1 or 10 Gigabit Ethernet interfaces should be named fsn#, where the first virtual interface is number zero (fsn0). Every additional interface should use the next sequential number. Virtual interfaces should be created in an active-standby configuration.

Below is a sample command line for virtual interface creation.

```
$ server_sysconfig server_2 -virtual -name fsn0 -create fsn -option "primary=cge2
device=cge2,cge3"
```

To view the fail-safe network device, type:

```
$ server_sysconfig server_2 -virtual -info fsn0
server_2 :
*** FSN fsn0: Link is Up ***
active=cge2 primary=cge2 standby=cge3
```

2.4 Aggregated (Trunked) Network Interfaces

Cisco Link Aggregation Control Protocol (LACP) is the primary means of trunking at Citi.

Link Aggregation Control Protocol – Resembles an Ethernet channel but uses the Link Aggregation Control Protocol (LACP). Link aggregation works with switches that support the IEEE 802.3ad Link Aggregation standard.

The following standards should be followed when implementing fail-safe networking or aggregated IP network ports:

- Additional Ethernet ports may be used in data movers where greater aggregate network bandwidth is required than offered by a single active/passive Ethernet pair. It is recommended that ports should be added in pairs, 1 port from each Flex I/O module, up to a total of 8 ports when using 2 x GigE modules or 4 ports with 2 x 10GigE modules.
- Each of the fail safe networks from the data mover **MUST** be connected to different network switches for redundancy in case of switch failure.
- All data movers in the VNX **MUST** be on the same network segment since in a data mover fail-over scenario; the standby datamover has to takeover the name and IP address of the failed data mover.

When creating the virtual interfaces on the VNX, the following naming standard should be followed:

- The letter at the end designates each port on system board.
- All 10/100Mb Ethernet virtual interfaces should be named ana#
- All Gigabit Ethernet virtual interfaces should be named gev#.
- All 10Gig Ethernet virtual interfaces should be named tgev#
- Trunks should be named trk#.
- Failsafe networks (failover) should be named fsn#.
- Every additional virtual interface (trunked or not) should use the next sequential number. In a cluster configuration all Ethernet virtual interfaces in the cluster pair **MUST** be unique.

Example:

```
$ server_sysconfig server_2 -virtual -name fsn0 -create fsn -option "primary=trk0_ec
device=trk1_ec"
```

2.5 Failsafe Network Planning

Failsafe networks should always span multiple I/O cards in a datamover. FSNs should also span network switches (when possible). This configuration will prevent network failures for switch failures, multiple port failures on the switch, multiple port failures on the I/O card, or I/O card failures.

Keep in mind that failsafe networks will go to the failover side only when all of the primary components have failed. Therefore, a trunk will failover only when the entire trunk has failed. Since failover network definitions are active/passive, the cost of the passive link should be considered. Here are the various configurations supported:

Orientation	Network Cards	FSN Configuration	Notes
Performance	2x 10GigE (2 ports each)	FSN (LACP (1,2) + LACP (1,2))	using all 4 ports, 2 from each card
Performance	2x 1GigE (4 ports each)	FSN (LACP (1,2,3,4) + LACP (1,2,3,4))	using all 8 ports, 4 from each card
Performance	2x 1GigE (4 ports each)	FSN (LACP (1,2) + LACP (1,2))	using only 4 of 8 ports, 2 from each card
Cost	2x 10GigE (2 ports each)	FSN (LACP (1,2) + 1)	using 3 of 4 ports, 2 from card 1 + 1 card 2
Cost	2x 10GigE (2 ports each)	FSN (1+1)	using 2 of 4 ports, 1 from each card
Cost	2x 1GigE (4 ports each)	FSN (LACP (1,2,3,4) + LACP(1,2))	using 6 of 8 ports, 4 from card 1 + 2 card 2
Cost	2x 1GigE (4 ports each)	FSN (LACP (1,2) + 1)	using 3 of 8 ports, 2 from card 1 + 1 card 2
Cost	2x 1GigE (4 ports each)	FSN (1 + 1)	using 2 of 8 ports, 1 from each card

2.6 Control Station Naming Standards

2.6.1 Primary (Production) Site

The following pattern should be used:

nas<2 letter country code><3 letter location code> <2 digit sequential number><cs><0or1> where 0=primary cs 1=secondary
Example: nasusmwd01cs0

NAS stands for Network Attached Storage. CS stands for control station. In a dual control station pair, the Control stations should be number on sequence i.e. 1 & 2, 3 & 4 etc..

The two-letter country code standard can be found at:

http://en.wikipedia.org/wiki/ISO_3166-1

2.6.2 Backup / Disaster Recovery Site

To distinguish primary control station from disaster recovery control station, “dr” will be appended at the end of the above naming convention. Therefore, the names will be:

nas<2 letter country code><3 letter location code> <2 digit sequential number><cs><0or1>dr
where 0=primary cs 1=secondary cs

Example:nasswd01cs0dr

The two-letter country code standard can be found at:

http://en.wikipedia.org/wiki/ISO_3166-1

2.7 VDM Naming Standards

VNX offers a virtual datamover that is more flexible than a physical datamover. These VDMs should be named with the control station and a vdm_# (unique integer) format so that they will be easily identified as a virtual rather than physical datamover.

<controlstationname>_vdm_<unique integer>

where control station base name is the control station name with cs0 removed

Example: nasmwd01cs0_vdm_1

2.8 CIFS Server Naming Standards

2.8.1 Primary (Production) Site

The following pattern should be used:

nas<3 letter location code> <2 digit sequential number><optional 3 letter business code>

Example: nasSWD01car

Note: This server may change locations (i.e. during a failover), but the name should reflect the location of the primary site.

2.8.2 Backup / Disaster Recovery Site

There should not normally be a backup CIFS server named. The only exception to this is a special active-passive configuration where the file system is replicated to a read-only remote CIFS server.

2.9 Domain Name and Host Resolution (DNS)

VNX deployments should maintain the default configurations for name resolution for NFS and CIFS authentication. The following is required:

- During the initial VNX setup the operator will be asked to enter the Domain and DNS server information. The VNX must be configured to use the global DNS settings as described in http://ipns.ny.smb.com/documents/Product-Development/CA-NTP/GNAST/NTP_design_RFP.doc

The order in which these IP addresses should be configured at each site can be found in Appendix C of the above document. If the site at which the VNX is installed is not listed in the aforementioned document, then use the default order of 192.193.215.65, 192.193.215.69, 192.193.215.73.

For more details on how to configure DNS for the control station, see [Control Station Setup of DNS and NTP](#).

- During CIFS setup the operator will be asked to optionally enter NETBIOS name for WINS resolution. WINS is a legacy process at Citi and should not be set unless required by a specific application

- Dynamic DNS is not supported at Citi. By default, datamovers attempt to make periodic updates to DNS (and during updates during failovers). To prevent this conflict the parameter “dns updateMode” should be set to 0 (off).

```
server_param server_2 -facility dns - modify updateMode -value 0
```

2.10 Time Synchronization and Time Zones

2.10.1 NTP

Time synchronization needs to be configured for the VNX to get updates from the appropriate NTP time source for the region. Time synchronization is important, especially for CIFS users in an Active Directory environment, where VNX's with an incorrect system time can actually cause a user denial of access due to the mismatch between the Kerberos ticket time-stamp and the system time.

To start time services on a Data Mover and immediately set the system time to the value returned by the specified time server, use this command syntax:

```
$ server_date <movername> timesvc start ntp [<host> [,<host>,....]]
```

To start time services on a Data Mover and immediately set the system time to the value returned by the specified time server as well as to specify the interval between polls of the time server, use this command syntax:

```
$ server_date <movername> timesvc start ntp [-interval<hh>[:<mm>]][<host> [,<host>,....]]
```

To start time services on a Data Mover and begin gradually modifying the system time to the value returned by the specified time server as well as specifying the interval between polls of the time server, use this command syntax:

```
$ server_date <movername> timesvc start ntp [-sync_delay] [-interval  
<hh>[:<mm>]][<host> [,<host>,....]]
```

where:

<movername> = name of the Data Mover

<hh> = hours

<mm> = minutes

<host> = IP address of NTP server

Example: To start the time service on server_2 and begin skewing the system time to match the first value returned by the time server, and set the interval between polls of the time server to one hour, 30 minutes, type:

```
$server_date <movername> timesvc start ntp -sync_delay -interval 01:30 192.193.215.65  
192.193.215.69 192.193.215.73
```

For a list of NTP servers, refer to Appendix C of:

http://ipns.ny.ssmc.com/documents/Product-Development/CA-NTP/GNAST/NTP_design_RFP.doc

For the order in which these IP addresses should be configured at each site. Use the example above if the site is not listed.

For more details on how to configure DNS for the control station, see [Control Station Setup of DNS and NTP](#).

2.10.2 Time Zones

CATE Storage Engineering recommends that the local time zone be used for the geographic location at which the VNX is installed. For example, VNXs in California should be configured

PST8PDT. If the time zone in which the VNX is located supports Daylight Savings Time (DST), then the VNX should be configured with the time zone information that supports DST.

To check your current time zone configuration, type `server_date <servername> timezone` at the command line. To update your time zone configuration, type `server_date <servername> timezone zone_name` at the command line (where `zone_name` is the name of your time zone).

For example, to set the time zone for Central Time Zone (including daylight savings adjustment):
`server_date <movernam> time zone -name America/Chicago`

For example, to set the time zone for Eastern Time Zone (including daylight savings adjustment):
`server_date <movernam> time zone -name America/New_York`

2.11 VNX Standby Operations (Local Failover)

2.11.1 Failover Concepts


Failover occurs when a standby component takes over for a failed primary component by immediately routing data to an alternate data path or device to avoid interrupting services during a failure. Data Movers and Control Stations support failover.

2.11.2 Data Mover Failover

Creating a standby Data Mover assures continuous access to file systems. When a primary Data Mover fails over to a standby, the standby Data Mover assumes the identity and functionality of the failed Data Mover.

To act as a standby server, a Data Mover must first be configured as a standby for one or more primary Data Movers. For example, you can have one standby Data Mover acting as a standby for three other active Data Movers. If one of the primary Data Movers fails over, the standby Data Mover assumes the IP and MAC addresses and functions of the failed Data Mover. The former standby is now a primary and no longer available in a standby capacity.

The standby Data Mover must have the same type of network interface cards (NICs) as the Data Mover with which it is associated. Any FTP, archive, or Network Data Management Protocol (NDMP) sessions active when the failure occurred are automatically disconnected and must be manually restarted. The following figure (Primary/Standby Relationship) shows a standby relationship.



A Data Mover restarts itself when it detects a software problem (that is, a software panic or an exception). Typically, the restart takes less than 100 seconds, and NFS applications and NFS clients do not see any interruption, except for a possible “server not responding” message during the restart.

If a Data Mover fails, the VNX Network Server clients retain normal NFS functions, but any archive or NDMP sessions are lost and not restarted. Connections between CIFS clients and the Data Mover are lost, but the redirector on the client will reconnect with the Data Mover after the failover. However, all data cached by the clients prior to failover are lost and data loss can occur. Applications on the client using the shares might not recover

2.11.3 How Standbys Work

When any failover condition occurs, you can transfer functionality from the primary to the standby Data Mover without disrupting file system availability. The standby Data Mover assumes the following identities from the faulted Data Mover:

- Network identity — IP and MAC addresses of all its NICs
- Storage identity — File systems controlled by the faulted Data Mover
- Service identity — Shares and exports controlled by the faulted Data Mover

The standby Data Mover assumes user file system services (if the policy is set to automatic) within a few seconds of the failure, transparently and without requiring users to unmount and remount the file system.

2.11.4 Choose Standby Data Movers

To configure a standby Data Mover, you must first identify and designate a standby, link the standby with a primary Data Mover, and define the failover policy for the primary Data Mover. This can be accomplished in the GUI or CLI.

To create a standby relationship between server_2 (primary) and server_3 (standby), type:

```
$ server_standby server_2 -create mover=server_3
server_2 : server_3 is rebooting as standby
```

Before any other actions can take place, the reboot must be complete. To activate the server_3 (standby) to take over for server_2 (primary), use the `server_standby` command:

```
$ server_standby server_2 -activate mover
server_2 :
server_2 : going offline
server_3 : going active
replace in progress ...done
failover activity complete
commit in progress (not interruptible)...done
server_2 : renamed as server_2.faulted.server_3
server_3 : renamed as server_2
```

2.11.5 Failover Policies

A standby Data Mover can take over functionality automatically or by manual intervention. A failover policy is a predetermined action the Control Station invokes when it detects a Data Mover failure condition. The failover policy you specify determines how the Control Station handles a Data Mover failure. The VNX Network Server offers three policies listed in the following table (Failover Standby Policy Types).

Policy	Action
Auto (Citi Standard)	The Control Station immediately activates the standby Data Mover.
Retry	The Control Station first tries to recover the primary Data Mover. If the recovery fails, the Control Station automatically activates the standby Data Mover.
Manual	The Control Station shuts down the primary Data Mover and takes no other action. The standby must be activated manually.

Auto is the CATE Storage and Backup Engineering recommended mode to prevent unnecessary downtime of CIFS services.

2.12 Control Station Failover Considerations (Dual Control Stations)

VNX Network Server software provides uninterrupted file access to users even when the Control Station is restarted, upgraded, or unavailable. All of the VNX Network Servers support configurations with dual, redundant Control Stations. The Control Station software used to

configure and manage the VNX Network Server, operates independently of the file-access operations and services provided by Data Movers.

If a Control Station fails:

- Individual Data Movers continue to respond to user requests.
- User access to data is uninterrupted.

If the primary Control Station goes offline for any reason:

- The secondary Control Station, (configured during install) automatically takes over all Control Station functions.
- The secondary Control Station executes any of the VNX Network Server commands you type.
- The VNX Network Server uses the E-mail Home utility to notify EMC Customer Service (or your service provider) of the event.

Under normal circumstances, after the primary Control Station fails over, you continue to use the secondary Control Station as the primary. After the primary Control Station is repaired or replaced and Control Stations are restarted, either directly or as a result of a power-down and restart cycle, the first Control Station to start is restored as the primary.

IP aliasing allows multiple IP addresses to point to the same network device. The VNX Network Server allows you to configure IP aliases for the Control Station network devices. You can configure an IP alias that connects to the active Control Station.

2.13 Forcing Failovers and Takeovers on Dual Control Stations

The command `cs_standby` initiates a Control Station takeover and failover. When a Control Station is activated, the name of the primary Control Station is activated.

2.13.1 Takeover

Executed only from the standby control station, this option initiates a reboot of the primary control station and then changes the state of the standby to be primary. To run this command, log onto the standby control station as root (go to `nasgcd/sbin` or `nas/sbin`):

```
./cs_standby -takeover
Taking over as Primary Control Station.....done
```

This may take several minutes to stop all of the current processes.

File system extends may be interrupted. If a primary control station fails over to a standby control station, the service continues to run, but the management capabilities are no longer available.

2.13.2 Failover

Executed only from the primary control station, this option initiates a reboot of the primary control station and then activates the standby to take over the role of the primary control station. To run this command, log onto the primary control station as root (go to `/nasgcd/sbin` or `/nas/sbin`):

```
./cs_standby -failover
The system will reboot, do you wish to continue [yes or no]: y
Failing over from Primary Control Station
```

This may take several minutes to stop all of the current processes.

Note: File system extends may be interrupted. If a primary control station fails over to a standby control station, the service continues to run, but the management capabilities are no longer available.

2.13.3 Status

To display the primary Control Station, run the following:

```
[nasadmin@c5-vnx5500-cs0 ~]$ /nas/sbin/getreason
10 - slot_0 primary control station
11 - slot_1 secondary control station
 5 - slot_2 contacted
 5 - slot_3 contacted
 5 - slot_4 contacted
```

2.14 VNX Database

In order to preserve the datamover configurations, the nasdb database needs to be backed up and copied to a remote location. Loss or corruption of this file may result in the failure of one or more datamovers. For details on how to backup this critical database, see [Backing Up the VNX Database](#).

2.15 Using Robocopy with VNX (FAT Timestamp)

The Celerra uses a FAT based timestamp even though the file system is declared as NTFS. This causes time comparison logic to fail. NTFS supports milliseconds in the timestamp. FAT only supports 2-second accuracy. For example, if you were to mirror a source (on a non-Celerra location) and destination (on a Celerra share), the timestamp comparison would always be different and you would end up copying the entire data set rather than just changed files.

Robocopy has a /FFT option/switch that accounts for this. Use the /FFT (Fat File Time) option to round time comparisons to compensate for the FAT time stamp.

```
Robocopy c:\test_data i:\test_data /MIR /Z /FFT
```

EMC Copy has similar options to Robocopy.

2.16 Quotas

Quotas are not a supported feature at Citi due to the operational burden that it presents. Use of Quota based storage management is Restricted and must have CSBE approval.


3 Storage Pools and Volume Management

The Storage Pools feature of Unisphere allows you to create and manage storage pools on the VNX Network Server.

A storage pool is an aggregation of disk storage created and managed by the Block subsystem that the VNX File sub-system uses to create and extend file systems. The VNX Network Server supports both Block Storage Pools and File Storage Pools. File Storage Pools are considered legacy and block-based pools are used from 2013 and later.

3.1 Storage Concepts: VNX vs. Filer

The following diagram shows how the EMC VNX and NetApp Filer compare in logical hierarchical concepts for storage and file systems.



From this perspective, the biggest difference between the Filer and VNX is that the VNX relies upon a separate storage system to manage and control the storage. This storage is delivered to the VNX datamovers via FC as LUNs. From that point forward, the high level concepts are very similar except for terminology and detailed component interaction/behavior. Storage pools are very similar to aggregates. Files systems are similar to volumes. This is also an oversimplification of the environments and is for comparative purposes only.

3.2 AVM (Automated Volume Management)

AVM is powerful system utility built into DART that provides for the management of storage below the file system. It handles all of the striping functions and volume management so the NAS administrator only assigns file systems to a pool. This makes for an efficient and well performing storage management. As you create file systems, AVM identifies the storage pool to use and the amount of storage to be allocated based on the size of the file system. The automatic allocation of disk volume storage to a storage pool and then to a file system eliminates the need for you to manually create stripe volumes, slice volumes, or metavolumes before creating a file system. The alternative MVM (Manual Volume Management) is not recommended as it requires a significant amount of manual planning and maintenance with little or no gain in performance.

In previous releases of this document the definition of AVM was not clear and may have been associated with the style of storage pool. CATE recommends the use of AVM on top of the pool, whether it is a File Pool or a Block Pool.

3.3 File Storage Pools (Legacy Storage Pools for File)

File Storage Pools will be supported on legacy configurations, but should not be used for new VNX configurations starting in 2013. Block Storage Pools should be used instead due to their additional features and scalability.

File storage pools (in conjunction with AVM) automatically create and manage usable file system storage and are designed to achieve a balance between performance, high availability, and ease of use. These pools manage their own expansion and shrinkage. File storage pools allow you to have control over how storage is aggregated by creating storage pools for particular needs, such as better performance characteristics or archive characteristics.

File storage pools are created by default during installation and setup with previous releases of DART. These pools are extended automatically by AVM as you create file systems from them if there are free volumes (space) available to add to the pool.

3.4 Block Storage Pools


Block storage pools are the primary method of storage on all new VNX systems starting in 2013.

Block Storage Pools are configured on the block side of the VNX and then assigned to the NAS system. They will show up as “Mapped Pools” and referred to in documentation as “Virtual Pools”, “Block Pools”, “Mapped Pools” or “Block Storage Pools”. LUNs created from these block pools are called “Pool LUNs”. It is noted that Thin Pools (with thinly provisioned LUNs) or “Fast-VP (Multi-tier) Pools” or “Mixed Pools” are not currently support for file use.


For an initial installation, part of the EMC setup will include configuration of a pool. However, it is very likely that the NAS Administrator will need to expand the pool or to create additional pools. There should be at least one pool per drive type and raid configuration. Having too many pools will lead to inefficiencies as there is likely to be unallocated space in a given pool. The following steps describe how to create a Block Pool and assign it to File subsystem:

3.4.1 Creating Block Storage Pools

1. To configure a Block Pool, go to Storage → Storage Configuration → Storage Pools.
This is different from the Storage Pools for File under Storage Configuration.



2. Select the Create button.




The pool should be named so that it describes the disks that are used. If it is a Unified system, then the assignment to NAS (or FILE) should also be included. For example:

p4_NAS_R6+2_2000GB_NLSAS_72k

(Pool id, dedicated for NAS, RAID Configuration, Disk Size, Disk Speed)

Manual or automatic disk selection can be used (automatic is the default). However, if specific disks are to be used or excluded, then manual is the best method. **If using Manual, it is important to hit the Select button (otherwise the wrong disks will likely be selected).** Automatic and Manual have three choices (Extreme Performance for SSDs/SATA Flash drives, Performance for SAS drives, or Capacity for NL-SAS drives) based on drive type. RAID 6(14+2) (preferred) or RAID 6(6+2) should be used for SAS and NL-SAS drives. SSDs are only used as FAST Cache for File, and are not configured under Pools. Unbounded disks are listed. When Manually selecting disks, all disks must be highlighted. If you select the wrong number of disks for the selected RAID type, you will get several warnings before the pool is created:



3. Hit *OK* or *Apply* to complete the pool creation. This will create the underlying RAID groups and bring them into the pool.
4. Once the pool has been created and is *Ready* on the Block Pool screen (see step 1), you will need to create (thick) LUNs and assign them to the file host group. The status may remain at *Initializing* for several minutes or hours depending on the size of the pool.
5. To Create the LUNs from the pool, to assign them to the File subsystem, go to Storage → LUNs. Select the raid type and storage pool. Calculate the number of LUNs as follows:


Number of LUNs = number of drives in pool / 4 → (rounded up to nearest 10)

For example:

#LUNs = 4x RAID Groups of R6 4+2 = 64 / 4 = 16 → rounded to 20

#LUNs = 2x RAID Groups of R6 14+2= 64 / 4 = 16 → rounded to 20


Make the user capacity MAX (end of the list), choose a LUN ID and enter the minimum of 10 (or the value obtained in the formula above) as the number of LUNs to create. Once the LUNs have been added to the file subsystem control, they will be renamed, so adding a name here is not effective.




Additional Notes:

- Use only thick (or traditional) LUNs.
- Make all LUNs the same size.
- Balance LUN ownership between SPA and SPB.
- Do not use thin LUNs.
- Do not use compressed LUNs.
- Do not have LUNs for Block mixed with LUNs for File.
- Performance oriented environments may require different RAID configurations or disk types allowable by CSBE approval only (i.e. R1+0).

6. Allow the LUNs to complete the prepare process (*Ready* state) before assigning them. To assign the LUNs to ~filestorage host group, select the LUNs and then right-click to select Add to Storage Group (or select the *Add to Storage Group* button).



Move the storage group “~filestorage” to the Selected Storage Groups list.



7. Rescan the storage system to bring new pool(s) into file control (Storage → Select Rescan Storage System from wizard list). Several rescans may be needed if not immediately present. The new pool will be listed as a mapped storage pool


3.4.2 Expanding Block Storage Pools

To expand a pool, the methodology is similar. The same RAID configuration and disk type must be used to expand an existing pool. The new LUNs will be added to the host group ~filestorage and a rescan is performed. The pool size should be increased. Several rescans may be needed.

3.4.3 Deleting Block Storage Pools

You may want to delete a misconfigured storage pool or to reuse the storage in a different pool. The following steps should be followed to remove a block storage pool:

1. Ensure there are no file systems, snapshots, vdms present on the mapped storage pool.
2. Go to Hosts → Storage Groups. Select the host storage group “~filestorage” and select the Properties button. Go to the LUNs tab. Select the LUNs that comprise the pool. Select the Remove button.



3. Select Apply or OK. You will have to confirm the action.
4. Delete the LUNs from the Storage Pool.
5. Delete the Storage Pool.
6. Rescan the storage (as in previous sections) to confirm the pool is no longer available. This may take several rescans and/or manual deletion of disks from the command line.

3.5 Supported RAID Configurations, Spares, and Vault

- If using the default number of disk volumes in a disk group would result in a LUN greater than 2TB the number of disk volumes will be incremented by the default number of disk volumes until the LUN size created is less than the 2TB.
- The following are the supported VNX RAID configurations. At Citi, the R6(14+2) is the primary recommended RAID configuration for general NAS systems. R6(6+2) is a secondary RAID configuration option. *R6 is highly recommended as it provides additional fault tolerance during the longer rebuild times associated with larger drive types.*

Disk Group Type	Disk Type	Storage Pool
14+2 RAID 6	SAS	Mapped Pool
14+2 RAID 6	NL-SAS	Mapped Pool
6+2 RAID 6	SAS	Mapped Pool
6+2 RAID 6	NL-SAS	Mapped Pool

- *Certain performance environments may require the use of different RAID configurations and or disk types with approval by CSBE on an exceptional basis.*

3.5.1 Spare Drives

The current Citi VNX for Block spare drive policy should be followed. This spare drive policy is:

- One spare drive for each disk type (larger drives of same type may be used for smaller drives)
- One spare for every 60 drives
- A minimum of two spare drives per array for each drive type

Under DART 7, the host spare drives must be defined individually as a spare drive type.

Under DART 8, a policy is set for each drive type/size. The default policy is 1 for 30 drives, so a custom policy of 1 for 60 drives must be set to comply with the CSBE standard. When using Unisphere to set the custom spare policy, 60 or more drives must be present in the array. Alternatively the command line can be used to set the correct policy regardless of the number of drives present in the array:

Gather the IP's for the SP's on the VNX

```
[nasadmin@rutctinas]$ more /etc/hosts |grep SP
10.49.85.229 A_APM00133904700 SPA # CLARiION SP
10.49.85.230 B_APM00133904700 SPB # CLARiION SP
```

Query the policy ID's to identify the Disk types to change

```
[nasadmin@rutctinas]$ nas/sbin/navicli -h 10.49.85.229 hotsparepolicy -list -
allentries
Policy ID: 2
Disk Type: SAS
Raw Capacity (MBs): 274845
Number of Disks: 4
Unused disks for hot spares: 0
Recommended Ratio of Keep Unused: 1/30
Ratio of Keep Unused: 1/30
Number to Keep Unused: 1
Recommended: Yes

Policy ID: 1
Disk Type: SAS
Raw Capacity (MBs): 549691
Number of Disks: 311
Unused disks for hot spares: 23
Recommended Ratio of Keep Unused: 1/30
Ratio of Keep Unused: 1/30
Number to Keep Unused: 11
Recommended: Yes
```

Change the policy for ID=1 to a 1/60 disk ratio

```
[nasadmin@rutctinas]$ nas/sbin/navicli -h 10.49.85.229 hotsparepolicy -set 1 -
keepunusedper 60
Done
```

Verify Change

```
[nasadmin@rutctinas]$ nas/sbin/navicli -h 10.49.85.229 hotsparepolicy -list -
allentries
Policy ID: 1
Disk Type: SAS
Raw Capacity (MBs): 549691
Number of Disks: 311
Unused disks for hot spares: 23
Recommended Ratio of Keep Unused: 1/30
Ratio of Keep Unused: 1/60
Number to Keep Unused: 6
Recommended: No
```

In addition, under DART 8, spares have become permanent and once invoked will remain a part of the pool (or RAID group) from the original spare drive location. The replacement for the failed drive will then become a spare for the next drive failure for the same disk type from the original failed drive location.

For more details around the setup of hot spares and the policy, see the [SG for VNX Block](#).

3.5.2 Vault Drives and VNX Private Space

On a VNX storage array, the first 4 drives (on the first DAE or DPE) are used by the array to hold the operating system/ related software and are called the Vault Drives. The VNX Vault Drives should be used only for the VNX vault space and the VNX File Control LUNs. If the individual vault disk size is greater than 300GB, the remaining disk space will remain unbound to avoid using the extra disk capacity for user data.

The remaining drive slots on the first DAE (or built-in slots on certain models) can be used for user data, spares, FAST Cache etc. For more information about configuring FAST Cache, see [FAST Cache Configuration](#).

The control LUNs have been moved to the VNX Private Space. Control LUNs are automatically added to a special system storage group named “~filestorage” (only visible from Unisphere). The system protects these LUNs by disallowing them from being deleted or removed from the “~filestorage” group in the GUI. Navisecli commands cannot alter these LUNs.

3.5.3 Supported Drive Types

The following are the Citi allowable drive types based on VNX storage:

Drive Size	Interface	Rotational Speed	Connectivity
300GB ³	SAS	10k,15k	6Gbps
600GB	SAS	10k,15k	6Gbps
900GB ¹	SAS	10k	6Gbps
2000GB ¹	NL-SAS	7.2k	6Gbps
3000GB ¹	NL-SAS	7.2k	6Gbps
100GB ²	SSD (SLC)	N/A	6Gbps
200GB ²	SSD (SLC)	N/A	6Gbps

Notes: ¹The use of NL-SAS drives and 900GB 10k are only supported for use as archive.

²The use of SSDs (SLC) are only supported for use as FAST Cache.

³The 300GB 10k drive is not available on the VNX NG models (5200-8000)

3.6 Storage Pool Selection

When non-gateway model is configured, the CE must be informed that we want the underlying VNX disk to be configured as R6 14+2. In every case, there will be 4x drives configured as R5. This is the space remaining on the vault drives that are used for the CLARiiON vault area and likely the location of the VNX Control LUNs. As we do not want user data to complete for I/O on these drives, any storage pool created for these drives should be avoided.

3.7 VNX LUNs on Gateway Models (VG2, VG8)

The following LUNS need to be allocated to VNXs running in a gateway (VG2,VG8) configuration. These LUNs also apply to an integrated (non-gateway) unit, but typically are configured by the CE doing initial configuration / wizard.

LUN Type	Size	Host LUN ID	SP Assignment
----------	------	-------------	---------------

Control LUNS (DART OS, etc.)	11GB	0	SPA
	11GB	1	SPA
	2GB	2	SPA
	2GB	3	SPA
	2GB	4	SPA
	32GB (VG2) or 64GB (VG8)	5	SPA
		5	SPA
Data LUNS (File Systems)	Divide RAID Group by 4	> 15	Auto (alternating)

The Control LUNS should reside on the vault drives that VNX uses for its vault area. You will need to create a RAID5 Raid Group on these drives to accommodate these LUNS. *This is the only exception to the Citi policy of using the vault drives.*

3.8 Zoning and Registering with VNX Block Storage for Gateway Models

Use of the full auto configuration of storage is not recommended, as the scripting will want to configure the zoning and storage for the VNX.

3.8.1 Zoning on Gateway Models

On VNX Block storage for Gateways, each DM fiber channel port (be0, be1 at a minimum) should be zoned to both SPs. For simplicity and to avoid overuse of zones, a single zone for each fabric is recommended.

For example:

```
Zone nsg8__cx4-480-SPA1-SPB1_A
CX4-480_SPA1 50:06:01:63:3c:e0:04:6a
CX4-480_SPB1 50:06:01:69:3c:e0:04:6a
nsg8-dm2_be0 50:06:01:60:3b:20:19:6c
nsg8-dm3_be0 50:06:01:68:3b:20:19:6c
nsg8-dm4_be0 50:06:01:60:3b:20:19:29
nsg8-dm5_be0 50:06:01:68:3b:20:19:29

Zone nsg8__cx4-480-SPA3-SPB3_B
CX4-480_SPA3 50:06:01:61:3c:e0:04:6a
CX4-480_SPB3 50:06:01:6b:3c:e0:04:6a
nsg8-dm2_be1 50:06:01:61:3b:20:19:6c
nsg8-dm3_be1 50:06:01:69:3b:20:19:6c
nsg8-dm4_be1 50:06:01:61:3b:20:19:29
nsg8-dm5_be1 50:06:01:69:3b:20:19:29
```

3.8.2 Registering the VNX within Unisphere on Gateway Models

Like any other host, the VNX File must be registered with the CLARiiON or VNX Block system. Typically this requires a manual registration of the HBAs. This "host" will have 2x ports per Data Mover. Each port is connected to 2x SPs each. Therefore you should have 4x the number of Data Movers in the host registration entry. The initiator type should be "CLARiiON Open" with a Failover Mode of 1.

For example:


```
4x Data Movers
    2x ports each
    2x SPs each
= 16 host entries in the VNX host
```

3.9 FAST Cache Configuration

The use of FASTCache is recommended when there is a need for additional peak IOPs from a VNX. This may stem from a performance sensitive environment (such as VMware). The configuration of FAST Cache from a VNX perspective is straight forward.

Unisphere allows you to choose the Flash drives to be used for FAST Cache manually or automatically. The best method is to choose these disks manually so you can distribute the Flash drives across the back-end buses. Avoid placing drives in the DPE or DAE-OS enclosure (0_0) that will be mirrored with drives in another enclosure. For example, do not mirror a disk in 0_0 with a disk in 1_0.

Once FAST Cache has been configured on the block environment, all storage pools created after that time will have it enabled by default. Therefore, if FAST Cache was added or configured after any pools were created for VNX, then these pools need to have it enabled (Storage → Storage Pools → Properties → <select pool> Advanced). FAST Cache will be a checkbox under FAST Cache.



A maximum configuration of FAST Cache for each array is listed below. The configuration of the SSDs as FAST Cache is a mirrored RAID group (R10) split into Read/Write logical partitions.

Model	Fast Cache Useable Capacity (R10)
5300	400GB
5500	1000GB
5700	1400GB
7500	2000GB
5200	600GB
5400	1000GB
5600	2000GB
5800	3000GB
7600	4200GB
8000	4200GB

FAST Cache requires that (64KB chunks of) data be accessed multiple times (approximately 3x) before being copied to the FAST Cache memory space. Therefore, seldom accessed data may not benefit from the feature. This would also not benefit very large I/O streams that are sequential, such as backups or checkpoints, because each 64KB chunk is accessed only once.

FAST Cache is most useful for handling spike in I/O, and it benefits applications that may have unpredictable I/O profiles.

FAST Cache currently supports on SLC SSD drives, so larger 400GB+ MLC SSD drives cannot be used. Under DART 8, if the utilization is < 80% then promotions occur for all requests until 80% is reached. Checkpoints are excluded from FC.

Note: The use of FASTCache with NL-SAS drives is not a recommended configuration as FASTCache is not sufficient for “upgrading” an entire tier / pool of storage.

3.10 Adding more storage to VNX

After you add disk volumes to the VNX, you must rescan the storage systems to make the additional storage available for use.

Rescanning disks is required after new storage devices are added to the configuration or when you add an entire storage system to the configuration. The simplest way to manage the rescan is to initiate a rescan from Unisphere (Storage → Rescan button). If you have configured the additional storage devices to the VNX by using the `setup_clariion` or `add_clariion` commands, the rescan is done implicitly.

3.11 File systems

A file system is a method of naming and logically organizing files and directories on a storage system. A file system on the VNX Network Server must be created and stored on a metavolume. The metavolume provides:

- Expandable storage capacity that might be needed to dynamically expand a file system

- The means to form a logical volume that is larger than a single disk

A metavolume can include disk volumes, slice volumes, stripe volumes, or other metavolumes.

The VNX Network Server creates different file systems based on how they are used. The Citi standard is UXFS.

3.12 File System Size Guidelines

The size of file systems depends on a variety of factors within a business organization. Follow these general guidelines when you create file systems in your environment:

- Consider the size of the PFS (Primary File System) and plan the backup and restore strategy. Larger file systems might take more time to back up and restore. If a file system becomes inconsistent, larger file systems might take longer to run a consistency check.
- Consider the rate at which the data grows and plan to increase the space as the data increases.
- On each Datamover, the total size of all file systems, the size of all SavVOLs used by SnapSure, and the size of all SavVOLs used by the EMC VNX Replicator™ feature must be less than the total supported capacity of the Datamover.

CATE Storage and Backup Engineering recommends keeping file systems at 650GB or smaller to keep backup and consistency checking windows to reasonable times. This is also a reflection of the maximum number of NDMP streams coming from a datamover (single stream per file system) and the shared backup infrastructure. In order to allow for

file systems up to 650GB, there is currently a backup restriction that should not be exceeded:

- **50TB per datamover if using 10GigE network**
- **32TB if using a slower network link**

3.13 File Systems and Automatically Extending File Systems

You can create or extend file systems with storage pools and configure the file system to automatically extend as needed. You can enable automatic file system extension on a file system when it is created, or you can enable and disable it at any later time by modifying the file system. The options that work with automatic file system extension are:

- HWM (High Water Mark)
- Maximum size
- Virtual provisioning

The HWM is the point at which the file system must be extended to meet the usage demand. The default HWM is 90 percent.

The default supported maximum size for any file system is 16 TB.

The maximum file system (or volume) size recommended by CATE Storage Engineering is 650GB. This allows for reasonable backup and consistency check times. An exception to this rule would be files systems that do not ever require backups. In order to support 650GB volumes, there is a backup restriction of 50TB per datamover (if using 10GigE networking).

With automatic file system extension, the maximum size is the size to which the file system could grow, up to the supported 16 TB. Setting the maximum size is optional with automatic file system extension, but mandatory with virtual provisioning. With virtual provisioning enabled, users and applications see the maximum size, while only a portion of that size is actually allocated to the file system.

Automatic file system extension allows the file system to grow as needed without system administrator intervention, making it easier to meet system operations requirements continuously, without interruptions.

3.14 CIFS

3.14.1 Share Naming Recommendation

Since DART can consolidate many CIFS server names on to one Datamover with multiple shares of the same type (data_grp, data_tmp, data_pub, apps etc.) can reside on the same Datamover, all shares will have to be suffixed with an abbreviated department name followed by an optional sequential number.

<Share type>_<Abbreviated Department Name>#

Examples: DATA_GRP_FIR, DATA_PUB_FIR, DATA_TMP_FIR, DATA_GRP_IBD3, DATA_PUB_IBD3, APPS_FIR, APPS_FIR2, APPS_EQTY1 etc.

Home shares are not an issue since all user ID's and home shares are unique and therefore the same existing standard will be used (<login id>\$).

It is both the responsibility of the NAS operations team, and the business aligned System Administrators who are creating these shares to follow these standards.

3.14.2 CIFS File Permissions

Directory and file level permissions are managed using Global Groups. File level permissions get replicated during IP Replication. If the primary and backup/DR VNX are in a different domain, it should be verified that the same trust(s) to the master user account domains exist.

Access controls are managed according to regional support structures. For example: in NA and EMEA directory permissions and access controls are managed by the respective business support system administrators, while in smaller regions Wintel Infrastructure manages both VNX operations and user access controls.

3.14.3 Using NFS and CIFS File Naming Conventions

The table below (NFS and CIFS File Naming Conventions) explains the differences in file naming conventions between the NFS and CIFS protocols.

NFS names	CIFS names
Case-sensitive.	Not case-sensitive, but they are case-preserving.
Allows a directory to hold files, names of which differ only in case.	Does not allow a directory to have two files, names of which differ only in case, because for CIFS these are duplicate names.
Note	
When creating UID and GID names for Windows clients, Windows names (usernames, domain names, and global group names) must be written in lowercase in the NIS, the password files, and the group UNIX files. You need to be careful when doing explicit user and group mapping; if spontaneously, names like Windows are written, the Celerra Network Server does not recognize them.	

3.14.4 Resolving File Naming Conflicts

Because file naming conventions differ, files created by an NFS user might occasionally present a conflict for a CIFS user. This occurs when an NFS user creates files in a single directory with names that differ only in case.

To resolve naming conflicts, the VNX Network Server distinguishes the CIFS filename by appending a numerical identifier. For example:

- An NFS file, `Aaa`, in an NFS and CIFS directory is identical to CIFS users, `Aaa`.
- A second NFS file, `aAA` in the same directory is unique to NFS users, but is presented for CIFS users for `aAA~1`.
- A third NFS file, `aAa`, is unique to NFS users, but is presented for CIFS users for `aAa~2`.

3.14.5 Creating CIFS Servers

Creating a CIFS server should not be a complicated process. Once you have confirmed that DNS, NTP are working, you should be able to create and then add a CIFS server assuming you have the rights to create objects in the AD OU for VNXs. The OU for VNXs can be almost anywhere in the domain, but should follow existing Citi AD guidelines. For example, all NAS servers may be under an OU called Servers that is part of a StorageEng directory. CIFS servers are the server names (plus share name) that users will be connecting to from their hosts.

3.14.5.1 Removing Comment for CIFS Servers ('EMC-SNAS:T7.0.14.0')

By default, there is a comment field for the CIFS server that is used when adding a server to the domain. This comment appears as:

sharename on 'EMC-SNAS:T7.0.14.0 (CIFSserver.domain.net)' (X:)

Note: This comment is based on the current version of DART on the VNX.

To remove this comment from the server before mapping any shares the following command can be used:

```
server_cifs server_2 -add compname=NS700DM2,domain=nas.net,interface=cge0 -
comment " " <--Requires single space
```

Alternatively, once a (end-user) host has already connected to a CIFS server with this comment field enabled, then the comment field can be removed from the host registry:

```
HKEY_CURRENT_USER>Software>Microsoft>Windows>CurrentVersion>Explorer>Co
mputerDescriptions
```

Edit each VNX Comp name/Netbios mapping and delete the Value data 'EMC-SN...'

3.14.5.2 Local Admins on CIFS Servers

When creating a CIFS server you will want to create a local administrator account when you add the server to the AD Domain.

Unfortunately, the VNX takes the default property of "user must change password at next login" when creating the user. To login as the local user to change the initial password, you can hit CTL+ALT+DEL (CTL+ALT+END via remote desktop) from a windows server and select change password.

Alternatively, there is an option override which allows non domain admin users to join the cifs server and be added as administrator to the server.

```
server_param server_2 -f cifs -m djAddAdminToLg -v 1
```

Without this option a non domain user is not added to the administrators of the CIFS servers. You should also remove domain admins from being local administrators on the CIFS servers. Ideally only storage administrators have access and control of the CIFS servers. Having domain administrators as local admins is a default behavior when creating the CIFS server.

3.14.5.3 Alternative CIFS Create Method: Creating the Server in AD

Note: A CIFS server add to a domain can be a challenge if the administrator does not have sufficient rights to the OU container that is being joined to. In addition, the complete OU path must be specified and in the right order.

If you cannot add the CIFS server to the domain and you have rights to do so, you may need to use the following procedure to get the server added to the domain. This method involves the pre-positioning of the server into the Active Directory (AD) structure.

1. Create the new machine account in AD (use the appropriate OU)
2. Allow for plenty of time for replication of new machine account within AD
3. Use the `-reuse` command when joining to reuse the existing entry, Here is an example for the Cards Alb, New Mexico cards site:

```
server_cifs server_2 -Join compname=nsusnewmxcea1,domain=cardsadmin.cards.citicorp.com,
admin=user123@cardsprod.cards.citicorp.com,ou="ou=CIFS-Fileers:ou=Servers" -o reuse
(You will want to replace user123 with your userid)
```

3.14.6 Using SMB Protocols (SMB 2.0, SMB 2.1, SMB 3.0)

With the deployment of Windows 2008 servers and Windows 7 desktops at Citi, the SMB 2.1 (aka SMB 2.0 with signing) protocol is strictly enforced. Windows 2012 servers are SMB 3.0 capable, but the VNX does not support this yet so the protocol will be negotiated down to 2.1. Windows Vista clients are still negotiable down to pre SMB 2.0 protocols. However, Windows 2008 Servers and Windows 7 desktops only follow SMB 2.0. All DART 7.X systems have SMB 2.1 enabled by default. The following are the commands can be used to verify this:

1. You can check the status of the max protocol allowed on the datamover by running the command:

```
[nasadmin@E8H5-vnx5700-01-cs0 ~]$ server_cifs server_2
server_2 :
256 Cifs threads started
Security mode = NT
Max protocol = SMB2
I18N mode = UNICODE
Home Directory Shares DISABLED
Usermapper auto broadcast enabled
```

2. Set the max protocol on the datamover to "SMB2" with the command (if it was set to a previous version).

```
[nasadmin@E8H5-vnx5700-01-cs0 ~]$ server_cifs server_2 -add
security=NT, dialect=SMB2
```

3. To verify the exact protocol being used by each client, use the audit option for `server_cifs`

```
[nasadmin@E8H5-vnx5700-01-cs0 ~]$ server_cifs server_2 -option
audit
server_2 :

|||| AUDIT Ctx=0x033a663c08, ref=2, W2K8 Client(10.40.59.178)
Port=60618/445
```

```

||| NASEMCTEST301[NAMDEV] on if=10-40-81-20
||| CurrentDC 0x0344f41408=NAMDEV002
||| Proto=SMB2.10, Arch=Win2K, MaxReadWriteSz=0x100000,
MaxTransactSz=0x100000, popupMsg=1
||| Client GUID=20378444-98ab-11e2-b20a-005056830595
||| SMB2 credits: Granted=31, Max=512
||| 0 FNN in FNNlist NbUsr=1 NbCnx=2
||| Uid=0x1 NTcred(0x0344d96808 RC=3 KERBEROS Capa=0x200002)
'NAMDEV\gj06529'
|| Cnxp(0x0341897c08), Name=fs_compress_quick, cUid=0x1 Tid=0x2,
Ref=1, Aborted=0
| readOnly=0, umask=22, opened files/dirs=1
| Absolute path of the share=\\fs_compress_quick\\compression_data
|| Cnxp(0x034112e008), Name=fs_compress_deep, cUid=0x1 Tid=0x3,
Ref=1, Aborted=0
| readOnly=0, umask=22, opened files/dirs=1
| Absolute path of the share=\\fs_compress_deep\\compression_data
total smb_streamCtx: 2

```

If a user on a SMB 2.0 enforced system tries to map to a CIFS server that has a max protocol set to NT1, then they will likely get an error stating “The account is not authorized to log in from this station” and will not be able to connect to the share. Once the setting has been changed on the datamover, the user will be able to connect to the share as normal.

3.14.7 Renaming a CIFS Server (Changing Compname)

CIFS servers can be renamed following a special procedure. Using this method (see emc156835 for more details) will preserve local groups, shares, and the file system.

1. Unjoin the original compname from the domain:

```
# server_cifs server_2 -Unjoin compname=W2k3,domain=abc.com,admin=Administrator
```

2. Delete the compname from the Data Mover's CIFS configuration:

```
# server_cifs server_2 -delete compname=W2k3
```

3. Add the name back to the Data Mover's CIFS configuration as a Netbios name:

```
# server_cifs server_2 -add netbios=W2k3,domain=abc,interface=fsn01
```

4. Rename the Netbios server to the new name:

```
# server_cifs server_2 -rename -netbios W2k3 W2k3prod
```

5. Delete the newly renamed NetBIOS name from the Data Mover's CIFS configuration:

```
# server_cifs server_2 -delete netbios=W2k3prod
```

6. Add and join the new Compname to the CIFS configuration and AD domain:

```
# server_cifs server_2 -add compname=W2k3prod,domain=abc.com,interface=fsn01
# server_cifs server_2 -Join compname=W2k3prod,domain=abc.com,admin=Administrator
```

The resulting CIFS configuration will be that the CIFS server compname has been changed but the ipaddress and existing interface will not change.

3.14.8 File Extension Filtering

DART 7.0, 8.0 (and above) provide storage system support for simple native file blocking functionality. This feature allows storage administrators to control the types of files being stored on the EMC NAS using file extension. This functionality is covered in this document to make storage administrators aware of this feature. It is intended for use by businesses who wish to use this functionality to control what file types are saved to their shared storage. This is an optional configuration.

3.14.8.1 Overview

File screening policies are used to specify files or directories with restrictions to be placed on them. Upon receiving a file operation request (such as open, write, create, or rename), DART checks its file screening policies and file access control list (ACLs) before permitting the operation. If the policy specifies screening for that file based on its extension, file screening takes place on the storage system with native file blocking. From a user perspective the request is denied and the file operation is blocked.

File extension filtering runs at the share level on each CIFS Server. Each CIFS Server share would need to be configured with file extension filtering as well as the proper ACLs to deny write access. File extension filtering policies enabled on a single share will not impact or be inherited by other shares running on the same system.

Note: File extension filtering only works for SMB/CIFS environments. For multiprotocol environments, NFS clients will still have the ability to write to the storage via an NFS mount.

3.14.8.2 Limitations of file extension filtering

Native file blocking is a simple approach to preventing unwanted file based data from being stored on an EMC NAS Device. It is a low cost solution as it does not require any additional infrastructure to implement however the tradeoff is that it lacks management and reporting functionality. If you have a requirement for the following features then native file blocking may not be the appropriate solution for your environment.

Native file blocking scans the extension of the file only. It does not scan the file header to determine if it should be blocked or if the file extension does not reflect the true content of the file. File extension filtering only works for SMB/CIFS access. If your environment is mixed CIFS and NFS, NFS users are not affected by file extension filtering and are able to store prohibited files on the file system.

3.14.8.3 Implementation for File Extension Filtering

Use the following steps to implement native file blocking on the EMC NAS platform.

In this example a policy will be created to stop txt files from being saved to a share "emctest96fs2" on CIFS server "nasemctest96".

1. Enable File Filtering parameter to enable the file screening process

Note: Datamover will need to be rebooted to load new parameters.

Verify existing configured value:

```
[E8H5-vnx5700-01-cs0 ~]$ server_param server_2 -facility cifs -info enableFileFiltering
server_2 :
name           = enableFileFiltering
facility_name    = cifs
default_value   = 0x00000007
current_value   = 0x00000007
configured_value =
user_action     = none
change_effective = immediate
range           = (0x00000000,0xffffffff)
```

description = Bit vector to enable various forms of file filtering

Change the value to “7” to enable filtering, pop-up messages, and auditing enabled.

```
[E8H5-vnx5700-01-cs0 ~]$ server_param server_2 -facility cifs -modify enableFileFiltering -
value 7
server_2 : done
```

Verify configured value has been applied.

```
[E8H5-vnx5700-01-cs0 ~]$ server_param server_2 -facility cifs -info enableFileFiltering
server_2 :
name           = enableFileFiltering
facility_name    = cifs
default_value   = 0x00000007
current_value   = 0x00000007
configured_value = 7
user_action     = none
change_effective = immediate
range          = (0x00000000,0xffffffff)
description     = Bit vector to enable various forms of file filtering
```

2. Connect to the root file system of CIFS Server “[\\cifs_server\c\\$](#)” and navigate to `\.etc\filefilter` folder
3. Create a text file with the restriction requirements without the .txt extension.


Create the filter files with the following syntax:

<extension_name>[@<sharename>][@<netbios_name>]

where:

<extension_name> = file extension that you want to filter.


<sharename> = name of the share to which you want to apply the filter. The <sharename> is an optional part of the filename. If you do not include a <sharename>, the filter is applied to all shares on the Data Mover.



Sample Filter Files


4. Configure exceptions to the filter policy by configuring the filter file's ACL via the file's Properties sheet.

Remove all group or users to apply a blanket restriction for the share.



Restriction for all users

To restrict only certain groups, the users or groups would need to be added as well as explicit security privileges would need to be applied for authenticated users.




Restriction for specific groups

3.14.8.4 End user experience

DART makes use of the access denied RPC to advise the end user that the file blocking process has occurred. This can be somewhat confusing for the end user as native file blocking only has a limited number of RPCs that can be used to communicate the denial of operation back to the client.

An example of the error message experienced by a Windows client when a user tries to copy an txt file (test.txt) to a share with the .txt file extension filter policy enabled (using the example from the previous section) is shown below:



File blocked client message

3.15 NFS

3.15.1 NFS Overview

The NFS protocol is used to provide access to data over an IP network, typically to Unix or Linux clients and is comparable to the functions provided by CIFS for Windows clients.

Summary Standards/Recommendations

- LDAP should be used for NFS configurations. The use of local files is not recommended.
- The initial certification of NFS on EMC NAS is limited to NFSv3. NFSv4 guidelines will be released at a later date.
- LDAP can now be configured on a VDM and is a requirement if the VDM is intended to be configured for an LDAP name space that is different from the Data Mover hosting it
- All file systems to be exported via NFS should be placed inside a VDM
- ALL NFS exports should be presented directly on the VDM.
- For NFS only VDM's, the associated interface should be attached to the VDM
- Operational procedures should be put in place to maintain consistency of NFS configuration files and settings between Production NAS and COB NAS.
- Operational procedures should be put in place to ensure that the usermapper database is backed up and securely stored offsite on at least a daily basis so that in the event of a disaster the usermapper database is recoverable.
- If requiring use of ntxmap, operational procedures should be put in place to ensure that the ntxmap.conf file is made available on the DR NAS to ensure correct multi-protocol mappings during a DR invocation

Known Limitations

- If a VDM containing NFS data is migrated to another physical DM then this will require manual reconfiguration of both the NAS and the clients.

Summary of Configuration Steps

- Configure LDAP name service on Data Mover and VDM (if required)
- Configure Name Service Switch (nsswitch.conf) on Data Mover and VDM (if required)
- Set highest NFS version to NFSv3
- If also providing access through CIFS then configure settings for multiprotocol access.
- Create VDM's as required to group NFS and Multi Protocol File Systems
- Create file system to be exported and set access mode as required
- Configure ntxmap (only complete in circumstances where additional mapping between Unix and Windows accounts is required)
- Export NFS File systems with appropriate access settings

3.15.2 LDAP name service

Before NFS can be used on EMC NAS a number of other facilities need to be configured as NFS is reliant on these functions to function in a Citi environment.

These facilities include:

- LDAP
- Naming service configuration (nsswitch.conf file)

This sub section describes how these facilities need to be configured to support NFS before moving on to the configuration of NFS itself.

The use of a local passwd file on the EMC NAS in place of LDAP connectivity is a non-standard CATE configuration. The document <https://catecollaboration.citigroup.net/domains/platstor/gdse/products/platforms/EMCCelerra/GDS E-11008-DG.pdf> details the use of local files in place of LDAP on EMC NAS **where LDAP is currently unavailable**. Approval from CATE Storage and backup Engineer as well as local BISO is required for this non-standard configuration.

When LDAP becomes available the EMC NAS should be reconfigured to use LDAP as soon as is practical.

LDAP is used to provide a central repository for netgroups (used to grant or restrict access to NFS exports) and also to lookup user and group information in a multi-protocol environment where Unix and Windows users must be translated to determine file access. Once certified, NFSv4 will also require the use of LDAP to lookup user and group data for NFS access as translation of UID and GID values take place in the storage rather than at the client in NFSv4.

With the latest code release, LDAP is now configurable on both datamover (DM) and VDM. Each VDM can only be configured with one set of LDAP configuration information, while each physical DM can contain multiple LDAP configuration entries. Before configuring LDAP on the VDM, the configuration must first be added into the DM hosting the VDM or the configuration will not be successful..

Citi has separate LDAP infrastructure for Windows (AD is based on LDAP) and Unix systems. This LDAP configuration discussion will focus solely on configuration of an EMC NAS into the Unix LDAP infrastructure. The Citi layout of the Unix LDAP is configured in each region as follows:




Figure 1 Citi LDAP layout

Typically there are separate infrastructures for different classes of server (prod, uat, dev) and within each of these infrastructures each businesses data is segregated into its own area. Each business has their own areas within the LDAP as some data is not unique between businesses (e.g. non-unique UID or GID values) and hence must be segregated to ensure unique answers during lookups. For this reason storage for a particular business must point to the top of the LDAP tree for that particular business rather than the overall top of the LDAP tree.

Citi Unix LDAP connections do not require ssl encryption as the user password data is held within Kerberos and a higher level of security is not required, hence no bind information is required to gain access to the Unix LDAP.

With the latest code release for EMC NAS, the Citi layout can now be configured on the NAS by configuring multiple LDAP locations on different VDMs hosted within a single physical datamover.

3.15.3 Configuring LDAP

LDAP cannot be configured from the EMC NAS management GUI and hence must be configured directly from the CLI. Not all settings can be set from the command line and several items such as the path of netgroup data and any specific search limit can only be configured via the `ldap.conf` file that must be copied to the physical datamover and VDM.

The following steps are required to configure LDAP on a per physical datamover or VDM basis.

- Copy `ldap.conf` configuration file to the physical datamover and VDM
- Configure LDAP on physical datamover using “`server_ldap`” and on VDM using “`server_nsddomains`”.
- Copy `nsswitch.conf` configuration file to the physical datamover and VDM

To configure LDAP the following mechanisms are used to set the information:

Information	Mechanism	Parameter name
LDAP servers and port (DM)	<code>server_ldap</code>	<code>-servers</code>

		<server_name:port#>:<server_name:port#>
LDAP base DN (DM)	server_ldap	-basedn <base_dn>
LDAP base DN (VDM)	server_nsdomains	-resolver LDAP=<base_dn>
Location of user data	ldap.conf	nss_base_passwd
Location of group data	ldap.conf	nss_base_group
Location of netgroup data	ldap.conf	nss_base_netgroup
Disable ssl encryption	server_ldap	-sslenabled n

3.15.3.1 ldap.conf configuration file

The ldap.conf file has the same format as a standard Unix ldap.conf file, any parameters in the file that are not understood by the EMC NAS will be ignored.

Below are the contents of an example ldap.conf file

```
nss_base_passwd ou=people,ou=infra,dc=wlab,dc=net?one
nss_base_group ou=group,ou=infra,dc=wlab,dc=net?one
nss_base_netgroup ou=netgroup,ou=infra,dc=wlab,dc=net?one
```

Each of the data paths has a search scope of “one” set, this limits the amount of searching that LDAP will perform to a single level in the tree and is the recommended setting.

To configure LDAP on physical DM:

Create an appropriate ldap.conf file and copy this file to the physical datamover. In the example below the file containing the ldap.conf data is ldap.conf.server_3 and is copied to datamover “server_3”.

```
[nasadmin@cs0]$ server_file server_3 -put ldap.conf.server_3 ldap.conf
server_3 : done
```

Once the ldap.conf file has been copied to the datamover the rest of the ldap information may be set but first the existing ldap data must be cleared in case there are any previous settings.

```
[nasadmin@cs0]$ server_ldap server_3 -clear
Server_3:done
```

If LDAP has never been configured then this will error as shown below, this is expected:

```
[nasadmin@cs0]$ server_ldap server_3 -clear
server_3 :
Error 4020: server_3 : failed to complete command
```

LDAP may be initialized with the rest of the LDAP information. Below is an example configuring LDAP. This configures a base LDAP value of “dc=infra,dc=wlab,dc=net”, servers of 10.40.48.154 and 1.2.3.4 both using port 389 with ssl encryption disabled.

```
[nasadmin@cs0]$ server_ldap server_3 -set -basedn dc=infra,dc=wlab,dc=net -
servers 10.40.48.154,1.2.3.4:389 -sslenabled n
server_3 : done
```

You may check that LDAP configuration via the ldap_server command:

```
[nasadmin@cs0]$ server_ldap server_3 -info
server_3 :
LDAP domain: infra.wlab.net
base DN: dc=infra,dc=wlab,dc=net
```

```
State: Configured - Connected
Configuration file - TTL: 1200 seconds
Next configuration update in 1175 seconds
DIT schema type: OPEN
Connected to LDAP server address: 10.40.48.154 - port 389
SSL enabled/disabled by Command line, cipher suites configured by default
```

Note: The normal status does not show the detailed LDAP information (e.g the settings found in the ldap.conf file) to return this information run “server_ldap server_3 -info -verbose”

In addition, if there is going to be one or more VDMs hosted on the DM which will be in a different LDAP namespace from the DM, then it is a requirement to **add** the LDAP configuration information into the DM.

Below is an example of ldap.conf file with new LDAP configuration (infra1.wlab.net) added:

```
nss_base_passwd ou=people,ou=infra,dc=wlab,dc=net?one
nss_base_group ou=group,ou=infra,dc=wlab,dc=net?one
nss_base_netgroup ou=netgroup,ou=infra,dc=wlab,dc=net?one
nss_base_passwd ou=people,ou=infra1,dc=wlab,dc=net?one
nss_base_group ou=group,ou=infra1,dc=wlab,dc=net?one
nss_base_netgroup ou=netgroup,ou=infra1,dc=wlab,dc=net?one
```

After copying the above ldap.conf file into the physical DM, the new LDAP configuration will need to be initialized. Below is an example using the new “server_ldap -add” option.

```
[nasadmin@c5-vnx5500-cs0 .etc]$ server_ldap server_2 -add -basedn
dc=infra1,dc=wlab,dc=net -servers 10.40.48.154:389 -sslenabled n
server_2 : done
```

You may check that the new LDAP configuration is in place via the ldap_server command:

```
[nasadmin@c5-vnx5500-cs0 .etc]$ server_ldap server_2 -info -all
server_2 :
LDAP domain:      infra.wlab.net
State:            Configured - Connected
Schema:           Sun Directory Server (iPlanet) (Sun-Java(tm)-System-
Directory/6.3)
Base dn:          dc=infra,dc=wlab,dc=net
Bind dn:          <anonymous>
Configuration:    File: ldap.conf - TTL: 1200 s
LDAP conf server: 10.40.48.154 - Port: 389
SSL:              Not enabled
LDAP domain:      infra1.wlab.net
State:            Configured - Connected
Schema:           Sun Directory Server (iPlanet) (Sun-Java(tm)-System-
Directory/6.3)
Base dn:          dc=infra1,dc=wlab,dc=net
Bind dn:          <anonymous>
Configuration:    File: ldap.conf - TTL: 1200 s
LDAP conf server: 10.40.48.154 - Port: 389
SSL:              Not enabled
```

To configure LDAP on VDM:

If the VDM is in the same LDAP namespace as its hosting DM, then it is optional to configure LDAP on the VDM. LDAP configuration is only required if the VDM will be in a different LDAP namespace from its hosting DM.

Copy the ldap.conf file containing the relevant LDAP information into the VDM as shown:

```
[nasadmin@c5-vnx5500-cs0 .etc]$ server_file emcnasvdm2 -put ldap.conf.emcnasvdm2
ldap.conf
emcnasvdm1 : done
```

Initialize the LDAP configuration on VDM using the `server_nsdomains` command:

```
server_nsdomains emcnasvdm2 -set -resolver ldap=infra1.wlab.net -resolver
dns=namdev.nsrootdev.net
emcnasvdm2 : done

[nasadmin@c5-vnx5500-cs0 .etc]$ server_nsdomains emcnasvdm2 -enable
emcnasvdm2 : done
```

You can check that LDAP is now configured on the VDM:

```
[nasadmin@c5-vnx5500-cs0 .etc]$ server_nsdomains emcnasvdm2
emcnasvdm2 :
NSDOMAINS CONFIGURATION = Enabled

NIS domain          : <not defined>
LDAP domain         : infra1.wlab.net
DNS domain          : namdev.nsrootdev.net
NFSV4 domain        : <not defined>
```

LDAP is now configured and running on the physical datamover and VDM. However the datamover and VDM must now be instructed to use LDAP for lookups. This is achieved via the `nsswitch.conf` file which again must be copied to the physical datamover and VDM.

3.15.3.2 Name Service Switch configuration file (`nsswitch.conf`)

The `nsswitch.conf` file should contain the following information:

```
passwd:    files ldap
group:     files ldap
hosts:     files dns
netgroup:  files ldap
```

This file instructs the data mover to search its local files first and then DNS for host information and LDAP for user, group and netgroup information. By default there is no `nsswitch.conf` and the DM will check the local files followed by NIS for passwd, group and netgroup entries, for hosts the DM queries local files, DNS and then NIS in that order. Therefore when using LDAP the `nsswitch.conf` file must be provided.

The `nsswitch.conf` file is copied using the same mechanism as for the `ldap.conf` file, and example of which is shown below:

```
[nasadmin@cs0]$ server_file server_3 -put nsswitch.conf.server_3 nsswitch.conf
server_3 : done
```

The `nsswitch.conf` file will be used automatically once it is written to the DM. If the VDM is also configured for LDAP, then the `nsswitch.conf` needs to be copied to the VDM using the same command above.

LDAP configuration and NFS exports information on VDM is replicated to the DR VDM as part of VDM replication. However the LDAP configuration on DM is not replicated to the DR NAS and hence LDAP on the DR NAS DM must be configured independently.

3.15.4 Debugging/Testing LDAP

Once LDAP has been configured then it is possible to test the LDAP configuration by directly querying LDAP using the “server_ldap” command. Below are examples of queries to LDAP from a DM for users, groups and netgroups for DM server_3. VDM configured with LDAP can also be queried using the same command.

```
[nasadmin@cs0]$ server_ldap server_3 -lookup -user sh939001
server_3 :
user: sh939001, uid: 939001, gid: 9001
[nasadmin@cs0]$ server_ldap server_3 -lookup -uid 939001
server_3 :
user: sh939001, uid: 939001, gid: 9001
[nasadmin@cs0]$ server_ldap server_3 -lookup -group nasgrp9000
server_3 :
Group name: nasgrp9000, gid: 9000
Group members:
[nasadmin@cs0]$ server_ldap server_3 -lookup -gid 9000
server_3 :
Group name: nasgrp9000, gid: 9000
Group members:
[nasadmin@cs0]$ server_ldap server_3 -lookup -netgroup filerldaptest1
server_3 :
Netgroup: filerldaptest1 - triples:
(nasldaptest01,-,)
[nasadmin@cs0]$ server_ldap server_3 -lookup -netgroup filerldaptest3
server_3 :
Netgroup: filerldaptest3 - triples:
(gdsenfsv4-1,-,)
(gdsenfsv4-2,-,)
```

Note: The nsswitch.conf file does not need be configured before testing LDAP. A successful lookup using server_ldap does not prove that the nsswitch.conf is configured correctly. The server_ldap command **does not** support the querying of nested netgroups but this will function correctly when used for NFS exports.

3.15.5 NFSv3 and NFSv4 configuration

EMC VNX storage systems supports and GDSE have certified both NFSv3 and NFSv4 protocol for use within Citi.

All NFSv3 exports will use the AUTH_SYS authentication mechanism, the use of Kerberos authentication for NFSv4 exports will be certified in a later phase of the current NFSv4 certification project.

NFSv4 has improvements in the following areas:

- Support for Access Control Lists (ACLs) – Each ACL being made up of one or more Access Control Entries (ACE)
- Mandatory file locking and automatic lock timeouts in the event of server failure

NFSv4 is currently being certified in a phased approach and the functionality detailed in this solutions guide relates to phase 1 of the NFSv4 certification project. Additional NFSv4 functionality will be added in later phases of the project and this solution guide will be updated as this work is completed.

The following NFSv4 functionality is currently certified for use within Citi:

- NFSv4 access to NFS only data areas.

- Concurrent access to the same data using both NFSv3 and NFSv4
- Support for NFSv4 ACLs
- Authentication of NFSv4 exports using AUTH_SYS, as used for existing NFSv3 data (Kerberos authentication is planned for later in the NFSv4 certification project)
- Support for NFSv4 read and write delegations

All volumes to be exported via NFS should be placed inside a VDM. NFS volumes may reside in the same VDM as existing CIFS shares, indeed this is required for multi-protocol access to the same data.

NFS exports information on VDM is stored in a vdm.cfg file which is replicated as part of VDM replication so that re-exporting is not required when failover to DR. Some NFS configuration files and settings must be configured on the physical DM which **are not** replicated to the DR NAS as part of any VDM replication and alternative procedures to synchronize NFS configuration files and settings must be put in place.

Warning: If a VDM containing NFS data is migrated to another physical DM then this data will become unavailable from NFS and will require manual reconfiguration of both the NAS and client to reconnect clients with their data.

It is possible that data needs to be migrated between physical datamovers to balance the load on an EMC NAS, however whilst moving a VDM containing purely CIFS data is a simple task (as all configuration files are contained within the VDM container) the same is not true when NFS is used. If a VDM is migrated to another physical DM in the same NAS then the following steps will be required to make that data available again.

- Dismount NFS exports on clients (The NFS exports will be exported from a different server/IP address after the migration)
- Check NAS parameter values on destination physical DM match source DM (i.e. parameters changed to support multi-protocol, LDAP configuration).
- Move VDM to new physical DM
- Recreate NFS exports on new physical DM
- Change client NFS mount entries to point to new mount point (at a minimum the server that will be providing the NFS exports will change)

During a local failover (due to the failure of a physical datamover) the standby datamover takes on the personality of the failed datamover (including all its configuration settings) and hence no manual changes are required in this scenario.

During a failover to a remote location (e.g. the invocation of COB) the NFS exports configuration stored in vdm.cfg is replicated as part of VDM replication and immediately available on failover. However, filesystems are not replicated and must be separately configured.

An example is shown below where failover of VDM and filesystem were done. Note the read-only permission on the filesystem and error when querying exports configuration on VDM before failover. After failover, the filesystem permission changed to read-write and exports configuration became available:

Before Failover:

```
[nasadmin@E8H5-vnx5700-01-cs0 .etc]$ server_mount emcnasvdm1
emcnasvdm1 :
nfs_unix3 on /nfs_unix3 ufs,perm,ro,accesspolicy=NATIVE,rwlock

[nasadmin@E8H5-vnx5700-01-cs0 .etc]$ server_export emcnasvdm1
emcnasvdm1 :
```

```
Error 4032: emcnasvdm1 : VDM is not loaded
```

After Failover:

```
[nasadmin@E8H5-vnx5700-01-cs0 .etc]$ server_mount emcnasvdm1
emcnasvdm1 :
nfs_unix3 on /nfs_unix3 udfs,perm,rw,accesspolicy=NATIVE,rwlock

[nasadmin@E8H5-vnx5700-01-cs0 .etc]$ server_export emcnasvdm1
emcnasvdm1 :
export "/nfs_unix3" sec=sys ro=filelrdaptest1 rw=filelrdaptest1
root=filelrdaptest1
```

Due to a change in IP address (even if 3DNS is used) any NFS mount points will need to be unmounted and remounted to regain access to the data on the COB NAS. Normal reconfiguration, such as making the COB volumes read-write may also be required. It is important to remember that while NFS configuration on VDM is replicated, all configuration on the physical DM is never replicated by the NAS and it is the storage administrator's responsibility to independently configure both the production and COB locations. Procedures must also be put in place to ensure that changes to the production DM NFS configuration are captured and re-created on the COB NAS.

3.15.5.1 NFSv4

NFSv4 is the latest standard for NFS and has improvements in the following areas:

- Support for Access Control Lists (ACLs) – Each ACL being made up of one or more Access Control Entries (ACE)
- Mandatory file locking and automatic lock timeouts in the event of server failure

If NFSv4 is enabled on a data mover or virtual data mover then all exports from that allocation can be mounted using either NFSv3 or NFSv4, it is not possible to permit or deny NFSv4 (or NFSv3) support on individual exports.

Conversion of existing physical data mover or virtual data mover to use NFSv4 must take the following items into account or loss of access to data may result:

- File system needs to be converted to Mixed Mode
- Data mover parameters need to be configured
- All storage providing NFSv4 must be configured to use LDAP
- The “root” user and groups must be defined in the local passwd and group files on the data mover/ virtual data mover or root ownership will not display correctly.

In addition to the above points NFSv4 has the following architectural differences with earlier versions of NFS:

- NFSv4 is now stateful (like CIFS) however the integration of NFS locking into NFSv4 actually improves the robustness of the protocol compared with NFSv3. It would also be incorrect to conclude that the move to a stateful protocol will adversely impact clients during failover as can be the case with in certain circumstances with CIFS.
- There are no longer ancillary protocols such as locking; all NFS related functions are now contained within the same protocol.
- IP port 2049 is used for all NFSv4 traffic.
- NFSv4 now makes used of compound operations with multiple operations being sent in

the same request, thus reducing traffic. It should be noted that whilst this may speed up some workloads, for those workloads with a high component of non-IO operations (i.e. meta-data operations) there may be a reduced performance compared with NFSv3.

- NFSv4 is now much closer in functionality to the CIFS protocol with the addition of extra operations such as the “open” call.
- UDP is not supported for NFSv4 hence only NFSv4 over TCP is now supported.

The standard for the NFSv4 ID domain name is to use the string “nfsdomain”, this value is to be used for all Citi locations. The same value will be set on all Citi OS platforms looking to connect using NFSv4.

NFSv4 delegations allow clients to cache file data. Whilst NFSv3 also had caching mechanisms this was based on polling of file last modification times. NFSv4 moves from this polling method to an interrupt based model where delegations are requested by clients for specific files and until these delegations are either given up by the client or revoked by the storage the cache can be guaranteed to be consistent without the need to poll the files status. Delegations are maintained similarly to NFSv4 locking with clients periodically requiring to renew their delegations. Clients must provide some form of callback daemon that is used by the storage to revoke delegations when required; if this callback daemon cannot be contacted then the delegation will not be granted. Delegations are granted on a per file basis and are requested at file open time via the NFSv4 OPEN system call.

Below is a summary of the two types of delegation available.

Read delegation

There may be multiple read delegations granted by the storage to the same file provided that no client either has the file open for write or “deny read access”. If a client holds a read delegation for a file then it knows that no other client can change the file and hence its cache is always valid. So unlike NFSv3 once a client holds a read delegation there is no requirement to check via `getattr` that the cache is still valid. If another client subsequently requests an open for write or “deny read access” (or non NFSv4 access to the file is attempted) then the read delegation is revoked by the storage from all clients with read delegations, the client has to comply and give back its delegation.

Write delegation

There can only be a single write delegation at a time and it cannot be granted if the file is open from read (or write) by another client. Normally (without delegations) once a file is closed any modified data must be flushed to the NFS server but this is not required by the NFS standard if write delegation is held; however the client must tell the server how much space to reserve for the file when the data is eventually flushed. Investigations with several Unix/Linux vendors have shown that they treat NFSv4 caching similar to NFSv3 caching with data being periodically written to the storage in the same way that normal block IO is flushed regularly even though this is not mandated by the NFS standard. If a write delegation is either voluntarily given up or revoked by the server then the client must flush any modified blocks to the server before handing back its delegation. If a write delegation is held for a file then any locking for that file is now managed by the client holding the write delegation.

If no delegations are held by any client then the client should work in a similar way to NFSv3 with regard to caching with the polling of modification times. It is also worth noting that if an application has a lock on a file when it gives up its lock it must also flush any modified data.

As delegations are functions of the file open call within NFSv4 they will only be used when specifically coded in an application. The way in which a file is being used should determine whether delegations are requested as it may be inappropriate for some patterns of file e.g. access from multiple clients. For critical data it is recommended that write delegation is not used to

ensure that minimal data is lost should a client crash before flushing data written to a file. Read delegation may not be appropriate where another client is periodically writing to the same file due to the need to revoke delegations.

In NFSv3 and earlier the client passed a user's credentials to the storage as numerical UID and GID values, with NFSv4 the client now provides the username and groups as strings. These strings must be looked up by the NFSv4 server (via LDAP or local files) to determine the resulting UID and GID values of the user and hence their access to file/directories.

As the same value for the NFSv4 id domain name must be set on both the server and client by using a derivation of the LDAP entry point this ensures that both client and server are also using the same LDAP and hence will both translate users, groups and netgroups correctly.

3.15.5.2 NFSv4 ACL

One of the most visible changes in NFSv4 is the availability of ACLs. Previously the NFSv3 standard did not support ACLs; although some vendors such as Sun produced their own non-standard extensions to NFSv3 to support ACLs. Vendor extensions to NFSv3 to support ACLs typically used Posix ACLs and these have a different syntax and granularity than NFSv4 ACLs, hence Posix ACLs do not have a one-to-one mapping with NFSv4.

If a file or directory does not have ACLs defined the client tools used to view ACLs will interpret the Unix permissions mask as an ACL. Below is an example of the ACL shown on a file (from Linux) with a Unix permissions mask of 330 but with no actual ACL set

```
$ nfs4_getfacl testfile
A::OWNER@:rwatTnNcCy
D::OWNER@:x
A:g:GROUP@:rwatTnNcCy
D:g:GROUP@:x
A::EVERYONE@:tc
D::EVERYONE@:rwaxTCy
```

3.15.5.3 NFSv4 and NFSv3 coexistence

It is not possible to limit a particular NFS export to either NFSv3 or NFSv4 as the only controls are at the data mover or virtual data mover level. Therefore concurrent access to the same data using NFSv3 and NFSv4 has been tested and is certified for use.

Below is the behavior of concurrent access to data using NFSv3 and NFSv4:

- NFSv3 access to files or directories with NFSv4 ACLs set will obey these ACLs but NFSv3 clients will not be able to view or change these ACLs.
- Creating files in sub-directories from NFSv3 in a parent directory with inherited NFSv4 ACLs defined will caused the created files to have the ACLs attached to the files/sub-directories in the same way as for an NFSv4 client file create.
- NFSv3 and NFSv4 locks are visible to each protocol and locks on one protocol are honored on the other.
- NFSv3 cannot perform byte range locking hence it will lock whole files, NFSv4 can perform byte range locking. Partial byte range locking in NFSv4 will prevent an NFSv3 gaining a lock, the reverse is also true.

- If an NFSv3 client crashes without releasing its locks then the lock will remain and must be cleared manually as for any other NFSv3 lock, this may impact NFSv4 clients. An NFSv4 clients locks will automatically be removed if an NFSv4 client crashes without releasing its locks.
- As files and directories have owners and groups defined on files using numeric UID/GID's both NFSv3 and NFSv4 can access the same files successfully.

3.15.5.4 Useful NFSv4 Host information

During the course of the NFSv4 certification testing the following information was gathered that may be useful to those using NFSv4.

RHEL Linux

- NFSv4 domain ID is set in the file "/etc/idmapd.conf"
- Requires module nfs4-acl-tools to manipulate NFSv4 ACLs
- Set NFSv4 ACLs using "nfs4_setfac"
- "nfs4_setfac" cannot set empty ACE's.
- nfs4_getfac and nfs4_setfac do not except wildcards for the filename.
- View NFSv4 ACLs using "nfs4_getfac"
- Using Posix "setfac" or "getfac" commands will return errors as NFSv4 does not support the Posix ACLs used by these commands.

3.15.5.5 Data mover configuration to enable NFSv4

The following options should be set on the data mover to be configured to use NFSv4

Configure NFSv4 domain name

```
[b1-vnx7600-cs0 ~]$ server_param server_4 -f nfsv4 -m domain -v nfsdomain
server_4 : done

[b1-vnx7600-cs0 ~]$ server_param server_4 -f nfsv4 -l
server_4 :
param_name          facility default  current  configured
numericId           nfsv4    1        1
IDCacheEnable       nfsv4    1        1
maxOwners            nfsv4    52428    52428
IDCacheSize         nfsv4    512      512
vnodePercent        nfsv4    80       80
leaseDuration       nfsv4    18       18
domain              nfsv4    nfsdomain nfsdomain
removeOpened        nfsv4    1        1
IDCacheTTL          nfsv4    10       10
advisoryLocking      nfsv4    0        0
ntAllSpecialUsers   nfsv4    0        0
```

Start NFSv4 Services

```
[b1-vnx7600-cs0 ~]$ server_nfs server_4 -v4
server_4 :
----- nfsv4 server status -----
      * Service Stopped *

[b1-vnx7600-cs0 ~]$ server_nfs server_4 -v4 -service -start
server_4 : done

[b1-vnx7600-cs0 ~]$ server_nfs server_4 -v4
server_4 :
----- nfsv4 server status -----
      * Service Started *
      * pNFS service Started *
      * (yet operating) *
----- NFSv4 Clients -----
Confirmed Clients      : 0
UnConfirmed Clients    : 0
Number of users        : 0
Number of lock owners  : 0
Longest List           : 0
Shortest List          : 0
Greatest depth to date : 0
Average List length    : 0.00
Domain Name            : nfsdomain
-----
----- NFSv4 State -----
Opens      : 0
Locks      : 0
Delegations: 0 (0 expired)
Layouts    : 0
OPEN stateCount = 0 max 0
LOCK stateCount = 0 max 0
DELEGATION stateCount = 0 max 0
DIRDELEGATION stateCount = 0 max 0
LAYOUT stateCount = 0 max 0
Free       : 0
-----
```

3.15.5.6 General configuration.

Before configuring NFS first configure LDAP as LDAP is used to lookup netgroup information in NFSv3 and to lookup netgroup, user and group data for NFSv4. LDAP is also used in multi-protocol deployments to map Windows users to Unix users.

To ensure that NFSv4 is not used edit the file `/nas/server/slot_<x>/config` where x is the DM slot number of the DM being configured.

Edit the line in the file

```
nfs config
```

and modify this to include "hivers=3"

```
nfs config hivers=3
```

This ensures that the maximum version of NFS that can be used on this datamover is NFSv3.

The DM must be rebooted for the change to take effect (this is disruptive to serving data hence schedule this action at a suitable time).

3.15.6 Access-security policy types

EMC NAS supports a number of access-security policies, these are:

Policy	Comment
NATIVE	- Two independent sets of file permissions - NFS access uses NFS file permissions - CIFS access uses CIFS file permissions - File ownership aligned between NFS/NFS during creation (if possible). - Permissions are not synchronized so changes to one set of file permissions are not reflected in the other - Changes to file owner are reflected in both permissions sets - Change to file group ownership only affects Unix NATIVE is the CATE recommended access-security policy.
UNIX	- As “NATIVE” with the addition of - CIFS access only granted if <u>both</u> CIFS and Unix permissions allow access to file. (This relies on Windows to Unix credential translations)
NT	- As “NATIVE” with the addition of - NFS access only granted if <u>both</u> CIFS and Unix permissions allow access to file. (This relies on Unix to Windows credential translations) NT is an alternative certified policy for use with Multi-protocol configurations where Windows ACLs are the primary permissions.
SECURE	- As “NATIVE” with the addition of - Access must be granted using <u>both</u> CIFS and NFS permission sets. (This relies on Unix to Windows and Windows to Unix credential translations)
MIXED	- Single permission set use by both NFS and CIFS - Unix permissions translated to/from Windows style ACLs - Unix umask essentially ignored during Unix file creation - Changes to inherited ACLs have major impact on Unix mode bits displayed
MIXED_COMPAT	- Single permission set use by both NFS and CIFS - Identical to NetApp “Mixed” security style - Security permissions based on last protocol to change permissions.

For NFSv3 only configurations the “NATIVE” access policy should be used (this is the EMC default). For Multi-protocol configurations the certified options are “NATIVE” and “NT” mode. See [Configuration to support concurrent Multi-protocol access to data](#). Note: NFSv4 requires “Mixed “ access policy.

The “MIXED” policy for multi-protocol access with NFSv3 should not be used for the following reasons:

- Challenges in reliably translating both Unix and Windows groups between these operating systems.

- There is a risk that a minor permissions change from Windows will dramatically affect the Unix permissions of a file and in many situations result in world read+write access being granted.
- Unix file permissions are based on inherited Windows permissions rather than the expected Unix umask.

For example:

- A file with Windows ACLs that allows two groups read/write access will result in any Unix user being granted read/write access to the file. This is because a Unix file can only be owned by a single group.
- A file with Windows ACLs that allows two users read/write access will result in any Unix user being granted read/write access to the file. This is because a Unix file can only be owned by a single user.
- Changes to inherited ACLs can dramatically and immediately change the Unix security permissions for all files that inherit the ACLs.
- A Unix user with a umask of 0022 has that umask ignored with the resulting security bits being dependent on the inherited ACLs.

There are also a number of important parameters that must be set on a per datamover basis before multi-protocol access may be used; this is detailed in [Configuration to support concurrent Multi-protocol access to data](#).

3.15.6.1 Exporting NFS file systems

Data volumes to be used for NFS exports should be mounted inside of a VDM (NFS volumes may be located inside existing VDM's used for CIFS storage). This provides a common way to export both NFS only and multi-protocol accessed data as only file systems mounted inside a VDM may be exported using CIFS from inside a VDM (the Citi standard). Hence locating all data volumes inside a VDM provides a common way to provision volumes and also makes it easier to enable CIFS access to the data for multi-protocol use at a later date.

Whilst the creation and mounting of the file system to be used for NFS may be created in the either the GUI or from the command line (using the “nas_fs”, “server_mountpoint” and “server_mount” commands) it is not possible to export a data volume located in a VDM via NFS from the GUI. The command line “server_export” command should therefore be used to export the data.

Previously when exporting a volume mounted within a VDM the path used needs modification to take into account the volumes location inside the VDM. It is now possible to export a NFS volume directly with its mount point without having to use the volume location. The NFS exports configuration is also replicated via VDM replication. This is the new recommended method to do NFS exports and example is given below:

```
[nasadmin@c5-vnx5500-cs0 ~]$ server_export emcnasvdm1 -P nfs -name nfs_unix2 -
option sec=sys,ro=filerldaptest1,rw=filerldaptest2 /nfs_unix2
emcnasvdm1 : done
```

The old method to export via the volume location while no longer recommended, is described as follow for compatibility with existing setup. Each VDM is allocated an identify number when it is created and it is this identify number that is required to create the path to be used for the NFS export. For example in the example below the identifier number of the VDM is 2.

```
[nasadmin@cs0]$ nas_server -info -vdm -all
id      = 2
name    = nfsv4_vdm_01
acl     = 0
type    = vdm
```

VDM id number

```
server      = server_3
rootfs      = root_fs_vdm_nfsv4_vdm_01
I18N mode   = UNICODE
mountedfs   = sh93900_nfsv4_vdm3
member_of   =
status      :
    defined = enabled
    actual   = loaded, active
Interfaces to services mapping:
```

Once the identify number of the VDM is known then the path to be used for the export can be created. The example below is for a file system mounted on VDM “vdm_1” with an id of “2” with a mount point inside the VDM of “/vdm_filesystem”. Using this information the export of this file system will be:

```
/root/vdm_<id>/<mount point>
or “/root_vdm_2/vdm_filesystem” using the values supplied.
```

Below is an example via the command line interface showing:

- Creation of a 100MB volume
- Creation of a mount point on a VDM
- Mounting volume on VDM mount point
- Exporting volume via NFS on physical DM “server_3” (VDM has id 2)

```
[nasadmin@cs0]$ nas_fs -name sh93900_nfsv4_vdm3 -create size=100M pool=clar_r6
id          = 418
name        = sh93900_nfsv4_vdm3
acl         = 0
in_use      = False
type        = uxfs
worm        = off
volume      = v748
pool        = clar_r6
member_of   = root_avm_fs_group_17
rw_servers=
ro_servers=
rw_vdms     =
ro_vdms     =
auto_ext    = no,virtual_provision=no
deduplication = unavailable
stor_devs   = APM00083300134-006C
disks       = d10

[nasadmin@cs0]$ server_mountpoint nfsv4_vdm_01 -create /sh93900_nfsv4_vdm3
nfsv4_vdm_01 : done
[nasadmin@cs0]$ server_mount nfsv4_vdm_01 sh93900_nfsv4_vdm3 /sh93900_nfsv4_vdm3
nfsv4_vdm_01 : done
[nasadmin@cs0]$ server_export server_3 -Protocol nfs -name nfsv4_vdm_3 -option
sec=sys,ro=filerldaptest1,rw=filerldaptest3,root=filerldaptest2
/root_vdm_2/sh93900_nfsv4_vdm3
server_3 : done
```

Once the file system has been successfully exported via the command line it is visible in both the CLI and GUI interfaces. Clients mount the NFS export by specifying the server_3 as the server and the full path as the mount point, hence in the example above the client will mount <servername>:/root_vdm_2/sh93900_nfsv4_vdm3.

3.15.6.2 NFS export settings

NFS export should be restricted so that only designated servers may mount the exports. Mounting of these exports should also be further restricted to allow each server read+write, read-only and allow access as the root user as required.

Note: Care should be taken with the access permissions set. If no servers are explicitly given read-only access then unless a list of servers is also explicitly set using the “Access Hosts” option then all servers will be granted read access to the export.

It is recommended that the list of servers used to restrict access to exports use LDAP netgroups wherever possible to simplify the management of the export configuration. When using netgroups EMC NAS **does not** permit them to start with the “@” sign as is normal for Unix or NetApp NAS storage. Therefore if the netgroup is identical to a DNS name then as DNS is queried first this DNS name will be used for access rather than the netgroup, for this reason netgroup names must be different from entries in DNS. Multiple servers or netgroups for the same access entry should be separated by a “:” in both the GUI and CLI.

As the volume being exported is then mounted inside a VDM the initial export must be via the CLI, however once it has been exported then access to the data may be modified via the GUI or CLI.

The example below exports the volume /nfs_unix3 using SYS authentication.

```
[nasadmin@c5-vnx5500-cs0 ~]$ server_export emcnasvdm1 -P nfs -name nfs_unix3 -
option sec=sys,ro=filerldaptest1,rw=filerldaptest3,root=filerldaptest2
/nfs_unix3
```

The command above grants all servers in the netgroup filerldaptest1 read-only access, all servers in the netgroup filerldaptest3 read+write access and in addition all servers in the netgroup filerldaptest2 are also permitted “root” access to the data.

Note: There is currently no way to restrict suid on files (via the nosuid option) as EMC NAS does not support this option. This configuration still complies with Citi security standards as checks for the nosuid flag are currently made on the client mount rather than the server export and hence this is an additional safeguard beyond the Citi security standards.

An example that uses the access permission to prevent all servers gaining read-only access to the export is shown below:

```
[nasadmin@c5-vnx5500-cs0 ~]$ server_export emcnasvdm1 -P nfs -name nfs_unix3 -
option sec=sys,rw=filerldaptest3,access=filerldaptest3 /nfs_unix3
```

In this case only those servers in the netgroup filerldaptest3 are permitted to mount the export and they gain read+write access. It is recommended that if there is any doubt as to the effective permissions that a list of access servers is specified containing all servers to access the export (read or read+write).

Only the AUTH_SYS security is currently certified for NFS use. (Kerberos authentication for NFSv4 will be certified as part of a later phase of the NFS certification project).

3.15.7 Configuration to support concurrent Multi-protocol access to data

Whilst the configuration of a new volume on a previously unused VDM for multi-protocol can be carried out by simply setting the correct configuration before it is used, adding a second access protocol to an existing volume, or to a VDM that is already used for CIFS data, requires additional steps.

Whilst “Native” mode is the only certified mode to be used for pure NFS data where a multi-protocol environment is being configured “NT” mode is also certified for use.

Native and NT mode access policies should be used in the following situations:

Native

- This is the default configuration choice when NFS is the sole or primary method of access in a multi-protocol configuration.
- Where only CIFS is used to access data (i.e. not multi-protocol)
- In a multi-protocol configuration Unix file permissions are the primary mechanism used to control access to data
- This provides the appropriate configuration for data being migrated from NetApp with the “Unix” security style (including multi-protocol areas with the “Unix” security style) or for CIFS only access data with the “ntfs security style”

NT

- CIFS is the primary access protocol to data in a multi-protocol configuration.
- Windows ACLs are the primary mechanism used to control access to data in a multi-protocol configuration
- This provides the closest match for data being migrated from multi-protocol areas configured on NetApp with the “ntfs security style”.

3.15.7.1 EMC NAS changes to support multi-protocol access

Before exporting a volume using both NFS and CIFS the following changes must be made, in addition to the steps required to export NFS data (e.g. the configuration of LDAP, nsswitch.conf etc):

- The advanced mount property “Multi-Protocol Locking Policy” should be set to “rwlock” (the default is “nolock”) on all volumes to be used for multi-protocol purposes.
- The following parameters must be set on the data mover to host the multi-protocol volume (parameters are changed on the physical DM not at the VDM level).
 - acl.useUnixGid – This causes Windows users to create files using the primary GID of the Unix user (found from LDAP) with the same username (or as modified by the ntxmap.conf file).
 - resolver – This removes the domain name from the group or username for lookups in LDAP, without this setting the Windows domain name is appended to items being searched for.
 - NTcred.winDomain – This sets the Windows domain used when creating nt credentials for Unix users (used for “NT” mode). This setting is common to all VDM’s hence if VDM’s connect to different Windows domains these VDM’s should be located on different physical Data Movers.

The changes above will allow a Windows user to create new files with the correct Unix user and primary group ownership. Changes to EMC NAS parameters are set on the physical DM and affect **all** VDM’s hosted on this physical DM. The “rwlock” option ensures that access via NFSv3 will obey CIFS and NFSv4 locking, without this flag set this is not the case and each protocol will ignore locks granted to other protocols.

The EMC NAS parameters may be set in either the GUI or via the CLI. The examples below show the commands required to configure these values on the DM. A reboot of the DM is not required to activate these values.

```
[nasadmin@cs0]$ server_param server_3 -facility cifs -modify acl.useUnixGid -
value 1
server_3 : done
```

```
[nasadmin@cs0]$ server_param server_3 -facility cifs -modify resolver -value 1
server_3 : done
[nasadmin@cs0 ~]$ server_param server_3 -facility nfs -modify NTcred.winDomain -
value namdev
server_3 : done
```

3.15.7.2 NT access policy mode

The “NT” mode access policy works differently than the default “Native” mode. As with the “Native” mode there are two sets of permissions, Unix and Windows, but in this case for a user to gain access to a file from Unix it must satisfy both sets of permissions. In order to achieve this the umask set on the CIFS share must be changed from the default value (“0022”) to “0000”, this ensures that all files created from CIFS have a Unix umask of “rw-rw-rw-” i.e. all users have access (read and write) to these files from a Unix perspective.

Normally Unix permissions of “rw-rw-rw-” would allow all users to modify and read these files this is **not true** when using NT mode. As the user must satisfy both the Unix and Windows permissions even though they may have read and write access from the Unix permissions they must also check the Windows ACL as just having access from one set of permissions is not sufficient. So a user will actually be granted (or denied) permissions based on the Windows ACL for the file, this is irrespective of whether it appears that a user has full access to a file from Unix. If however access is denied (or partially denied e.g. only read-only granted) to the Unix user by the Unix permissions then the ACL will not be checked, hence the need for the “0000” umask.

Files created by either a Windows or Unix user on a NT mode volume will have the Unix and Windows ownership and group ownership set identically to volumes using “Native” mode, it is only the mechanism used to determine access that changes.

In order to use NT mode access policy the following is required:

When mounting the underlying volume:

- Select NT mode as the Access-checking policy
- Ensure that “Use NT credentials” is selected.
- As with other multi-protocol modes ensure that rwlock is also enabled.

An example of mounting the volume from the command line using the required options is shown below:

```
[nasadmin@cs0 ~]$ server_mount nfsv4_vdm_01 -option
accesspolicy=NT,rwlock,ntcredential sh93900_nfsv4_nt1 /sh93900_nfsv4_nt1
nfsv4_vdm_01 : done
```

When creating the CIFS share to present the NT mode volume to users the umask must be set correctly i.e. to “0000”.

An example setting the umask on the CIFS share is shown below:

```
[nasadmin@cs0 ~]$ server_export nfsv4_vdm_01 -Protocol cifs -name nfsv4_nt1 -
options umask=0000,ntcredential,netbios=NAMECTEST35 /sh93900_nfsv4_nt1
nfsv4_vdm_01 : done
```

It is not possible to set the umask on a CIFS share from the GUI interface hence this option must be set via the command line.

3.15.7.3 Windows and Unix file creation in a multiprotocol environment

As a Windows user has no Unix credentials the EMC NAS has to translate the existing Windows credentials into the Unix credentials and apply these to the Unix file permissions on each file as they are created. This translation is done as follows:

1. The Windows user SID is looked up in the “secmap” database on the VDM. From this database a corresponding Unix UID is obtained. This UID will be used as the owner of the file for the Unix file permissions. The Unix UID in “secmap” may be shown as being sourced from a number of places and this is indicated by the “Origin” column when viewing the “secmap” entries:
 - a. ldap – These entries have been resolved in the Unix LDAP and show the Unix username associated with the Windows user.
 - b. usermapper – If a matching username cannot be found in the Unix LDAP then the usermapper will create an arbitrary, and currently unused (by usermapper), UID value for this user in the range of UID numbers used by usermapper.
 - c. ntxmap – This does not necessarily show the actual source of any UID but simply that the mapping was made with the help of the ntxmap mapping file. The UID will display a Unix username if the a Unix user was resolved in LDAP as a result of the ntxmap or a numeric UID value if usermapper was used.

There will always be an entry in the secmap relating to a Windows user as if a record does not exist when the user connects to the share then a record will be created.

2. The EMC NAS parameters “acl.useUnixGid” and “resolver” cause the username of the Windows user to be queried in the Unix LDAP to obtain the primary group of the user that matches the Windows username. This GID is then used for group ownership of the file for the Unix file permissions.
 - a. The Unix name looked up in the Unix LDAP can be changed by mapping the Windows name to a different Unix name in the ntxmap.conf file.
 - b. If no match for the user is found in LDAP then a GID will be provided by the usermapper.
3. The actual Unix permissions set will be dependent on the default umask set on the CIFS share, the default is 0022. Hence the usual permissions on a file will be “-rw-r--r--”. The umask used by the CIFS share may be changed if required via the “server_export” command and using the “umask=” option. [A umask of “0000” is required for NT mode shares]

A Unix user must also be given Windows credentials when creating a file however the way in which Windows file permissions work means that when creating a file from Unix the process is different.

1. The Unix username (the UID is translated to a username by a LDAP lookup) is looked up in the “secmap” database, if a corresponding match is found then the matching domain user will be used as the owner of the file for the Windows permissions. If no match is found in secmap but the UID supplied by NFS can be translated to a username then the ownership will be set to “UID=0x123 ‘user’” where “123” is the Unix users UID and “user” their Unix username. If there is no secmap match and the user cannot be translated to a username then the ownership will be set to “UID=0x123” where “123” is the Unix users UID. Note: If a matching secmap entry is later added then an ownership of the form UID=0x123 ‘user’ will now become shown as a domain user. Conversely if a secmap record is later removed a domain user ownership will revert to the “UID=0x123 ‘user’” form.
2. The resulting file will automatically inherit ACLs from the directory in which it resides.

Note: The incorrect setting of Windows file ownership for NFS created files will impact any chargeback processing that walks the file system and relies on the correct CIFS file ownership to apportion storage costs. It will not however affect file access via CIFS as the permissions are based on the CIFS ACLs rather than file ownership.

It should be recalled that each file has two sets of file permissions (Unix and Windows) in “NATIVE” and NT modes. Whilst these permissions will be set as mentioned above when a file is created the permissions will diverge if permission changes are made within Windows or Unix.

The following is a brief summary of the changes and their affects between the two sets of permissions:

Change	Unix permission effect	Windows permission effect
File ownership changed from Unix or Windows	Ownership changes (using translation rules used during file creation)	Ownership changes (using translation rules used during file creation)
Group ownership changed from Unix	Group ownership changes	No effect. Windows does not have the concept of group ownership.
Unix permissions changes	Unix permissions change	No effect
Windows ACLs added or removed (directly or via inherited ACLs)	No effect	Added or removed ACLs take effect

3.15.7.4 New NFS or CIFS share permissions

When a new CIFS share is created a default inherited ACL is created that allows the windows group “Everyone” full control of all files created beneath that share. New NFS volumes have the top level directory of the export owned by the root user and groups with the permissions set to “drwxr-xr-x”. Both these default values should be changed to values that comply with the Citi security requirements when the exports/shares are created.

3.15.7.5 Discussion of usermapper, secmap databases and ntxmap.conf configuration file

The role of the usermapper and secmap databases in conjunction with the ntxmap.conf file in multi-protocol environments cannot be overstated. Below is a brief summary of each of these components and commands that may be useful in manipulating them.

secmap database

The first place that will be interrogated for SID to UID/GID or UID/GID to SID mapping information is the “secmap” database. There is a separate secmap database on each DM or VDM; this database should be looked on as a local cache of credential mappings and contains data found elsewhere (either in LDAP or the usermapper database). However unlike most other caches an entry in the secmap database is **permanent** and will only be removed or updated due to the manual intervention by a storage admin.

New secmap records are automatically added when a user connects to a CIFS share if a corresponding secmap entry does not already exist. User lookups that are used to create the secmap entry use the search order defined in nsswitch.conf (after any ntxmap.conf mappings) to resolve information and only when these options result in no matches will the usermapper be queried. A new secmap entry will not be created due to NFS access to data only Windows share mounts trigger the creation of new secmap entries.

Below is a partial listing of a secmap database showing the database contents:

```
[nasadmin@cs0]$ server_cifssupport server_3 -secmap -list
server_3 : done
```

SECMAP USER MAPPING TABLE

UID	Origin	Date of creation	Name	SID
939001	ntxmap	Wed Dec 15 04:58:33 2010	NAMDEV\sh93900	S-1-5-15-
ce1979d7-b16b6410-a0bb877c-50013				
939001	ldap	Fri Dec 10 06:12:34 2010	NAMDEV\sh939001	S-1-5-15-
ce1979d7-b16b6410-a0bb877c-ab443				
939002	ldap	Fri Dec 10 04:49:29 2010	NAMDEV\sh939002	S-1-5-15-
ce1979d7-b16b6410-a0bb877c-ab444				
32776	usermapper	Fri Dec 10 09:02:00 2010	NAMDEV\sh939003	S-1-5-15-
ce1979d7-b16b6410-a0bb877c-ab445				
32775	usermapper	Fri Dec 10 09:01:59 2010	NAMDEV\root	S-1-5-15-
ce1979d7-b16b6410-a0bb877c-43f64				
32768	usermapper	Fri Dec 10 09:13:21 2010	NAMDEV\gj06529	S-1-5-15-
ce1979d7-b16b6410-a0bb877c-6c188				
32770	usermapper	Sun Dec 12 03:56:11 2010	NAMDEV\Administrator	S-1-5-15-
ce1979d7-b16b6410-a0bb877c-1f4				

SECMAP GROUP MAPPING TABLE

GID	Origin	Date of creation	Name	SID
32780	usermapper	Fri Dec 10 04:49:29 2010	_HISTORY_SID_RANGE_\Domain Users	S-1-5-15-
52482b5c-677262ba-beb1293-201				
32781	usermapper	Fri Dec 10 04:49:29 2010	_HISTORY_SID_RANGE_\Domain Users	S-1-5-15-
193a4d69-782a51d5-12df346e-201				
32769	usermapper	Fri Dec 10 04:49:08 2010	NAMDEV\Domain Admins	S-1-5-15-
ce1979d7-b16b6410-a0bb877c-200				
32774	usermapper	Fri Dec 10 04:49:29 2010	NAMDEV\Domain Users	S-1-5-15-
ce1979d7-b16b6410-a0bb877c-201				
32786	usermapper	Fri Dec 10 04:49:29 2010	NAMDEV\Domain Guests	S-1-5-15-
ce1979d7-b16b6410-a0bb877c-202				
32787	usermapper	Fri Dec 10 04:49:38 2010	NAMDEV\Domain Computers	S-1-5-15-
ce1979d7-b16b6410-a0bb877c-203				
32788	usermapper	Sun Dec 12 03:56:11 2010	NAMDEV\Group Policy Creator Ownr	S-1-5-15-
ce1979d7-b16b6410-a0bb877c-208				

...

To update an existing entry, delete an existing entry or create a new secmap entry the following commands are used:

```
server_cifssupport <data mover> -secmap -update -name <username> -domain <domain name>
server_cifssupport <data mover> -secmap -delete -name <username> -domain <domain name>
server_cifssupport <data mover> -secmap -create -name <username> -domain <domain name>
```

It is also possible to display the reverse lookup and hence determine the username that a particular Unix UID or GID will map to (if a mapping is defined).

```
server_cifssupport <data mover> -secmap -list -uid <numeric uid>
```

If required the whole secmap database may be deleted but this requires that CIFS be stopped.
Note: CIFS's can only be stopped on a physical DM and this will also stop all VDM CIFS server instances residing on the same physical DM as well.

```
server_setup <movername> -Protocol cifs -option stop
server_dbms <movername> -db -delete Secmap (you must log in as root on CS)
server_setup <movername> -Protocol cifs -option start
```

usermapper database

The usermapper database holds a mapping of UID/GID/SID and user/group names. It is very similar in format to the secmap database but the usermapper database does not hold information gathered from LDAP or ntxmap and can be looked on as something like a local EMC NAS specific user and group database. This database is used not only for multi-protocol access but also for pure CIFS use as the underlying NAS file system (uxfs) requires UID and GID values set on each file created and the usermapper is the way in which these are created in a Windows only environment.

A new user (or Windows group) that cannot be found in LDAP will have a record created in the usermapper database, this record will contain a UID, GID, SID and name. This entry is permanent (unless manually removed or updated, which is not recommended nor documented by EMC) and will be used by all EMC NAS clusters referencing this usermapper instance.

The UID and GID values used by the usermapper are simply a set of sequential integers with each new user or group simply being allocated the next unassigned value.

The usermapper database is used as the source of last resort for UID/GID information with the local secmap database being queried first and usermapper only being queried to add new secmap entries if a secmap record is not available.

Below is a partial list of a usermapper database.

```
[nasadmin@cs0]$ server_cifssupport server_2 -usrmap -list
server_2 : done
```

USRMAP USER MAPPING TABLE			
SID	UID	GID	Name
S-1-5-15-ce1979d7-b16b6410-a0bb877c-1f4	32770	32768	
namdev\administrator			
S-1-5-15-ce1979d7-b16b6410-a0bb877c-43f64	32775	32768	namdev\root
S-1-5-15-ce1979d7-b16b6410-a0bb877c-50013	32772	32768	namdev\xxx123
S-1-5-15-ce1979d7-b16b6410-a0bb877c-6c188	32768	32768	namdev\gj06529
S-1-5-15-ce1979d7-b16b6410-a0bb877c-ab443	32774	32768	namdev\sh939001
S-1-5-15-ce1979d7-b16b6410-a0bb877c-ab444	32773	32768	namdev\sh939002
S-1-5-15-ce1979d7-b16b6410-a0bb877c-ab445	32776	32768	namdev\sh939003
S-1-5-15-ce1979d7-b16b6410-a0bb877c-bc879	32769	32768	
namdev\srvc_elerra_av			
S-1-5-15-ce1979d7-b16b6410-a0bb877c-bc87b	32771	32768	namdev\besadmin5

USRMAP GROUP MAPPING TABLE			
SID	UID	GID	Name
S-1-5-15-193a4d69-782a51d5-12df346e-201		32781	
_history_sid_range_\domain=20users			
S-1-5-15-52482b5c-677262ba-beb1293-201		32780	
_history_sid_range_\domain=20users			
S-1-5-15-864dc38-d0e266d0-56af0938-21a1		32778	
nsrootdev\exchange=20view-only=20administrators			
S-1-5-15-ce1979d7-b16b6410-a0bb877c-200		32769	
namdev\domain=20admins			
S-1-5-15-ce1979d7-b16b6410-a0bb877c-201		32774	namdev\domain=20users
S-1-5-15-ce1979d7-b16b6410-a0bb877c-202		32786	
namdev\domain=20guests			
S-1-5-15-ce1979d7-b16b6410-a0bb877c-203		32787	
namdev\domain=20computers			
S-1-5-15-ce1979d7-b16b6410-a0bb877c-208		32788	
namdev\group=20policy=20creator=20owners			
...			

As these usermapper generated UID/GID values are used whenever a user cannot be found in LDAP it is important that where these files may also be accessed via NFS that the generated UID/GID values do not overlap with Unix UID/GID ranges used by the real Unix users otherwise users may be granted access to files to which they should not be.

To overcome this issue the minimum and maximum values for the usermapper UID/GID ranges may be set using the commands shown below:

```
server_param <movername> -facility usrmap -modify minuid -value <minimum value>
server_param <movername> -facility usrmap -modify mingid -value <minimum value>
server_param <movername> -facility usrmap -modify maxuid -value <maximum value>
server_param <movername> -facility usrmap -modify maxgid -value <maximum value>
```

Once the changes have been made, the DM hosting the usermapper must be rebooted to commit the settings. Any existing usermapper records will not be affected by this change and only newly created records will be affected by the new minimum and maximum values.

Warning: As the usermapper database is not replicated and resides in a single site it is recommended that operational procedures be put in place to ensure that the usermapper database is backed up and securely stored offsite on at least a daily basis so that in the event of a disaster the usermapper database is not lost.

The usermapper database is backed up as part of the automatic hourly control station backups but this backup is for CS use only. The usermapper database may be backed up using the CLI “server_usermapper” command and copying certain configuration files from the DM.

Back up Usermapper (on primary Usermapper datamover)

1. As root, dump the password and group files to a specified directory by typing:
\$ server_usermapper server_2 -Export -user /home/nasadmin/backup.passwd
\$ server_usermapper server_2 -Export -group /home/nasadmin/backup.group

2. Make a backup copy of the current usrmap.cfg file (if one is in use) by typing:
\$ cp /nas/rootfs/slot_2/etc/usermapper/usrmap.cfg /home/nasadmin/usrmap.cfg

3. Make a backup copy of the usrmap.settings file by typing:
\$ cp /nas/rootfs/slot_2/etc/usermapper/usrmap.settings
/home/nasadmin/usrmap.settings

All of these files should be copied daily to a central management server.

ntxmap.conf configuration file

The ntxmap.conf file provides a mechanism to manually map Windows and Unix users. The ntxmap.conf file is consulted before a new secmap entry is created (or updated) to locate the appropriate entry to search for in either LDAP or the usermapper database. The file is processed sequentially so EMC recommends that not more than 1,000 mapping entries are created in the file to prevent user logins being slowed.

The ntxmap.conf does not exist by default and if required must be created and edited manually. Each record in the file is of the format:

```
<Windows Domain>:<Windows username>::=<Unix username>
```

The Windows domain may be replaced by “*” to indicate a match against all Windows domains. Below is an example entry mapping the Windows user “ab12345” to the Unix user “xy67890” for all windows domains. Note: this is a bi-directional mapping and also maps the Unix user to the Windows user.

```
*:ab12345::=xy67890
```

Unix to Windows mapping occurs in the same manner but as NFS only supplies the UID the NAS must be able to translate the UID to a username, using either LDAP or a local passwd file, before the mapping will function correctly.

The ntxmap.conf file contains a mapping entry per line of the file. There is a ntxmap.conf file per datamover. The ntxmap.conf file may be copied to the datamover using the command as shown in the example below:

```
[nasadmin@cs0]$ server_file server_3 -put ntxmap.conf.server_3 ntxmap.conf
server_3 : done
```

The ntxmap.conf file may also be tested to ensure that the format is correct, as shown in the example below:

```
[nasadmin@cs0]$ server_checkup server_3 -test CIFS -subtest ntxmap
server_3 :
-----Checks-----

Component CIFS :

Ntxmap      : Checking the ntxmap configuration file..... Pass
-----
```

Note: If the ntxmap.conf file is used then operational processes should be put in place to replicate this file to the DR NAS to ensure correct multi-protocol mappings during a DR invocation.

If the ntxmap.conf file is no longer required it may be disabled by simply copying an empty ntxmap.conf file to the datamover.

3.15.7.6 Adding multi-protocol access to existing volumes

If the access to the volume is primarily for Windows systems and the Windows ACL is to be the primary controller of file access then the recommended access policy is "NT" mode, changes must be made to the existing file system mount to change the access policy, add ntcredentials to the mount and change the default umask on the CIFS share.

As has already been stated, file permissions for a file system with NATIVE access-policy are set when the file is created and after creation they will diverge if changes are made to them. As it's likely that a CIFS only volume will not have been configured with the appropriate LDAP configuration to match the expected Unix access, the Unix permissions on all existing files will therefore not have the correct Unix UID/GID set on them but will instead have values created by the usermapper.

The entries in the secmap database on the VDM (or DM if a VDM is not used) will also hold mappings from the Windows SID to the UID/GID values generated by usermapper rather than the values that the Unix account holds found from LDAP. For this reason the following steps are required will be required to add a second protocol to existing data volumes:

Adding NFS to an existing CIFS volume

- Configure LDAP on the DM and VDM to support NFS (as detailed previously)
- Either delete the secmap database (on the DM or VDM hosting the CIFS service) thus causing the entries to be re-created or update each secmap record individually (as detailed previously)
- Mount the volume on a Unix/Linux server and reset the user and group ownership of the existing files to their correct values. This will require root access to the export.

Resetting the file ownership will affect both Windows and Unix ownership. The same rules are used when changing a files ownership as are applied when files are created, see [Configuration to support concurrent Multi-protocol to data](#).

If the permissions on the existing files are not modified then the current file permissions will not be correct and either allow or disallow access to the wrong users. New files will however be created with the correct permissions (after the secmap entries have been re-created/refreshed).

Adding CIFS to an existing NFS volume

- Create CIFS shares
- Mount CIFS share on Windows and set inherited or specific ACLs on existing files.

Providing inherited ACLs are used the existing files will automatically acquire the required Windows permissions.

3.15.8 Use of Local Files

During NAS certification work the CATE storage team have historically expected and tested the use of LDAP (prior to LDAP NIS was tested) to provide the Unix UID/GID value translations for concurrent multi-protocol environments. However CATE have been made aware that some historical operational deployments have been using local passwd files to provide similar functionality. The use of local files to provide this functionality has not been part of the CATE testing and it is strongly recommended by CATE that LDAP be used in these environments to provide this functionality.

Below are some reasons for that using local files instead of LDAP is not recommended:

- The local passwd files must be manually edited as new users are added thus adding operational work to manage these files. Once configured LDAP will automatically find new users (and groups) as they are added into LDAP.
- The local files on each NAS storage devices must be managed individually with the same change potentially being required on many storage devices. With LDAP once a user is added to LDAP this information is available to all servers configured to use LDAP.
- There is the potential for errors being introduced to the translation of Unix to Windows and Windows to Unix mappings due to typing errors in the local files.
- The Citi security standards apply a segregation of roles such that storage administrators do not have access to data. However by manually maintaining the passwd files as translations between Unix and Windows (and the reverse) the storage administrators are in effect controlling which files can be accessed by which users. By using LDAP the control of these mappings are no longer maintained by the storage team.
- NFSv3-only environments do not require the translation of usernames and UID/GID's and hence historically have not needed LDAP configured unless LDAP netgroups were being used. With the introduction of NFSv4 all user information provided by clients to the storage is provided as strings, these strings must be translated to UID/GID values by the storage and hence name services are now required even for NFSv4 only deployments.
- The use of NFSv4 Kerberos authentication (the subject of a current CATE certification project) requires that the storage is configured to connect to the Unix Kerberos key servers. The use of LDAP, which is used in conjunction with these key servers, positions these NAS devices for the use of NFSv4.

In general the use of LDAP for the Unix component of NAS storage should be looked on as equivalent to the use of AD naming services for CIFS data (AD is actually based on LDAP) and

hence the equivalent to using the local passwd and group files for Unix data would be the use of local Windows accounts in a CIFS environment.

3.15.9 Multi-protocol permission and access checking

The EMC NAS has a useful command that can help when troubleshooting multi-protocol permissions and access to files, this tool is the “server_cifssupport” command. This is particularly helpful when troubleshooting the NT mode access policy.

Previously in the solutions guide we have used this command to display contents of the secmap and usmap databases, you can also use it to display user credentials and access to files.

The examples given we use shall use the Unix/Windows user sh939001, to help interpret the EMC NAS information below are shown the permissions as viewed from Unix and Windows:

From Unix

```
$ id
uid=939001(sh939001) gid=9001(nasgrp9001) groups=9001(nasgrp9001),9101(nasgrp9101)
```

From Windows

```
V:\>whoami
namdev\sh939001

V:\>whoami -groups

GROUP INFORMATION
-----

Group Name                Type                SID
-----
Attributes
=====
Everyone                  Well-known group S-1-1-0
Mandatory group, Enabled by default, Enabled group
BUILTIN\Remote Desktop Users Alias                S-1-5-32-555
Mandatory group, Enabled by default, Enabled group
BUILTIN\Users             Alias                S-1-5-32-545
Mandatory group, Enabled by default, Enabled group
NT AUTHORITY\REMOTE INTERACTIVE LOGON Well-known group S-1-5-14
Mandatory group, Enabled by default, Enabled group
NT AUTHORITY\INTERACTIVE Well-known group S-1-5-4
Mandatory group, Enabled by default, Enabled group
NT AUTHORITY\Authenticated Users Well-known group S-1-5-11
Mandatory group, Enabled by default, Enabled group
NT AUTHORITY\This Organization Well-known group S-1-5-15
Mandatory group, Enabled by default, Enabled group
LOCAL                     Well-known group S-1-2-0
Mandatory group, Enabled by default, Enabled group
NAMDEV\nfsv4t1            Group                S-1-5-21-3457776087-29766
05200-2696644476-879962 Mandatory group, Enabled by default, Enabled group
NAMDEV\nfsv4t3            Group                S-1-5-21-3457776087-29766
05200-2696644476-879964 Mandatory group, Enabled by default, Enabled group
NAMDEV\MOM_Operators_DG   Group                S-1-5-21-3457776087-29766
05200-2696644476-398616 Mandatory group, Enabled by default, Enabled group
```

Below is an example showing the credentials granted to the Unix user sh939001 when accessing a volume with the “ntcredential” option set.

```
[nasadmin@cs0 ~]$ server_cifssupport nfsv4_vdm_01 -cred -uname sh939001
nfsv4_vdm_01 : done
```

```
ACCOUNT GENERAL INFORMATION
Name                : sh939001
Domain              : NAMDEV
Primary SID         : S-1-5-15-ce1979d7-b16b6410-a0bb877c-ab443
UID                : 939001
GID                : 9001
Authentication      : UNIX
Privileges          : 0x0
System privileges   : 0x0
Extra credential data : 0x0
NT administrator    : False
NT credential capability : 0x0

ACCOUNT GROUPS INFORMATION
Type UNIX ID      Name                Domain                SID
NT 32804          b16b6410-a0bb877c-d6d5a S-1-5-15-ce1979d7-
NT 32774          b16b6410-a0bb877c-201  S-1-5-15-ce1979d7-
NT 32803          b16b6410-a0bb877c-d6d5c S-1-5-15-ce1979d7-
NT 9101           NT 9101                 S-1-5-12-2-238d
NT 9001           NT 9001                 S-1-5-12-2-2329
NT 9001           NT 9001                 S-1-5-12-2-2329
UNIX 9101         UNIX 9101
UNIX 32804        UNIX 32804
UNIX 32774        UNIX 32774
UNIX 32803        UNIX 32803
[nasadmin@cs0 ~]$
```

These "NT" type credentials allow the Unix user to use these SID's to access files via the Windows ACL. These SID's are retrieved from the Windows user of the same name.

A GID value will be displayed for all Windows native groups (e.g. SID S-1-5-15-ce1979d7-b16b6410-a0bb877c-d6d5a), this value is supplied by the NAS as there is no equivalent group name in ldap, this translation cannot be disabled but will not affect access to files in "NT" or "Native" modes..

There is also an equivalent command to display the credentials for the equivalent Windows user, as shown below:

```
[nasadmin@cs0 ~]$ server_cifssupport nfsv4_vdm_01 -cred -name sh939001 -domain
namdev
nfsv4_vdm_01 : done

ACCOUNT GENERAL INFORMATION
Name                : sh939001
Domain              : NAMDEV
Server              : NASEMCTEST35
Primary SID         : S-1-5-15-ce1979d7-b16b6410-a0bb877c-ab443
UID                : 939001
GID                : 9001
Authentication      : KERBEROS
Privileges          : 0x8
                    : 0x00008 => SeChangeNotify
System privileges   : 0x2
                    : 0x00002 => SysAccessNetworkLogon
Extra credential data : 0x2
                    : 0x00002 => Bypass Traverse Checking
NT administrator    : False
NT credential capability : 0x2
                    : 0x00002 => Kerberos Auth Used
```

UID and GID values are obtained from Unix (ldap)

```
ACCOUNT GROUPS INFORMATION

Type UNIX ID      Name                Domain                SID
NT 32804          nfsv4t1             NAMDEV                S-1-5-15-ce1979d7-
b16b6410-a0bb877c-d6d5a
NT 32774          Domain Users        NAMDEV                S-1-5-15-ce1979d7-
b16b6410-a0bb877c-201
NT 32803          nfsv4t3             NAMDEV                S-1-5-15-ce1979d7-
b16b6410-a0bb877c-d6d5c
NT 32775          MOM_Operators_DG    NAMDEV                S-1-5-15-ce1979d7-
b16b6410-a0bb877c-61518
NT 32780          Domain Users        NAMDEV                S-1-5-15-52482b5c-
677262ba-beb1293-201
NT 32781          Domain Users        NAMDEV                S-1-5-15-193a4d69-
782a51d5-12df346e-201
NT 4294967294     Everyone            S-1-1-0
NT 4294967294     NETWORK             NT AUTHORITY          S-1-5-2
NT 4294967294     Authenticated Us    NT AUTHORITY          S-1-5-b
NT 2151678497     Users               BUILTIN               S-1-5-20-221
NT 9001           UNIX GID=0x2329     S-1-5-12-2-2329
UNIX 32804
UNIX 32774
UNIX 32803
UNIX 32775
UNIX 32780
UNIX 32781
UNIX 4294967294
UNIX 2151678497
[nasadmin@cs0 ~]$
```

The “server_cifssupport –cred” command displays information in the credential cache, if the cache has no entries for this user (entries have a time-to-live of 20 minutes by default) then the command will fail as shown below:

```
[nasadmin@cs0 ~]$ server_cifssupport nfsv4_vdm_01 -cred -name sh939002 -domain
namdev
nfsv4_vdm_01 :
Error 0: nfsv4_vdm_01 : Invalid error code [0 The user sh939002 in the domain
namdev is not connected.
]
[nasadmin@cs0 ~]$ server_cifssupport nfsv4_vdm_01 -cred -uname sh939002
nfsv4_vdm_01 :
Error 0: nfsv4_vdm_01 : Invalid error code [0 The UNIX user sh939002 has no
credential in the cache.
]
[nasadmin@cs0 ~]$
```

If this occurs it is possible to build and display the credential as shown below:

```
[nasadmin@cs0 ~]$ server_cifssupport nfsv4_vdm_01 -cred -uname sh939002 -build
nfsv4_vdm_01 : done

ACCOUNT GENERAL INFORMATION

Name                : sh939002
Domain              : NAMDEV
Server              : NASEMCTEST35
Primary SID         : S-1-5-15-ce1979d7-b16b6410-a0bb877c-ab444
UID                 : 939002
GID                 : 9002
Authentication      : UNIX
```

```
Privileges           : 0x0
System privileges   : 0x0
Extra credential data : 0x0
NT administrator    : False
NT credential capability : 0x0

ACCOUNT GROUPS INFORMATION

Type UNIX ID      Name                               Domain                               SID
NT 32802          b16b6410-a0bb877c-d6d5b                          S-1-5-15-ce1979d7-
NT 32774          b16b6410-a0bb877c-201                             S-1-5-15-ce1979d7-
NT 32803          b16b6410-a0bb877c-d6d5c                          S-1-5-15-ce1979d7-
NT 9102           S-1-5-12-2-238e
NT 9002           S-1-5-12-2-232a
UNIX 32802
UNIX 32774
UNIX 32803
UNIX 9102
[nasadmin@cs0 ~]$
```

The command is similar for Windows users with “-build” simply being appended to the query command line.

The command can also be used to show the permissions on or access to a specific file. The example below shows the access to the file “sh939002-rw-nfsv4t1-rw.txt”.

Firstly a view of the file from Unix and Windows to aid interpretation of the NAS information.

From Unix

```
$ ls -l sh939002-rw-nfsv4t1-rw.txt
-rw-rw-rw- 1 sh939002 nasgrp9002 72 Apr 15 08:36 sh939002-rw-nfsv4t1-rw.txt
```

From Windows:

```
V:\>cacls sh939002-rw-nfsv4t1-rw.txt
V:\sh939002-rw-nfsv4t1-rw.txt NAMDEV\nfsv4t1:F
                             NAMDEV\sh939002:F
```

We can now use the command to show the ACLs as viewed within the EMC NAS:

```
[nasadmin@cs0 ~]$ server_cifssupport nfsv4_vdm_01 -acl -path
/sh93900_nfsv4_nt1/sh939002-rw-nfsv4t1-rw.txt
nfsv4_vdm_01 : done

ACL DUMP REPORT

Path           : /sh93900_nfsv4_nt1/sh939002-rw-nfsv4t1-rw.txt
UID            : 939002
GID            : 9002
Rights         : rw-rw-rw-
owner SID      : S-1-5-15-ce1979d7-b16b6410-a0bb877c-ab444
group SID      : S-1-5-12-2-232a

DACL

GROUP 32804 S-1-5-15-ce1979d7-b16b6410-a0bb877c-d6d5a
ALLOWED 0x0 0x1f01ff RWXPD0
```

```
USER 939002 S-1-5-15-ce1979d7-b16b6410-a0bb877c-ab444
ALLOWED 0x0 0x1f01ff RWXPDO
```

We may now show the expected access for user sh939001 (Unix and Windows) to this particular file.

```
[nasadmin@cs0 ~]$ server_cifssupport nfsv4_vdm_01 -accessright -uname sh939001 -
path /sh93900_nfsv4_nt1/sh939002-rw-nfsv4t1-rw.txt
nfsv4_vdm_01 : done

ACCOUNT GENERAL INFORMATION

Name                : sh939001
Path                : /sh93900_nfsv4_nt1/sh939002-rw-nfsv4t1-rw.txt
Allowed mask        : 0x301df
Action              : List Folder / Read data
Action              : Create Files / Write data
Action              : Create Folders / Append Data
Action              : Read Extended Attributes
Action              : Write Extended Attributes
Action              : Delete Subfolders and Files
Action              : Read Attributes
Action              : Write Attributes
Action              : Delete
Action              : Read Permissions

[nasadmin@cs0 ~]$ server_cifssupport nfsv4_vdm_01 -accessright -name sh939001
-domain namdev -path /sh93900_nfsv4_nt1/sh939002-rw-nfsv4t1-rw.txt
nfsv4_vdm_01 : done

ACCOUNT GENERAL INFORMATION

Name                : sh939001
Domain              : namdev
Path                : /sh93900_nfsv4_nt1/sh939002-rw-nfsv4t1-rw.txt
Allowed mask        : 0x301ff
Action              : List Folder / Read data
Action              : Create Files / Write data
Action              : Create Folders / Append Data
Action              : Read Extended Attributes
Action              : Write Extended Attributes
Action              : Traverse Folder / Execute File
Action              : Delete Subfolders and Files
Action              : Read Attributes
Action              : Write Attributes
Action              : Delete
Action              : Read Permissions

[nasadmin@cs0 ~]$
```

As expected both the Unix and Windows user are shown as having identical access rights to this file.

If the Unix permissions are changed to remove write access to “others” and the test repeated.

```
[root@nasldapttest01 nt1]# ls -l sh939002-rw-nfsv4t1-rw.txt
-rw-rw-rw- 1 sh939002 nasgrp9002 72 Apr 15 08:36 sh939002-rw-nfsv4t1-rw.txt
[root@nasldapttest01 nt1]# chmod o-w sh939002-rw-nfsv4t1-rw.txt
[root@nasldapttest01 nt1]# ls -l sh939002-rw-nfsv4t1-rw.txt
-rw-rw-r-- 1 sh939002 nasgrp9002 72 Apr 15 08:36 sh939002-rw-nfsv4t1-rw.txt
```

From NAS

```
[nasadmin@cs0 ~]$ server_cifssupport nfsv4_vdm_01 -accessright -uname sh939001 -
path /sh93900_nfsv4_nt1/sh939002-rw-nfsv4t1-rw.txt
nfsv4_vdm_01 : done

ACCOUNT GENERAL INFORMATION

Name                : sh939001
Path                : /sh93900_nfsv4_nt1/sh939002-rw-nfsv4t1-rw.txt
Allowed mask        : 0x20089
Action              : List Folder / Read data
Action              : Read Extended Attributes
Action              : Read Attributes
Action              : Read Permissions
[nasadmin@cs0 ~]$
```

As can be seen the Unix users now loses write permissions to the file (as expected). The Windows permissions remain unaffected by changes to the Unix permissions.

3.15.10 Assign an interface to a VDM

One limitation in previous Dart codes was the association of a network interface for NFS only VDM's. You are now able to assign one or more interfaces to a VDM for hosting NFS only endpoints.

This process can be applied to any loaded VDM as long as the interface is not already assigned to another VDM or the interface is not first allocated to a CIFS Server on the same VDM.

To attach one or more interfaces to a VDM, use this command syntax:

```
$ nas_server -vdm <vdm_name> -attach <if_name>[,...]
```

where:

<vdm_name> = name of the VDM

<if_name> = name of the interface

Example:

To attach network interfaces to vdm1, type:

```
$ nas_server -vdm vdm1 -attach vdm1if1,vdm1if2
```

4 Alerting

CATE Storage Engineering recommends configuring the VNX to send SNMP traps directly to the enterprise alerting system.

All production VNXs should have a basic level of alerting configured. This alerting is a combination of several different processes within the VNX, and thusly has several different areas for configuration.

4.1 System Alerting via SNMP and Email

This alerting encompasses critical system alerts. These alerts should be configured to use both SNMP and email for notification for redundancy.

4.1.1 Alerting with SNMP Traps

The VNX can be configured to send SNMP based traps to the enterprise alerting system (ESS). In several regions this is a Tivoli system. CATE Storage and Backup Engineering recommends configuring the VNX to send SNMP traps directly to the enterprise alerting system.

4.1.1.1 Configure VNX to send SNMP traps to Tivoli

Do the following to configure VNX to send SNMP traps to Tivoli (or similar ESS environment):

1. The SNMP trap daemon (snmptrapd) is already running. To prevent two sets of logs being generated during testing, you can stop the existing daemon by issuing the following command (as root or similar):

```
killall snmptrapd
```

To restart the daemon after testing issue the following command (as root or similar):

```
/usr/sbin/snmptrapd -c /nas/sys/snmptrapd.conf -p 162 -u /var/run/snmptrapd.pid  
>/dev/null 2>&1 &
```

2. Modify the /nas/site/trap.cfg file so that there is an entry for every trap destination that you want to send to. Typically this is a single trap server, but there may be more than one entry:

```
#snmp trap configuration file
#example:
#snmpmanager 128.154.11.20 ; communityname public
#snmpmanager host1 ; communityname public
#
#gdse-ecc06.nam.nsroot.net (ecc v6.0)
snmpmanager 10.40.43.170 ; communityname public
#
#gdse-ecc02.nam.nsroot.net (ecc v6.1)
snmpmanager 10.40.43.153 ; communityname public
```

Note: Avoid the use of the “public” community name. Examples provided in this document are “public” for simplicity. Citi Information Security Standards recommend the use of a more complex community name. This should be addressed with the regional enterprise alerting (ESS) teams for specific implementations.

3. Start the SNMP trap daemon on the control station by issuing the following command (as root or similar):

```
/usr/sbin/snmptrapd -f -Le -Lf /nbsnas/var/log/my_eventlog_messages &
```

4. Send a test trap by issuing the following command:

```
/nas/sbin/nas_snmptrap <config_file_path> -m /nas/sys/emcVNX.mib -r <trap_number> -f <facility_id>
```

Note: Up to 255 characters are supported in <description> length.

Where:

<config_file_path>= path of the trap configuration file

/nas/sys/emccelerra.mib= VNX MIB file

<trap_number>= unique trap number for the event

<facility_id>= ID number of the facility generating the event

<event_id>= event ID number

<severity_level>= severity level of the event

<description>= description of the trap

Example:

```
/nas/sbin/nas_snmptrap /nas/site/trap.cfg -m /nas/sys/emcVNX.mib -r 1 -f 64 -i 5 -s 7 -d "This is a test SNMP trap."
```

This command will send the test trap to all of the servers listed in trap.cfg. The following is how the raw trap should be seen:

Date	Time	Priority	Hostname	Message
05-29-2009	10:54:44	Local7.Debug	10.40.42.201	community=public, enterprise=1.3.6.1.4.1.1139.2, enterprise_mib_name=emcVNX, uptime=69430662, agent_ip=10.40.42.201, generic_num=6, specificTrap_num=1, specificTrap_name=emcVNX.1, version=Ver1, celEvent.1=64, celEvent.2=5, celEvent.3=7, celEvent.4="test SNMP trap from c5-ns480-c1. Have a nice day. "

4.1.1.2 Additional Configuration/Customization of SNMP Traps (optional)

Enter the following command to check the events that are bound to sending traps from the VNX:

```
/nas/bin/nas_event -list -a trap
```

If an event is not in the list, edit /nas/sys/nas_eventlog.cfg to add an event, adding a line under the correct facility policy as follows, where ID1 and ID2 are event ID values:

1. disposition range=<ID1>-<ID2>, trap "/nas/site/trap.cfg 2"Load the changes to the nas_eventlog.cfg file:
2. /nas/bin/nas_event -Load /nas/sys/nas_eventlog.cfgVerify that the configuration file is loaded:

```
/nas/bin/nas_event -Load -info5. Verify the events are bound to traps:/nas/bin/nas_event -list -a trap
```

4.1.1.3 Configuring SNMP for the datamovers

There are three settings on the dtamover that need to be verified / set. These settings are:

- community (public) Sets or displays community name being used by the server SNMP agent. The default is **public**. The community name must be alphanumeric characters only and cannot contain any symbols or slashes.
- syscontact (nasadmin) Sets or displays syscontact value being used by the server SNMP agent. The default is **nasadmin**.
- location Sets location or facility name

To query the settings on the data movers issue the following commands:

```
server_snmp ALL --community
server_snmp ALL --syscontact
server_snmp ALL --location
```

To set the value of each snmp setting on all data movers, issue the following commands:

```
server_snmp ALL --community public
server_snmp ALL --syscontact nasadmin
server_snmp ALL --location here
```

You should also confirm that the active data movers (and control station) has SNMP enabled. (Unisphere → Security → Network Services). If there are any that are disabled, go ahead and enable them.

4.1.2 System Alerting with Email

The VNX can be configured to send email alerts just as it can be configured to send SNMP traps.

4.1.2.1 Configure DNS on the Control Station

By using DNS services, it is easier to administer hostname resolution. Typically the DNS configuration is done at the time of VNX Network Server implementation.

4.1.2.2 Confirm Email is Working

This procedure combines the commands echo, mail, and hostname. The results of the echo command are passed to the mail command and the hostname (Control Station name) is sent to the recipient.

Action

To test if email is working properly before configuring notifications, use this command syntax:

```
$ echo test from 'hostname' | mail -v -s "'hostname' subject" recipient email address
```

Be sure to enclose hostname in single, backward quotes; otherwise, the command fails. The man pages for echo, mail, and hostname commands provide more information.

where:

-v = verbose output option of mail command

-s = subject option of the message

'hostname' = result of hostname command that is passed to the requested commands (echo, and mail)

subject = email subject

recipient email address = recipient's email address

Example:

To send a test email to user1, type:

```
$ echo test from 'hostname' | mail -v -s "'hostname' test message" user1@nasdocs.emc.com
```

Output

The following output provides information about the state of email service on the Control Station.

```
user1@nasdocs.emc.com... Connecting to eagle.lss.nasdocs.emc.com. via esmtp...
```

```
220 igw ESMTP Sendmail 8.10.1/8.10.1; Tue, 22 Jun 2004 11:14:19 -0400 (EDT)
```

```
>>> EHLO manatee6.pag.nasdocs.emc.com
```

```
250-igw.nasdocs.emc.com Hello [192.168.80.140], pleased to meet you
```

```
.
```

```
.
```

```
.
```

```
>>> MAIL From:<nasadmin@manatee6.pag.nasdocs.emc.com> SIZE=104
```

```
250 2.1.0 <nasadmin@manatee6.pag.nasdocs.emc.com>... Sender ok
```

```
>>> RCPT To:<user1@nasdocs.emc.com>
```

```
250 2.1.5 <user1@nasdocs.emc.com>... Recipient ok
```

```
>>> DATA
```

```
354 Enter mail, end with "." on a line by itself
```

```
>>> .
```

```
250 2.0.0 i5MFEJr00252 Message accepted for delivery
```

```
user1@nasdocs.emc.com... Sent (i5MFEJr00252 Message accepted for delivery)
```

```
Closing connection to eagle.lss.nasdocs.emc.com.
```

```
>>> QUIT
```

```
221 2.0.0 igw.nasdocs.emc.com closing connection
```

This test only confirms that the control station can send messages. Next we configure VNX with email capabilities.

4.1.2.3 Configuring Email from Unisphere

You should also configure and test email notifications by using Unisphere or from the CLI. This step describes the setup in Unisphere. To configure via the CLI go to [Configuring Email Notifications via CLI](#). In Unisphere, go to System, then select Manage Email User from the Service Tasks list.

The following fields should be filled out in order to send/receive email messages. This is independent of the remote

- *Service Enabled* should be checked.
- *Recipient Email Addresses* should always include emcalert@emc.com . Use commas to separate values.
- *CC Email Addresses* can be distributions lists (external email address required) or individuals. Use commas to separate the values.
- *Email Server* should be the primary email server for the site/region. Typically this is rotated among servers via the network.DNS layer.
- *Subject Prefix* is used to distinguish the emails from other email messages.
- *Sender Email Address* is the where the emails will be coming from and must be a valid Citi (or Citigroup) address or they may be rejected.

To test the email facility, press the *Test* button at the bottom of the page.

The following is a sample test email:

```
From: gdse@citigroup.com [mailto:gdse@citigroup.com]
Sent: Monday, November 01, 2010 5:09 PM
To: emcalert@emc.com
Cc: Johnson, Gary G [CCC-OT_IT]
Subject: VNX Notification - c5-nsg8-cs0.namdev.nsrootdev.net, INFO, EmailUser

Event Time:          Nov  1 18:08:59 2010
```

Brief Description: Test event for Email User.

Full Description:
Test event for Email User.

Recommended Action:
No action is required.

VNX Name: c5-nsg8-cs0.namdev.nsrootdev.net (10.40.42.211)
VNX Model: NS-G8
VNX Serial #: APM00091600918
NAS Version: 5.6.50-2

Severity: INFO
Component: CS_PLATFORM
Facility: EmailUser
Message ID: 96109330532

System Configuration Snapshot: Mon Nov 1 17:15:01 EDT 2010

Control Station Slot No: 0

/nas/sbin/model:

NS-G8

/nas/bin/nas_inventory -list:

Component	Type	Status
System ID		
Battery A	Battery	OK
CLARiiON CX4-480 APM00083300134		
Battery B	Battery	OK
CLARiiON CX4-480 APM00083300134		
VNX NS-G8 APM000833001340064	VNX	OK
VNX NS-G8 APM000833001340064		
CLARiiON CX4-480 APM00083300134	CLARiiON	OK
CLARiiON CX4-480 APM00083300134		
Control Station 0	Control Station	OK
VNX NS-G8 APM000833001340064		
Control Station 1	Control Station	OK
VNX NS-G8 APM000833001340064		
DME 0 Blower A	Blower	OK
VNX NS-G8 APM000833001340064		
DME 0 Blower B	Blower	OK
VNX NS-G8 APM000833001340064		
DME 0 Blower C	Blower	OK
VNX NS-G8 APM000833001340064		
DME 0 Blower D	Blower	OK
VNX NS-G8 APM000833001340064		
DME 0 Data Mover 2	Data Mover	OK
VNX NS-G8 APM000833001340064		
DME 0 Data Mover 2 IO Module 0	IO Module	OK
VNX NS-G8 APM000833001340064		
DME 0 Data Mover 2 IO Module 0 SFP 0	SFP	OK
VNX NS-G8 APM000833001340064		
DME 0 Data Mover 2 IO Module 0 SFP 1	SFP	OK
VNX NS-G8 APM000833001340064		
DME 0 Data Mover 2 IO Module 0 SFP 2	SFP	OK
VNX NS-G8 APM000833001340064		
DME 0 Data Mover 2 IO Module 0 SFP 3	SFP	OK
VNX NS-G8 APM000833001340064		
DME 0 Data Mover 2 IO Module 1	IO Module	OK
VNX NS-G8 APM000833001340064		
DME 0 Data Mover 2 IO Module 2	IO Module	OK
VNX NS-G8 APM000833001340064		
DME 0 Data Mover 3	Data Mover	OK
VNX NS-G8 APM000833001340064		
DME 0 Data Mover 3 IO Module 0	IO Module	OK
VNX NS-G8 APM000833001340064		
DME 0 Data Mover 3 IO Module 0 SFP 0	SFP	OK
VNX NS-G8 APM000833001340064		
DME 0 Data Mover 3 IO Module 0 SFP 1	SFP	OK
VNX NS-G8 APM000833001340064		
DME 0 Data Mover 3 IO Module 0 SFP 2	SFP	OK
VNX NS-G8 APM000833001340064		

DME 0 Data Mover 3 IO Module 0 SFP 3	SFP	OK
VNX NS-G8 APM000833001340064		
DME 0 Data Mover 3 IO Module 1	IO Module	OK
VNX NS-G8 APM000833001340064		
DME 0 Data Mover 3 IO Module 2	IO Module	OK
VNX NS-G8 APM000833001340064		
DME 0 DME 0	Enclosure	OK
VNX NS-G8 APM000833001340064		
DME 0 I/O Annex A	I/O Annex	OK
VNX NS-G8 APM000833001340064		
DME 0 I/O Annex B	I/O Annex	OK
VNX NS-G8 APM000833001340064		
DME 0 Management Switch A	Management Switch	OK
VNX NS-G8 APM000833001340064		
DME 0 Management Switch B	Management Switch	OK
VNX NS-G8 APM000833001340064		
DME 0 Power Supply A	Power Supply	OK
VNX NS-G8 APM000833001340064		
DME 0 Power Supply B	Power Supply	OK
VNX NS-G8 APM000833001340064		
DME 1 Blower A	Blower	OK
VNX NS-G8 APM000833001340064		
DME 1 Blower B	Blower	OK
VNX NS-G8 APM000833001340064		
DME 1 Blower C	Blower	OK
VNX NS-G8 APM000833001340064		
DME 1 Blower D	Blower	OK
VNX NS-G8 APM000833001340064		
DME 1 Data Mover 4	Data Mover	OK
VNX NS-G8 APM000833001340064		
DME 1 Data Mover 4 IO Module 0	IO Module	OK
VNX NS-G8 APM000833001340064		
DME 1 Data Mover 4 IO Module 0 SFP 0	SFP	OK
VNX NS-G8 APM000833001340064		
DME 1 Data Mover 4 IO Module 0 SFP 1	SFP	OK
VNX NS-G8 APM000833001340064		
DME 1 Data Mover 4 IO Module 0 SFP 2	SFP	OK
VNX NS-G8 APM000833001340064		
DME 1 Data Mover 4 IO Module 0 SFP 3	SFP	OK
VNX NS-G8 APM000833001340064		
DME 1 Data Mover 4 IO Module 1	IO Module	OK
VNX NS-G8 APM000833001340064		
DME 1 Data Mover 4 IO Module 2	IO Module	OK
VNX NS-G8 APM000833001340064		
DME 1 Data Mover 5	Data Mover	OK
VNX NS-G8 APM000833001340064		
DME 1 Data Mover 5 IO Module 0	IO Module	OK
VNX NS-G8 APM000833001340064		
DME 1 Data Mover 5 IO Module 0 SFP 0	SFP	OK
VNX NS-G8 APM000833001340064		
DME 1 Data Mover 5 IO Module 0 SFP 1	SFP	OK
VNX NS-G8 APM000833001340064		
DME 1 Data Mover 5 IO Module 0 SFP 2	SFP	OK
VNX NS-G8 APM000833001340064		
DME 1 Data Mover 5 IO Module 0 SFP 3	SFP	OK
VNX NS-G8 APM000833001340064		
DME 1 Data Mover 5 IO Module 1	IO Module	OK
VNX NS-G8 APM000833001340064		
DME 1 Data Mover 5 IO Module 2	IO Module	OK
VNX NS-G8 APM000833001340064		
DME 1 DME 1	Enclosure	OK
VNX NS-G8 APM000833001340064		
DME 1 I/O Annex A	I/O Annex	OK
VNX NS-G8 APM000833001340064		
DME 1 I/O Annex B	I/O Annex	OK
VNX NS-G8 APM000833001340064		
DME 1 Management Switch A	Management Switch	OK
VNX NS-G8 APM000833001340064		
DME 1 Management Switch B	Management Switch	OK
VNX NS-G8 APM000833001340064		
DME 1 Power Supply A	Power Supply	OK
VNX NS-G8 APM000833001340064		
DME 1 Power Supply B	Power Supply	OK
VNX NS-G8 APM000833001340064		

IO Module A0	IO Module	OK
CLARiion CX4-480 APM00083300134		
IO Module A1	IO Module	OK
CLARiion CX4-480 APM00083300134		
IO Module A2	IO Module	OK
CLARiion CX4-480 APM00083300134		
IO Module A3	IO Module	Empty.
CLARiion CX4-480 APM00083300134		
IO Module A4	IO Module	Empty.
CLARiion CX4-480 APM00083300134		
IO Module B0	IO Module	OK
CLARiion CX4-480 APM00083300134		
IO Module B1	IO Module	OK
CLARiion CX4-480 APM00083300134		
IO Module B2	IO Module	OK
CLARiion CX4-480 APM00083300134		
IO Module B3	IO Module	Empty.
CLARiion CX4-480 APM00083300134		
IO Module B4	IO Module	Empty.
CLARiion CX4-480 APM00083300134		
Power Supply A0	Power Supply	OK
CLARiion CX4-480 APM00083300134		
Power Supply A1	Power Supply	OK
CLARiion CX4-480 APM00083300134		
Power Supply B0	Power Supply	OK
CLARiion CX4-480 APM00083300134		
Power Supply B1	Power Supply	OK
CLARiion CX4-480 APM00083300134		
Shelf 0/0 Disk 0	Disk	OK
CLARiion CX4-480 APM00083300134		
Shelf 0/0 Disk 1	Disk	OK
CLARiion CX4-480 APM00083300134		
Shelf 0/0 Disk 2	Disk	OK
CLARiion CX4-480 APM00083300134		
Shelf 0/0 Disk 3	Disk	OK
CLARiion CX4-480 APM00083300134		
Shelf 0/0 Disk 4	Disk	OK
CLARiion CX4-480 APM00083300134		
Shelf 0/0 Fan A	Fan	OK
CLARiion CX4-480 APM00083300134		
Shelf 0/0 Fan B	Fan	OK
CLARiion CX4-480 APM00083300134		
Shelf 0/0 LCC A	LCC	OK
CLARiion CX4-480 APM00083300134		
Shelf 0/0 LCC B	LCC	OK
CLARiion CX4-480 APM00083300134		
Shelf 0/0 Power Supply A	Power Supply	OK
CLARiion CX4-480 APM00083300134		
Shelf 0/0 Power Supply B	Power Supply	OK
CLARiion CX4-480 APM00083300134		
Shelf 0/0 Shelf 0/0	Enclosure	OK
CLARiion CX4-480 APM00083300134		
Shelf 1/0 Disk 0	Disk	OK
CLARiion CX4-480 APM00083300134		
Shelf 1/0 Fan A	Fan	OK
CLARiion CX4-480 APM00083300134		
Shelf 1/0 Fan B	Fan	OK
CLARiion CX4-480 APM00083300134		
Shelf 1/0 LCC A	LCC	OK
CLARiion CX4-480 APM00083300134		
Shelf 1/0 LCC B	LCC	OK
CLARiion CX4-480 APM00083300134		
Shelf 1/0 Power Supply A	Power Supply	OK
CLARiion CX4-480 APM00083300134		
Shelf 1/0 Power Supply B	Power Supply	OK
CLARiion CX4-480 APM00083300134		
Shelf 1/0 Shelf 1/0	Enclosure	OK
CLARiion CX4-480 APM00083300134		
Shelf 2/0 Disk 1	Disk	OK
CLARiion CX4-480 APM00083300134		
Shelf 2/0 Disk 10	Disk	OK
CLARiion CX4-480 APM00083300134		
Shelf 2/0 Disk 11	Disk	OK
CLARiion CX4-480 APM00083300134		

Shelf 2/0 Disk 12	Disk	OK
CLARiION CX4-480 APM00083300134		
Shelf 2/0 Disk 13	Disk	OK
CLARiION CX4-480 APM00083300134		
Shelf 2/0 Disk 14	Disk	OK
CLARiION CX4-480 APM00083300134		
Shelf 2/0 Disk 2	Disk	OK
CLARiION CX4-480 APM00083300134		
Shelf 2/0 Disk 3	Disk	OK
CLARiION CX4-480 APM00083300134		
Shelf 2/0 Disk 4	Disk	OK
CLARiION CX4-480 APM00083300134		
Shelf 2/0 Disk 5	Disk	OK
CLARiION CX4-480 APM00083300134		
Shelf 2/0 Disk 6	Disk	OK
CLARiION CX4-480 APM00083300134		
Shelf 2/0 Disk 7	Disk	OK
CLARiION CX4-480 APM00083300134		
Shelf 2/0 Disk 8	Disk	OK
CLARiION CX4-480 APM00083300134		
Shelf 2/0 Disk 9	Disk	OK
CLARiION CX4-480 APM00083300134		
Shelf 2/0 Fan A	Fan	OK
CLARiION CX4-480 APM00083300134		
Shelf 2/0 Fan B	Fan	OK
CLARiION CX4-480 APM00083300134		
Shelf 2/0 LCC A	LCC	OK
CLARiION CX4-480 APM00083300134		
Shelf 2/0 LCC B	LCC	OK
CLARiION CX4-480 APM00083300134		
Shelf 2/0 Power Supply A	Power Supply	OK
CLARiION CX4-480 APM00083300134		
Shelf 2/0 Power Supply B	Power Supply	OK
CLARiION CX4-480 APM00083300134		
Shelf 2/0 Shelf 2/0	Enclosure	OK
CLARiION CX4-480 APM00083300134		
Shelf 3/0 Disk 0	Disk	OK
CLARiION CX4-480 APM00083300134		
Shelf 3/0 Fan A	Fan	OK
CLARiION CX4-480 APM00083300134		
Shelf 3/0 Fan B	Fan	OK
CLARiION CX4-480 APM00083300134		
Shelf 3/0 LCC A	LCC	OK
CLARiION CX4-480 APM00083300134		
Shelf 3/0 LCC B	LCC	OK
CLARiION CX4-480 APM00083300134		
Shelf 3/0 Power Supply A	Power Supply	OK
CLARiION CX4-480 APM00083300134		
Shelf 3/0 Power Supply B	Power Supply	OK
CLARiION CX4-480 APM00083300134		
Shelf 3/0 Shelf 3/0	Enclosure	OK
CLARiION CX4-480 APM00083300134		
SP A	SP	OK
CLARiION CX4-480 APM00083300134		
SP B	SP	OK
CLARiION CX4-480 APM00083300134		
Write Cache	Cache	OK
CLARiION CX4-480 APM00083300134		

Depending on your service level agreement, you can contact your service provider to resolve the problem or follow the instructions in the recommended action.

4.1.2.4 Configuring Email Notifications via CLI

Email notifications allow you to configure email for multiple, serious system events in one place, eliminating the need to configure notifications with the same email recipients. To configure email notifications:

Enable Email Notifications

- Configure Email Notifications
- Display an Email Notifications Configuration
- Test an Email Notification Configuration
- Re-Initialize Email Notifications
- Disable Email Notifications

4.1.2.4.1 Enable Email Notifications

Action
To enable email notification, use this command syntax: <pre>\$ nas_emailuser -modify -enabled {yes no}</pre> where: yes = enabled. Example: To enable email notification, type: <pre>\$ nas_emailuser -modify -enabled yes</pre>
Output
OK

4.1.2.4.2 Configure Email Notifications

Use this procedure to configure an email notification.

Note: If DNS is not configured on the Control Station, email notifications will not be delivered.

The mailbox portion of the fully qualified domain name uses the following ASCII characters: a–z, A–Z, 0–9, !, #, \$, %, &, *, +, -, /, =, ?, ^, _, ` , {, |, }, ., ' , and ~. Periods cannot be the first or last character in the mailbox. Alphanumeric strings are accepted. Enclose email addresses in single quotes. For example: 'storage_admin1@nasdocs.emc.com'. You can use single quotes within an email address, but you must escape them correctly. The first example shows how to use a single quote on its own in an address, while the second shows how to use a single quote within a single-quoted address:

```
admin'email@nasdocs.emc.com
```

```
'first\'admin@nasdocs.emc.com,second\'admin@nasdocs.emc.com'
```

You can customize the subject prefix to meet your specific requirements, such as email filtering.

Action
To configure email notifications, use this command syntax: <pre>\$ nas_emailuser -modify -to {<email_addr>,...} -cc {<email_addr>,...} -email_server <email_server> -subject_prefix <email_subject> -from <email_addr></pre> where: <email_addr> = for -to, one or more comma-separated, email addresses, from 3 to 255 characters, in the format 'mailbox@fully-qualified-domain-name'. Each individual email address can be a maximum of 63 characters. <email_addr> = for -cc, one or more optional, comma-separated, email addresses, from 3 to 255 characters, in the format 'mailbox@fully-qualified-domain-name'. Each individual email address can be a maximum of 63 characters. <email_addr> = for -from, sender's optional email address, from 3 to 63 characters, in the format 'mailbox@fully-qualified-domain-name'. If you do not specify a sender, an address in the format root@<hostname> is used by the system. For example: 'root@doc.emc.com'. <email_server> = the optional email server that accepts and routes the email notifications. Specifies an optional IP address or the fully qualified domain name, from 1 to 63 characters, of the email server. The IPv4 addresses 0.0.0.0 and 255.255.255.255 are not allowed. <email_subject> = the optional subject prefix of the email notification, from 1 to 63 printable ASCII characters. The default subject prefix is Celerra Notification. Example: To configure email notifications using email server 10.6.50.122 from administrator to support, while copying engineering and documentation, type: <pre>\$ nas_emailuser -modify -to 'support1@nasdocs.emc.com','support2@nasdocs.emc.com' -cc 'engineering@nasdocs.emc.com','documentation@nasdocs.emc.com' - email_server 10.6.50.122 -subject_prefix "Celerra Notification -- May 10, 2008" -from 'administrator@nasdocs.emc.com'</pre>
Output
OK

4.1.2.4.3 Display an Email Notifications Configuration

Action
To display the email notifications configuration, type: <pre>\$ nas_emailuser -info</pre>
Output
<pre>Service Enabled = Yes Recipient Address(es) = xu2018support1@nasdocs.emc.comxu2019,su2019support2@nasdocs.emc.comxu20 19 Carbon copy Address(es) = xu2018engineering@nasdocs.emc.comxu2019,xu2019documentation@nasdocs.emc .comxu2019 Email Server = 10.6.50.122 Subject Prefix = Celerra Notification -- May 10, 2008 Sender Address = xu2018administrator@nasdocs.emc.comxu2019</pre>

4.1.2.4.4 Sample Email Notification

The following output shows an example of an email notification based on the configuration information you supplied and generated when a BoxMonitor error occurs:

Output
<pre> From: administrator@nasdocs.emc.com To: support1@nasdocs.emc.com,support2@nasdocs.emc.com CC: engineering@nasdocs.emc.com,documentation@nasdocs.emc.com Subject: Celerra Notification -- May 10, 2008 Event Time: May 10 03:14:07 2038 Brief Description: Slot 2: Enclosure 0 blade 1 I2C PSA bus error. Full Description: This is a management switch internal error. The error is logged when the management switch is not able to control or query the status of enclosure components because of a hardware problem that may be caused by a component or connection outside the management switch. Recommended Action: Examine the /nas/log/sys_log to locate the cause of the management switch event. This may be caused by bad I2C connections on the motherboard or the enclosure. This can also happen when the two management switches in the enclosure have conflicting views of the system. Replacing the enclosure hardware or one of the management switches may be necessary. Celerra Name: csdev-45(1.2.3.4) Celerra Model: NS40 Celerra Serial #: APM123455445 NAS Version: 5.6.38.0 Severity: CRITICAL Component: CS_PLATFORM Facility: BoxMonitor Message ID: 78928609795 Depending on your service level agreement, you can contact your service provider to resolve the problem or follow the instructions in the recommended actions. </pre>

4.1.2.4.5 Test an Email Notification Configuration

Test an email notification posts a test event, which delivers an email notification to the configured recipients.

Action
To test email notifications, type: <pre>\$ nas_emailuser -test</pre>
Output
OK

After Testing an Email Notification Configuration

Confirm with the configured recipients that they received the test email with the correct system identification information.

4.1.2.4.6 Disable email notifications

Action
To disable email notification, use this command syntax: <pre>\$ nas_emailuser -modify -enabled {yes no}</pre> where: no = disabled. Example: To disable email notification, type: <pre>\$ nas_emailuser -modify -enabled no</pre>
Output
OK

4.1.2.4.7 Re-initialize Email Notifications

If the configuration file becomes corrupt or is deleted, you can use the -init option to re-initialize the email notification feature; however, you should only use this option if you are directed to run it by a system message.

Action
To re-create the configuration file, type: <pre>\$ nas_emailuser -init</pre>
Output
OK

4.2 File System Capacity Alerting

The most effective alerting of the file system capacity is configured as a resource notification. These notifications are only configurable via the Unisphere GUI. Go to the Storage Usage tab under Notifications (Unisphere → Notifications → Storage Usage).

The following are the recommended thresholds to configure for file system alerting.

Notification Condition	Storage Type	Notify On	Storage Resource	Frequency (hours)*	Email Addresses (comma separated)	SNMP Destination (IP1;public,IP2;public)
90% Used	File System	Current Size	All File Systems	24	NAS-distrib-list@citi.com	SNMPserver; public
95% Used	File System	Current Size	All File Systems	24	NAS-distrib-list@citi.com	SNMPserver; public
97% Used	File System	Current Size	All File Systems	12	NAS-distrib-list@citi.com	SNMPserver; public
99% Used	File System	Current Size	All File Systems	12	NAS-distrib-list@citi.com	SNMPserver; public

* In the case of DART 5.6.43, the frequency is fixed at 24 hours.

Note: SNMP Traps must be already configured. See [Alerting with SNMP Traps](#) for details.
Email configuration must already be defined. See [System Alerting with Email](#) for details.

The same screen is used to configure file system and storage pool notifications. Once file system utilization has reached a configured threshold, the notifications (email and SNMP) will be sent from the control station to the defined destinations. For email, an external distribution list is recommended. For SNMP, one or more trap servers are recommended. These notifications will repeat according to the frequency specified (for all notification conditions met). On email alerts, if there are more than one resource notification that has the same time (i.e. same setup time), then there will be only a single email like in the example below.

4.2.1 Example SNMP Traps

Date	Time	Priority	Hostname	Message
11-01-2010	18:05:58	Local7.Debug	10.40.42.211	community=public, enterprise=1.3.6.1.4.1.1139.2, enterprise_mib_name=emcVNX, uptime=391435374, agent_ip=10.40.42.211, generic_num=6, specificTrap_num=9, specificTrap_name=emcVNX.9, version=Ver1, celEvent.1=135, celEvent.2=210, celEvent.3=4, celEvent.4="Nov 1 18:05:57 2010 JServer:4:210 Alert W0: File system ndmp_fs2 is using 87.1% of its current space capacity."
11-01-2010	18:05:58	Local7.Debug	10.40.42.211	community=public, enterprise=1.3.6.1.4.1.1139.2, enterprise_mib_name=emcVNX, uptime=391435374, agent_ip=10.40.42.211, generic_num=6, specificTrap_num=9, specificTrap_name=emcVNX.9, version=Ver1, celEvent.1=135, celEvent.2=210, celEvent.3=4, celEvent.4="Nov 1 18:05:57 2010 JServer:4:210 Alert W0: File system ndmp_fs5 is using 72.6% of its current space capacity."

4.2.2 Example Email Alert

From: root [mailto:root@c5-vnx5500-cs0.namdev.nsrootdev.net]
 Sent: Wednesday, August 31, 2011 6:02 PM
 To: Johnson, Gary G [CCC-OT_IT]
 Subject: Celerra Notification - c5-vnx5500-cs1 - WARNING: Resource Usage

```
#
# Celerra celerraSerial:APM00111601868, Control Station: c5-vnx5500-cs1
#
***** Aug 31, 2011 7:02:19 PM *****
Alert W0:
File system 'fs33' is using 46.3% of its current space capacity.
```

For more detailed information, Please use nas_message with associated Message ID 87518806226.

```
***** Aug 31, 2011 7:02:19 PM *****
Alert W0:
File system 'fs5' is using 49.7% of its current space capacity.
```

For more detailed information, Please use nas_message with associated Message ID 87518806226.

4.2.3 System Warning Alert

Notifications will also generate a  system warning alert from the JServer facility. These alerts will need to be acknowledged.

4.3 Storage Pool Utilization Alerting

Like with file system notifications, the most effective alerting of the storage pool capacity is configured as a resource notification. These notifications are only configurable via the Unisphere GUI. Go to the Storage Pool Notifications.
 Unisphere → Notifications → Storage Pool

The following are the recommended thresholds to configure for file system alerting.

Notification Condition	Storage Type	Notify On	Storage Resource	Frequency (hours)*	Email Addresses (comma separated)	SNMP Destination (IP1;public,IP2;public)
90% Used	Storage Pool	Current Size	All Storage Pools	24	NAS-distrib-list@citi.com	SNMPserver; public
95% Used	Storage Pool	Current Size	All Storage Pools	24	NAS-distrib-list@citi.com	SNMPserver; public
97% Used	Storage Pool	Current Size	All Storage Pools	24	NAS-distrib-list@citi.com	SNMPserver; public
99% Used	Storage Pool	Current Size	All Storage Pools	24	NAS-distrib-list@citi.com	SNMPserver; public

Note: SNMP Traps must be already configured. See [Alerting with SNMP Traps](#) for details.
Email configuration must already be defined. See [System Alerting with Email](#) for details.

The same screen is used to configure file system and storage pool notifications. Once storage pool utilization has reached a configured threshold, the notifications (email and SNMP) will be sent from the control station to the defined destinations. For email, an external distribution list is recommended. For SNMP, one or more trap servers are recommended. These notifications will repeat according to the frequency specified (for all notification conditions met). On email alerts, if there are more than one resource notification that has the same time (i.e. same setup time), then there will be only a single email.

4.3.1 System Warning Alert

Notifications will also generate a  system warning alert from the JServer facility. These alerts will need to be acknowledged.

4.3.2 Email Examples for Utilization

The following are example email messages for these event alerts. Note that multiple notifications (i.e. different thresholds) may be included in the same email alert. There are only minor cosmetic differences between the versions.

```
From: root [mailto:root@c5-vnx5500-cs0.namdev.nsrootdev.net]
Sent: Wednesday, August 31, 2011 6:02 PM
To: Johnson, Gary G [CCC-OT_IT]
Subject: Celerra Notification - c5-vnx5500-cs1 - WARNING: Resource Usage
```

```
#
# Celerra celerraSerial:APM00111601868, Control Station: c5-vnx5500-cs1
#
***** Aug 31, 2011 7:02:19 PM *****
Alert W0:
File system 'fs33' is using 46.3% of its current space capacity.
```

```
For more detailed information, Please use nas_message with associated
Message ID 87518806226.
```

```
***** Aug 31, 2011 7:02:19 PM *****
Alert W0:
File system 'fs5' is using 49.7% of its current space capacity.
```

```
For more detailed information, Please use nas_message with associated
Message ID 87518806226.
```

4.3.3 Trap Examples for Utilization

The following are example SNMP Trap messages for these event alerts. Note that each file system for each threshold sends out a different trap as highly utilized file systems will generate multiple traps.

Date	Time	Priority	Hostname	Message
08-31-2011	17:35:57	Local7.D ebug	10.40.42.211	community=public, enterprise=1.3.6.1.4.1.1139.2, enterprise_mib_name=emcCelerra,

				uptime=104250238, agent_ip=10.40.42.211, generic_num=6, specificTrap_num=9, specificTrap_name=emcCelerra.9, version=Ver1, celEvent.1=135, celEvent.2=210, celEvent.3=4, celEvent.4="Aug 31 17:35:55 2011 JServer:4:210 Alert W0: File system ndmp_fs13 is using 69.2% of its current space capacity."
08-31-2011	17:35:56	Local7.D ebug	10.40.42.211	community=public, enterprise=1.3.6.1.4.1.1139.2, enterprise_mib_name=emcCelerra, uptime=104250189, agent_ip=10.40.42.211, generic_num=6, specificTrap_num=9, specificTrap_name=emcCelerra.9, version=Ver1, celEvent.1=135, celEvent.2=210, celEvent.3=4, celEvent.4="Aug 31 17:35:54 2011 JServer:4:210 Alert W0: File system ndmp_fs14 is using 67.6% of its current space capacity."
08-31-2011	17:35:56	Local7.D ebug	10.40.42.211	community=public, enterprise=1.3.6.1.4.1.1139.2, enterprise_mib_name=emcCelerra, uptime=104250180, agent_ip=10.40.42.211, generic_num=6, specificTrap_num=9, specificTrap_name=emcCelerra.9, version=Ver1, celEvent.1=135, celEvent.2=210, celEvent.3=4, celEvent.4="Aug 31 17:35:53 2011 JServer:4:210 Alert W0: File system ndmp_fs17 is using 67.6% of its current space capacity."

4.4 SavVol Utilization Alerting

SavVol alerting requires the modification of the MIB file and notifications configuration file for SNMP and email alerting. There is only a single threshold available for alerting. This threshold is defined by the high water mark for the SavVol. The default value for this parameter is 90 (or 90%).

Note: An alternative method for alerting on SavVol utilization is under development.

The following are the steps needed to enable this alerting.

1. Create a new notification configuration file named `/nas/site/my_eventlog.cfg` by copying `/nas/sys/nas_eventlog.cfg` to `my_eventlog.cfg`.
2. Use a text editor to add the following statements to `my_eventlog.cfg`:

```
# SVFS High Water Mark Event Control
# notification for the SavVol high water mark usage reached
#
facilitypolicy 1:70,7
disposition range=1-1, logfile "/nas/log/sys_log"
disposition range=1-1, mail "stor_admin_01@nasdocs.emc.com"
disposition range=1-1, mail "stor_admin_02@nasdocs.emc.com"
disposition range=1-1, mail "stor_admin_03@nasdocs.emc.com"
disposition range=1-1, trap "/nas/site/trap.cfg 14"
where:
facilitypolicy 1:70 = dispositions that need to be applied to events in the SVFS
facility (facility ID 70) of the DART component (component ID 1).
7 = severity-level threshold (captures events of severity 7 or lower).
disposition range=1-1 = events 1 through 1, which equates to the SVFS event
"high water mark of SavVol reached." Event id 1 is the event id of the High Water Mark
event.
logfile = entry added to the /nas/log/sys_log file for events 1 through 1.
stor_admin_01@nasdocs.emc.com, stor_admin_02@nasdocs.emc.com, and
stor_admin_03@nasdocs.emc.com = email addresses that receive email notification
Replace these with one or more email addresses at company.
```

14 = arbitrary, yet unique, trap number you associate with the event in the VNX MIB (code shown in step 5).

3. Verify the SNMP Manager by confirming the following line to the file /nas/site/trap.cfg:
snmpmanager <a.b.c.d> ; communityname <public>
4. Modify the VNX MIB, found at /nas/sys/emcVNX.mib, to include the following lines:
celSVFS TRAP-TYPE
ENTERPRISE emcVNX
VARIABLES {celEvent}
DESCRIPTION
"Trap message will be sent when a SavVol high water mark reached."
::=14
where:
14 = arbitrary, yet unique, trap number for the event you specified in the my_eventlog.cfg file
5. Load the new notification configuration file using the following command:
\$ nas_event -Load /nas/site/my_eventlog.cfg

Output:
EventLog : will load /nas/site/my_eventlog.cfg...done

Note: Be sure to use the full path name of the configuration file in the -Load command.

6. Verify that the new notifications (email and traps) have been accepted using the following commands:
\$ nas_event -l -a mail
\$ nas_event -l -a trap
where:
-l = list
-a = action

Output:
For nas_event -l -a mail, the output should at least display:
DART(1)
|--> SVFS(70)
BaseID Severity Brief_Description
1 WARNING(4) FSID:\${id,5,%u} SavVol:\${vol,8,%s}
MaxSize:\${maxsize,5,%u} MB %Full(hwm=\${hwm,2,%d}) reached
(t:\${ticks,3,%q})
For nas_event -l -a trap, the output should at least display:
DART(1)
|--> SVFS(70)
BaseID Severity Brief_Description
1 WARNING(4) FSID:\${id,5,%u} SavVol:\${vol,8,%s}
MaxSize:\${maxsize,5,%u} MB %Full(hwm=\${hwm,2,%d}) reached
(t:\${ticks,3,%q})

4.5 CAVA and DEDUPE Alerting (Event Notifications)

Like with file system notifications, the most effective alerting of the storage pool capacity is configured under the notifications are within Unisphere.

Event Notifications can be configured to send out email or traps when a certain severity level is reached for a given facility. For example, if the CAVA service fails on a CAVA server, an ERROR will be generated under the VC facility. If there is a notification entry for the VC facility like below, then the action will be taken (emails, traps).

The following are the recommended Event Notifications:

Facility	Severity	Action	Destination (email addresses with commas or single trap ip;public)
VC	Error	mail	NAS-distrib-list@citi.com
VC	Error	trap	SNMPserver; public
DEDUPE	Error	mail	NAS-distrib-list@citi.com
DEDUPE	Error	trap	SNMPserver; public

There will need to be a separate entry for each destination on these notifications. The following is an example of multiple traps and email destinations:

Notifications nasadmin@10.40.42.201

Events Storage Usage Storage Projections Data Mover Load

Facility ▲	Severity	Action	Destination
 DEDUPE	Error	mail	gary.g.johnson@citigroup.com
 DEDUPE	Error	trap	10.40.43.169;public
 DEDUPE	Error	trap	153.40.226.68;public
 DEDUPE	Error	trap	169.185.204.123;public
 VC	Error	mail	gary.g.johnson@citigroup.com
 VC	Error	trap	10.40.43.169;public
 VC	Error	trap	153.40.226.68;public
 VC	Error	trap	169.185.204.123;public

Notes: Setting a severity level lower than the event will not trigger an alert. Therefore if warnings and errors are desired then there would need to be two entries for that facility.

Unlike resource notifications, there is only a single trap available for event alerts. To have multiple trap destinations, multiple event alerts will have to be configured (each with a different trap destination).

With CAVA (VC) these events are generated when one or more CAVA servers are disabled or not communicating.

For DEDUPE events, if there is a failure reported (like cannot complete due to space limitations), the Error event will be generated.

4.5.1 Email Examples for CAVA and Dedupe

The following are example email messages for these event alerts:

CAVA server offline

From: root [mailto:root@c5-nsg8-cs0.namdev.nsrootdev.net]
 Sent: Tuesday, December 28, 2010 1:03 PM
 To: Johnson, Gary G [CCC-OT_IT]
 Subject: EMCServer c5-nsg8-cs0, Component - DART, Facility - VC, Severity - ERROR

Dec 28 14:03:17 2010 DART:VC:ERROR:30 Slot 2: 1293562992: Error on CAVA server 10.40.43.180: OFFLINE, ntStatus: SUCCESS.

CAVA both servers offline (example is after each CAVA server offline reported)

From: root [mailto:root@c5-nsg8-cs0.namdev.nsrootdev.net]
 Sent: Tuesday, December 28, 2010 1:03 PM
 To: Johnson, Gary G [CCC-OT_IT]
 Subject: EMCServer c5-nsg8-cs0, Component - DART, Facility - VC, Severity - ERROR

Dec 28 14:03:19 2010 DART:VC:ERROR:6 Slot 2: 1293562990: No virus checker server is available. Virus checking is stopped.

CAVA High Watermark reached

From: root [mailto:root@c5-nsg8-cs0.namdev.nsrootdev.net]
Sent: Tuesday, December 28, 2010 1:06 PM
To: Johnson, Gary G [CCC-OT_IT]
Subject: EMCServer c5-nsg8-cs0, Component - DART, Facility - VC, Severity - ERROR

Dec 28 14:05:37 2010 DART:VC:ERROR:2 Slot 2: 1293563133: high water mark reached.

DEDUPE Failure

From: root [mailto:root@c5-ns480-cs0.namdev.nsrootdev.net]
Sent: Tuesday, December 28, 2010 1:53 PM
To: Johnson, Gary G [CCC-OT_IT]
Subject: EMCServer c5-ns480-cs0, Component - DART, Facility - DEDUPE, Severity - ERROR

Dec 28 14:52:39 2010 DART:DEDUPE:ERROR:2 Slot 2: 1293565951: Deduplication of file system id 57 was aborted after 10 failure(s), 2400 KB available on FS. Last reported error was: NoSpace.

4.5.2 SNMP Trap Examples for CAVA and DEDUPE

The following are example CAVA trap messages for these event alerts (in reverse order):

Date	Time	Priority	Hostname	Message
12-28-2010	14:03:19	Local7.Debug	10.40.42.201	community=public, enterprise=1.3.6.1.4.1.1139.2, enterprise_mib_name=emcCelerra, uptime=129104443, agent_ip=10.40.42.201, generic_num=6, specificTrap_num=10, specificTrap_name=emcVNX.10, version=Ver1, celEvent.1=81, celEvent.2=6, celEvent.3=3, celEvent.4="Dec 28 14:03:19 2010 VC:3:6 Slot 2: 1293562990: No virus checker server is available. Virus checking is stopped. "
12-28-2010	14:03:19	Local7.Debug	10.40.42.201	community=public, enterprise=1.3.6.1.4.1.1139.2, enterprise_mib_name= emcCelerra, uptime=129104434, agent_ip=10.40.42.201, generic_num=6, specificTrap_num=10, specificTrap_name=emcVNX.10, version=Ver1, celEvent.1=81, celEvent.2=9, celEvent.3=3, celEvent.4="Dec 28 14:03:19 2010 VC:3:9 Slot 2: 1293562990: Error on server 10.40.43.181: OFFLINE, RPC program version 0, Pre 3.0 CAVA. "
12-28-2010	14:03:19	Local7.Debug	10.40.42.201	community=public, enterprise=1.3.6.1.4.1.1139.2, enterprise_mib_name= emcCelerra, uptime=129104398, agent_ip=10.40.42.201, generic_num=6, specificTrap_num=10, specificTrap_name=emcVNX.10, version=Ver1, celEvent.1=81, celEvent.2=9, celEvent.3=3, celEvent.4="Dec 28 14:03:19 2010 VC:3:9 Slot 2: 1293562990: Error on server 10.40.43.180: OFFLINE, RPC program version 0, Pre 3.0 CAVA. "
12-28-2010	14:03:17	Local7.Debug	10.40.42.211	community=public, enterprise=1.3.6.1.4.1.1139.2, enterprise_mib_name= emcCelerra, uptime=882819336, agent_ip=10.40.42.211, generic_num=6, specificTrap_num=10, specificTrap_name=emcVNX.10, version=Ver1, celEvent.1=81, celEvent.2=6, celEvent.3=3, celEvent.4="Dec 28 14:03:17 2010 VC:3:6 Slot 2: 1293562992: No virus checker server is available. Virus checking is stopped. "
12-28-2010	14:03:17	Local7.Debug	10.40.42.211	community=public, enterprise=1.3.6.1.4.1.1139.2, enterprise_mib_name= emcCelerra, uptime=882819316, agent_ip=10.40.42.211, generic_num=6, specificTrap_num=10, specificTrap_name=emcVNX.10, version=Ver1, celEvent.1=81, celEvent.2=30, celEvent.3=3,

				celEvent.4="Dec 28 14:03:17 2010 VC:3:30 Slot 2: 1293562992: Error on CAVA server 10.40.43.180: OFFLINE, ntStatus: SUCCESS. "
12-28-2010	14:03:07	Local7.Debug	10.40.42.211	community=public, enterprise=1.3.6.1.4.1.1139.2, enterprise_mib_name= emcCelerra, uptime=882818311, agent_ip=10.40.42.211, generic_num=6, specificTrap_num=10, specificTrap_name=emcVNX.10, version=Ver1, celEvent.1=81, celEvent.2=30, celEvent.3=3, celEvent.4="Dec 28 14:03:07 2010 VC:3:30 Slot 2: 1293562982: Error on CAVA server 10.40.43.181: OFFLINE, ntStatus: SUCCESS. "
12-28-2010	11:29:59	Local7.Debug	10.40.42.201	community=public, enterprise=1.3.6.1.4.1.1139.2, enterprise_mib_name= emcCelerra, uptime=128184392, agent_ip=10.40.42.201, generic_num=6, specificTrap_num=10, specificTrap_name=emcVNX.10, version=Ver1, celEvent.1=81, celEvent.2=2, celEvent.3=3, celEvent.4="Dec 28 11:29:59 2010 VC:3:2 Slot 2: 1293553795: high water mark reached. "

In this case, two CAVA servers when offline soon after each other and caused a shutdown of virus checking. There is also an example of a high water mark event.

The following is an example DEDUPE trap message for these event alerts:

Date	Time	Priority	Hostname	Message
12-28-2010	14:52:39	Local7.Debug	10.40.42.201	community=public, enterprise=1.3.6.1.4.1.1139.2, enterprise_mib_name= emcCelerra, uptime=129400399, agent_ip=10.40.42.201, generic_num=6, specificTrap_num=10, specificTrap_name=emcVNX.10, version=Ver1, celEvent.1=148, celEvent.2=2, celEvent.3=3, celEvent.4="Dec 28 14:52:39 2010 DEDUPE:3:2 Slot 2: 1293565951: Deduplication of file system id 57 was aborted after 10 failure(s), 2400 KB available on FS. Last reported error was: NoSpace. "

In this case you are not getting the file system name, but rather the file system id.

4.6 DataMover CPU/Memory Alerting

In order to be notified of an overburdened datamover, it is recommended that the following DataMover Load alerting be added to the system.

Computation Method	Data Mover	Interval Start	Interval Length	Memory / CPU Limit	Email Addresses (comma separated)	SNMP Destination (IP1;public,IP2;public)
Averages Method	All	0:00	23:59	75% / 75%	NAS-distrib-list@citi.com	SNMPserver; public
Averages Method	All	0:00	23:59	85% / 85%	NAS-distrib-list@citi.com	SNMPserver; public

These alerts will be triggered when the average memory or cpu utilization reaches 75% or 85%. Reaching these averages would indicate a heavily loaded datamover.

5 Remote Support via ESRS

The EMC VNX uses the ESRS for remote support. The VNX sends home alerts via an email process.

5.1 EMC Secure Remote Support Gateway (ESRS)

The EMC VNX utilizes the ESRS Gateway for remote support activities. EMC will connect to any VNX via the ESRS Gateway.

5.1.1 EMC VNX Remote Support Access

EMC Support staff may expect to have access to the root and nasadmin users in order to diagnose potential issues. Alternatively, root or nasadmin level access could be provided by logging into Unisphere and/or a SSH session during a WebEx shared session.

5.1.2 Gateway Server Agent

The Gateway server agent is an HTTP handler. The agent functions as the communications broker between the Gateway-managed devices, the Policy Manager, and the EMC® Device Relationship Manager (DRM). All messages are encoded using standard XML and SOAP application protocols. Agent message types include:

- Device state heartbeat polling
- Data file transfer
- Remote access session initiation
- User authentication requests
- Device management synchronization

The Gateway agent acts as a proxy, carrying information to and from the Gateway-managed devices. To maximize remote support availability, EMC configures all Gateway agents to employ built-in failover to redundant EMC remote-support enterprise systems in the event that access to the primary site is unavailable. The Gateway agent can also queue session requests in the event of a temporary local network failure.

Network traffic can be configured to route from the Gateway through proxy servers to the Internet. Such configurations include support for auto-configuration, HTTP, and SOCKS proxy standards. The agent does not have its own user interface, but the customer can run it as either a Windows service or a console application. All agent actions are logged to a local runtime file.

5.1.3 Gateway to EMC communication

All communication between the customer's site and EMC is initiated by the Gateway server agent within Citi. Using industry standard Secure Sockets Layer (SSL) encryption over the Internet and EMC-signed digital certificate authentication, the Gateway creates a communication tunnel. The Secure Remote Support Gateway uses an industry-accepted bilateral authentication for the EMC servers and the Gateway Agent. Each Gateway has a unique digital certificate that is verified by EMC whenever a Gateway makes a connection attempt. The Gateway then verifies EMC's server certificate. Only when the mutual SSL authentication passes and the client and server negotiate a shared secret does the Gateway transmit messages to EMC, securing the connection against spoofing and man-in-the-middle attacks.

If EMC Customer Support needs to work on a VNX problem, a support engineer needs to remotely access these VNX applications utilizing the local VNX administrator account.

- Unisphere
- CLI via SSH

Although other applications are denied access, those specific applications are set at **Always Allow**

The VNX does not currently support the use of the ESRS Gateway as its email router. This is due to a Citi standard on ESRS to change the email port. Until this is resolved, the VNXs will use Citi email servers to route email based alerts to EMC.

5.1.4 Device Setting

From the group hierarchy, select the group for a particular (VNX) device. It is represented below the device type name by a serial number — for example, ML2805000499. You now see policy settings for that device only. Some may be inherited from the Global settings, some from the VNX settings, and some may be specific to that device.

5.1.5 Access Rights

The policy for each new device registering with the Policy Manager is inherited from the device type. Device type policy is preset by EMC, but can be edited. Policy for a particular group consists of a set of permissions (action-parameter combination), each with an associated access right. For a particular permission, one of three allowed access right options is set:

- **Always Allow (current Citi ESRS standard)**
- Ask for Approval
- Never Allow

5.2 Sending Alerts to EMC

The ESRS component on the VNX (called connectEMC) does not support the use of non-standard ports for email alerts. The ESRS Gateway at Citi, has been implemented using port 3325 to avoid traffic conflicts with the standard port 25 traffic. Therefore, we configure the VNXs to send email alerts to EMC via Citi email servers (primary and secondary). This also allows us to use the cc field to send a copy to Citi NAS Support Teams distribution list (this is not supported if we use the ESRS gateway).

6 Audit and System Logging

6.1 Audit Logging via SYSLOG

A new audit logging facility is available for improved audit logging capabilities. This facility combines audit records from two sources (/nas/log/cmd_log and /var/log/messages), adds user information, and places them into the system Syslog message file and/or any external Syslog server. This new method of audit logging replaces the existing manual processes of log archiving.

The VNX Audit tool consists of several scripts running on the VNX Control station. Each script is responsible for monitoring a specific log file for events to be sent to a remote log collector via the SYSlog protocol. These scripts are designed so that they constantly monitor an assigned log and when one of the listed events occurs a standardized message is posted using the Linux Syslog functionality.

The scripts map the User ID number to the LDAP name or local name as defined in the VNX admin database. Each instance of the script creates a internal mapping at startup and updates it periodically while running

The scripts are added to inittab. This ensures that the scripts are restarted in case of a system restart. The scripts do not add new events to the logs, instead they monitor for events posted in existing logs and then send the appropriate events to the local Syslog service. The local Syslog service is, in turn, configured to forward the event records to a remote Syslog server.

6.1.1 Message Format

When a script detects a detected message it reformats the message into the following format.

Date/time stamp, Time zone, System ID (Serial number), EVENTID, Operation, Username, User Id, Log name, Original Celera Log Entry

Field	Description
Date/Time Stamp	Time on the VNX Control Station
Time zone	Time zone in standard format from the Linux time zone functionality
System ID	Serial number of the VNX as reported by /nas/sbin/serial
EVENTID	Unique Number to facilitate quick identification of events (A list of the EVENTID's is provided as part to the tool delivery)
Operation	Text describing the event
Username	User name of the individual that executed the operation (Using the name in the "/nas/site/user_db" file on the VNX.)
User Id	User Id as defined on the VNX
Log Name	Name of the log on the VNX where the event was posted
Original Celera Log Entry	Text of the original Celera Log Entry.

Examples of the AUDIT tool output from a remote Syslog when the scripts are running:

```
May 18 18:32:14 nasdev244cs0 AUDIT_Messages.pl:
05/18/2010,18:32:14,EDT,ABC12345678901 ,1101,Succesfull
Login,root(uid=0)@local,0,/var/log/messages,May 18 18:32:02 nasdev244cs0
sshd(pam_unix)[16132]: session opened for user root by root(uid=0)
```

```
May 18 18:32:32 nasdev244cs0 sshd[16348]: Accepted password for nasadmin from
128.222.7.47 port 2221 ssh2
```

```
May 18 18:32:32 nasdev244cs0 sshd(pam_unix)[16493]: session opened for user
nasadmin by (uid=0)
```

```
May 18 18:32:39 nasdev244cs0 AUDIT_Messages.pl:
05/18/2010,18:32:39,EDT,ABC12345678901 ,1103,Password for session
accepted,root(uid=0)@local,0,/var/log/messages,May 18 18:32:32 nasdev244cs0
sshd[16348]: Accepted password for nasadmin from 128.222.7.47 port 2221 ssh2
```

6.1.2 Audit Scripting Setup

The administrator will need to copy the AUDIT_tool directory to the /etc/tools directory on the Control Station. RCP via an SSH session is the preferred method as FTP is disabled.

6.1.2.1 Copying Files and Configuring Inittab

To add the scripts to the /etc/inittab, perform the following actions.

1. Log in to the VNX as root.
2. Go to the AUDIT_tool directory . Then “ls” the directory to verify the files.

List of files :

```
AUDIT_cmd_bs
AUDIT_messages.pl
custom_cmd_log.csv
AUDIT_cmd.pl
AUDIT_ms_bs
cmd_log.csv
custom_cmd_log.csv
Custom_Log_message_files.xls
custom_messages_log.csv
messages_log.csv
ReadMe.doc
ReadMe.txt
```

List of directories:

```
auto
File
Time
```

3. Copy the file AUDIT_ms_bs to /etc/.AUDIT_ms_bs (the '.' makes it a hidden file)

```
cp -f AUDIT_ms_bs /etc/.AUDIT_ms_bs
```

4. Copy the file AUDIT_cmd_bs to /etc/.AUDIT_cmd_bs (the '.' makes it a hidden file)

```
cp -f AUDIT_cmd_bs /etc/.AUDIT_cmd_bs
```

5. Make sure both scripts are executable by running the following commands in the /etc directory

```
chmod u+x .AUDIT_ms_bs
chmod u+x .AUDIT_cmd_bs
```

6. Run the .AUDIT_ms_bs in the original console to make sure it executes.

```
/etc/.AUDIT_ms_bs
```

7. Tail the log from a different console window to verify the tool is operating.

```
tail -f /var/log/messages
```

8. Add the following 3 lines to the /etc/inittab file on the VNX (or uncomment existing lines)

```
# Run the Audit tool shell scripts
aml:3:respawn:/etc/.AUDIT_ms_bs
acl:3:respawn:/etc/.AUDIT_cmd_bs
```

9. Restart the inittab using the command below.

```
telinit q
```

6.1.2.2 Sending Messages to Syslog

To see the message on a remote Syslog server the Administrator will need to perform the following actions.

1. Add the loghost information to the /etc/hosts file on the VNX Control Station

For Example :

```
#log host entries for remote syslog destinations
#Ipaddress (###.###.###) -- Fully qualified DNS name -- "loghost"
10.11.12.13    syslog_host.nsroot.net    loghost1
10.11.12.130  syslog_host.nsroot.net    loghost2
```

2. Verify that you can ping the loghost entries added to /etc/hosts
3. Add the following lines to the /etc/syslog.conf file on the VNX

```
# write audit to remote log
auth.notice          @loghost1
auth.notice          @loghost2
```

4. Restart the Syslog service on the VNX.

```
/sbin/service syslog restart
```

Note: There are typically two SYSLOG destinations in an audit repository system such as SEMS. There would likely be two entries one for the primary system and another for the secondary system as they are not splitting the messages on a network layer.

6.1.3 Syslog Server Setup

In order for Syslog messages to be stored on a local server and routed to a log messaging repository such as SEMS, a server must have its standard Citi configuration modified. This modification will configure the Syslog services to receive external messages. The default behavior of a Citi server is to reject external Syslog messages. **In most cases, this is not a necessary step as the SEMS environment will be the Syslog destination(s).**

Note: Windows/AIX/Linux/HP Syslog platforms are not discussed here. Only Solaris platforms are discussed.

As Syslog is a common service among Solaris platforms, it is a simple process to configure a server to receive Syslog messages from an HDS array.

6.1.3.1 Syslog on Solaris 10

There are two configuration files that need to be modified (on Solaris 10) to enable it for remote logging. In addition, there are SVC area values to modify.

Solaris 10

/etc/default/syslogd

```
LOG_FROM_REMOTE=YES (all other lines should be commented out)
```

/etc/syslog.conf

Add a line for "user.info" to point to a "VNX_messages" file. This file requires the use of tab characters. Note this will require additional scripting to handle file management (prevent file from growing too large).

```
auth.notice                               /var/adm/VNX_messages
```

To modify the SVC area values and restart the syslog daemon in Solaris 10, use the following commands:

```
svccfg -s svc:/system/system-log setprop config/log_from_remote = true
svcadm restart svc:/system/system-log
```

6.1.4 Audit Log Scripting Options

The following are options that can be specified. Changes to these options are not recommended at this time.

AUDIT_messages [options]

Options:

- help Print a brief help message and exits.
- scaninterval=<filename> Default = 20, Min = 10, Max = 300
- userfile=<filename> Default => 'custom_messages_log.csv'
- logfile=<filename> Default => '/var/log/messages'
- man Prints the manual page and exits.
- debug Display debug information during operation

6.2 System Log Files

6.2.1 System Log Files Overview

The EMC VNX Network Server generates events to record errors, commands, and other information you might need to know. As described in the previous section, these events are sent to the SYSLOG facility. It is important to know where these events originate and how to interpret the messages.

6.2.2 Interpreting System Log Files

Log files contain messages to assist you in determining what is happening with the VNX Network Server. Some logs have information about the state of the VNX Network Server, while others have information about errors. As of VNX version 5.6, the system log (sys_log) must be viewed using the nas_logviewer command when viewing the log from the CLI.

Log	Command line or log location	Description
Command	<code>\$ /nas/log/cmd_log</code>	Displays a list of all successful commands executed for the Celerra Network Server. Commands displaying information only and most commands in the /nas/sbin directory are not logged.
Command error	<code>\$ /nas/log/cmd_log.err</code>	Lists all the failed commands issued by the administrator. Commands displaying information only and most commands in the /nas/sbin directory are not logged.
Data Mover or blade	<code>\$ server_log <movernam></code>	Displays the current log updates from the most recent restart.
Data Mover or blade	<code>\$ server_log <movernam> -a</code> Note: Use the -s option to display logs with a user-friendly timestamp.	Displays the available history of logs for a Data Mover or blade.
System	<code>\$ /nas/log/sys_log</code>	Displays a cumulative list of system event and log messages from the most recent Control Station restarts. Note: Use this command to format system logs: <code>nas_logviewer -f /nas/log/sys_log</code>
OS system	<code>\$ /nas/log/osmlog</code>	Lists events related to the Control Station operating-system operation.
ConnectHome	<code>\$ /nas/log/ConnectHome</code>	Stores each CallHome pending in XML format.
SIB POST	<code>\$ /nas/log/sibpost_log</code>	Shows results from Control Station hardware tests performed at boot time.

7 Reporting

7.1 Reporting via ECC

The ECC Storage Agent for NAS agent collects VNX device information as required by the NAS object model, and monitors those devices for “alertable” conditions. The ECC engine must have fully discovered the VNX in order for Storage Scope to be able to report on the NAS array. For more information on NAS array discovery in ECC see the [CATE Storage Engineering ECC Website](#) for the latest solutions guide and related documentation.

7.1.1.1 Functional ID for ECC

A local administrative user should be created in order for ECC to discover and monitor each VNX. This local user should have administrative privileges similar to nasadmin.

7.1.1.2 Add ECC server to ACL (IPTables FW)

ECC uses SSH, HTTPS, and HTTP to communicate with the VNX. Therefore, the ECC server running the Storage agent for NAS should be added to the ACL on the VNX control station.

Within Storage Scope (under ECC 6.1) the following data sets are available:

- NAS Servers
- Data Movers
- NAS File System
- IP Interfaces
- Exports
- Storage Device
- Mount Points

These reports can be displayed as HTML or can be exported via a comma delimited text file.

7.2 Reporting using Unisphere

Unisphere does provide some real-time and historical data trends for the datamovers and other major components.

7.3 EERS Entitlement Reporting

Active Directory accounts should be used for management and client access of shares for the VNX. There should already be an EERS export of active directory accounts. All accounts in the active directory group defined on the VNX for control station access will have the same administrator privileges.

Alternatively, there is a manual entitlement reporting mechanism that could be sent to the EERS system for centralized entitlement confirmation. Until scripting has been developed, the following is the base command (nas_acl) that will be used:

```
[nasadmin@c5-vnx5500-cs0 ~]$ nas_acl -list
index      type      level      num_id      name
1           user      admin      201          nasadmin
2           user      observer   500          sysadmin
3           user      observer   501          nasadmin1
4           user      observer   502          admin
5           user      observer   503          root1
7           user      observer   504          gj06529
8           user      observer   505          gj06529_admin
9           user      observer   506          ql68074
10          user      observer   507          ta13615
11          user      observer   508          jt51235
12          user      observer   509          emc
13          user      observer   510          ta136151
14          user      observer   511          jt512351
15          user      observer   512          eccuser
16          user      observer   513          sansscreen
19          user      observer   515          yt38855
20          user      observer   516          yt388551
21          user      observer   517          ql68074_mon01
22          user      observer   518          ql68074_mon02
23          user      observer   519          ql68074_sec01
25          user      observer   520          ql68074_stg01
26          user      observer   521          sansscreen2
29          user      observer   522          srv_ecc
```

Note that the level given in this export does not directly correspond to roles defined in the VNX.

7.4 Performance Reporting

Performance data on the VNX can be obtained by the CLI interface. This performance information is obtained via the server_stats command. This command is limited but should be familiar to a Unix administrator.

The command structure is as follows:

```
server_stats <movername>
  [{ -summary {basic|cifs|nfs|caches}{,...}
    | -table {net|dvol|fsvol|cifs|nfs}{,...}
    | -table {net|dvol|fsvol|cifs|nfs}
      [-sort <field_name>
        [-order {asc|desc}]
        [-lines <lines_of_output>]] [...]
    }...]
  [-count <count>]
  [-port <port>]
  [-interval <seconds>]
  [-terminationsummary {no|yes|only}]
  [-format {text|csv}]
  [-type {rate|diff|accu}]
  [-titles {never|once|repeat}]
```

7.4.1 Summary Views

The summary views give an overview of performance statistics and for the requested VNX component. The different Summary Views are:

- basic – Displays SPU, network, and volume statistics. This is the recommended view as it is a good general view of the system activity.
- cifs – Displays CIFS protocol summary (read ops, read throughput, ave read size, write ops, write throughput, average write request, other cifs ops, number cifs connections) nfs – Displays NFS protocol summary (read ops, read throughput, ave read size, write ops, write throughput, ave write request, other nfs ops, number nfs connections)
- caches – Displays cache-related summary

Summary View Sample Output

```
server_stats server_2 -summary basic -interval 3 -count 5
server_2 CPU Network Network dVol dVol
Timestamp Util In Out Read Write
          % KiB/s KiB/s KiB/s KiB/s
11:40:55 1 3751 175 35 4602
11:40:58 1 2535 113 11 3510
11:41:01 0 1802 87 21 2220
11:41:04 1 1762 97 32 2686
11:41:07 1 1908 94 8 2784

server_2 CPU Network Network dVol dVol
Summary Util In Out Read Write
          % KiB/s KiB/s KiB/s KiB/s
Minimum 0 1762 87 8 2220
Average 1 2352 113 21 3160
Maximum 1 3751 175 35 4602
```

7.4.2 Table Views

The table views allow for more detailed information, to sort the results, and to truncate the results (say top 10). The different table views are:

- net – Displays the network device information
- dvol – Displays basic volume statistics
- fsvol – Displays statistics for mounted file systems
- cifs – Displays CIFS protocol details (minimum, average, maximum of SMB1_Mkdir, SMB1_Rmdir, SMB1_Open, SMB1_Create, SMB1_Close, SMB1_Flush, SMB1_Unlink, SMB1_Rename, SMB1_GetAttr, SMB1_SetAttr, SMB1_Read, SMB1_Write, SMB1_Lock, SMB1_Unlock, SMB1_CreateTmp, SMB1_MkNew, SMB1_ChkPath, SMB1_Exit, SMB1_Lseek, SMB1_LockRead, SMB1_WriteUnlock, SMB1_ReadBlockRaw, SMB1_WriteBlockRaw, SMB1_SetAttrExp, SMB1_GetAttrExp, SMB1_LockingX,

SMB1_Trans, SMB1_TransSec, SMB1_Copy, SMB1_Move, SMB1_Echo, SMB1_WriteClose, SMB1_OpenX, SMB1_ReadX, SMB1_WriteX, SMB1_CloseTreeDisco, SMB1_Trans2Prim, SMB1_Trans2Secd, SMB1_FindClose2, SMB1_FindNotifyClose, SMB1_TreeConnect, SMB1_TreeDisco, SMB1_NegProt, SMB1_SessSetupX, SMB1_UserLogoffX, SMB1_TreeConnectX, SMB1_DiskAttr, SMB1_Search, SMB1_FindFirst, SMB1_FindUnique, SMB1_FindClose, SMB1_TransNT, SMB1_TransNTSecd, SMB1_CreateNTX, SMB1_CancelNT, SMB1_SendMessage, SMB1_BeginMessage, SMB1_EndMessage, SMB1_MessageText, SMB1_TRANS2_Open, SMB1_TRANS2_FindFirst, SMB1_TRANS2_FindNext, SMB1_TRANS2_QFsInfo, SMB1_TRANS2_SetFileSystemInfo, SMB1_TRANS2_QPathInfo, SMB1_TRANS2_SetPathInfo, SMB1_TRANS2_QFileInfo, SMB1_TRANS2_SetFileInfo, SMB1_TRANS2_Trans2-09, SMB1_TRANS2_Trans2-0A, SMB1_TRANS2_Trans2-0B, SMB1_TRANS2_Trans2-0C, SMB1_TRANS2_Mkdir, SMB1_TRANS2_Trans2-0E, SMB1_TRANS2_Trans2-0F, SMB1_TRANS2_GetDfsReferral, SMB1_TRANS2_Trans2-11, SMB1_TRANSNT_TransNT-00, SMB1_TRANSNT_Create, SMB1_TRANSNT_IOCTL, SMB1_TRANSNT_SetSD, SMB1_TRANSNT_NotifyChange, SMB1_TRANSNT_Rename, SMB1_TRANSNT_QuerySD, SMB1_TRANSNT_GetQuotas, SMB1_TRANSNT_SetQuotas, SMB2_Negotiate, SMB2_SessionSetup, SMB2_Logoff, SMB2_TreeConnect, SMB2_TreeDisConnect, SMB2_Create, SMB2_Close, SMB2_Flush, SMB2_Read, SMB2_Write, SMB2_Lock, SMB2_IOCTL, SMB2_Cancel, SMB2_Echo, SMB2_QueryDirectory, SMB2_ChangeNotify, SMB2_QueryInfo, SMB2_SetInfo, SMB2_OplockBreak, SMB2_IOCTL_GetDfsReferral, SMB2_IOCTL_IsVolumeMounted, SMB2_IOCTL_FindFilesBySID, SMB2_IOCTL_GetObjectID, SMB2_IOCTL_SetReparsePoint, SMB2_IOCTL_GetReparsePoint, SMB2_IOCTL_CreateGetObjID, SMB2_IOCTL_SetSparse, SMB2_IOCTL_SetZeroData, SMB2_IOCTL_QueryAllRanges, SMB2_IOCTL_SrvEnumSnapshot, SMB2_IOCTL_GetNTFSVolData, SMB2_IOCTL_SrvReqResumeKey, SMB2_IOCTL_SrvCopyChunk, SMB2_IOCTL_PipePeek, SMB2_IOCTL_PipeTransceive, SMB2_IOCTL_PipeWait, SMB2_IOCTL_Unknown140198, SMB2_IOCTL_IsPathnameValid, SMB2_IOCTL_SetCompression, SMB2_IOCTL_EMCCopyChunk, SMB2_IOCTL_GetCompression, SMB2_QUERYINFO_InfoFile, SMB2_QUERYINFO_InfoFileSystem, SMB2_QUERYINFO_InfoSecurity, SMB2_QUERYINFO_InfoQuota, SMB2_SETINFO_InfoFile, SMB2_SETINFO_InfoFileSystem, SMB2_SETINFO_InfoSecurity, SMB2_SETINFO_InfoQuota)

- nfs – Displays NFS protocol details (minimum, average, maximum of various NFS statistics)

7.4.2.1 Table View Sample Output (net - Network)

```
server_stats server_2 -table net -interval 3 -count 5
```

server_2 Timestamp	device	Network In Pkts/s	Network In Errors diff	Network In KiB/s	Network Out Pkts/s	Network Out Errors diff	Network Out KiB/s
11:47:52	cge-2-0	169	0	176	143	0	56
	mge0	55	0	5	52	0	9
11:47:55	cge-2-0	469	0	572	341	0	79
	mge0	0	0	0	0	0	0
	mge1	0	0	0	0	0	0
11:47:58	cge-2-0	317	0	415	214	0	32
	cge-2-1	0	0	0	0	0	0
	mge0	2	0	0	2	0	0
11:48:01	cge-2-0	619	0	821	384	0	47
	cge-2-1	0	0	0	0	0	0
	mge0	1	0	0	1	0	0
	mge1	0	0	0	0	0	0
11:48:04	cge-2-0	355	0	436	258	0	34
	cge-2-1	0	0	0	0	0	0
	mge0	1	0	0	1	0	1
	mge1	0	0	0	0	0	0

server_2 Summary	device	Network In Pkts/s	Network In Errors diff	Network In KiB/s	Network Out Pkts/s	Network Out Errors diff	Network Out KiB/s
Minimum	cge-2-0	169	0	176	143	0	32
	cge-2-1	0	0	0	0	0	0
	cge-2-2	0	0	0	0	0	0
	cge-2-3	0	0	0	0	0	0
	mge0	0	0	0	0	0	0
	mge1	0	0	0	0	0	0
Average	cge-2-0	386	0	484	268	0	50
	cge-2-1	0	0	0	0	0	0
	cge-2-2	0	0	0	0	0	0
	cge-2-3	0	0	0	0	0	0
	mge0	12	0	1	11	0	2
	mge1	0	0	0	0	0	0
Maximum	cge-2-0	619	0	821	384	0	79
	cge-2-1	0	0	0	0	0	0
	cge-2-2	0	0	0	0	0	0
	cge-2-3	0	0	0	0	0	0
	mge0	55	0	5	52	0	9
	mge1	0	0	0	0	0	0

7.4.2.2 Table View Output (data volume - DVOL)

server_stats server_2 -table dvol -interval 3 -count 5

server_2 Timestamp	dVol	Total IO Ops/s	Queue Depth	Read Ops/s	Read KiB/s	Avg Read Size Bytes	Write Ops/s	Write KiB/s	Avg Write Size Bytes	Util %	IO Retries diff	Avg Service uSec/call	Avg uSec/call
11:50:10	NBS2	43	0	0	0	-	43	28	679	1	0	236	237
	d11	17	0	2	13	8192	16	357	23356	1	0	433	433
	d12	0	0	2	16	8192	24	371	15815	1	0	-	-
	d13	0	0	0	0	-	18	179	10164	0	0	-	-
	d14	0	0	0	0	-	17	219	13171	0	0	-	-
11:50:13	NBS2	44	0	0	0	-	44	49	1144	3	0	588	590
	d11	42	0	1	8	8192	41	477	11922	2	0	2402	2397
	d12	171	0	2	13	8192	44	485	11210	3	0	504	504
	d13	0	0	0	3	8192	2	19	11469	0	0	-	-
	d14	0	0	0	0	-	4	35	8192	0	0	-	-
11:50:16	NBS2	36	0	0	0	-	36	24	689	1	0	296	297
	d11	13	0	2	16	8192	11	123	11776	0	0	202	216
	d12	0	0	0	3	8192	28	288	10409	1	0	-	-
	d14	0	0	0	0	-	0	3	8192	0	0	-	-
11:50:19	NBS2	66	0	0	0	-	66	39	605	1	0	224	225
	d11	30	0	2	13	8192	29	387	13812	1	0	203	203
	d12	0	0	1	5	8192	37	357	9979	1	0	-	-
	d13	0	0	0	0	-	17	245	15073	0	0	-	-
	d14	0	0	0	0	-	12	131	11150	0	0	-	-
11:50:22	NBS2	56	0	0	0	-	56	52	960	1	0	236	237
	d7	0	0	0	0	-	0	3	8192	0	0	6744	6746
	d8	1	0	0	0	-	1	5	8192	0	0	10356	10359
	d11	50	0	2	16	8192	48	539	11492	1	0	389	396
	d12	0	0	1	5	8192	58	563	9934	1	0	-	-
	d13	0	0	0	0	-	1	11	8192	0	0	-	-
	d14	0	0	0	0	-	6	45	8192	0	0	-	-
server_2 Summary	dVol	Total IO Ops/s	Queue Depth	Read Ops/s	Read KiB/s	Avg Read Size Bytes	Write Ops/s	Write KiB/s	Avg Write Size Bytes	Util %	IO Retries diff	Avg Service uSec/call	Avg uSec/call
Minimum	NBS1	0	0	0	0	-	0	0	-	0	0	-	-
	NBS2	36	0	0	0	-	36	24	605	1	0	224	225
	NBS3	0	0	0	0	-	0	0	-	0	0	-	-
	NBS4	0	0	0	0	-	0	0	-	0	0	-	-
	NBS5	0	0	0	0	-	0	0	-	0	0	-	-
	NBS6	0	0	0	0	-	0	0	-	0	0	-	-
	d7	0	0	0	0	-	0	0	8192	0	0	6744	6746
	d8	0	0	0	0	-	0	0	8192	0	0	10356	10359
	d11	13	0	1	8	8192	11	123	11492	0	0	202	203
	d12	0	0	0	3	8192	24	288	9934	1	0	504	504
	d13	0	0	0	0	8192	0	0	8192	0	0	-	-
	d14	0	0	0	0	-	0	3	8192	0	0	-	-
Average	NBS1	0	0	0	0	-	0	0	-	0	0	-	-
	NBS2	49	0	0	0	-	49	39	815	1	0	316	317
	NBS3	0	0	0	0	-	0	0	-	0	0	-	-
	NBS4	0	0	0	0	-	0	0	-	0	0	-	-
	NBS5	0	0	0	0	-	0	0	-	0	0	-	-
	NBS6	0	0	0	0	-	0	0	-	0	0	-	-
	d7	0	0	0	0	-	0	1	8192	0	0	6744	6746
	d8	0	0	0	0	-	0	1	8192	0	0	10357	10359
	d11	30	0	2	13	8192	29	377	14471	1	0	726	729
	d12	34	0	1	9	8192	38	413	11469	1	0	504	504
	d13	0	0	0	1	8192	8	91	11225	0	0	-	-
	d14	0	0	0	0	-	8	86	9780	0	0	-	-
Maximum	NBS1	0	0	0	0	-	0	0	-	0	0	-	-
	NBS2	66	0	0	0	-	66	52	1144	3	0	588	590
	NBS3	0	0	0	0	-	0	0	-	0	0	-	-
	NBS4	0	0	0	0	-	0	0	-	0	0	-	-
	NBS5	0	0	0	0	-	0	0	-	0	0	-	-
	NBS6	0	0	0	0	-	0	0	-	0	0	-	-
	d7	0	0	0	0	-	0	3	8192	0	0	6744	6746
	d8	1	0	0	0	-	1	5	8192	0	0	10356	10359
	d11	50	0	2	16	8192	48	539	23356	2	0	2402	2397
	d12	171	0	2	16	8192	58	563	15815	3	0	504	504
	d13	0	0	0	3	8192	18	245	15073	0	0	-	-
	d14	0	0	0	0	-	17	219	13171	0	0	-	-

7.4.2.1 Table View Output (CIFS)

server_stats server_2 -table cifs -interval 3 -count 5

server_2 Timestamp	SMB Operation	Op Calls/s	Min uSec	Max uSec	Avg uSec/call	SMB Op %
11:58:57	SMB2_Create	21	7	34011	637	1
	SMB2_Close	8	3	53239	72	1
	SMB2_Read	664	1	361307	8788	45
	SMB2_Write	682	8	297989	524	46
	SMB2_QueryInfo	81	1	41	9	6
	SMB2_SetInfo	17	4	46905	860	1
11:59:00	SMB2_Create	5	7	34011	493	0
	SMB2_Close	32	3	53239	1243	2
	SMB2_Read	627	1	361307	9671	48
	SMB2_Write	623	8	297989	444	48
	SMB2_QueryInfo	17	1	41	10	1
	SMB2_SetInfo	4	4	46905	903	0
11:59:03	SMB2_Create	3	7	34011	464	0
	SMB2_Close	27	3	53239	1770	2
	SMB2_Read	838	1	361307	7208	49
	SMB2_Write	845	8	297989	495	49
	SMB2_QueryInfo	10	1	41	9	1
	SMB2_SetInfo	2	4	46905	888	0
11:59:06	SMB2_Close	17	3	53514	3019	1
	SMB2_Read	816	1	361307	10866	51
	SMB2_Write	773	8	297989	930	48
11:59:09	SMB2_Create	12	7	34011	523	1
	SMB2_Close	32	3	53514	913	3
	SMB2_Read	400	1	361307	55908	43
	SMB2_Write	428	8	297989	437	46
	SMB2_QueryInfo	51	1	41	9	6
	SMB2_SetInfo	10	4	46905	1337	1

server_2 Summary	SMB Operation	Op Calls/s	Min uSec	Max uSec	Avg uSec/call	SMB Op %
Minimum	SMB1_Mkdir	0	0	0	-	0
	SMB1_Rmdir	0	321	4811	-	0
	SMB1_Open	0	0	0	-	0
	SMB1_Create	0	0	0	-	0
	SMB1_Close	0	5	354131	-	0
	SMB1_Flush	0	0	0	-	0
	SMB1_Unlink	0	0	0	-	0
	SMB1_Rename	0	205	812	-	0
	SMB1_GetAttr	0	0	0	-	0
	SMB1_SetAttr	0	0	0	-	0
	SMB1_Read	0	0	0	-	0
	SMB1_Write	0	13	15	-	0
	SMB1_Lock	0	0	0	-	0
	SMB1_Unlock	0	0	0	-	0
	SMB1_CreateTmp	0	0	0	-	0
	SMB1_MkNew	0	0	0	-	0
	SMB1_ChkPath	0	0	0	-	0
	SMB1_Exit	0	0	0	-	0
	SMB1_Lseek	0	0	0	-	0
	SMB1_LockRead	0	0	0	-	0
	SMB1_WriteUnlock	0	0	0	-	0
	SMB1_ReadBlockRaw	0	0	0	-	0
	SMB1_WriteBlockRaw	0	0	0	-	0
	SMB1_SetAttrExp	0	0	0	-	0
	SMB1_GetAttrExp	0	0	0	-	0
	SMB1_LockingX	0	2	20	-	0
	SMB1_Trans	0	4	69	-	0
	SMB1_TransSec	0	0	0	-	0
	SMB1_Copy	0	0	0	-	0
	SMB1_Move	0	0	0	-	0
	SMB1_Echo	0	4	18	-	0
	SMB1_WriteClose	0	0	0	-	0
	SMB1_OpenX	0	0	0	-	0
	SMB1_ReadX	0	2	566251	-	0
	SMB1_WriteX	0	4	282927	-	0

	SMB1_CloseTreeDisco	0	0	0	-	0
	SMB1_Trans2Prim	0	2	306745	-	0
	SMB1_Trans2Secd	0	0	0	-	0
	SMB1_FindClose2	0	2	5	-	0
	SMB1_FindNotifyClose	0	0	0	-	0
	SMB1_TreeConnect	0	0	0	-	0
	SMB1_TreeDisco	0	3	37	-	0
	SMB1_NegProt	0	37	21932	-	0
	SMB1_SessSetupX	0	17	251564	-	0
	SMB1_UserLogoffX	0	3	28	-	0
	SMB1_TreeConnectX	0	23	7254	-	0
	SMB1_DiskAttr	0	0	0	-	0
	SMB1_Search	0	0	0	-	0
	SMB1_FindFirst	0	0	0	-	0
	SMB1_FindUnique	0	0	0	-	0
	SMB1_FindClose	0	0	0	-	0
	SMB1_TransNT	0	1	68	-	0
	SMB1_TransNTSecd	0	0	0	-	0
	SMB1_CreateNTX	0	6	58262	-	0
	SMB1_CancelNT	0	3	17	-	0
	SMB1_NTRename	0	0	0	-	0
	SMB1_SendMessage	0	0	0	-	0
	SMB1_BeginMessage	0	0	0	-	0
	SMB1_EndMessage	0	0	0	-	0
	SMB1_MessageText	0	0	0	-	0
	SMB2_Negotiate	0	22	11526	-	0
	SMB2_SessionSetup	0	11	77909	-	0
	SMB2_Logoff	0	1	15	-	0
	SMB2_TreeConnect	0	36	8190	-	0
	SMB2_TreeDisConnect	0	6	50	-	0
	SMB2_Create	0	7	34011	464	0
	SMB2_Close	8	3	53239	72	1
	SMB2_Flush	0	0	0	-	0
	SMB2_Read	400	1	361307	7208	43
	SMB2_Write	428	8	297989	437	46
	SMB2_Lock	0	0	0	-	0
	SMB2_Ioctl	0	1	4256	-	0
	SMB2_Cancel	0	0	0	-	0
	SMB2_Echo	0	0	0	-	0
	SMB2_QueryDirectory	0	3	1348	-	0
	SMB2_ChangeNotify	0	1	28	-	0
	SMB2_QueryInfo	0	1	41	9	0
	SMB2_SetInfo	0	4	46905	860	0
	SMB2_OplockBreak	0	10	10	-	0
Average	SMB1_Mkdir	0	0	0	-	0
	SMB1_Rmdir	0	321	4811	-	0
	SMB1_Open	0	0	0	-	0
	SMB1_Create	0	0	0	-	0
	SMB1_Close	0	5	354131	-	0
	SMB1_Flush	0	0	0	-	0
	SMB1_Unlink	0	0	0	-	0
	SMB1_Rename	0	205	812	-	0
	SMB1_GetAttr	0	0	0	-	0
	SMB1_SetAttr	0	0	0	-	0
	SMB1_Read	0	0	0	-	0
	SMB1_Write	0	13	15	-	0
	SMB1_Lock	0	0	0	-	0
	SMB1_Unlock	0	0	0	-	0
	SMB1_CreateTmp	0	0	0	-	0
	SMB1_MkNew	0	0	0	-	0
	SMB1_ChkPath	0	0	0	-	0
	SMB1_Exit	0	0	0	-	0
	SMB1_Lseek	0	0	0	-	0
	SMB1_LockRead	0	0	0	-	0
	SMB1_WriteUnlock	0	0	0	-	0
	SMB1_ReadBlockRaw	0	0	0	-	0
	SMB1_WriteBlockRaw	0	0	0	-	0
	SMB1_SetAttrExp	0	0	0	-	0
	SMB1_GetAttrExp	0	0	0	-	0
	SMB1_LockingX	0	2	20	-	0
	SMB1_Trans	0	4	69	-	0

	SMB1_TransSec	0	0	0	-	0
	SMB1_Copy	0	0	0	-	0
	SMB1_Move	0	0	0	-	0
	SMB1_Echo	0	4	18	-	0
	SMB1_WriteClose	0	0	0	-	0
	SMB1_OpenX	0	0	0	-	0
	SMB1_ReadX	0	2	566251	-	0
	SMB1_WriteX	0	4	282927	-	0
	SMB1_CloseTreeDisco	0	0	0	-	0
	SMB1_Trans2Prim	0	2	306745	-	0
	SMB1_Trans2Secd	0	0	0	-	0
	SMB1_FindClose2	0	2	5	-	0
	SMB1_FindNotifyClose	0	0	0	-	0
	SMB1_TreeConnect	0	0	0	-	0
	SMB1_TreeDisco	0	3	37	-	0
	SMB1_NegProt	0	37	21932	-	0
	SMB1_SessSetupX	0	17	251564	-	0
	SMB1_UserLogoffX	0	3	28	-	0
	SMB1_TreeConnectX	0	23	7254	-	0
	SMB1_DiskAttr	0	0	0	-	0
	SMB1_Search	0	0	0	-	0
	SMB1_FindFirst	0	0	0	-	0
	SMB1_FindUnique	0	0	0	-	0
	SMB1_FindClose	0	0	0	-	0
	SMB1_TransNT	0	1	68	-	0
	SMB1_TransNTSecd	0	0	0	-	0
	SMB1_CreateNTX	0	6	58262	-	0
	SMB1_CancelNT	0	3	17	-	0
	SMB1_NTRename	0	0	0	-	0
	SMB1_SendMessage	0	0	0	-	0
	SMB1_BeginMessage	0	0	0	-	0
	SMB1_EndMessage	0	0	0	-	0
	SMB1_MessageText	0	0	0	-	0
	SMB2_Negotiate	0	22	11526	-	0
	SMB2_SessionSetup	0	11	77909	-	0
	SMB2_Logoff	0	1	15	-	0
	SMB2_TreeConnect	0	36	8190	-	0
	SMB2_TreeDisConnect	0	6	50	-	0
	SMB2_Create	8	7	34011	529	1
	SMB2_Close	23	3	53349	1403	2
	SMB2_Flush	0	0	0	-	0
	SMB2_Read	669	1	361307	18488	47
	SMB2_Write	670	8	297989	566	47
	SMB2_Lock	0	0	0	-	0
	SMB2_Ioctl	0	1	4256	-	0
	SMB2_Cancel	0	0	0	-	0
	SMB2_Echo	0	0	0	-	0
	SMB2_QueryDirectory	0	3	1348	-	0
	SMB2_ChangeNotify	0	1	28	-	0
	SMB2_QueryInfo	32	1	41	9	3
	SMB2_SetInfo	7	4	46905	997	1
	SMB2_OplockBreak	0	10	10	-	0
Maximum	SMB1_Mkdir	0	0	0	-	0
	SMB1_Rmdir	0	321	4811	-	0
	SMB1_Open	0	0	0	-	0
	SMB1_Create	0	0	0	-	0
	SMB1_Close	0	5	354131	-	0
	SMB1_Flush	0	0	0	-	0
	SMB1_Unlink	0	0	0	-	0
	SMB1_Rename	0	205	812	-	0
	SMB1_GetAttr	0	0	0	-	0
	SMB1_SetAttr	0	0	0	-	0
	SMB1_Read	0	0	0	-	0
	SMB1_Write	0	13	15	-	0
	SMB1_Lock	0	0	0	-	0
	SMB1_Unlock	0	0	0	-	0
	SMB1_CreateTmp	0	0	0	-	0
	SMB1_MkNew	0	0	0	-	0
	SMB1_ChkPath	0	0	0	-	0
	SMB1_Exit	0	0	0	-	0
	SMB1_Lseek	0	0	0	-	0

SMB1_LockRead	0	0	0	-	0
SMB1_WriteUnlock	0	0	0	-	0
SMB1_ReadBlockRaw	0	0	0	-	0
SMB1_WriteBlockRaw	0	0	0	-	0
SMB1_SetAttrExp	0	0	0	-	0
SMB1_GetAttrExp	0	0	0	-	0
SMB1_LockingX	0	2	20	-	0
SMB1_Trans	0	4	69	-	0
SMB1_TransSec	0	0	0	-	0
SMB1_Copy	0	0	0	-	0
SMB1_Move	0	0	0	-	0
SMB1_Echo	0	4	18	-	0
SMB1_WriteClose	0	0	0	-	0
SMB1_OpenX	0	0	0	-	0
SMB1_ReadX	0	2	566251	-	0
SMB1_WriteX	0	4	282927	-	0
SMB1_CloseTreeDisco	0	0	0	-	0
SMB1_Trans2Prim	0	2	306745	-	0
SMB1_Trans2Secd	0	0	0	-	0
SMB1_FindClose2	0	2	5	-	0
SMB1_FindNotifyClose	0	0	0	-	0
SMB1_TreeConnect	0	0	0	-	0
SMB1_TreeDisco	0	3	37	-	0
SMB1_NegProt	0	37	21932	-	0
SMB1_SessSetupX	0	17	251564	-	0
SMB1_UserLogoffX	0	3	28	-	0
SMB1_TreeConnectX	0	23	7254	-	0
SMB1_DiskAttr	0	0	0	-	0
SMB1_Search	0	0	0	-	0
SMB1_FindFirst	0	0	0	-	0
SMB1_FindUnique	0	0	0	-	0
SMB1_FindClose	0	0	0	-	0
SMB1_TransNT	0	1	68	-	0
SMB1_TransNTSecd	0	0	0	-	0
SMB1_CreateNTX	0	6	58262	-	0
SMB1_CancelNT	0	3	17	-	0
SMB1_NTRename	0	0	0	-	0
SMB1_SendMessage	0	0	0	-	0
SMB1_BeginMessage	0	0	0	-	0
SMB1_EndMessage	0	0	0	-	0
SMB1_MessageText	0	0	0	-	0
SMB2_Negotiate	0	22	11526	-	0
SMB2_SessionSetup	0	11	77909	-	0
SMB2_Logoff	0	1	15	-	0
SMB2_TreeConnect	0	36	8190	-	0
SMB2_TreeDisConnect	0	6	50	-	0
SMB2_Create	21	7	34011	637	1
SMB2_Close	32	3	53514	3019	3
SMB2_Flush	0	0	0	-	0
SMB2_Read	838	1	361307	55908	51
SMB2_Write	845	8	297989	930	49
SMB2_Lock	0	0	0	-	0
SMB2_Ioctl	0	1	4256	-	0
SMB2_Cancel	0	0	0	-	0
SMB2_Echo	0	0	0	-	0
SMB2_QueryDirectory	0	3	1348	-	0
SMB2_ChangeNotify	0	1	28	-	0
SMB2_QueryInfo	81	1	41	10	6
SMB2_SetInfo	17	4	46905	1337	1
SMB2_OplockBreak	0	10	10	-	0

7.4.2.2 Table View Output (NFS)

server_stats server_2 -table nfs -interval 3 -count 5					
server_2	NFS Op	NFS Op	NFS Op	NFS Op %	
Timestamp		Op	Errors	Op	
		Calls/s	diff	uSec/Call	
11:53:42	v3GetAttr	1	0	12	1
	v3Read	1	0	17244	1
	v3Write	90	0	1238	98
11:53:45	v3GetAttr	0	0	11	1
	v3Read	1	0	14	3

	v3Write	51	0	652	97
11:53:48	v3GetAttr	2	0	14	2
	v3Write	84	0	690	97
	v3Commit	0	0	489	0
11:53:51	v3Write	75	0	519	100
11:53:54	v3Write	108	0	1139	100

server_2	NFS Op	NFS Op	NFS Op	NFS Op	NFS Op %
Summary		Op	Errors	Op	
		Calls/s	diff	uSec/Call	
Minimum	v2Null	0	0	-	0
	v2GetAttr	0	0	-	0
	v2SetAttr	0	0	-	0
	v2Root	0	0	-	0
	v2Lookup	0	0	-	0
	v2ReadLink	0	0	-	0
	v2Read	0	0	-	0
	v2WriteCache	0	0	-	0
	v2Write	0	0	-	0
	v2Create	0	0	-	0
	v2Remove	0	0	-	0
	v2Rename	0	0	-	0
	v2Link	0	0	-	0
	v2Symlink	0	0	-	0
	v2Mkdir	0	0	-	0
	v2Rmdir	0	0	-	0
	v2ReadDir	0	0	-	0
	v2StatFs	0	0	-	0
	v3Null	0	0	-	0
	v3GetAttr	0	0	11	0
	v3SetAttr	0	0	-	0
	v3Lookup	0	0	-	0
	v3Access	0	0	-	0
	v3ReadLink	0	0	-	0
	v3Read	0	0	14	0
	v3Write	51	0	519	97
	v3Create	0	0	-	0
	v3Mkdir	0	0	-	0
	v3Symlink	0	0	-	0
	v3Mknod	0	0	-	0
	v3Remove	0	0	-	0
	v3Rmdir	0	0	-	0
	v3Rename	0	0	-	0
	v3Link	0	0	-	0
	v3ReadDir	0	0	-	0
	v3ReadDirPlus	0	0	-	0
	v3FsStat	0	0	-	0
	v3FsInfo	0	0	-	0
	v3PathConf	0	0	-	0
	v3Commit	0	0	489	0
	VAAI	0	0	-	0
	v4Null	0	0	-	0
	v4Compound	0	0	-	0
	v4Reserved	0	0	-	0
	v4Access	0	0	-	0
	v4Close	0	0	-	0
	v4Commit	0	0	-	0
	v4Create	0	0	-	0
	v4DelegPrg	0	0	-	0
	v4DelegRet	0	0	-	0
	v4GetAttr	0	0	-	0
	v4GetFh	0	0	-	0
	v4Link	0	0	-	0
	v4Lock	0	0	-	0
	v4LockT	0	0	-	0
	v4LockU	0	0	-	0
	v4Lookup	0	0	-	0
	v4Lookupp	0	0	-	0
	v4NVerify	0	0	-	0
	v4Open	0	0	-	0
	v4OpenAttr	0	0	-	0

	v4Open_Conf	0	0	-	0
	v4Open_DG	0	0	-	0
	v4PutFh	0	0	-	0
	v4PutpubFh	0	0	-	0
	v4PutrootFh	0	0	-	0
	v4Read	0	0	-	0
	v4ReadDir	0	0	-	0
	v4ReadLink	0	0	-	0
	v4Remove	0	0	-	0
	v4Rename	0	0	-	0
	v4Renew	0	0	-	0
	v4RestoreFh	0	0	-	0
	v4SaveFh	0	0	-	0
	v4SecInfo	0	0	-	0
	v4SetAttr	0	0	-	0
	v4SetClntid	0	0	-	0
	v4Clntid_Conf	0	0	-	0
	v4Verify	0	0	-	0
	v4Write	0	0	-	0
	v4Rel_Lockown	0	0	-	0
	v4BackChanCtl	0	0	-	0
	v4BindConn	0	0	-	0
	v4ExchangeId	0	0	-	0
	v4CreateSess	0	0	-	0
	v4DestroySess	0	0	-	0
	v4FreeStateid	0	0	-	0
	v4GetDirDeleg	0	0	-	0
	v4GetDevInfo	0	0	-	0
	v4GetDevList	0	0	-	0
	v4LayoutCmmt	0	0	-	0
	v4LayoutGet	0	0	-	0
	v4LayoutRet	0	0	-	0
	v4SecinfoNoName	0	0	-	0
	v4Sequence	0	0	-	0
	v4SetSsv	0	0	-	0
	v4TestStateid	0	0	-	0
	v4WantDeleg	0	0	-	0
	v4DestroyClid	0	0	-	0
	v4ReclaimCmpl	0	0	-	0
	v4Illegal	0	0	-	0
Average	v2Null	0	0	-	0
	v2GetAttr	0	0	-	0
	v2SetAttr	0	0	-	0
	v2Root	0	0	-	0
	v2Lookup	0	0	-	0
	v2ReadLink	0	0	-	0
	v2Read	0	0	-	0
	v2WriteCache	0	0	-	0
	v2Write	0	0	-	0
	v2Create	0	0	-	0
	v2Remove	0	0	-	0
	v2Rename	0	0	-	0
	v2Link	0	0	-	0
	v2Symlink	0	0	-	0
	v2Mkdir	0	0	-	0
	v2Rmdir	0	0	-	0
	v2ReadDir	0	0	-	0
	v2StatFs	0	0	-	0
	v3Null	0	0	-	0
	v3GetAttr	1	0	12	1
	v3SetAttr	0	0	-	0
	v3Lookup	0	0	-	0
	v3Access	0	0	-	0
	v3ReadLink	0	0	-	0
	v3Read	0	0	8629	1
	v3Write	81	0	848	98
	v3Create	0	0	-	0
	v3Mkdir	0	0	-	0
	v3Symlink	0	0	-	0
	v3Mknod	0	0	-	0
	v3Remove	0	0	-	0

v3Rmdir	0	0	-	0
v3Rename	0	0	-	0
v3Link	0	0	-	0
v3ReadDir	0	0	-	0
v3ReadDirPlus	0	0	-	0
v3FsStat	0	0	-	0
v3FsInfo	0	0	-	0
v3PathConf	0	0	-	0
v3Commit	0	0	489	0
VAAI	0	0	-	0
v4Null	0	0	-	0
v4Compound	0	0	-	0
v4Reserved	0	0	-	0
v4Access	0	0	-	0
v4Close	0	0	-	0
v4Commit	0	0	-	0
v4Create	0	0	-	0
v4DelegPrg	0	0	-	0
v4DelegRet	0	0	-	0
v4GetAttr	0	0	-	0
v4GetFh	0	0	-	0
v4Link	0	0	-	0
v4Lock	0	0	-	0
v4LockT	0	0	-	0
v4LockU	0	0	-	0
v4Lookup	0	0	-	0
v4Lookupp	0	0	-	0
v4NVerify	0	0	-	0
v4Open	0	0	-	0
v4OpenAttr	0	0	-	0
v4Open_Conf	0	0	-	0
v4Open_DG	0	0	-	0
v4PutFh	0	0	-	0
v4PutpubFh	0	0	-	0
v4PutrootFh	0	0	-	0
v4Read	0	0	-	0
v4ReadDir	0	0	-	0
v4ReadLink	0	0	-	0
v4Remove	0	0	-	0
v4Rename	0	0	-	0
v4Renew	0	0	-	0
v4RestoreFh	0	0	-	0
v4SaveFh	0	0	-	0
v4SecInfo	0	0	-	0
v4SetAttr	0	0	-	0
v4SetClntid	0	0	-	0
v4Clntid_Conf	0	0	-	0
v4Verify	0	0	-	0
v4Write	0	0	-	0
v4Rel_Lockown	0	0	-	0
v4BackChanCtl	0	0	-	0
v4BindConn	0	0	-	0
v4ExchangeId	0	0	-	0
v4CreateSess	0	0	-	0
v4DestroySess	0	0	-	0
v4FreeStateid	0	0	-	0
v4GetDirDeleg	0	0	-	0
v4GetDevInfo	0	0	-	0
v4GetDevList	0	0	-	0
v4LayoutCmmt	0	0	-	0
v4LayoutGet	0	0	-	0
v4LayoutRet	0	0	-	0
v4SecinfoNoName	0	0	-	0
v4Sequence	0	0	-	0
v4SetSsv	0	0	-	0
v4TestStateid	0	0	-	0
v4WantDeleg	0	0	-	0
v4DestroyClid	0	0	-	0
v4ReclaimCmpl	0	0	-	0
v4Illegal	0	0	-	0
Maximum	0	0	-	0
v2Null	0	0	-	0

v2GetAttr	0	0	-	0
v2SetAttr	0	0	-	0
v2Root	0	0	-	0
v2Lookup	0	0	-	0
v2ReadLink	0	0	-	0
v2Read	0	0	-	0
v2WriteCache	0	0	-	0
v2Write	0	0	-	0
v2Create	0	0	-	0
v2Remove	0	0	-	0
v2Rename	0	0	-	0
v2Link	0	0	-	0
v2Symlink	0	0	-	0
v2Mkdir	0	0	-	0
v2Rmdir	0	0	-	0
v2ReadDir	0	0	-	0
v2StatFs	0	0	-	0
v3Null	0	0	-	0
v3GetAttr	2	0	14	2
v3SetAttr	0	0	-	0
v3Lookup	0	0	-	0
v3Access	0	0	-	0
v3ReadLink	0	0	-	0
v3Read	1	0	17244	3
v3Write	108	0	1238	100
v3Create	0	0	-	0
v3Mkdir	0	0	-	0
v3Symlink	0	0	-	0
v3Mknod	0	0	-	0
v3Remove	0	0	-	0
v3Rmdir	0	0	-	0
v3Rename	0	0	-	0
v3Link	0	0	-	0
v3ReadDir	0	0	-	0
v3ReadDirPlus	0	0	-	0
v3FsStat	0	0	-	0
v3FsInfo	0	0	-	0
v3PathConf	0	0	-	0
v3Commit	0	0	489	0
VAAI	0	0	-	0
v4Null	0	0	-	0
v4Compound	0	0	-	0
v4Reserved	0	0	-	0
v4Access	0	0	-	0
v4Close	0	0	-	0
v4Commit	0	0	-	0
v4Create	0	0	-	0
v4DelegPrg	0	0	-	0
v4DelegRet	0	0	-	0
v4GetAttr	0	0	-	0
v4GetFh	0	0	-	0
v4Link	0	0	-	0
v4Lock	0	0	-	0
v4LockT	0	0	-	0
v4LockU	0	0	-	0
v4Lookup	0	0	-	0
v4Lookupp	0	0	-	0
v4NVerify	0	0	-	0
v4Open	0	0	-	0
v4OpenAttr	0	0	-	0
v4Open_Conf	0	0	-	0
v4Open_DG	0	0	-	0
v4PutFh	0	0	-	0
v4PutpubFh	0	0	-	0
v4PutrootFh	0	0	-	0
v4Read	0	0	-	0
v4ReadDir	0	0	-	0
v4ReadLink	0	0	-	0
v4Remove	0	0	-	0
v4Rename	0	0	-	0
v4Renew	0	0	-	0

v4RestoreFh	0	0	-	0
v4SaveFh	0	0	-	0
v4SecInfo	0	0	-	0
v4SetAttr	0	0	-	0
v4SetClntid	0	0	-	0
v4Clntid_Conf	0	0	-	0
v4Verify	0	0	-	0
v4Write	0	0	-	0
v4Rel_Lockown	0	0	-	0
v4BackChanCtl	0	0	-	0
v4BindConn	0	0	-	0
v4ExchangeId	0	0	-	0
v4CreateSess	0	0	-	0
v4DestroySess	0	0	-	0
v4FreeStateid	0	0	-	0
v4GetDirDeleg	0	0	-	0
v4GetDevInfo	0	0	-	0
v4GetDevList	0	0	-	0
v4LayoutCmmt	0	0	-	0
v4LayoutGet	0	0	-	0
v4LayoutRet	0	0	-	0
v4SecinfoNoName	0	0	-	0
v4Sequence	0	0	-	0
v4SetSsv	0	0	-	0
v4TestStateid	0	0	-	0
v4WantDeleg	0	0	-	0
v4DestroyClid	0	0	-	0
v4ReclaimCmpl	0	0	-	0
v4Illegal	0	0	-	0

8 Virtual Data Mover (VDM)

A Virtual Data Mover (VDM) is an EMC VNX Network Server software feature that enables the grouping of file systems and CIFS servers into virtual containers. Each VDM contains all the data necessary to support one or more CIFS servers and their file systems. The servers in a VDM store their dynamic configuration information (such as local groups, shares, security credentials, audit logs, and so on) in a configuration file system. A VDM can then be loaded and unloaded, moved from Data Mover to Data Mover, or replicated to a remote Data Mover as an autonomous unit. The servers, their file systems, and configuration data are available in one virtual container.

This technical module is part of the VNX Network Server information set and is intended primarily for VNX system administrators.

8.1 Terminology

This section defines terms important to understanding Virtual Data Mover capabilities on the VNX Network Server. The VNX Glossary provides a complete list of VNX terminology.

- **CIFS server:** A logical server using the CIFS protocol to transfer files. A Data Mover can host many instances of a CIFS server. Each instance is referred to as a CIFS server.
- **CIFS service:** A CIFS server process running on the Data Mover that presents shares on a network as well as on Microsoft Windows-based computers.
- **Control Station:** A hardware and software component of the VNX Network Server that manages the system and provides the user interface to all VNX components.
- **Virtual Data Mover (VDM):** A VNX software feature enabling users to administratively separate CIFS servers, replicate their CIFS environments, and move servers from Data Mover to Data Mover.

8.2 Restrictions

These are the restrictions for VDMs:

- In addition to CIFS servers created within a VDM, a global CIFS server is required for antivirus (CAVA) functionality. A global CIFS server is a CIFS server created at the physical Data Mover level.
- A default CIFS server and CIFS servers within a VDM cannot coexist on the same Data Mover. A default CIFS server is a global CIFS server assigned to all interfaces, and CIFS Servers within a VDM require specified interfaces. If a VDM exists on a Data Mover, a default CIFS server cannot be created.
- VDM supports CIFS and Security=NT mode only.
- The VDM feature does not support NFS, UNIX and SHARE mode, iSCSI, VNX Data Migration Service (migration file systems mounted within VDMs), or resource allocation (CPU, memory, and so on) for a VDM.
- A full path is required to back up VDM file systems with NDMP backup and server_archive. A server archive example is `server_archive <movername> -w -f /dev/cit4IO/ -J /root_vdm1/ufs1`, and an NDMP example is `/root_vdm1/ufs1`.
- IP replication failover support for local groups must include VDMs.

CAUTION: This section lists the cautions for using this feature on the VNX Network Server: The convert operation, using the `nas_server -vdm -setstate <setstate> <movername> -Convert18N` option, is irreversible. After a VDM is converted to Unicode mode, it cannot be converted to ASCII mode.

Converting the VDM using the `-Convert18N` option does not trigger the file or directory name conversion (from non-ASCII to Unicode) of user file systems mounted on the VDM. You might

have to convert the file or directory names separately using the `uc_config -convert` command if you have non-ASCII file/directory names in the user file systems.

The root file system of a Data Mover or a VDM does not support automatic file system extension.

8.3 Virtual Data Mover Concepts

Virtual Data Movers enable system administrators to group file systems and CIFS servers. Each VDM contains the necessary information to support one or more CIFS servers. Each VDM has access only to the file systems mounted to that VDM, providing a logical isolation between the VDM and the CIFS servers it contains.

Each VDM stores its configuration information in a VDM root file system, which is a directory within the root file system of the physical Data Mover. No data file systems are stored within the VDM root file system and all user data is kept in user file systems.

Use VDMs for the following:

- To partition or group file systems and CIFS servers. Users see only the file systems on the VDM.
- To easily move VDMs from one to another physical Data Mover.
- As part of a replication data-recovery strategy.

When you create a VDM, a root file system is created for the VDM. This file system stores the CIFS server configuration information for the CIFS servers created within the VDM. The VDM root file system stores the majority of the CIFS servers' dynamic data, which includes:

- Local group database for servers in the VDM
- Share database for servers in the VDM
- CIFS server configuration (computer names, interface names, and so on)
- VNX home directory information for servers in the VDM
- Auditing and Event Log information
- Kerberos information for servers in the VDM

The following figure shows the VDM container.




Figure 8-1: The VDM Container

Having the file systems and the configuration information contained in a VDM:

- Enables system administrators to separate CIFS servers and shares
- Allows replication of the CIFS environment from a source to a destination site

- Enables system administrators to move CIFS servers within a VDM from one to another physical Data Mover

8.4 Creating VDMs

The following is an example of how to create a VDM. **Please note, the examples in this section do not follow the current naming standards and are for demonstration purposes only.**

1. Create a VDM on the appropriate storage pool (in the default loaded state).

```
[nasadmin@c5-nsg8-cs0 ~]$ nas_server -name gj_vdm_02 -type vdm -create server_2 pool=clar_r6
id          = 6
name        = gj_vdm_02
acl         = 0
type        = vdm
server      = server_2
rootfs      = root_fs_vdm_gj_vdm_02
I18N mode   = UNICODE
mountedfs   =
member_of   =
status      :
  defined    = enabled
  actual     = loaded, ready
Interfaces to services mapping:
```

2. Create a user file system on the appropriate storage pool. Then create a mount point for the file system on the vdm and mount the file system.

```
[nasadmin@c5-nsg8-cs0 ~]$ nas_fs -name fs91 -create size=10240M pool=clar_r6
id          = 485
name        = fs91
acl         = 0
in_use      = False
type        = uxfs
worm         = off
volume      = v900
pool        = clar_r6
member_of   = root_avm_fs_group_17
rw_servers=
ro_servers=
rw_vdms     =
ro_vdms     =
auto_ext    = no,virtual_provision=no
deduplication = unavailable
stor_devs   = APM00083300134-006A
disks       = d9

[nasadmin@c5-nsg8-cs0 ~]$ server_mountpoint gj_vdm_02 -create /fs91
gj_vdm_02 : done

[nasadmin@c5-nsg8-cs0 ~]$ server_mount gj_vdm_02 fs91 /fs91
gj_vdm_02 : done
```

3. Configure an interface for the CIFS server. The interface will need to be named identically on any datamover hosting the VDM (i.e. interface on secondary = interface on primary).

```
[nasadmin@c5-nsg8-cs0 ~]$ server_ifconfig server_2 -create -Device fsn0_2 -name ifvdm1-3 -protocol
IP 10.40.42.219 255.255.255.0 10.40.42.255

server_2 : done
```

4. Create a CIFS server using a network interface that will be identically named at the secondary site. The secondary site IP address will be different, but the interface will be named the same. Join the CIFS server to the domain.

```
[nasadmin@c5-nsg8-cs0 ~]$ server_cifs gj_vdm_02 -add
compname=nasemctest91_A,domain=namdev.nsrootdev.net,interface=ifvdm1-3
gj_vdm_02 : done

[nasadmin@c5-nsg8-cs0 ~]$ server_cifs gj_vdm_02 -Join compname=nasemctest91_A,domain=namdev.nsr
ootdev.net,admin=gj06529,ou="ou=Servers:ou=StorageEng"
gj_vdm_02 : Enter Password:*****
done
```

8.5 Replicating and Failing Over VDMs

To setup VDM replication, a replication interconnect must be created and in working order. This interconnect needs to be between the source datamover and the destination datamover. It is assumed that the interconnect is already present.

1. Validate the interconnect:

```
[nasadmin@c5-nsg8-cs0 ~]$ nas_cel -interconnect -validate nsg8-s2_vnx5500-s2
nsg8-s2_vnx5500-s2: has 1 source and 1 destination interface(s); validating - please wait...ok
```

2. You will also need to create a VDM at the secondary site. This VDM will need identically named network interfaces to be able to come up correctly during failover. The IP address will be different, but named the same on the target datamover as it is named on the primary datamover. It is assumed that the target datamover at the secondary site has also been set up for DNS and Time services similar to the primary site. The VDM at the secondary state should be created and placed into the mounted state. This will keep the VDM inactive to users, but capable of replication. The interface for the vdm (same name from primary datamover) should also be created.

```
[nasadmin@c5-vnx5500-cs0 ~]$ nas_server -name gj_vdm_02_cob -type vdm -create se
rver_2 -setstate mounted pool=clarsas_r6
id          = 8
name        = gj_vdm_02_cob
acl         = 0
type        = vdm
server      = server_2
rootfs      = root_fs_vdm_gj_vdm_02_cob
I18N mode   = UNICODE
mountedfs   =
member_of   =
status      :
    defined = enabled
    actual  = mounted
Interfaces to services mapping:

[nasadmin@c5-vnx5500-cs0 ~]$ server_ifconfig server_2 -create -Device fsn2_0 -name ifvdm1-3 -
protocol IP 10.40.42.251 255.255.255.0 10.40.42.255
server_2 : done
```

3. Once the interconnect and secondary VDM have been established, the replication sessions can be defined. There is one session for the VDM (including its root file system), and one session for each of the file systems.

```
[nasadmin@c5-nsg8-cs0 ~]$ nas_replicate -create repl_gj_vdm_02 -source -vdm gj_vdm_02 -destination
-vdm gj_vdm_02_cob -interconnect nsg8-s2_vnx5500-s2 -max_time_out_of_sync 10
OK

[nasadmin@c5-nsg8-cs0 ~]$ nas_replicate -create repl_fs91 -source -fs fs91 -destination -pool
clarsas_r6 -vdm gj_vdm_02_cob -interconnect nsg8-s2_vnx5500-s2 -max_time_out_of_sync 10
OK
```

4. Verify the replication sessions.

```
[nasadmin@c5-nsg8-cs0 ~]$ nas_replicate -list
Name                Type          Local Mover          Interconnect          Celerra
Status
repl_gj_vdm_02      vdm           server_2             -->nsg8-s2_vnx5500-+  c5-vnx5500-+
OK
repl_fs91           filesystem    server_2             -->nsg8-s2_vnx5500-+  c5-vnx5500-+
OK

[nasadmin@c5-nsg8-cs0 ~]$ nas_replicate -info repl_gj_vdm_02
ID                  = 898_APM00083300134_0064_278_APM00111601868_2007
Name                = repl_gj_vdm_02
Source Status       = OK
Network Status      = OK
Destination Status  = OK
Last Sync Time      = Fri Dec 09 13:09:20 EST 2011
Type                = vdm
Celerra Network Server = c5-vnx5500-cs0
Dart Interconnect    = nsg8-s2_vnx5500-s2
Peer Dart Interconnect = nsg8-s2_vnx5500-s2
Replication Role     = source
Source VDM           = gj_vdm_02
Source Data Mover     = server_2
Source Interface      = 10.40.42.213
Source Control Port   = 0
Source Current Data Port = 0
Destination VDM       = gj_vdm_02_cob
Destination Data Mover = server_2
Destination Interface = 10.40.42.204
Destination Control Port = 5085
Destination Data Port = 8888
Max Out of Sync Time (minutes) = 10
Current Transfer Size (KB) = 0
Current Transfer Remain (KB) = 0
Estimated Completion Time =
Current Transfer is Full Copy = No
Current Transfer Rate (KB/s) = 0
Current Read Rate (KB/s) = 0
Current Write Rate (KB/s) = 0
Previous Transfer Rate (KB/s) = 6795
Previous Read Rate (KB/s) = 1894
Previous Write Rate (KB/s) = 499
Average Transfer Rate (KB/s) = 6795
Average Read Rate (KB/s) = 1894
Average Write Rate (KB/s) = 499

[nasadmin@c5-nsg8-cs0 ~]$ nas_replicate -info repl_fs91
ID                  = 900_APM00083300134_0064_283_APM00111601868_2007
Name                = repl_fs91
Source Status       = OK
Network Status      = OK
Destination Status  = OK
Last Sync Time      = Fri Dec 09 13:14:28 EST 2011
Type                = filesystem
Celerra Network Server = c5-vnx5500-cs0
Dart Interconnect    = nsg8-s2_vnx5500-s2
Peer Dart Interconnect = nsg8-s2_vnx5500-s2
Replication Role     = source
Source Filesystem    = fs91
Source Data Mover     = server_2
Source Interface      = 10.40.42.213
Source Control Port   = 0
Source Current Data Port = 0
Destination Filesystem = fs91_replica1
Destination Data Mover = server_2
```

```

Destination Interface      = 10.40.42.204
Destination Control Port   = 5085
Destination Data Port      = 8888
Max Out of Sync Time (minutes) = 10
Current Transfer Size (KB)  = 0
Current Transfer Remain (KB) = 0
Estimated Completion Time  =
Current Transfer is Full Copy = No
Current Transfer Rate (KB/s) = 0
Current Read Rate (KB/s)    = 0
Current Write Rate (KB/s)   = 0
Previous Transfer Rate (KB/s) = 5309
Previous Read Rate (KB/s)   = 400
Previous Write Rate (KB/s)  = 371
Average Transfer Rate (KB/s) = 5309
Average Read Rate (KB/s)   = 400
Average Write Rate (KB/s)  = 371

```

5. To initiate the failover of a VDM from the primary site to the secondary site, the failover command is invoked. Verify the status of replication.

```

[nasadmin@c5-vnx5500-cs0 ~]$ nas_replicate -failover repl_gj_vdm_02
OK

[nasadmin@c5-vnx5500-cs0 ~]$ nas_replicate -failover repl_fs91
OK

[nasadmin@c5-vnx5500-cs0 ~]$ nas_replicate -list

```

Name	Type	Local Mover	Interconnect	Celerra	Status
repl_gj_vdm_02	vdm	server_2	<--nsg8-s2_vnx5500--	c5-nsg8-cs0	Info 26045317486: Failed Over
repl_fs91	filesystem	server_2	<--nsg8-s2_vnx5500--	c5-nsg8-cs0	Info 26045317486: Failed Over

8.6 Using VDMs with 3DNS

3DNS is a type of dynamic DNS that will redirect users to an alternative IP address when the primary site IP (and port 445 for CIFS) fail.

When using 3DNS, there is a separation between the primary CIFS server that a VDM hosts and the name of the virtual CIFS server that clients attempt to connect to. These are tied together via DNS and the alternative machine names associated with the primary CIFS server.

For example, the user may expect to use a CIFS server named “nasemctest90”. This is the virtual CIFS server name that will float between the primary and secondary sites by redirecting users via IP. The CIFS server actually hosted by the VDM will have a slightly different name of “nasemctest90_A”. The name format is not important so long as it is similar to the virtual name, but not the same name.

There are a few steps which need to be taken to make CIFS shares on Celerra compatible with 3DNS failover.

8.6.1 Setting the Aliases (SPNs) for 3DNS

Note: There is a setspn command that is used to configure active directory. This tool is not a part of DART.

The following is a Microsoft article on its use. It is included with 2008:
[http://technet.microsoft.com/en-us/library/cc731241\(v=ws.10\).aspx](http://technet.microsoft.com/en-us/library/cc731241(v=ws.10).aspx)

For 2003, it is included in the Windows Server 2003 Support tools.
<http://support.microsoft.com/kb/970536>

After a CIFS server has been created for the VDM, there are several SPNs that need to be added to AD to allow for proper Kerberos security. Failure to do this may result making the share inaccessible by users or force them to use the less secure and less reliable NTLM. In order to keep clients connecting with Kerberos certificates, the following setspn commands need to be run

against the active directory domain that contains the primary CIFS server(s). This CIFS server will not be accessed directly by users.

```
C:\>setspn -A HOST/nasemctest91 nasemctest91_A
Registering ServicePrincipalNames for CN=nasemctest91_A,OU=Servers,OU=StorageEng
,DC=namdev,DC=nsrootdev,DC=net
      HOST/nasemctest91
Updated object

C:\>setspn -A HOST/nasemctest91.namdev.nsrootdev.net nasemctest91_A
Registering ServicePrincipalNames for CN=nasemctest91_A,OU=Servers,OU=StorageEng
,DC=namdev,DC=nsrootdev,DC=net
      HOST/nasemctest91.namdev.nsrootdev.net
Updated object

C:\>setspn -A HOST/nasemctest91.wlb2.namdev.nsrootdev.net nasemctest91_A
Registering ServicePrincipalNames for CN=nasemctest91_A,OU=Servers,OU=StorageEng
,DC=namdev,DC=nsrootdev,DC=net
      HOST/nasemctest91.wlb2.namdev.nsrootdev.net
Updated object

C:\>setspn -A HOST/nasemctest91_A.namdev.nsrootdev.net nasemctest91_A
Registering ServicePrincipalNames for CN=nasemctest91_A,OU=Servers,OU=StorageEng
,DC=namdev,DC=nsrootdev,DC=net
      HOST/nasemctest91_A.namdev.nsrootdev.net
Updated object
```

This covers the combination of the simple server names and FQDN as well as the 3DNS sub-domain entry (wlb or wlb2).

To check the current machine names for a AD server, run the list command:

```
C:\>setspn -L nasemctest91_A
Registered ServicePrincipalNames for CN=nasemctest91_A,OU=Servers,OU=StorageEng,
DC=namdev,DC=nsrootdev,DC=net:
      HOST/nasemctest91.wlb2.namdev.nsrootdev.net
      HOST/nasemctest91.namdev.nsrootdev.net
      HOST/nasemctest91
      host/nasemctest91_A.namdev.nsrootdev.net
      host/nasemctest91_A
```

8.6.2 Disable Name Checking on Datamover for 3DNS

You will need to set a parameter change on the datamover to disable name checking. Without this parameter change, access to the share will be prevented.

Set the boot parameter value:

```
[nasadmin@c5-nsg8-cs0 ~]$ server_param server_2 -f cifs -m LanmanServer.disableNameChecking -v 1
server_2 : done

Warning 17716815750: server_2 : You must reboot server_2 for LanmanServer.disableNameChecking
changes to take effect.
```

Verify the changed parameter value:

```
[nasadmin@c5-nsg8-cs0 ~]$ server_param server_2 -f cifs -info LanmanServer.disableNameChecking -v
server_2 :
name                = LanmanServer.disableNameChecking
facility_name        = cifs
default_value        = 0
current_value        = 1
configured_value     = 1
user_action          = reboot DataMover
change_effective     = reboot DataMover
range                = (0,1)
description          = Disables checking of the server's principal name of the client's
kerberos ticket.

detailed_description
```

When set to 1, this parameter disables the control of the server's principal name of the client's kerberos ticket. The client is then allowed to connect with a DNS alias (please refer to Microsoft article 281308). When set to 0, the client is only allowed to connect using the primary computer name.

When set to 1, this parameter disables the control of the server's principal name of the client's Kerberos ticket. The client is then allowed to connect with a DNS alias (please refer to Microsoft article 281308). When set to 0, the client is only allowed to connect using the primary computer name. A reboot is required to make this parameter change and is the recommended method.

Under certain conditions, it may not be possible to reboot the datamover. To make the parameter change without a reboot of the datamovers, the following command can be run:

```
[nasadmin@c5-nsg8-cs0 ~]$ server_param server_2 -f cifs -m LanmanServer.disableNameChecking -v 1
server_2 : done

Warning 17716815750: server_2 : You must reboot server_2 for LanmanServer.disableNameChecking
changes to take effect.
[nasadmin@c5-nsg8-cs0 ~]$
[nasadmin@c5-nsg8-cs0 ~]$ .server_config server_2 -v "param cifs
LanmanServer.disableNameChecking=1"
server_2 : commands processed: 1
command(s) succeeded
output is complete

1323458685: ADMIN: 6: Command succeeded: param cifs LanmanServer.disableNameChecking=1
```

To verify the current running value:

```
[nasadmin@c5-nsg8-cs0 ~]$ .server_config server_2 -v "param cifs LanmanServer.disableNameChecking"
server_2 : commands processed: 1
command(s) succeeded
output is complete

cifs.LanmanServer.disableNameChecking          INT 0x0004317150 1 0 (0,1) TRUE REBOOT
'Disables checking of the server's principal name of the client's kerberos ticket.'
1323458696: ADMIN: 6: Command succeeded: param cifs LanmanServer.disableNameChecking
```

The value 1 0 above indicates the current value is 1 or disable checking is turned on. A value of 0 0 would indicate that disable checking is turned off.

8.6.3 3DNS Port/IP Monitoring

3DNS should be configured as port based checking, so that both the port and IP address need to be alive to direct users to that site. The primary site is heavily favored in the decision process. Within seconds of losing the primary site (port or IP address) users will be redirected to the secondary IP address. Any active session would need to be re-requested to see the updated DNS entry. Unfortunately, due to a limitation in DART, the Celerra does not fully close the port (responds to a connect rather than connection refused) after a CIFS server has been used on the IP address even if the VDM has been failed over or if the interface was removed from the CIFS server directly. Therefore, as a part of the failover procedures, the NAS administrators will need to disable the primary IP interface to force 3DNS to redirect users.

```
[nasadmin@c5-nsg8-cs0 ~]$ server_ifconfig server_2 ifvdm1-3 down
server_2 : done
```

8.6.4 Verify the User Connections (CIFS / SMB Connections)

Verify the users are connecting to the CIFS share via Kerberos using the audit option on server_cifs. There will be a record for each users' access to the shares on the VDM. Kerberos will be listed on a correctly configured connection. NTLMSSP will be listed on an incorrectly configured connection (i.e. using IP address or if the setspn commands are not set correctly). If users cannot connect to the share at all, then likely the parameter changes have not been set on the datamover.

```
[nasadmin@cs5-vnx5500-cs0 ~]$ server_cifs gj_vdm_01_cob -o audit
gj_vdm_01_cob :

||| AUDIT Ctx=0x000aee9408, ref=2, W2K3_R2 Client(10.40.43.182) Port=1729/445
||| NASEMCTEST90_A[NAMDEV] on if=ifvdm1-2
||| CurrentDC 0x0070827408=NAMDEV001
||| Proto=NT1, Arch=Win2K, RemBufsz=0xffff, LocBufsz=0xffff, popupMsg=1
||| 0 FNN in FNNlist NbUsr=1 NbCnx=1
||| Uid=0x3f NTcred(0x02b75e1c08 RC=2 KERBEROS Capa=0x2) 'NAMDEV\gj06529'
||| Cnxp(0x0306efa008), Name=fs90, cUid=0x3f Tid=0x3f, Ref=1, Aborted=0
||| readOnly=0, umask=22, opened files/dirs=1
||| Absolute path of the share=\root_vdm_7\fs90
total smb_streamCtx: 1

||| AUDIT Ctx=0x032f4b7008, ref=2, W2K3 Client(GDSEHDS01) Port=1479/445
||| NASEMCTEST90_A[NAMDEV] on if=ifvdm1-2
||| CurrentDC 0x0070827408=NAMDEV001
||| Proto=NT1, Arch=Win2K, RemBufsz=0xffff, LocBufsz=0xffff, popupMsg=1
||| 0 FNN in FNNlist NbUsr=1 NbCnx=2
||| Uid=0x3f NTcred(0x0064d0c408 RC=3 NTLMSSP Capa=0x4001) 'NAMDEV\gj06529'
||| Cnxp(0x01942d1808), Name=IPC$, cUid=0x3f Tid=0x3f, Ref=1, Aborted=0
||| readOnly=0, umask=22, opened files/dirs=0
||| Absolute path of the share=\root_vdm_7\etc
||| Cnxp(0x00647b0008), Name=fs90, cUid=0x3f Tid=0x40, Ref=1, Aborted=0
||| readOnly=0, umask=22, opened files/dirs=1
||| Absolute path of the share=\root_vdm_7\fs90
total smb_streamCtx: 2
```

Note: In the examples above, the Kerberos (correct) connection is highlighted in green and the NTLM (incorrect) connection is highlighted in red.

8.6.5 Verifying the 3DNS Failover

To initiate the 3DNS failover, the following steps should be taken.

- Run the failover commands for the VDM
- Set the interface on the primary site as down

To verify the 3DNS configuration is working, the IP address (via ping or nslookup) of the virtual CIFS server (nasemctest91) and the real CIFS server (nasemctest91_A) should be the same value and at the remote site. Users should be able to connect to the share and be using Kerberos certificate security.

8.6.6 Troubleshooting

If SPNs have been generated but the parameter has not been applied users will experience access problems. It is recommended to apply the parameter and reboot the datamover. If the parameter cannot be applied then the following workarounds can also be used to resolve access problems.

- Removing the SPN from AD by running setspn -D on a domain controller. The CIFS client accesses the Data Mover through its DNS alias, without SPN. The clients will not be able to obtain Kerberos tickets for the CIFS server's service principals, but NTLM or NTLMV2 can be used to connect to the Data Mover's shares. This may take some time to propagate in AD. This should only be used temporarily.
- Users could access the real CIFS server (nasemctest90_A)

8.6.7 Preventing Client Redirection during DM Reboots or Panics

The 3DNS environment should allow for both unexpected panics of datamovers and upgrades (planned) due to the manual control of the secondary site IP interface. The manual control will prevent user redirects from the primary site until that is a desired condition. Due to the current limitation in the DART that allows the DM to respond to port 445 even if there is not an active

CIFS server, the IP interface at the secondary site should be disabled. This interface needs to be turned back on before failing over the VDM to the secondary site.

```
[nasadmin@c5-vnx5500-cs0 ~]$ server_ifconfig server_2 ifvdm1-3 down  
server_2 : done
```

9 File-Level Retention (FLR / WORM)

File Level Retention (FLR) is software that runs on the EMC NAS platform to prevent the modification or deletion of data until a set retention period is reached. In essence it provides WORM (Write Once Read Many) storage on NAS.

There are two variants of FLR:

- Enterprise – Whilst this prevents changes or modification of data from an application it can be overridden by storage administrators.
- Compliance – This is fully compliant with SEC rule 17a4 and prevents all changes or modification of data.

Both forms of FLR are licensed software; these licenses must be added to the NAS storage before using FLR.

FLR is applied to volumes at volume creation time and each volume has its own compliance clock, it is a comparison with this clock that is used to determine whether a file has reached its expiration. The use of an independent and non-modifiable compliance clock prevents a storage administrator setting the data movers own clock forward and hence allowing deletion of files before they have reached their retention time.

Each volume has its own FLR compliance clock; this clock has the following properties:

- The compliance clock is set using DM clock on the volumes first mount
- The clock is only updated when the volume is mounted for read/write (or for a replica when the compliance clock is replicated from the primary copy).
- The clock on volumes not mounted are frozen, when mounted again they will begin incrementing again from the time when they were last mounted.
- The clock is replicated to remote copy via normal synchronization updates.
- The clock can catch up to the DM clock by 138 seconds per hour e.g. after the clock has not been updated for a period after either being dismounted or mounted read-only .

Whilst an individual files state cannot be read directly from the NAS array its state can be implied from the files meta-data. The possible file states are as follows:

- Not locked - Normal file access (i.e. no retention set)
- Locked – Retention has been applied. No modification, deletes etc are permitted until the file has reached its retention time.
- Append-only – Data can only be written to the end of the file. Existing contents cannot be modified or deleted. While an append-only file does not have a retention time, it is usual to convert an append-only file to a normal locked FLR file when no further updates are required, which will set a retention time.
- Expired - The retention period on the file has been reached and the file may now be deleted or modified again, just like a file that is not locked.

A file is locked by setting a retention date on the file. From CIFS, this is achieved by setting the file's "Last Accessed Time" (LAT) to the date/time when the retention will expire and then setting the read-only bit on the file. On Windows setting the LAT can be either achieved by using standard file system API calls or by using the FLR tools supplied by EMC (these EMC tools have not been certified by CATE), you cannot use Windows Explorer to set the file access time directly. From Unix the procedure is similar with the LAT being set from the command line and the file being set to read-only by removing all write permissions and setting the read bits. Whilst the time used to set the LAT is the local time of the server setting the retention time, all retention times are set and compared as the equivalent GMT/UTC time. For example setting a retention time of 12pm on a server using EST will set a time of 5pm GMT/UTC on the file itself, it is this GMT/UTC form of the retention time that is quoted in the FLR security audit logs.

In either case it is not possible to set a retention time behind the current compliance clock, **if a time is set earlier than the current compliance clock time the retention time is set to infinite**. It is however possible to change a file retention period whilst it is locked by setting a new LAT, this new LAT must be a later date than the original retention period, it is not possible to set an earlier date than the existing retention time for locked files.

If no LAT is set prior to setting the file to read-only the retention date defaults to infinite, an infinite retention time will display as “Dec 31 1969”. A file with an infinite retention time can, as the name implies, never be deleted.

All files should be set with a specific retention time before they are changed to read-only to prevent files having infinite retention times.

The contents of a locked file cannot be changed, neither can the ownership or file permissions. It is however possible to change the ownership of append-only files as for a normal file.

EMC compliance storage should be allocated on a dedicated array and isolated from non-compliant workloads.

9.1 Compliance version of File Level Retention (FLR-C)

To be SEC Rule 17a4(f) compliant the FLR-C variant of FLR is required but to achieve full compliance it also requires that write verification to be enabled on volume. Write verification reads data back from disk after it has been written and compares it to the original data to ensure that the data has been written correctly. When reading data this data if after 3 attempts the verification still shows an error this is logged and the DM panics to ensure that the data corruption event is investigated.

Write verification is enabled on the volume using the command shown below:

```
$ server_param <movername> -facility FLRCompliance -modify writeverify -value 1
```

The value “1” enables write verification whilst “0” disables verification, “0” is the default. EMC performance information states that enabling write verification reduces write performance by between 20-40%, write verification is also enabled for replicated data and hence this will impact on replication speeds.

9.2 Enterprise version of File Level Retention (FLR-E)

The FLR-E variant of FLR does **not** meet SEC Rule 17a4(f) but does provide WORM storage for other uses.

The main difference between FLR-C and FLR-E are as follows:

- A file system containing locked files **cannot** be deleted when using FLR-C but **can** be deleted if the file system uses FLR-E. This means that using FLR-E locked files can be removed by a storage administrator by deleting the whole file system. (Neither FLR-C nor FLR-E can protect against the underlying storage being deleted or reconfigured at the underlying block storage layer).
- A file protected with an infinite retention period may have this retention period reduced when using an FLR-E system, this is not possible for FLR-C systems.
- Whilst it is possible to create a writeable checkpoint on a FLR-E protected file system it is not possible to create a writeable checkpoint of a FLR-C protected file system.

Retention times are set and enabled in the same way for both FLR-C and FLR-E thus from an application perspective they act identically.

If SEC or a similar level of compliance is **not** required then it is recommended that the Enterprise version of FLR is used.

9.3 Configuration

FLR is enabled at volume time using the command

```
$ nas_fs -name <name> -create [size=<integer>[T|G|M]] pool=<pool> [worm={enterprise|compliance|off}]
```

If a value for the worm flag is not provided the volume will **not** be configured with FLR, this is the default.

The FLR status and the FLR clock for a file systems can be displayed using the following command.

```
$nas_fs -info <fs_name>
```

The field “worm=” shows the compliance mode (enterprise or compliance) and the “worm_clock=” shows the current compliance clock for the volume. The “worm max retention date=” shows the most distant retention date set, after this date the volume may be deleted (FLR-E allows deletion before this date has been reached) provided that no new files of a later retention period are written.

9.4 Replication

Replication of volumes is permissible and indeed required for Citi systems to protect the data.

- Both the source and destination file systems must have the same FLR type enabled
- The destination cannot contain existing FLR protected files when the replication is enabled.
- It is not possible to reverse the replication session if unique changes have been made to the original source volume that is about to become the destination.

The source volumes compliance clock is replicated along with the data to the target volume, as replication is asynchronous the compliance clock of the target volume will lag the primary volume and remains at the time when the last replication cycle completed. If the target volume is brought online its compliance clock will start incrementing as normal.

9.5 FLR Activity log

The EMC NAS logs its FLR based activity in a file located in the root of the FLR filesystem. The maximum size of an individual log file is 10MB and once a file is full a new log is created and used. This log is an append-only file (i.e. it cannot be deleted) and once a log file is full it is converted to a normal FLR protected file with a retention period set based on the datamover setting (see below). Whilst the system will automatically create new log files it is the administrator's responsibility to delete log files that are no longer required once the retention period expires.

The default activity log retention period is 7 days but this should be changed to a value the same as the maximum retention period for the data held on this volume as shown below (the example shows a figure for retention of 3 years).

```
$ server_param {<movername>|ALL} -facility FLRcompliance -modify ActivityLogRP -value 1095
```

Log files have the following filename format:

```
flrLogyyyyymmddhhmm
```

where the yyyyymmddhhmm timestamp suffix is the FLR clock time when the file was created

FLR logs the following actions:

- An append-only file is created
- A file is set to the locked state
- A retention date is extended for a locked or extended file
- An operation is attempted on a locked or expire file to make it writeable
- An operation deletes or attempts to delete a protected file (locked, expired, append-only).

The activity log records events against inode numbers not against filenames. Log entries also show the UID/GID of the users performing action hence for CIFS users this must be translated to a specific CIFS user using the `server_cifssupport` command e.g.

```
$ server_cifssupport <movername> -list -secmap -uid <uid>
```

or

```
$ server_cifssupport <movername> -list -secmap -gid <gid>
```

Below are some examples of audit log event for typical events:

```
Tue May 31 12:52:42 2011 : Worm commit clean file : Inode No = 23 : Generation No
= 1306846283 : Uid = 939001 : Gid = 9001 : FileMode = 444 : FileSize = 4 : RP =
Tue May 31 17:00:00 2011 : Passed
```

```
Tue May 31 14:54:13 2011 : Make writeable worm committed file : Inode No = 22 :
Generation No = 1306844640 : Uid = 939001 : Gid = 9001 : FileMode = 644 : FileSize
= 17 : RP = Tue May 31 20:00:00 2011 : Failed
```

```
Wed Jun 1 13:40:30 2011 : Create a worm append-only file : Inode No = 25 :
Generation No = 1306917771 : Uid = 939001 : Gid = 9001 : FileMode = 644 : FileSize
= 0 : RP = Wed Jun 1 14:00:00 2011 : Passed
```

```
Wed Jun 1 13:56:25 2011 : Make writeable worm committed file : Inode No = 25 :
Generation No = 1306935741 : Uid = 939001 : Gid = 9001 : FileMode = 644 : FileSize
= 19 : RP = Wed Jun 1 14:00:00 2011 : Failed
```

```
Wed Jun 1 13:41:17 2011 : Delete worm append-only file : Inode No = 25 :
Generation No = 1306917771 : Uid = 939001 : Gid = 9001 : FileMode = 644 : FileSize
= 0 : RP = Wed Jun 1 14:00:00 2011 : Passed
```

```
Thu Jun 2 15:37:28 2011 : Extend Retention Period on worm committed file : Inode
No = 30 : Generation No = 1307028990 : Uid = 939001 : Gid = 9001 : FileMode = 444
: FileSize = 0 : Old RP = Thu Jun 2 16:00:00 2011 : New RP = Thu Jun 2 20:00:00
2011 : Passed
```

```
Thu Jun 2 15:37:37 2011 : Delete worm committed file : Inode No = 30 : Generation
No = 1307028990 : Uid = 939001 : Gid = 9001 : FileMode = 444 : FileSize = 0 : RP =
Thu Jun 2 20:00:00 2011 : Failed
```

```
Thu Jun 2 15:39:52 2011 : Shorten Retention Period on worm committed file : Inode
No = 30 : Generation No = 1307028990 : Uid = 939001 : Gid = 9001 : FileMode = 444
: FileSize = 0 : Old RP = Thu Jun 2 20:00:00 2011 : New RP = Thu Jun 2 18:00:00
2011 : Failed
```

The activity log records events against inode numbers not against filenames. Log entries also show the UID/GID of the users performing action hence for CIFS users this must be translated to a specific CIFS user using the `server_cifssupport` command e.g.

```
$ server_cifssupport <movername> -list -secmap -uid <uid>
```

or

```
$ server_cifssupport <movername> -list -secmap -gid <gid>
```

It is possible that the mapping of UID/GID to CIFS users may change over time, for example if multi-protocol access to data is configured as the UID/GID used may be obtained from LDAP, local mapping files or the usermapper database.

The audit log entry timestamps are based on the current DM time even though checking is performed against the compliance clock. It should also be noted that the retention times are shown in UTC.

9.6 Anti-Virus

Whilst FLR protected files may be scanned it is not possible to clean or delete such files if a virus is discovered. However, a file is only protected by FLR after it has been first written to a volume, its retention period set and finally its permissions set to read-only. This means that when a file is first written to the EMC NAS, and before it is protected by FLR, it will still be scanned and cleaned if deemed infected. Hence the issue of cleaning infected FLR protected files is mainly limited to files containing viruses that were not identified on the initial scan but are later identified when the virus definitions on the AV scanners are updated. The other main avenue for AV scanning issues for FLR protected files is for append-only files as the contents of these files are protected as soon as the data is written.

9.7 Deduplication

Both deduplication and compression are supported on both FLR-E and FLR-C volumes.

10 NDMP Backups

Due to the different backup products used globally within Citi, this section will only include configuration information that is specific to the NAS device itself, or backup infrastructure details that are product neutral. Support groups are advised to read through the section to familiarize themselves with the topology and NDMP requirements. NAS operations groups should document their specific backup implementations as part of their CITMP level 3 documentation.

There is also a CATE Storage and Backup Engineering Solutions Guide for NAS backups that covers the EMC and NetApp NAS. The latest version of this solutions guide can be found at: <https://catecollaboration.citigroup.net/domains/platstor/gdse/products/backup/Netbackup/CSBE-10047-SG-v01.pdf>

The website that includes previous versions and related documentation can be found at: <https://catecollaboration.citigroup.net/domains/platstor/gdse/products/techstdshmpgs/netbackup.aspx>

10.1 NDMP Overview

Network Data Management Protocol (NDMP) controls backups and restores on the VNX File Server. NDMP provides a standard protocol by which network data can be backed up and restored, and by which network storage devices such as tapes and libraries can be controlled.

The Data Mover maintains a state machine for each NDMP client connection that executes and maintains the backup and restore processes. Each Data Mover supports 8 concurrent NDMP sessions at one time. For example, server_2 can run seven backup sessions and one restore session, simultaneously. For more details see [NDMP Datastreams](#).

10.2 Authorizing Access to the NDMP Host

The server controlling the backup job will need to be authorized for NDMP backups of the target VNX. Refer to the backup software vendor's documentation for this step.

NDMP configuration for the backup software and VNX Network Server requires setting a trusted username and password for each NDMP-Host Data Mover.

10.2.1 Authorization with NIS or LDAP

For Network Information Service (NIS) or Lightweight Directory Access Protocol (LDAP) installations, the username and password must match the username and password in the NIS database.

Note: NIS or LDAP must be enabled on every NDMP-Host Data Mover.

The same username and password should be typed when you configure the backup software.

10.2.2 Authorization without NIS or LDAP

In this case, a local account must be created and used for the NDMP account. Use the server_user command to configure a local account on each datamover that will be used to manage the NDMP backups.

```
[nasadmin@c5-vnx5500-cs0 ~]$ /nas/sbin/server_user server_2 -add -md5 -password
ndmp_user2
Creating new user ndmp_user2
User ID: 1001
Group ID: 1001
Home directory:
Changing password for user ndmp_user2
New passwd:
Retype new passwd:
```

Note: In the output, the two mandatory fields, User ID (UID) and Group ID (GID), are integers. The VNX Network Server uses UNIX-style UIDs and GIDs to record the ownership of files and directories. The UID of the root user is 0 (do not use). The password you assign to the Data Mover can contain between six and eight characters.

10.3 Backup of Checkpoints

Regular backups will be made from a snapshot/checkpoint on a COB VNX. The backup schedule will typically include a monthly full backup, weekly cumulative incremental, and daily differential incremental backups.

To integrate a checkpoint into the NDMP backup process and automate the checkpoint creation, management, and deletion tasks, upon initiation of backups, set the “**SNAPSURE = y**” Environment variable, in the NetBackup Policy’s Backup Selection List. For example, the backup a snapshot of volume, /fs1, the Backup Selection list would contain the following.

```
set SNAPSURE = y
/fs1
```

Or, in the case of backing up a snapshot that has already been created, use the snapshot path, in the Backup Selection List.

Note the following restrictions, when adding volumes to the Backup Selection List.

- You cannot back up files where the path length is greater than 1024 characters.
- In the file list for an NDMP Policy, you can include only directory paths and cannot use wildcards.

The checkpoint that is automatically created in the backup process is given a unique, system-defined name that appears in the automaticNDMPckpts directory, and in the SnapSure checkpoints list, as follows:

```
automaticTempNDMPckpt<id>-<srcFsid>-<timestamp>
where:
  <id> = the checkpoint ID
  <srcFsid> = the ID of the source file system (being backed up using the
  checkpoint)
  <timestamp> = the date and time of checkpoint creation
  Checkpoints with this name are for NDMP use only and must not be unmounted,
  refreshed, or deleted
```

10.4 Backup configurations


There are several configurations of backups that are possible with NDMP. These configurations include:

- Local NDMP – Tape drives are SAN attached to fiber on each datamover
- Remote NDMP to Tape – Media servers are SAN attached to tape drives
- Remote NDMP to Disk – Media servers are SAN attached to disk-based virtual tape
- Three Way NDMP – Redirection of SAN attached tape drives

The recommended configuration for Citi is to use Remote NDMP to Tape. Local NDMP is an alternative configuration allowed.

10.4.1 Remote NDMP to Tape


Under this configuration, the NAS data is backed up over the network to a NetBackup Enterprise Media server.



10.4.2 Remote NDMP to Tape

The most basic is known as local NDMP, where tape drives are attached to NAS Filers either using SCSI or SAN. This approach works well for NAS systems with large amounts of data because it keeps data from flowing across the network, preventing network congestion. Either physical or virtual tape drives may be used. This essentially makes the NAS Array a NetBackup Media Server.

NetApp Filer and EMC Celerra backing up, over SAN, using NDMP



10.5 NDMP Operating Notes and Restrictions

The following are restrictions to NDMP backups:

- A tape created on an NDMP storage unit is in backup format and cannot be restored from a non-NDMP storage unit.
- You cannot back up files where the path length is greater than 1024 characters.
- In the file list for an NDMP class, you can include only directory paths and cannot use wildcards.
- You should turn off (disable) CAVA on the VNX during restores
- Checkpoints should be deleted or disabled before restores
- When restoring a single file of a volume backup, it will take at least the amount of time it took to back up the entire volume.

10.6 NDMP Volume Backup (NVB)

NDMP Volume Backup (NVB) is an EMC NDMP backup mechanism that backs up data blocks at a volume level, rather than at a file level. NVB reads a set of disk data blocks in an efficient manner compared to the method used for traditional, file-based backups. NVB is commonly known as VBB.

VBB can be enabled on NDMP clients in two different ways.

- Setting 'vbb' as the backup type. In Legato NetWorker, this means using the command line '`nsrndmp_save -Type vbb`' instead of '-Type tar' or dump.
- Setting the environment variable vbb to y.

VBB does not support

- File Filtering for backups
- 3-Way NDMP
- Filemover

Note: VBB is not a CATE supported backup configuration at this time.

10.7 NDMP Datastreams

The datamovers on most VNX models support up to 8x datastreams (simultaneous data sessions) for NDMP backups. There is a minimum of 8GB of memory on the datamover for 8x datastreams. The 6GB memory models only support 4x datastreams.

VNX Model	Number of Data Streams for NDMP
5300	4 (due to 6GB RAM)
5500	8
5700	8
7500	8
VG2	4 (due to 6GB RAM)
VG8	8
5200	4 (due to 6GB RAM)
5400	4 (due to 6GB RAM)
5600	8
5800	8
7600	8
8000	8

The default configuration is 4x datastreams, but is easily changed to 8x by modifying the NDMP parameter. This change requires a datamover reboot to be effective.

Note: The backup jobs will also need to be modified to take advantage of this. Changing the backup jobs only without the parameter change will result in failed jobs.

```
[nasadmin@c5-vnx5500-cs0 ~]$ server_param server_2 -facility NDMP -info
concurrentDataStreams
server_2 :
name                = concurrentDataStreams
facility_name        = NDMP
default_value        = 4
current_value        = 4
configured_value     =
user_action          = reboot DataMover
change_effective     = reboot DataMover
range                = (1,8)
description          = Number of concurrent data streams

[nasadmin@c5-vnx5500-cs0 ~]$ server_param server_2 -facility NDMP -modify concurrentDataStreams -value 8
server_2 : done
Warning 17716815750: server_2 : You must reboot server_2 for concurrentDataStreams changes to take effect.

[nasadmin@c5-vnx5500-cs0 ~]$ server_cpu server_2 -reboot -monitor now
server_2 : reboot in progress 3.3.3.3.3.3.3.3.3.4.4.done

[nasadmin@c5-vnx5500-cs0 ~]$ server_param server_2 -facility NDMP -info
concurrentDataStreams
server_2 :
name                = concurrentDataStreams
facility_name        = NDMP
default_value        = 4
current_value        = 8
configured_value     = 8
user_action          = reboot DataMover
change_effective     = reboot DataMover
```

range	= (1,8)
description	= Number of concurrent data streams

10.8 Average File Size

There is a well known issue with backing up NAS volumes that contain many small files, where the backup takes a very long time to complete, and uses tape drives at a very slow pace. This is a big problem, since, not only does it slow down the speed of the tape drives and contribute to their wear (i.e. shoe shining the drives), it also causes other backups to have to wait and often miss their backup window.

Both the number and size of files contribute to this problem. For example, backing up 2 million 512 KB files may run at a reasonable speed, but backing up 2 million 10 KB files may run at a very low speed.

For VNX Data Movers, the `server_df` command can be used to determine the average size and the number of files on a Data Mover. For example, to determine this information for volume `ndmp_fs1`, on Data Mover `server_2`, run the following commands, then divide used Kbytes by used inodes:

```
[nasadmin@c5-vnx5500-cs0 ~]$ server_df server_2 ndmp_fs1
server_2 :
Filesystem      kbytes      used      avail capacity Mounted on
ndmp_fs1        10325744    3117088    7208656   30%    /ndmp_fs1
[nasadmin@c5-vnx5500-cs0 ~]$ server_df server_2 -i ndmp_fs1
server_2 :
Filesystem      inodes      used      avail capacity Mounted on
ndmp_fs1        1259518    1470    1258048    0%    /ndmp_fs1
```

$$\text{Average file size} = \frac{3117088 \text{ Kb}}{1470 \text{ inodes}} = 2120 \text{ Kb per file}$$

10.9 Exclusions Lists for Backups

To exclude files or directories from a backup, the following methods can be used. This is typical for static or archived data that no longer requires regular backups.

10.9.1 Directory Exclusions

Action
To exclude a specific directory from the NDMP backup, type: <pre>EMC_EDIR<nn>=<valid directory pattern></pre> where: <code>nn</code> = an ASCII number from 01 to 50 that helps you keep track of the total number of variable entries. By default, five variable entries are allowed per NDMP backup session for directory filtering. <code>valid directory pattern</code> = a pattern that identifies the specific type of directory to be excluded from the backup. Note: You can increase the number of directory filters, up to 50, using the <code>filter.numDirFilter</code> parameter. "Modify parameters" on page 44 explains how. Example: To exclude the directory <code>ufs1/dir1</code> from the NDMP backup, type: <pre>EMC_EDIR01=/ufs1/dir1</pre> To also exclude all directories that begin with "d" under <code>/ufs2</code>, type: <pre>EMC_EDIR02=/ufs2/d*</pre> Note: If you are using the Celerra internationalization feature with a character set other than UTF-8, use the <code>filter.dialect</code> parameter to translate the pattern received to UTF-8. The default value of the parameter is <code>.</code>. After the pattern is converted to a UTF-8 pattern, the valid path rules are applied, and if the criteria are met, the pattern is accepted.

10.9.2 File Exclusions for Backups

Action
To exclude a specific file from the NDMP backup, type: <pre>EMC_EFILE<nn>=<valid file pattern></pre> where: <code>nn</code> = an ASCII number from 01 to 50 that helps you keep track of the total number of variable entries. By default, five variable entries are allowed per NDMP backup session for file filtering. Note: You can increase the number of file filters, up to 50, using the <code>filter.numFileFilter</code> parameter. "Modify parameters" on page 44 explains how. <code>valid file pattern</code> = a pattern that identifies the specific type of file to be excluded from the backup. Example: To exclude all MP3 files from the NDMP backup, type: <pre>EMC_EFILE01=*.mp3</pre> To also exclude all files that begin with "temp" from the backup, type: <pre>EMC_EFILE02=temp*</pre> To also exclude all WMA files from the backup, type: <pre>EMC_EFILE03=*.wma</pre> Note: If you are using the Celerra internationalization feature with a character set other than UTF-8, use the <code>filter.dialect</code> parameter to translate the pattern received to UTF-8. The default value of the parameter is <code>.</code>. After the pattern is converted to a UTF-8 pattern, the valid path rules are applied, and if the criteria are met, the pattern is accepted.

10.9.3 Valid path rules for excluding directories from backup

Follow these rules to specify the valid path:

- The path must start with a forward slash (/) and the file system name.
- The path can contain a maximum of 1,023 ASCII characters.
- Supported wildcards are: question mark (?), which matches any single character, and asterisk (*), which matches any sequence of zero or more characters.
- Users cannot use the question mark and asterisk wildcards on the same pattern.
- For directory filtering, the question mark and asterisk wildcards can appear only at the end of the string.

Note: The question mark wildcard is not accepted when VNX international character sets are enabled. A message is sent to the server log when this is attempted.

The pattern-match algorithm is case-sensitive, by default. To make it case insensitive, use the pax param filter.caseSensitive. The default value 1 should be changed to 0.

11 SnapSure Checkpoints (Snapshots)

11.1 Planning Considerations for SnapSure

With SnapSure Checkpoint, you can perform administrative tasks such as creating, listing, refreshing, or deleting checkpoints (or file system snapshots), and employing them for such tasks as restoring a PFS or copying to a backup device. Clients can use checkpoint data to independently restore files and directories or for business modeling.

11.2 Use of NL-SAS

SavVOL checkpoint specific storage should be primarily used on lower tiers of storage. In most cases this would be a NL-SAS (similar to SATA) based raid group. However, this tier of storage may not be available to all systems, especially smaller VNXs. In those cases, additional SAS based storage should be used (largest drive capacity).

11.3 Checkpoint Operations

SnapSure supports the following checkpoint operations:

Create — Creates a read-only checkpoint from a PFS, or a writeable checkpoint from a baseline read-only checkpoint.

Note: SnapSure can create 96 read-only checkpoints on a PFS, and one writeable checkpoint per read-only checkpoint. A maximum of 16 writeable checkpoints can exist concurrently on a PFS.

Refresh — Recycles the SavVol space by deleting the data in the checkpoint of a PFS and creating a new checkpoint by using the same checkpoint filename, file system identification number (ID), and mount state. Baseline and writeable checkpoints cannot be refreshed.

Restore — Restores a PFS to a point in time by using a read-only or writeable checkpoint.

Mount — Mounts a read-only checkpoint as a read-only file system, or a writeable checkpoint as a read-only or read/write file system. Both read-only and writeable checkpoints are mounted automatically upon creation.

Note: The number of writeable checkpoints per PFS does not count towards the total number of read-only checkpoints per PFS. However, mounted writeable checkpoint file systems count towards the 2048 maximum number of file systems per Data Mover and the 4096 maximum number of file systems per VNX cabinet. Also, writeable checkpoints might increase the size of the SavVol as changed blocks are written to it directly.

Delete — Deletes a read-only or writeable checkpoint. A read-only checkpoint which is a baseline checkpoint cannot be deleted unless the associated writeable checkpoint is deleted.

Export — Makes available a read-only or writeable checkpoint to NFS users on a network.

Share — Makes available a read-only or writeable checkpoint to CIFS users on a network.

Schedule — Automates read-only checkpoint creation and refresh operations. Writeable checkpoints cannot be scheduled.

11.4 Disk Space for Creating and Extending SavVols (Automatic)

If SavVols will be using the same storage pool as the primary file system, then a maximum size of the SavVol is required to prevent an overrun of the storage pool.

SnapSure requires a SavVol, or checkpoint-specific volume in which to hold data when you create the first Production File System (PFS) checkpoint. When you create more than one checkpoint of the same PFS, SnapSure uses the same SavVol, but logically separates the point-in-time data by using unique checkpoint names.

When you create the first checkpoint, you can choose to let SnapSure automatically create and manage a SavVol by using its default settings, or you can create and manage SavVol activity yourself by using the SnapSure volume and HWM options.

You can create the first checkpoint by using the SnapSure default values. For example, if you use the CLI command `fs_ckpt <fsname> -Create` and specify no `volume_name`, `size=`, or `%full=` option, SnapSure automatically:

Creates a SavVol by using this criterion:

If PFS > 10 GB, then SavVol = 10 GB.

If (PFS < 10 GB) and (PFS > 64 MB), then SavVol = PFS size.

If PFS < 64 MB, then SavVol = 64 MB, which is the minimum SavVol size.

Monitors the SavVol and extends it by 10 GB if a checkpoint reaches a HWM of 90 percent to keep the checkpoint active or readable.

Note: By default, EMC Automatic Volume Manager (AVM) algorithms determine the selection of disks used for a SavVol. AVM tries to match the storage pool for the SavVol with that of the PFS whenever possible. However, SnapSure provides the option of specifying any available pool for the SavVol when you create the first PFS checkpoint if you do not want to use the AVM selection method. The SavVol can be autoextended as long as space is available on the same disk type. For example, if the SavVol is built on an EMC CLARiiON® standard (CLSTD) disk, then autoextend can occur only if space is available on another CLSTD disk. The autoextend feature cannot use a CLARiiON ATA (CLATA) disk for the autoextend. Use the `nas_disk -list` command to display the disk types. You can also specify a limit to which SavVol can automatically extend by using the `fs_ckpt -modify` command. An event message is sent to the VNX system log each time an automatic SavVol extension succeeds or fails.

11.5 SavVol Size Customization (Manual)

If SavVols will be using the same storage pool as the primary file system, then a maximum size of the SavVol is required to prevent an overrun of the storage pool.

A large and busy PFS requires a bigger SavVol than that of a smaller, less active PFS. The SnapSure options allow you to customize the SavVol size proportional to the PFS size and the activity level you anticipate the checkpoints will need to capture. *You can set the maximum size of the SavVol once the initial checkpoint has been created.*

You can extend the SavVol size by using the `nas_fs -xtend` command. You can contract the SavVol size as follows:

Back up the existing checkpoint data to tape or available disks.

Delete all checkpoints. SnapSure automatically deletes associated SavVol.

Create a new checkpoint of the PFS, specifying a custom-size SavVol.

Note: If you use only the size option, for example, `size = 100M`, when creating the first checkpoint of a PFS, VNX first selects the same pool and storage system as the PFS when creating the SavVol. If the PFS spans multiple storage systems, the system picks the same set of storage

systems to find space for the SavVol after it ensures there is at least the specified 100 MB amount of space in the pool on the storage system. If not, the system extends the storage pool by adding unused disk volumes, and then creates the SavVol. If the pool cannot be extended, the QoS method of disk selection is used.

11.18 Recommended SavVol Sizes

The following are the recommendations for SavVOL storage on the VNX.

SavVOL storage can be defined with a maximum storage to limit their overall growth. With autoextend in place and a low initial storage size, the SavVol storage can grow according to need based on the behavior of the file system. In addition, SavVols should be implemented (with autoextend) on a lower tier of storage (i.e. SATA) or at least a different storage pool from primary data. These policies avoid the scenario of massive file deletes causing a large uptake of primary file system storage. With autoextend on, you can choose a much lower ADCR and have the system adapt to the characteristics of the files system.

Assumptions

- 7x snapshots active on any given day
- Volume sizes are 500GB-650GB

Average Daily Change Rate (ADCR)	Checkpoint sizes	Dedupe Buffer	Total Size of SavVol (percent of File System size)
5%	$7 \times 5 = 35\%$	5%	36.75 or 37%
4%	$7 \times 4 = 28\%$	5%	29.4 or 29%
3%	$7 \times 3 = 21\%$	5%	22.05 or 22%
2%	$7 \times 2 = 14\%$	5%	14.7 or 15%
1%	$7 \times 1 = 7\%$	5%	7.35 or 7%

For a file system that does not have good historical change rate history, 3% should be assumed. It is noted that EMC and CATE Backup engineering have used a 5% assumption but this would create significant wasted SavVOL space for many of the current file systems*. Ideally, setting the initial SavVOL space to a small size (say 5% of the file system space) and letting it autogrow to the needed size is the most efficient model. This may force initial deduplication to run slower, but they will eventually catch up.

**an analysis of the current file systems and SavVOLs was performed by NA NAS INT and CATE CSBE on 09/28/2010. Thusly these recommendations may change over time.*

11.19 Performance Considerations

SnapSure checkpoint operations affect the system as follows:

Creating a checkpoint requires the PFS to be paused. Therefore, PFS write activity suspends but read activity continues while the system creates the checkpoint. The pause time depends on the amount of data in the cache, but is typically a few seconds. EMC recommends a 10-minute minimum interval between the creation or refresh of checkpoints of the same PFS.

Restoring a PFS from a checkpoint requires the PFS to be frozen. Therefore, all PFS activities are suspended while the system restores the PFS from the selected checkpoint. The pause time depends on the amount of data in the cache, but is typically a few seconds.

Note: When read activity is suspended during a freeze, connections to CIFS users are broken. However, this is not the case when write activity is suspended.

Refreshing a checkpoint requires the checkpoint to be frozen. Therefore, checkpoint read activity suspends while the system refreshes the checkpoint. Clients attempting to access the checkpoint during a refresh process experience the following:

NFS — The system continuously tries to connect indefinitely. When the system thaws, the file system automatically remounts.

CIFS — Depending on the application running on Windows, or if the system freezes for more than 45 seconds, the Windows application might drop the link. The share might need to be remounted and remapped.

Deleting a checkpoint requires the PFS to be paused. Therefore, PFS write activity suspends momentarily but read activity continues while the system deletes the checkpoint.

If a checkpoint becomes inactive for any reason, read/write activity on the PFS continues uninterrupted.

11.20 Guidelines

Follow these guidelines when using SnapSure in your business environment:

Configure VNX to provide email or SNMP trap notification of the following events:

A SavVol or PFS becomes full or reaches its percent-full threshold, also known as HWM.

A scheduled checkpoint refresh attempt fails or the schedule fails to run.

11.21 Recommended Schedule of Checkpoints

The following chart is a summary of the recommended checkpoints that should be implemented. There may be regional requirements that exceed this schedule. However, this schedule is a good balance between having additional copies of changed data and the overhead associated with those copies. None of the recommended checkpoints occur at the zero minute due to potential system level schedule conflicts (system processes are scheduled at :00 of each hour).

Checkpoint	Schedule	Retention	Comments
Nightly	Sunday through Saturday at 00:10	7	This retention, in combination with the weekly checkpoints, will provide 7 days worth of data to be available on-line for any data recovery. Restores beyond that period will need to be obtained from backup/restore facilities.

11.22 Restrictions

Review the following restrictions before using SnapSure with other VNX features

11.22.1 SnapSure and VDMs

If a PFS is mounted on a Virtual Data Mover (VDM), the checkpoint must be mounted on the same VDM.

If a checkpoint is mounted on a VDM and you move the VDM to a different Data Mover, active checkpoints will persist. Checkpoint schedules resume at the next scheduled runtime after the move occurs. The CVFS name parameter will be set to that of the new Data Mover where the VDM is loaded. The CVFS timestamp will be adjusted to the new Data Mover's time and time zone. The SCSF timestamp will also be adjusted to the new Data Mover's time, but not the new time zone.

11.22.2 SnapSure and Replication

Writeable checkpoints cannot be created on replication V2 internal checkpoints. These checkpoints are identified by the following naming conventions:

Replication V2 internal checkpoints are identified by the name prefixed with the string
 “root_rep_ckpt.”


11.22.3 Checkpoints and Backup Restores

Checkpoints should be deleted or disabled before restoring backup data as large changes to the file system will result in a significant increase in the use of the checkpoint file system (increased storage use).

11.22.4 User Access for checkpoints (User Managed Restores)


Windows 2003 (or Older) Clients

Similar to the NetApp hidden directory structure for snapshots, checkpoints are user accessible from a hidden directory called “.ckpt”. There will be directories for each active checkpoint under this hidden directory that the user can access to self-administer file or directory level restores. The following example demonstrates this directory structure.



Vista, Windows 7, Windows 2008 Clients

Use of these hidden directories is discouraged with the newer clients and the user should instead use the Previous Versions tab from Windows Explorer. From Windows Explorer, right-clicking the directory / file system / network drive and select properties. Select the previous versions tab.



You can scroll through the list of checkpoints and find the one that represents the time from which you want to recover a file. You will be able to Open (explore), Copy (to a new location), or Restore (replace the directory/share/drive). Be careful using the restore function and be sure you are restoring the right location as this action cannot be undone and may fill up the file system. It is safer to use the Open and Copy functions.

12 VNX IP Replication

12.1 How VNX Replication sessions work

VNX Replicator uses internal checkpoints to ensure availability of the most recent point-in-time copy. These internal checkpoints are based on SnapSure Checkpoint technology and are visible to replication users, but are not available for user management. This means you cannot perform operations on the checkpoint such as refresh, delete, and restore.

12.2 Space requirements

Setting up a new replication session can require up to double the amount of space on the destination site as on the source site. VNX Replicator creates a destination file system that is the same size as the primary file system. It also creates a SavVol file system the size of the amount of data on the source file system. The SavVol cannot be reduced or re-created. This means that if the source file system is full, double the space will be required on the destination when setting up a new replication session.

For example, if you have a 2 TB file system that is half-full, your new replication session will require a 2 TB file system on the destination site plus a 1 TB SavVol, for a total of 3 TB.

12.3 Required Communication Ports

VNX Replicator requires the following additional unblocked, unfiltered ports:

- 443 - HTTPS. Required between control stations.
- 5081 - DART Interconnect Server: DIC facility for Data Mover exchange
- 5083 - VNX Interconnect Server (S): CIC facility for short command support
- 5084 - VNX Interconnect Server (L): CIC facility for long command support
- 5085 - VNX Interconnect Server (D): CIC facility for DART-DART command support

12.4 Replicator/SnapSure

By using VNX Replicator, you can enable automatic file system extension on the source file system. When the source file system hits its high water mark (HWM), the destination file system automatically extends first. Then the source file system automatically extends. If the extension of the destination file system succeeds but the source file extension fails, the file systems differ in size which causes replication failure. Use the `nas_fs -xtend <fs_name> -option src_only` command to manually adjust the size of the source file system.

There must be sufficient file system space and disk storage available to support Replicator and SnapSure operations. To review the entire file system size, use `nas_fs -list`. To calculate the SavVol file size, use `nas_disk -size`. The EMC VNX Network Server Command Reference Manual provides a detailed synopsis of the commands associated with SnapSure and VNX Replicator.

12.5 Controlling replication sessions with replication policies

Most of these policies can be established for one replication session by using the `nas_replicate` command or all replication sessions on a Data Mover by using `nas_cel` or a parameter.

VNX Replicator has policies to:

- Control how often the destination object is refreshed by using the `max_time_out_of_sync` setting.
- Control throttle bandwidth by specifying bandwidth limits on specific days, specific hours, or both.
- Set the data amount to be sent across the IP network before an acknowledgement is required from the receiving side. The TCP window is automatically sized.

12.6 Update policy

The update policy set for a replication session determines how frequently the destination is updated with source changes. You can define a `max_time_out_of_sync` value for a replication session or you can perform on demand, manual updates.

12.6.1 Max time out of sync

The `max_time_out_of_sync` value represents the elapsed time window within which the system attempts to keep the data on the destination synchronized with the data on the source. For example, if the `max_time_out_of_sync` is set to 10 minutes, then the content of, and latest checkpoint on, the destination should not be older than 10 minutes and that checkpoint time is identical to the time of the corresponding checkpoint on the source. The destination could be updated sooner than this value. The source write rate, network link speed, and the interconnect throttle, when set, determine when and how often data is sent. Large bursts of writes on any file system that is sharing an interconnect, or physical connection, or network problems between the source and destination can cause the `max_time_out_of_sync` setting to be exceeded.

CSBE recommends using the lowest value that meets the business requirements associated with acceptable data loss ranging from 10 minutes up to 60 minutes.

12.6.2 Manual refresh

The manual refresh option allows you to perform on-demand updates of the destination with the data on the source. When you change a replication session from a `max_time_out_of_sync` policy to a manual refresh, the system suspends the session. To resume the session, reset the `max_time_out_of_sync` value.

12.7 Bandwidth schedule

Each side of a Data Mover interconnect can define a bandwidth schedule for all replication sessions using that interconnect. A bandwidth schedule allocates the interconnect bandwidth used on specific days, specific times, or both, instead of using all available bandwidth at all times for the interconnect. By default, an interconnect provides all available bandwidth at all times for the interconnect.

Note: A bandwidth limit of 0 means no data is sent over the interconnect.

You use the `nas_cel -interconnect -bandwidth` option to configure a bandwidth schedule for a specified interconnect. You can set overlapping bandwidth schedules to specify smaller time intervals first and bigger time intervals next. The limit set for the first applicable time period is taken into account, so the order in which the schedule is created is important.

The following schedule syntax means that Monday at 9:30 A.M., the limit is 5000 Kb/s but Monday at 8:00 A.M., the limit is 7000 Kb/s:

`Mo09:00-12:00/5000,MoTu08:00-17:00/7000,MoTuWeThFr06:00-18:00/10000`

Example 1

The following bandwidth schedule modifies interconnect `s2_s3` to use a limit of 2000 Kb/s from 7 A.M. to 6 P.M. Monday through Friday; otherwise use a bandwidth schedule of 8000 Kb/s:

`nas_cel -interconnect -modify s2_s3 -bandwidth MoTuWeThFr07:00-18:00/2000,/8000`

Example 2

The following bandwidth schedule modifies interconnect `s2_s3` to use a limit of 2000 Kb/s from 9 A.M. to 2 P.M. Monday and Tuesday; and a limit of 5000 Kb/s Wednesday through Friday from 6 A.M. to 2 P.M.; otherwise use a bandwidth schedule of 8000 Kb/s:

`nas_cel -interconnect -modify s2_s3 -bandwidth MoTu09:00- 14:00/2000,/8000,WeThFr06:00-14:00/5000,/8000`

Example 3

The following bandwidth schedule modifies interconnect s2_s3 to use a limit of 2000 Kb/s from 9 A.M. to 2 P.M. Monday and Tuesday; a limit of 5000 Kb/s Wednesday and Thursday from 6 A.M. to 2 P.M.; a limit of 8000 Kbit/sec all day Friday; otherwise use a bandwidth schedule of 8000 Kb/s:

```
nas_cel -interconnect -modify s2_s3 -bandwidth MoTu09:00- 14:00/2000,/8000,WeTh06:00-
14:00/5000,/8000,Fr00:00- 23:00/8000
```

Example 4

The following bandwidth schedule modifies interconnect s2_s3 to use no bandwidth limit, which in effect stops all data transfer over the interconnect:

```
nas_cel -interconnect -modify s2_s3 -bandwidth MoTuFr08:00- 22:00/0
```

12.8 General Restrictions

- The maximum number of sessions per Data Mover is 1024 (the total number of configured replication sessions and copy sessions). This enforced limit includes all configured file system, VDM, and iSCSI replication sessions as well as copy sessions.
- The maximum number of active sessions is 256 per Data Mover. This limit includes all active file system, VDM, and iSCSI replication sessions as well as copy sessions. Any additional sessions are queued until one of the current sessions completes.
- The maximum number of initial copies that can be in progress at one time is 16. This limit includes all active file system, VDM, and iSCSI replication sessions as well as copy sessions.

12.8.1 Filesystem Restrictions

- A source file system supports up to four replication sessions (one-time copy or ongoing replication sessions).
- You cannot unmount a file system if replication is configured. Replication has to be stopped for an unmount of the file system to succeed. When the file system is unmounted, the internal checkpoints will be automatically unmounted and when mounted, the internal checkpoints will be automatically mounted.
- You cannot move a file system to another Data Mover if replication is configured on the file system. You must stop replication before the file system is unmounted. After the file system is mounted on the new Data Mover, start the replication session.
- When replicating databases, additional application-specific considerations maybe necessary to bring the database to a consistent state (for example, quiescing the database).
- If you plan to enable international character sets (Unicode) on your source and destination sites, you must first set up translation files on both sites before starting Unicode conversion on the source site. *Using International Character Sets with VNX* describes this action in detail.
- Replication failover is not supported for CIFS local users and local groups unless you use Virtual Data Movers.

12.8.2 VDM Replication Restrictions

- A given VDM can only be replicated once.
- VDM replication does not support a cascade configuration.
- VDM replication does not support a one-to-many configuration.
- Unloading a VDM with a replication configured is not supported. Replication has to be stopped to unload a VDM.
- Replication of a physical Data Mover root file system is not supported. VDM root file systems are supported.

12.9 Read-Only COB Optional Replication

In order to preserve existing processes for replication testing (COB Testing), CSBE allows for a read-only (or read-write) share of the destination copy of the PFS on a different CIFS server. This can be implemented as an alternative to traditional VDM replication as it establishes another active CIFS server on the secondary VNX that users can access. Normally the file system on the remote system will be read-only. However, if a disaster is declared, the replication session can be halted on the secondary control station and the file system can be remounted as read-write. Alternatively, if the business needs to test via a read-write share for application testing, the COB test file system can be made read-write.

This requires that all CIFS server settings be manually reproduced by system administrator so that the remote CIFS server will behave identically to the primary CIFS server. Unlike a VDM replication, the CIFS settings will not be replicated.

12.9.1 Read-Only COB Impact to VDMs

With the current VNX DART code, any file systems owned by VDMs must have snapshots that are also owned by that VDM. Therefore it is difficult to maintain a normal VDM replication and this read-only COB at the same time. There are two options here.

1. Stop replicating with the VDM and switch to file-system only replication. The destination file system needs to be located in an active VDM or a physical DM that can host the different CIFS server. This replication would be instead of normal VDM replication. *This is the recommended solution when the read-only file copy needs to be accessible all of the time.* This should be discouraged as it changes the replication environment significantly.
2. Keep replicating the VDM and use checkpoints for the read-only (or read-write) COB copy. Checkpoints can be used to facilitate a series of copies of the replicated file system as read-only or read-write, but based on a point in time. With the use of the MMC, we can then create a share based on the non-rooted directory. This is the least intrusive to the VDM environment, though more work for the storage admin. *This is the recommended solution when the COB test volume needs to be accessible for a COB test and can be manually synchronized.* An example is provided below.
3. Other options that include secondary replication sessions and cascaded replications are discouraged because they require additional storage.

12.9.2 Configuring the Read-Only Replication (VDM replication independent)

1. Establish the connectivity between the local and remote control stations and datamovers. This requires root access. The secondary (remote) control station will need to be added to the iptables ACL on the primary control station. Likewise the primary control station will need to be added to the secondary (remote) control station if no previous replication sessions are in place. Use the `nas_cel` command to create the control station relationships and the interconnects and then to test the interconnects.

```
From Primary Control Station
[root@c5-nsg8-cs0 nasadmin]# nas_cel -create cs5-vnx5500-cs0 -ip 10.40.42.201 -passphrase
Password123
operation in progress (not interruptible)...
id          = 3
name        = c5-vnx5500-cs0
owner       = 0
device      =
channel     =
net_path    = 10.40.42.201
VNX_id      = APM001116018682007
passphrase  = Password123

[root@c5-nsg8-cs0 nasadmin] nas_cel -interconnect -create nsg8-s3_vnx5500-s2 -source_s
erver server_3 -destination_system c5-vnx5500-cs0 -destination_server server_2 -
source_interfaces 10.40.42.216 -destination_interfaces 10.40.42.204
operation in progress (not interruptible)...
```

```

id = 30003
name = nsg8-s3_vnx5500-s2
source_server = server_3
source_interfaces = 10.40.42.216
destination_system = c5-vnx5500-cs0
destination_server = server_2
destination_interfaces = 10.40.42.204
bandwidth schedule = uses available bandwidth
crc enabled = yes
number of configured replications = 0
number of replications in transfer = 0
status = The interconnect is OK.
[root@c5-nsg8-cs0 ~]# nas_cel -interconnect -list
id name source_server destination_system destination_server
20001 loopback server_2 c5-nsg8-cs0 server_2
30001 loopback server_3 c5-nsg8-cs0 server_3
30003 nsg8-s3_vnx5500-s2 server_3 c5-vnx5500-cs0 server_2
40001 loopback server_4 c5-nsg8-cs0 server_4
[root@c5-nsg8-cs0 ~]# nas_cel -interconnect -validate nsg8-s3_vnx5500-s2
nsg8-s3_vnx5500-s2: has 1 source and 1 destination interface(s); validating - please wait...ok

From Secondary/Remote Control Station
[root@c5-vnx5500-cs0 nasadmin]# nas_cel -create cs5-nsg8-cs0 -ip 10.40.42.211 -passphrase
Password123
operation in progress (not interruptible)...
id = 1
name = cs5-nsg8-cs0
owner = 0
device =
channel =
net_path = 10.40.42.211
VNX_id = APM000833001340064
passphrase = Password123

[root@c5-vnx5500-cs0 nasadmin]# server_ping server_2 -interface ifdm2-0 10.40.42.216
server_2 : 10.40.42.216 is alive, time= 0 ms
[root@c5-vnx5500-cs0 nasadmin]# nas_cel -interconnect -create nsg8-s3_vnx5500-s2 -
source_server server_2 -destination_system c5-nsg8-cs0 -destination_server server_3 -
source_interfaces 10.40.42.204 -destination_interfaces 10.40.42.216
operation in progress (not interruptible)...
id = 20003
name = nsg8-s3_vnx5500-s2
source_server = server_2
source_interfaces = 10.40.42.204
destination_system = c5-nsg8-cs0
destination_server = server_3
destination_interfaces = 10.40.42.216
bandwidth schedule = uses available bandwidth
crc enabled = yes
number of configured replications = 0
number of replications in transfer = 0
status = The interconnect is OK.
[root@c5-vnx5500-cs0 nasadmin]# nas_cel -interconnect -list
id name source_server destination_system destination_server
20001 loopback server_2 c5-vnx5500-cs0 server_2
20003 nsg8-s3_vnx5500-s2 server_2 c5-nsg8-cs0 server_3
30001 loopback server_3 c5-vnx5500-cs0 server_3
[root@c5-vnx5500-cs0 nasadmin]# nas_cel -interconnect -validate nsg8-s3_vnx5500-s2
nsg8-s3_vnx5500-s2: has 1 source and 1 destination interface(s); validating - please wait...ok

```

2. You can create the file system on the remote datamover during the replication session or you can create it before you establish the communication session. Updating existing files systems may have a negative impact to the checkpoint storage, so creating the file system (specify a pool to use) is recommended. In this example, the file system was created before.

```

From the Primary Control Station
[root@c5-nsg8-cs0 nasadmin]# nas_replicate -create fs33_repl -source -fs fs33 -destination
-fs fs33 -interconnect nsg8-s3_vnx5500-s2 -max_time_out_of_sync 5 -overwrite_destination
OK

```

3. Create an identical share for the replicated file system. It should be already mounted as read-only (auto-mounted in previous step). This requires an active CIFS server joined to the domain, but the CIFS server could be created at anytime. The CIFS server should be named similar to the production CIFS server so that it is easy to associate them together for maintenance and easy to tell apart for production vs. COB status.

12.9.3 Configuring the Read-Only or Read-Write COB Test (with VDM replication)

It is assumed that since VDM replication is already present, we will focus on the checkpoints and sharing of those checkpoints here.

1. You should create a CIFS server on the COB system and join it to the appropriate domain. The CIFS server should be named similar to the production CIFS server so that it is easy to associate them together for maintenance and easy to tell apart for production vs COB status.
2. Create a baseline (read-only) checkpoint of the replicated file system (currently owned by the inactive VDM).

```
[nasadmin@IRNT1C01 ~]$ fs_ckpt PST_FLR_Test -name baseline_ckpt_PST_FLR_Test -C operation in
progress (not interruptible)...

id          = 450
name        = PST_FLR_Test
acl         = 0
in_use      = True
type        = uxfs
worm        = enterprise with no protected files
worm_clock= Mon Jul 25 10:59:25 CDT 2011 worm Max Retention Date= No protected files created
volume      = v577
pool        = clar_r5_performance
member_of   = root_avm_fs_group_3
rw_servers=
ro_servers= server_2
rw_vdms     =
ro_vdms     = jvnt1c01_ap_vdm
auto_ext    = no,virtual_provision=no
deduplication = Off
ckpts       = ckpnt-
PST_FLR_TEST,PST_FLR_Test_ckpt01,root_rep_ckpt_450_52602_2,root_rep_ckpt_450_52602_1,baseline_ckpt
_PST_FLR_Test
rep_sess    = 480_APM00060703453_0000_577_APM00060604036_0000(ckpts: root_rep_ckpt_450_52602_1,
root_rep_ckpt_450_52602_2) stor_devs = APM00060604036-0015
disks       = d15
disk=d15    stor_dev=APM00060604036-0015 addr=c16t115      server=server_2
disk=d15    stor_dev=APM00060604036-0015 addr=c0t115      server=server_2

id          = 462
name        = baseline_ckpt_PST_FLR_Test
acl         = 0
in_use      = True
type        = ckpt
worm        = enterprise with no protected files
volume      = vp580
pool        = clar_r5_performance
member_of   =
rw_servers=
ro_servers= server_2
rw_vdms     =
ro_vdms     = jvnt1c01_ap_vdm
ckpt_of= PST_FLR_Test Mon Jul 25 11:29:03 CDT 2011
deduplication = Off
used        = 4%
full(mark)= 90%
stor_devs   = APM00060604036-0010
disks       = d7
disk=d7     stor_dev=APM00060604036-0010 addr=c0t110      server=server_2
disk=d7     stor_dev=APM00060604036-0010 addr=c16t110      server=server_2
```

```
[nasadmin@IRNT1C01 ~]$ nas_fs -i baseline_ckpt_PST_FLR_Test
id          = 462
name        = baseline_ckpt_PST_FLR_Test
acl         = 0
in_use      = True
type        = ckpt
worm         = enterprise
volume      = vp580
pool        = clar_r5_performance
member_of   =
rw_servers=
ro_servers= server_2
rw_vdms     =
ro_vdms     = jvnt1c01_ap_vdm
ckpt_of= PST_FLR_Test Mon Jul 25 11:29:03 CDT 2011
deduplication = Off
used        = 4%
full(mark)= 90%
stor_devs   = APM00060604036-0010
disks       = d7
  disk=d7    stor_dev=APM00060604036-0010 addr=c0t110      server=server_2
  disk=d7    stor_dev=APM00060604036-0010 addr=c16t110     server=server_2
[nasadmin@IRNT1C01 ~]$ fs_ckpt baseline_ckpt_PST_FLR_Test -name PST_FLR_Test_DR_RW -C -readonly n
operation in progress (not interruptible)...id          = 462
name        = baseline_ckpt_PST_FLR_Test
acl         = 0
in_use      = True
type        = ckpt
worm         = enterprise
volume      = vp580
pool        = clar_r5_performance
member_of   =
rw_servers=
ro_servers= server_2
rw_vdms     =
ro_vdms     = jvnt1c01_ap_vdm
ckpt_of= PST_FLR_Test Mon Jul 25 11:29:03 CDT 2011
deduplication = Off
ckpts       = PST_FLR_Test_DR_RW
used        = 6%
full(mark)= 90%
stor_devs   = APM00060604036-0010
disks       = d7
  disk=d7    stor_dev=APM00060604036-0010 addr=c0t110      server=server_2
  disk=d7    stor_dev=APM00060604036-0010 addr=c16t110     server=server_2

id          = 463
name        = PST_FLR_Test_DR_RW
acl         = 0
in_use      = True
type        = wckpt
worm         = enterprise
volume      = vp580
pool        = clar_r5_performance
member_of   =
rw_servers= server_2
ro_servers=
rw_vdms     = jvnt1c01_ap_vdm
ro_vdms     =
ckpt_of= PST_FLR_Test
baseline_ckpt = baseline_ckpt_PST_FLR_Test Mon Jul 25 11:29:03 CDT 2011
deduplication = Off
used        = 6%
full(mark)= 90%
stor_devs   = APM00060604036-0010
disks       = d7
  disk=d7    stor_dev=APM00060604036-0010 addr=c0t110      server=server_2
  disk=d7    stor_dev=APM00060604036-0010 addr=c16t110     server=server_2
```

3. Create a writeable checkpoint based on the baseline checkpoint.

```
[nasadmin@IRNT1C01 ~]$ server_mount server_2 | grep -i RW
root_fs_2 on / uxfs,perm,rw
larry on /larry uxfs,perm,rw
steve_nfs on /steve_nfs uxfs,perm,rw
filecen_test1 on /root_vdm_5/filecen_test_rep uxfs,perm,rw
filecen_test2 on /root_vdm_5/filecen_test2 uxfs,perm,rw
pm01 on /pm01 uxfs,perm,rw,noprefetch,nonotify,accesspolicy=NATIVE
PST_FLR_Test_DR_RW on /root_vdm_8/PST_FLR_Test_DR_RW ckpt,perm,rw
[nasadmin@IRNT1C01 ~]$ server_mount server_2 | grep -i DR_RW PST_FLR_Test_DR_RW
/root_vdm_8/PST_FLR_Test_DR_RW ckpt,perm,rw
[nasadmin@IRNT1C01 ~]$ server_mount server_2 | grep -i "DR_RW"
PST_FLR_Test_DR_RW on /root_vdm_8/PST_FLR_Test_DR_RW ckpt,perm,rw
```

If the business needs the COB test volume to be read-only, then skip the next step and proceed to step 5.

4. Only if a read-write COB test volume is desired, determine the amount of SAVVOL space that is currently available. This is important because a read-write file system will write changes to the SAVVOL area of the PFS on the COB system. The amount of data expected to be written to the read-write file system must be available in the SAVVOL unless the existing SAVVOL space has autoextend turned on (by default).

Note: There will need to be a clean-up of the SAV-VOL space once the COB test has been completed. This only applies to the read-write option.

5. If read-only is desired, then un-mount the second checkpoint and remount it as read-only. Use the -F option to force the action.

```
[nasadmin@IRNT1C01 ~]$ server_umount server_2 -t -F PST_FLR_Test_DR_RW
server_2 : done
[nasadmin@IRNT1C01 ~]$ server_mount server_2 -o ro PST_FLR_Test_DR_RW
server_2 : done
[nasadmin@IRNT1C01 ~]$ server_mount server_2 | grep -i "DR_RW"
PST_FLR_Test_DR_RW on /root_vdm_8/PST_FLR_Test_DR_RW ckpt,perm,ro
```

6. Create the same shares as the production site on the DR site using the MMC system administrator tool from a windows environment. This requires an administrative share to the CIFS server (i.e. C\$). Follow the share wizard to create the new share, following the path to /root_vdm_x/checkpoint_name/folder_directory.

Alternatively, this can be accomplished with the RMTSHARE utility from a windows environment (with an administrative share). This utility was provided in the NT 4.0 Server Resource Kit. However, it has been briefly tested with Vista, 2000 Server and 2003 Server on DART 5 and DART 6. It has not been tested with 2008. The syntax of this utility is:

```
RMTSHARE \\server
\\server\sharename
\\server\sharename=drive:path [/USERS:number | /UNLIMITED]
    [/REMARK:"text"]
    [/GRANT [user[:perm]][ /GRANT user[:perm]]]
    [/REMOVE user]
\\server\sharename=printrname /PRINTER [/USERS:number | /UNLIMITED]
    [/REMARK:"text"]
    [/GRANT [user[:perm]][ /GRANT user[:perm]]]
    [/REMOVE user]
\\server\sharename [/USERS:number | /UNLIMITED]
    [/REMARK:"text"]
    [/GRANT [user[:perm]][ /GRANT user[:perm]]]
    [/REMOVE user]
\\server\sharename /DELETE
if a sharename or path contains spaces, it should be enclosed in quotes:
\\server\ "with space" = "c:\with space"

rmtshare \\CIFS_SERVER\share_name=c:\mountpoint
```

```
C:\temp>net view \\nsustex31ct01
Shared resources at \\nsustex31ct01
EMC-SNAS:T5.6.50.209
Share name      Type  Used as  Comment
-----
PST_FLR_Test   Disk
The command completed successfully.

C:\temp>net use q: \\nsustex31ct01\c$ /user:Administrator
Enter the password for 'Administrator' to connect to 'nsustex31ct01':
The command completed successfully.

C:\temp>rmtshare \\nsustex31ct01\gary_share=c:\root_vdm_8\PST_FLR_Test_DR_RW
The command completed successfully.

C:\temp>net view \\nsustex31ct01
Shared resources at \\nsustex31ct01
EMC-SNAS:T5.6.50.209
Share name      Type  Used as  Comment
-----
gary_share      Disk           Shared by remote command.
PST_FLR_Test   Disk
The command completed successfully.
```

7. To update the secondary and baseline checkpoints, the secondary checkpoint must first be deleted. Then, the baseline checkpoint can be refreshed.

```
[nasadmin@IRNT1C01 ~]$ fs_ckpt baseline_ckpt_PST_FLR_Test -refresh
```

8. If the COB test volume was kept at read-write, then it is critical that the SAV-VOL space be reclaimed once the COB testing has been completed. This is because any data written to the read-write share will be written to the SAVVOL area for the primary file system on the COB array. Clean up of the SAVVOL space while keeping replication alive is a undocumented feature and subsequently you will need to contact EMC Support for this process. This process will manually delete all of the replication and non-replication checkpoints that have been written to the SAVVOL area for the PFS.

13 VNX Anti-Virus Agent (CAVA)

13.1 Integration with Citi AV

Only one option exists today for virus scanning on the CIFS space for VNX (FILE):

On-access virus scanning using 3rd-party scan-engines acting as gateways/proxies.

CTI employs McAfee AV Suite of products as its SOE standard for virus scanning. With the availability of CAVA agent on-access scanning, the only certified method for scanning network drives is via the external scanning engine. CAVA for VNX protects from viruses by running the anti-virus program on a separate scanner-server. When a client system attempts to access or store a file, the operating system (DART) virus scan API determines (in real-time) if the file should be scanned. If the file needs to be scanned, the VNX sends a request to the scanner-server. The CAVA for VNX scanner-server scans the file and takes appropriate action, if necessary.

CAVA is only for CIFS servers and does not include files managed by NFS.

13.2 CAVA System Requirements

The following table details the minimum system requirements recommended for each AV server. The CAVA server may be a VMWare host (if using VSE 8.7i or later). CAVA servers are required to be deployed in the same location of the NAS array with a recommendation of 2ms response time. A higher response time of 2ms would still function however would cause additional latency for the scan process and potentially increase scan queues.

Subject	Minimum Requirement
Hardware	1.2 GHz 64-bit processor 4 CPU/core 2 GB RAM 80 MB of free disk space VMWare or physical server
Operating System	Windows 2008 x64 Server or Windows 2012 x64 Server
AD and Network	The Windows network must contain a domain controller with Active Directory and DNS enabled.

Note: This server must be dedicated and is not to be used for any other function (i.e. scripting or otherwise).

*** Appendix G references the previous hardware requirements and ratios based on dual core server requirements.

13.3 Supported 3rd Party AV Engines

- McAfee Virus Scan Enterprise versions: 8.8P4
- McAfee Agent versions: 4.8

For the latest versions please refer to the following:

- [ePO Product Page](#)
- [McAfee VSE Page](#)

13.4 VNX to AV scanner ratio

This table is only a guideline based upon best practices.

This guideline provides the flexibility to add additional CAVA servers to maximize scanner utilization as needed. CAVA server performance should be closely monitored via BMC agents to ensure VNX responsiveness to end-user requests. AV scan pools are recommended and all members must be similar in performance characteristics in hardware (cpu, memory, network) since round-robin load balancing does not have insight into individual server utilization.

VMware is a candidate to host AV scanners, provided vendor support is available. McAfee supports VMware for version 8.7i and later.

For reference, please consider the follow resources as a minimum requirement for an AV Pool:

VNX Active Datamovers	CAVA Scanners
1-3	2+
4-7	3+
8-11	4+
12-15	5+

13.5 CAVA File Exclusions

CATE Host Security Engineering has authorized the NAS OPS teams to use the following exclusions. These exclusions are for compressed file that would likely be skipped by the current McAfee scanning policy.

```
excl=*.arc:*.arj:*.cab:*.dat:*.gz:*.iso:*.jar:*.pst:*.tar:*.tgz:*.z:*.zip
```

For more on CAVA file exclusions, please refer to [CAVA Configuration file](#)

13.6 CAVA Restrictions

CAVA should be disabled during backup restores to prevent an unnecessary constraint on the restore and to prevent the CAVA servers from being overloaded.

CAVA Servers should be dedicated for use with McAfee Virus Scanning and CAVA. No other applications or processes should be running on the CAVA servers.

13.7 CAVA Setup Procedures

The following steps are required to set up CAVA at Citi:

- Create a domain functional user in the appropriate domain service accounts area.
- Setup CAVA servers (minimum 2 physical or Virtual Machines per VNX) and install CAVA software on the CAVA servers
- Create dedicated CIFS servers for CAVA on each active datamover
- Create a local group on each CAVA CIFS server and assign virus checking right
- Create a local user on each DM that correlates to the domain user
- Configure the viruschecker configuration file on each active datamover
- Start the CAVA services on each datamover

13.8 CAVA Domain Functional ID

You must create a domain functional account that is the CAVA domain user. This account should be named according to current domain standards and will be placed in the service accounts area. For example: "namdev\svr_emc_celerra_av". This user's password should not expire. Otherwise,

if local security restrictions require this, then the password would need to be updated on each DM and CAVA server. This user will be assigned to a viruschecker group on each CAVA server (minimum 2 CAVA servers per Celerra) in a later step.

13.9 Create CIFS Server for CAVA

Create a dedicated CIFS server for CAVA on each active datamover. This requires the use of an additional virtual IP interface. The new CIFS server is recommended to be part of the same domain that the CAVA domain user was defined (though not required). Follow the standard CIFS server creation procedures. This CIFS server should be named so that it is identified as a CIFS server for the CAVA interface and not a normal file-sharing CIFS server. This CIFS server should not be used for any file sharing or related CIFS services.

Note: To setup the password for the local administrator, set an initial password when creating the CIFS server from Celerra Manager (or CLI). You will need to update this password before remotely administering the server. To update the password go to the windows security screen (CTL+ALT+DEL, or CTL+ALT+END if using terminals services). Select Change password and fill out the information for the CIFS_CAVA_Server\Administrator local administrator account.

13.10 Create Local User Group on CAVA CIFS Servers

For each of the CIFS servers that were created for CAVA above, you will need to create a new local group called "viruscheckers" (or similar). You will then need to add the CAVA domain user to this group.


Note: To administratively control a newly created CIFS server, you will need to first map an administrative interface (such as c\$), using a local administrator account on the CIFS server. Then you will be able to manage the computer via Computer Manager (Computer Manager → Connect to another computer) and Local Policy Setting Tool.

13.11 Assign EMC Virus-Checking Right to Local Group

You will need to assign the viruschecking right to the local virus checking group on each of the CAVA CIFS servers. In order to distinguish the CAVA domain user from all other domain users, you will be assigning the group the EMC virus-checking right. This right is not a domain privilege, but rather it exists locally in the Data Mover.

Note: You cannot you cannot use Microsoft Local Policy Settings tool for this step. You must use the Celerra Management Console (part of EMC Celerra Management Applications). This is because the Microsoft tool does not allow you to remotely manage user rights assignments.

Under the Celerra Management Console (Control Pane → Administrative Tools → Celerra Management if installed), go to Data Mover Management, and select Connect to Data Mover. Type in the name of the CAVA CIFS server created above. Double click the Data Mover Security Settings. Click User Rights Assignment. Double-click the EMC Virus Checking and then click Add. Select the local AV group created above and then click ok to assign the special right to the group. Once added, there should be an entry on the screen similar to below:



13.12 Assign Local Admin Rights to CAVA Domain User on CAVA Server

The CAVA domain user needs to be a local administrator on all CAVA servers for a given datamover. This allows the user to start and stop processes as needed.

Note: If the AV server is a domain controller (not typical at Citi), the virus-checking user account should join the domain Administrator group instead of the local administrator group. This is because the local administrator group is not managed on a domain controller.

13.13 CAVA Setup on a scanning server

There is a minimum of two servers required for each VNX FILE to be used as external virus scanning servers. These servers should be dedicated to the scanning processes. These servers should also be a member of the domain that contains the CAVA AV user. These servers can be physical or VMWare hosted servers.

The CAVA AV user that was created above should be set as a local administrator on each CAVA server.

The latest CAVA package (see [Current Software Packages](#)) should be installed. The CAVA package is a silent installation. Once the package has installed, the service must be configured. Initially the service is not started, but is set to start automatically (verify this). Modify the service properties and use the CAVA domain user that was created above. Change the actions to restart the service after 1-5 minutes on each error (1st time, 2nd time, etc.) Restart the system after the install and setup is completed.

13.14 CAVA Server Network Configuration


The CAVA server should be configured so that any dual NICs are in a Network Fault Tolerance Only configuration sometimes referred to as Active/Passive. This may be different than a default configuration of Load Balancing (with Fault Tolerance) sometimes referred to as Active/Active. Both configurations should be available to the system administrators setting up and managing the CAVA servers.

This configuration is required to prevent CAVA communication issues when dealing with multiple MAC addresses. The default configuration of AV on the Datamover uses MAC addressing and cannot handle multiple MAC addresses. Therefore a single MAC address is required and thus the use of a Fault Tolerant configuration of any dual NICs. This is also described frequently as ACTIVE/PASSIVE.

As with all fault tolerant configurations, each NIC should be connected to separate physical switches to ensure availability of the CAVA server. In addition, there should be at least two CAVA servers defined for every datamover.

13.15 ePO Policy on CAVA Server

The scanning server running CAVA should be part of the local ePO group managed by the [EMC NAS ePO Configuration Standards](#). This assignment is controlled by the ePO environment. The EMC NAS ePO policy is very similar to the default Citi policy except for the configuration difference of scanning network drives. This is important as all scan requests from CAVA are provided to the AV software as UNC formatted network file locations. This setting (if using the correct ePO policy) can be recognized by right-clicking on the McAfee OAS icon → On-Access Scan Properties → Default Processes → Scan items Tab.



13.16 CAVA Configuration file

The viruschecker.conf file is the configuration file that defines the IP address of the CAVA servers and any exclusion patterns. Additional exclusions other than listed below should be avoided if possible as the Citi exclusion policy will be followed (defined by ePO / Virus Scan Enterprise exclusion policies).

Configure the /etc/viruschecker.conf file (run vi or similar typically as root)

Note: if file does not exist copy template from /nas/sys directory

```
masks=*.*
excl=*.arc:*.arj:*.cab:*.dat:*.gz:*.iso:*.jar:*.pst:*.tar:*.tgz:*.z:*.zip
addr=nasmgmtgdse04dr.namdev.nsrootdev.net:nasmgmtgdse05.namdev.nsrootdev.net
CIFSServer=NASEMCCA02
shutdown=viruschecking
RPCRetryTimeout=30000
RPCRequestTimeout=90000
```

Note: The *addr* field is not followed by a colon or blank line after the last virus scanning server fully qualified name (or ip). There should not be any blank lines in this file.

The masks should be set to a simple *.* to allow for the ePO policy to control exclusions. Compressed files are allowed to be added to the exclusion list as above.

The CIFSserver field is the NetBIOS (or ip) of the CAVA CIFS server setup above. This entry will make the file unique for each DM.

The excl (exclusion list) will prevent compressed files from being sent to the CAVA server. Normally only file pointers are sent over the network to the CAVA server for scanning. Compressed files are handled differently. Compressed files are copied over the network and extracted locally by the CAVA server and then scanned before a write is allowed. Placing these compressed file exclusions here will prevent these files from being copied or scanned. This specific exclusion is acceptable to CATE Host Security Engineering as virus scan is configured to skip scanning these files by policy.

The addr field is the list of CAVA servers for this datamover. There should be at least two CAVA servers listed here. **The CAVA servers should be listed with fully qualified names to allow for the best protocol search. This is especially important when using 64-bit CAVA.** Listing CAVA servers by IP is allowed for 32-bit CAVA, but FQN is recommended. In order to use FQN, reverse DNS (i.e. lookup of IP by FQN is required to be working correctly).

13.16.1 RPCTimeOut and RPCRetryTimeout

In order to allow for some network and response delay between the CAVA server and the DM, the following timeout values should be set:

RPCRetryTimeout=30000 (default is 500)

RPCRequestTimeout=90000 (default is 25000)

When a RPC is sent to the CAVA server, if the CAVA does answer after "RPCRetryTimeout", the Data Mover makes retries until "RPCRequestTimeout" is reached.

13.16.2 Shutdown Parameter

The shutdown field is used to modify the behavior of the system if no CAVA scanning servers are available.

- Viruschecking: Stops the virus checking if no AV server is available. This is the recommended setting.
- No: Continues retrying list of AV servers – users delayed until resolved
- CIFS: Stops CIFS if no AV servers are available – no user access

13.16.3 Parameter File Location

This parameter file should be in /etc/viruschecker.conf. If the file does not exist, then you may have to start with a sample copy from /nas/site/viruscheckers.conf

13.16.4 Unique File for each Datamover

Because each datamover will have its own unique version (due to CIFSserver field), you should keep a local copy of each file named according to the data mover. For example, have

etc/viruschecker.conf.server_2

etc/viruschecker.conf.server_3

etc/viruschecker.conf.server_4

...

However, this file must be called "viruschecker.conf" locally on each datamover in order to be read. To distribute this file to the datamovers, the server_file command can be used:

(server_file server_n [-put | -get] sourcefile destinationfile)

(as root)

```
# server_file server_2 -put viruschecker.conf.server_2 viruschecker.conf
server_2 : done
```

13.17 Create Local User on VNX

Before the virus-checking services are started, a local user that corresponds to the CAVA domain user must be created on each datamover. The purpose of this is to retain the CAVA domain user password. The groupid and userid information (when prompted) are not critical. The following is the command used to create this local user.

Note: You will be prompted for the group and uid values (these are not important, but should be unique).

```
[nasadmin@c5-ns480-cs0 ~]$ /nas/sbin/server_user server_4 -add -md5 -passwd
namdev.srv_emc_celerra_av
Creating new user namdev.srv_emc_celerra_av
User ID: 903
Group ID: 903
Home directory:
Changing password for user namdev.srv_emc_celerra_av
New passwd:
Retype new passwd:
```

To list the current local users on the DM, use the following command:

```
[nasadmin@c5-ns480-cs0 ~]$ /nas/sbin/server_user server_4 -list
server_4:
namdev.srv_emc_celerra_av:33smmmja0VJqE:903:903:f7mfBL321draowdbsQ4saPbgEH::ndmp_md5
```

13.18 Starting and Stopping CAVA

Before starting:

- Make sure that the CIFS services are configured and started.
- Make sure that CAVA is running on all of the AV servers.

To start the CAVA process on a datamover, the start option is used.

```
[root@cs-ns480-cs1 etc]# server_setup server_2 -Protocol viruschk -option start
server_2 : done
```

Use *ALL* for all datamovers (standby DMs will not be affected).

To stop the CAVA process on a datamover, the stop option is used.

```
[root@cs-ns480-cs1 etc]# server_setup server_2 -Protocol viruschk -option stop
server_2 : done
```

Use *ALL* for all datamovers (standby DMs will not be affected).

13.19 Updating Configuration File

To load the updated viruscheckers.conf into memory file without having to stop and start the services, use the following command:

```
#server_viruschk server_2 -update
server_2: done
```

Use *ALL* for all datamovers (standby DMs will not be affected).

13.20 Deleting CAVA

To delete the CAVA configuration on a datamover, the following command can be used:

```
[root@cs-ns480-cs1 etc]# server_setup server_2 -Protocol viruschk -option delete
server_2 : done
```

Use *ALL* for all datamovers (standby DMs will not be affected).

13.21 Checking the status of CAVA

To check the configuration and overall status of CAVA, the following command can be used:

```
[root@c5-vnx5500-cs0 nasadmin]# server_viruschk server_2

server_2 :
10 threads started.
2 Checker IP Address(es):
10.40.59.215                                ONLINE at Fri Feb 27 03:16:19 2015 (GMT-05:00)
                                           MS-RPC over SMB, CAVA version: 6.5.0.0, ntStatus:
SUCCESS
                                           AV Engine: Network Associates
                                           Server Name: icl-cee650-vm01.namdev.nsrootdev.net
                                           Last time signature updated: Thu Feb 26 16:00:11
2015 (GMT-05:00)
10.40.59.48                                ONLINE at Fri Feb 27 03:16:19 2015 (GMT-05:00)
                                           MS-RPC over SMB, CAVA version: 6.5.0.0, ntStatus:
SUCCESS
                                           AV Engine: Network Associates
                                           Server Name: i10-cava-vm01.namdev.nsrootdev.net
                                           Last time signature updated: Thu Feb 26 17:07:43
2015 (GMT-05:00)

1 File Mask(s):
*.*
12 Excluded File(s):
*.ARC *.ARJ *.CAB *.DAT *.GZ *.ISO *.JAR *.PST *.TAR *.TGZ *.Z *.ZIP
Share \\10.40.42.206\CHECK$.
RPC request timeout=90000 milliseconds.
RPC retry timeout=30000 milliseconds.
High water mark=200.
Low water mark=50.
Scan all virus checkers every 10 seconds.
When all virus checkers are offline:
Shutdown Virus Checking.
Scan on read disable.
Panic handler registered for 65 chunks.
MS-RPC User: NASEMCCAVA04$
MS-RPC ClientName: nasemccava04.NAMDEV.NSROOTDEV.NET
```

This status shows the current configuration including values from the active viruschecker.conf file. To check the counters and brief status of CAVA, the following commands can be used:

```
[nasadmin@c5-ns480-cs0 ~]$ server_viruschk server_2 -a
server_2 :
Total Requests: 117664.
Requests in progress: 0.

NO ANSWER from the Virus Checker Servers: 0.
ERROR_SETUP: 0.
FAIL: 0.
TIMEOUT: 0.
min=2071 uS, max=2238148 uS, average=3851 uS

0 files in the collector queue.
0 files processed by the AV threads.
```

Note: The total requests will be an ever increasing counter for every file that needs to be scanned. As users update files or copy files to the CIFS share, the counter will increase. The Requests in Progress counter may hover around 0-10 depending on how quickly the requests are being processed. An increasing number of in-progress files indicate that the virus scanning is falling behind.

For a detailed status of the CAVA servers, use the `–audit` option.

```
[nasadmin@cs5-ns480-cs0 ~]$ server_viruschk server_2 -audit
server_2 :
Total Requests                : 138
Requests in progress          : 25

NO ANSWER from the Virus Checker Servers: 0
ERROR_SETUP                   : 0
FILE_NOT_FOUND                 : 0
ACCESS_DENIED                  : 0
FAIL                           : 0
TIMEOUT                        : 0
Total Infected Files           : 875
Deleted Infected Files         : 64
Renamed Infected Files         : 0
Modified Infected Files        : 811
min=70915 uS, max=1164891 uS, average=439708 uS

15 File(s) in the collector queue
10 File(s) processed by the AV threads
Read file '/.etc/viruschecker.audit' to display the list of pending requests
```

13.22 Alerting on CAVA events

It is important to be notified when CAVA servers have failed or if virus scanning has been disabled if all CAVA servers have failed. This can be configured under the notifications within Celerra Manager. For information on how to configure these alerts see [CAVA and DEDUPE Alerting \(Event Notifications\)](#).

13.23 Confirmation of Virus Scanning Protection

You can examine the log files (for a passive approach) or you can actively test the virus scanning protection.

13.23.1 Test Virus Scanning

The following are procedures to test that Virus Scanning is in place and working. It is not intended to be used on frequent basis as it is a manual process, but more of an end-to-end check to confirm CAVA is working as intended. You will need a client that will have read/write access to a mapped CIFS share that is presumed to be protected by CAVA and VSE.

1. Disable the VSE software on the client that is going to write the file, otherwise the client VSE software will pick up the file.
2. Go to the EICAR (European Institute for Computer Antivirus Research) at: <http://eicar.org/86-0-Intended-use.html> to obtain the anti-malware test file – approximately 68 characters long. This is a string that is indented to identify a file as a virus, though it is not a harmful piece of code. Copy the string into the buffer (select the string and copy it from IE).
3. Open an editor such as WordPad and paste the EICAR string into a new file.
4. Save the new file as a typical file extension (such as .doc) on the mapped CIFS share.
5. If VSE is working successfully, you will get a virus found message in the AV log file of the CAVA server (see example below).

-
6. On the NAS share you will either see the file has been deleted or similar (depending on the action policy defined in ePO).

13.23.2 Log Files

Log entries will only appear for infected files detected or unsuccessful scans. The log file for McAfee Virus Scans is:

C:\Documents and Settings\All Users\Application Data\McAfee\DesktopProtection\OnAccessScanLog.txt

13.24 CAVA Sizing Tool

The CAVA sizing tool is available as an interactive menu and command line script. It can be used for generating CAVA stats and basic sizing purposes.


When estimating whether the current CAVA setup has enough capacity, it is recommended to analyze the “saturation” level (output from the tool). An optimal saturation level should be around 60%. If the saturation level is consistently above (or below) the optimal level, consideration should be given to completing a capacity planning exercise to determine whether more (or less) CAVA servers are required to meet the current requirements. For comprehensive capacity analysis, data from regular wintel performance monitoring of the CAVA servers should also be included to provide better understanding of the server resource utilization.

It is noted that output from the tool on older CAVA versions does not give an accurate measurement of the capacity because files are scanned serially. Newer versions of CAVA (above 4.6.7) are able to do multiple file scanning and the tool is more accurate for capacity measurement. Hence it is recommended to upgrade older CAVA setups to the latest CAVA versions for the multiple file scanning feature and accurate measurement by the tool. This also means that newer CAVA versions are more likely to become saturated, so it is recommended that the CAVA servers should be deployed for CAVA only and not be used as AV server for other purposes. No other applications or processes should reside on the CAVA server.

13.24.1 Enable the CAVA Sizing Tool Registry Option

Before you can run the tool, on each CAVA server that you wish to monitor, you will have to change a registry key to enable monitoring.

1. Run the registry editor (regedit.exe) on the CAVA server.
2. Locate the Sizing entry in
HKEY_LOCAL_MACHINE\SOFTWARE\EMC\Celerra Event Enabler\CAVA\Sizing
3. Double-click on Sizing entry and the Edit DWORD Value dialog box appears.
4. Change the Value data field to 1 to enable the tool (0 to disable).




5. Close the registry editor.
6. Restart the CAVA service.
7. Repeat for any related CAVA servers.


13.24.2 Interactive CAVA Sizing Tool

The simplest way to run the tool is to run cavamon.exe directly. This method requires that you create a dat file and it has an interactive display for the results.


1. Create a text file named `C:\Program Files\EMC\CAVA\cavamon.dat`
2. Add a line with the IP address (or FQDN) for each CAVA server you want to monitor.
3. Save and close the file.
4. Run `Program Files\EMC\CAVA\cavamon.exe`




5. Select the Get Stats button to display the CAVA server stats. It will periodically refresh to display the latest stats. Take note of the Saturation % value if you wish to estimate on the CAVA server(s) capacity. The Saturation % of any CAVA server should be below 60%. Above that value, additional CAVA servers or CPU/Mem resources should be considered.



If not all of the CAVA servers in the cavamon.dat file have the registry entry enabled or if there are newer versions than you are running the tool from, then you may get an error message:



6. Select the Clear Output button.
7. Select the size button.



The size output will never be less than the number of configured CAVA servers. Therefore, the saturation percentage (below 60%) for each server is a better indication of the number of CAVA servers needed.

13.24.3 Scripting CAVA Sizing Tool

The command line (scripting) version of the tool only shows stats of the CAVA servers. It does not make the overall recommendation (that has issues described above). The Visual Basic scripting file is located in the same folder as the interactive menu.

1. Open up Command Prompt
2. Navigate to the folder containing the tool (C:\Program Files\ EMC\CAVA)
3. The format of the command is:

```
csript cavamon.vbs <hostname1> <hostname2>...
```

<hostname> is the domain name or IP address of the CAVA server(s) you want to examine with a space between multiple servers.

4. Take note of the Saturation % value if you wish to estimate the CAVA server(s) capacity. The Saturation % of any CAVA server should be below 60%. Above that value, additional CAVA servers or CPU/Mem resources should be considered.

```
C:\WINNT\system32\cmd.exe
C:\Program Files\EMC\CAVA>cscript cavamon.vbs nasmgtgdse05
Microsoft (R) Windows Script Host Version 5.6
Copyright (C) Microsoft Corporation 1996-2001. All rights reserved.

Server:                \\ NASMGMTGDSE05 \
AU Engine State:       Up
AU Engine Type:        Network Associates
Files Scanned:         2156
Health:                Good
Msec Per Scan:         85.48
Saturation %:          53.00
Scans Per Second:      6.2
CAVA State:            NORMAL
CAVA Version:          4.6.8
```

14 DeDuplication (Dedupe) and Compression

For each file system, file-level deduplication gives the Data Mover the ability to process files in order to compress them, as well as the ability to share the same instance of the data only if they happen to be identical. Deduplication functionality operates on whole files and is applicable to files that are static or nearly static. Enabling deduplication also enables compression.

During deduplication, each deduplication-enabled file system on the Data Mover is scanned for files that match specific criteria, such as last access time and a modification time older than a certain date. Once a file is found that matches the criteria, the file data is deduplicated—compressed if appropriate and single instanced. Different instances of the file can have different names, security attributes, and timestamps. None of the metadata is affected by deduplication. When a user reads a file that was deduplicated, the VNX retrieves the data so that the NAS clients are unaware that the data was deduplicated.

If there are 100 unique files in a file system that are eligible to be deduplicated, the 100 unique files will remain but the data will be compressed, yielding space savings up to 50%. If there are 50 identical copies of a presentation on a file system, the 50 files will still exist, but all will share the same file data. In addition, the single file data would also be compressed (saving more space).

Deduplication /Compression is highly recommended as it requires low system impact and a high degree of storage savings.

14.1 Configuration

To configure Dedupe and Compression, go to Storage → Storage Configuration → File Systems. The file system properties page has a simple Off / On/ Suspended state. To enable deduplication and compression with default settings, select the On

To make changes to the Deduplication of Compression settings, select the properties of any file system. Then select the Deduplication Settings tab.

14.2 Compression

In addition to removing extra copies of data (deduplication), there is a compression function that is also enabled when you enable deduplication. Read operations decompress the portion of the file requested on-the-fly, in memory. Read operations do not cause the file to be decompressed on disk. When a user tries to write to a deduplicated file, the entire file is decompressed and reduplicated on disk to allow the write operation to be performed.

Quick Compression

Quick Compression is the recommended compression algorithm to use.

Deep Compression

14.3 Enabling Deduplication

There are three possible states for deduplication: on, off, or suspended. To enable deduplication, go to the file system properties (Storage → Storage Configuration → File Systems). This can be done for on existing files systems with data or on files systems that are new and empty. The deduplication states are:

- On = Enable deduplication/compression processing on a file system. Deduplication is the process used to compress redundant data, allowing space to be saved on a file system. Setting the state to on schedules the file system to be the next file system that is scanned. If there are no active scans, it starts immediately. Compression is also turned on. Deduplication and compression cannot be separated.
- Off = Undo all deduplication/compression processing. Do not perform any new space reduction. Any data that was deduplicated is now reduplicated, which is the process used to restore a file that was deduplicated to its original condition. Any compressed files are uncompressed. This process may take some time. If reduplication fails, such as when the system detects that the file system is too small to contain the reduplicated data, then the state transitions to the suspended state, the file system is left in a consistent, usable state, and a CCMD message will be sent to the server's event log. If reduplication succeeds, then it remains in the off state.
- Suspended = Suspend deduplication processing on a file system. The state transitions to the suspended state. No new space reduction is performed, but existing space-reduced files (compressed or dedupe) remain as-is.

To enable deduplication use this command syntax for the command line:

```
$ fs_dedupe -modify { <fs_name> | id=<fs_id> } -state on
```

where:

<fs_name> = name of the file system

<fs_id> = identifier of the file system

Example:

To enable deduplication space-reduction processing on file system fs_ufs1, type:

```
$ fs_dedupe -modify fs_ufs1 -state on
```

To list all deduplication-enabled file systems on the VNX, type:

```
$ fs_dedupe -list
```

Output:

```
fs_dedupe -list
```

```
id name state status time_of_last_scan original_data_size usage space_saved
401 fs_ufs1 On Idle Wed Nov 12 13:54:29 EST 2008 352 MB 119% 93 MB (26%)
```

14.4 Processing Time

If more than 1 file system has been enabled for deduplication, then only the first file system will be deduplicated and compressed immediately. The remaining file systems will be queued until the first file systems have been completed. The GUI deduplication statistics may not be updated frequently during the processing. Processing time for each file system depends on the number and size of the existing files. The exact timing for any file system can be observed in the server_log listing. For example:

```
2010-10-04 17:16:48: DEDUPE: 6: PE: Task 1 starting scan on fsid 323
2010-10-04 17:16:59: DEDUPE: 6: PE: Task 1 finished scan on fsid 323, duration: 11
seconds
```

It may take several seconds to hours depending on the number of files and if the deduplication process has been run before. Subsequent scans against unchanged portions of file systems will be significantly faster than the original scan.

As an approximate guideline the following example is provided:

File System Scan	Space Before Dedupe / Compression	Space After Dedupe / Compression	Time To Complete
1 st scan of 600GB FS with 31,079 files	76,584 MB	60,624 MB	36 min
2 nd scan of 600GB FS with 31,079 files	<i>Dedupe/compression already on and has run successfully</i>	60,573 MB	46 sec

14.5 Performance

The performance of the scan and ingest process depends on the system load on the datamover. At this level, the deduplication process:

- Scans up to 3.6 billion files per week at an average rate of 6,000 files per second.
- Uses approximately 5 percent of CPU processing capacity

14.5.1 Read Access

Read access to a deduplicated file is a pass-through operation so users will see little impact in terms of read-performance. Random reads of deduplicated files may appear to users in the same way as reads of normal files appear to them or slightly faster. Large sequential file reads may take longer than reads of files in their uncompressed form. Read performance depends on the number of streams and compressibility of the files. Reading deduplicated files does require more CPU cycles than reading uncompressed files.

- Random reads: 54% to 139% compared to uncompressed files
- Sequential reads: 68% to 107% compared to uncompressed files

14.5.2 Write Access

Writing to a deduplicated file prompts a reduplication of the requested portion of the compressed file data. This effectively removes the limit on the maximum size of a file that can be deduplicated and the performance impact.

14.6 File System Selection

All file systems should be set with deduplication turned on by default. Deduplication can be suspended during restores or large file moves as necessary. High CPU DMs may not be good candidates as the deduplication process takes additional resources to compress/decompress files in memory and to deduplicate files. The throttle thresholds for deduplication are 40% and 75% (full throttle enabled at 40% and below, low throttle is enabled at 75%). Low throttle will result in longer deduplication times.

File systems with high write performance requirements may also be excluded.

14.7 Additional Configuration and Deduplication Parameters

To make changes to the Deduplication of Compression settings, select the properties of any file system. Then select the Deduplication Settings tab.

Changes are effective by selecting the *OK* or *Apply* buttons. The default values are recommended with the exception of the cycle time (minimum scan interval). Normally you want the process to run once per week (7 days), but when converting new file systems, it may be beneficial to temporarily change this value to once per day (1 day). For more information about conversions see [Methodology for Initial Deduplication](#) below.

The parameters are specified on each file system (or datamover):

Parameter	Description
Case Sensitive	Select to perform case sensitive file name comparisons for NFS environments. Leave this checkbox unselected (default and recommend value) to perform case insensitive comparisons for CIFS environments.
Compression Method	Choose the compression method to use for this file system: fast (default) or deep. The fast option is the (default and) recommended compression method that is optimized for speed rather than space efficiency. It achieves the original compression ratios and performance. The deep option is the compression method that is optimized for space efficiency rather than speed. It achieves space savings up to 30% greater than the fast method. For example, if a file is 50% compressible then the deep algorithm can compress the same file up to 65%. However, the compression and decompression time when using this deep option is longer than when using the fast option. You obtain more storage space at the cost of slower access. Selecting this deep compression method applies only to new files that are subsequently compressed, and not to existing compressed files. When using VNX Replicator, VNX systems that use version 7.0 and earlier cannot read the deep compression format and will return an I/O error if a read operation is attempted. Select the deep compression format only if downstream replication sessions are using compatible software or are scheduled to be upgraded soon

CIFS Compression Enabled	Select to enable CIFS compression (default and recommended value) . When enabled and the deduplication state of the file system is either on or suspended, then CIFS compression is allowed. If the deduplication state is either off or in the process of being turned off, then CIFS compression is not allowed, regardless of whether this CIFS Compression Enabled checkbox is selected.
Duplicate Detection Method	Choose the duplicate detection method capability on a Data Mover: sha1 (default value), byte, or off: • sha1 identifies that a SHA-1 hash is used to detect duplicate data. This is the default and recommended method. • byte identifies that a byte-to-byte comparison is used to detect duplicate data. This adds considerable overhead, especially for large files. • off identifies that the duplicate detection method is disabled. With this setting, every deduplicated file is considered unique and the only space savings made are accomplished with compression.
Access Time	Enter the minimum required file age in days based on read access time. Files that have been read within the specified number of days will not be deduplicated. This option does not apply to files with an FLR locked state. Setting this value to 0 disables this option. 0-365 is the range of values and the default value is 15 days and is the desired setting after dedupe can run without issue . Temporarily setting this to 365 may help run dedupe initially.
Modification Time	Enter the minimum required file age in days based on modification time. Files updated within the specified number of days will not be deduplicated. Setting this value to 0 disables this option. 0-365 is the range of values and the default value is 15 days and is the desired setting after dedupe can run without issue . Temporarily setting this to 365 may help run dedupe initially.
Minimum Size	Enter the file size in KB that limits deduplication. File sizes equal to this value or smaller will not be deduplicated. File sizes greater than this value will be candidates for deduplication. This value should not be set lower than 24 KB. Setting this value to 0 disables this option. 0-1000 is the range of values and the default and recommended value is 24 KB .
Maximum Size	Enter the file size in MB, GB, or TB of the largest file to be processed for deduplication. Files larger than this size will not be deduplicated. Setting this value too high may affect system write performance as the first write operation reduplicates the file in its entirety. Setting this value to 0 disables this option. 0-8388608 MB is the range of values and the default and recommended value is 8388608 MB (or 8 TB) .
File Extensions Excluded	Specify a colon-delimited list of filename extensions to be excluded from deduplication. Filename extensions can be any valid alphanumeric character. Each extension must include the leading dot. The default value is empty and is the desired setting after dedupe can run without issue . Temporarily putting certain exclusions may help run dedupe initially.
Minimum Scan Interval	Enter the minimum number of days between completing one scan of a file system and before scanning the same file system again. 1-365 is the range of values and the default value is 7 days and is the desired setting after dedupe can run without issue . Temporarily setting this to 1 (run every day) may help run dedupe initially for conversions.
SavVol High Water Mark	Enter the percentage of the configured save volume (SavVol) auto extension threshold that can be used during deduplication. Once the specified amount of SavVol is used, deduplication stops on this file system. By default, this value is 90 percent and the SavVol auto-extension is also 90 percent. The recommended setting is 90 . This option will apply when the SavVol is 81 percent full (90 * 90). Setting this value to 0 disables this option. 0-100 is the

	range of values. Temporarily setting this value to 97 will take the SavVol usage to 87 and may help run dedupe initially.
Backup Data High Water Mark	Enter the full percentage that a deduplicated file has to be below in order to trigger space-reduced backups for NDMP. For example, when set to 90, any deduplicated file that has 90 percent or more of its blocks uncompressed will have the entire file data backed up without attempting to back it up in a space-reduced format. Setting this value to 0 disables this option. 0-200 is the range of values and the default and recommended value is 90 percent.
CPU % Low Water Mark	Enter the average percent of CPU usage that can be used during the deduplication process at which full throttle mode is re-entered. 0-100 is the range of values and the default and recommended value is 40 percent.
CPU % High Water Mark	Enter the average percent of CPU usage that can be used during the deduplication process which should trigger a slow throttle mode. The system starts in full throttle mode. 0-100 is the range of values and the default and recommended value is 75 percent.

Default values are recommended for all of the parameters as they reflect the capabilities of the DART code. Exceptions are listed below in [Methodology for Initial Deduplication](#) below.

14.8 File System Space Usage

14.8.1 Auto-extend file systems

The deduplication process aborts if the file system is within 1% of the auto-extend threshold. This prevents the file system from unexpectedly auto-extending. For example if the auto-extend HWM is set to 90%, the deduplication process aborts after the file system reaches 89% full.

14.8.2 Non auto-extend file systems

The deduplication process aborts if the file system is within 99% of its capacity.

14.8.3 Working space from File System

During the deduplication process, the file system must have enough space available that is equivalent to the size of the original file to be deduplicated plus the size of the compressed version of the file to be stored. This is in addition to the 1%/5% (or 1%/5% below autoextend) listed above.

DART Version + FileSystem	Freespace to Run Dedupe Process	Working Space**	Total Freespace
5.6.47 or later + autoextend	1% below extension → (typically 89%)	largest_file + .80 * largest_file	.11*FS-(1.8*LF)
5.6.47 or later	1%	largest_file + .80 * largest_file	.1*FS-(1.8*LF)

****The calculations above assume a 20% savings on compression.**

If there is not enough file system space you will see errors in the server_log listing regarding NoSpace. For example:

```
2010-10-04 16:14:47: DEDUPE: 3: PE: ProcessFile failed on file
/dedupe_fs2/TestFolderNumber2/hbsmupgrade.b_i_n.txt with error: NoSpace
```

Once the dedupe process encounters 10 errors (i.e. 10 files that are too big for the available freespace) the dedupe process halts.

14.9 SavVol Space Usage

The SavVol high water mark threshold is the percentage of configured SavVol auto-extension threshold that can be used during deduplication. After the specified amount of SavVol is used, deduplication stops on the specified file system. By default, this value is 90% of the SavVol auto-extension is also 90% (.9 * .9 = .81) or 81%.

Deduplication will grow the checkpoints when duplicate files are removed from the primary file system. The growth is approximately the size of a compressed copy of the files. Once the dedupe process encounters 10 errors (i.e. 10 files that will not fit into the checkpoint) the dedupe process halts.

The SavVol space is compared against the max savvol space value. Therefore if the max space is less than the current space, dedupe will not be successful. You will see SavVolAtCapacity errors in the server_log listing. For example:

```
2010-10-04 16:56:30: DEDUPE: 3: PE: ProcessFile failed on file
/dedupe_fs1/TestDedupeData/Example_SharedDisk_Configuration.doc with error:
SavVolAtCapacity
```

These same errors occur if there is simply not enough space to delete the duplicates from the file system and to place a compressed copy of the deleted files in the latest checkpoint without triggering an expansion.

The natural cycle of deduplication freeing up space for future deduplication works best with SavVOL storage set to autoextend. Autoextension of the SavVOL is recommended especially if a different tier (or storage pool) of storage is used for SavVOLs. This is ideal because undesirable behavior (such as massive file deletes) can force SavVOL space to grow unexpectedly. If SavVOLs are sharing space with primary file systems, then this could put those primary file systems at risk for running out of storage. *However, a maximum size can be set on a SavVol to prevent such a freespace robbery.* For more information see [Disk Space for Creating and Extending SavVols \(Automatic\)](#) or [SavVol Size Customization \(Manual\)](#).

14.10 Methodology for Initial Deduplication (Due to Space Limitations)

There is a certain methodology that can be followed to successfully deduplicate an existing file system with limited SavVol space or File System space. This methodology is to alter some of the parameters as described below and then to reattempt dedupe. Getting the deduplication to work at least one time is a critical step to saving space. No space will be saved until it completes at least one successful run. Once the deduplication has successfully run and space is recovered, then parameters can be adjusted slowly (with repeated manual or automatic dedupe cycles) back down to the recommended levels.

If the file system is tight on space (primary file system or SavVol), the following are the recommended steps.

14.10.1 File System Workarounds

1. Change the Minimum Scan Interval to be a much shorter period (for example every 1-2 days). This will help recover storage more quickly and will free up more primary file system space.
2. Adjust the parameters for AccessTime and ModificationTime to 365.
`fs_dedupe -modify <fs_name> -modificationTime 365`
`fs_dedupe -modify <fs_name> -accessTime 365`
 Rereun dedupe. If successful, then modify these values to 90 days less and repeat this parameter modification/dedupe cycle until all parameters are back to recommended values (15 days for each).

If after rerunning dedupe again without success, the following more aggressive steps should be considered:

3. Exclude certain file types (such as zip or cab files) that may help trigger space limitations. Examine the sever log to see what files are causing the error to prevent dedupe from succeeding. Reattempt dedupe.
4. Turn off deduplication and only keep the compression turned on. Turn back on deduplication when enough space has been released.
5. Remove older files /directories temporarily. Re-enable dedupe and then replace the files/directories after more space is released.
6. Extend the primary file system manually. This should be the last option as it is somewhat defeating the purpose of deduplication (saving space). This may the option of last resort.

Note: The File System space is non-reclaimable once extended.

14.10.2 SavVol Workarounds

1. Change the Minimum Scan Interval to be a much shorter period (like every 1-2 days). This will help recover SavVOL storage more quickly. However, checkpoints that contain previous deletes must roll off before the SavVol space is recovered.
2. Change the SavVol High Water Mark to 97. This will allow up to 97*90 or 87% of the SavVol to be used during dedupe.
`fs_dedupe -modify <fs_name> -savvol_threshold 97`
3. Adjust the parameters for AccessTime and ModificationTime to 365.
`fs_dedupe -modify <fs_name> -modificationTime 365`
`fs_dedupe -modify <fs_name> -accessTime 365`
 Rerun dedupe. If successful, then modify these values to 90 days less and repeat this parameter modification/dedupe cycle until all parameters are back to recommended values.

If after rerunning dedupe again without success, the following more aggressive steps should be considered:

4. Exclude certain large file types (such as .zip,.cab,or .pst files) that may be triggering space limitations. The server_log may indicate this. Rerun dedupe.
5. Turn off deduplication and only keep the compression turned on. Turn back on deduplication when enough space has been released.
6. Delete the oldest checkpoints and force deduplication to re-run by re- enabling deduplication on the file system. You could also look at the size of each checkpoint and selectively delete a checkpoint that is taking up a significant amount of the SavVOL space. Rerun dedupe.
7. Manually extend the SavVol space and then immediately re-enable deduplication on the file system. This should be the option of last resort as it defeats the purpose of deduplication to save/recover space.

Note: The SavVol space is non-reclaimable once extended unless all of the checkpoints are deleted.

14.11 Controlling When Deduplication is Run

Currently, there is no definition for when deduplication is run. The schedule time is based on the last time it was run for a file system plus the parameter Minimum Scan Interval. However, if deduplication is running successfully, you can effectively schedule it by running it manually at a selected time. This is desirable in an environment when Dedupe should be run in off-peak periods of time to minimize impact. For example, you could suspend deduplication for all file systems and then enable it for all file systems at a good start time.

```
$ fs_dedupe -modify xzy -state suspended
Wait until desired time range has started...
```

```
$ fs_dedupe -modify xyz -state on
```

14.12 Deduplication and Backups

There are several aspects to deduplication that carry onto backups. Most backups will take advantage of compression and some will take advantage of both compression and deduplication.

14.12.1 Network Share Based Backups

When a VNX file system is backed up through CIFS or NFS, space-reduced files are not reduplicated to their original size for transfer to the backup application. They are also not reduplicated on disk. However, the space saving and storage efficiency savings realized in the production file system do not flow through to backups when using network-based backups. This backup method is typically not used at Citi.

14.12.2 NDMP Based Backups

When backing up a deduplication-enabled file system with NDMP PAX (tar and dump), the compressed version of the file is backed up and all instances of the deduplicated file appear in the backup stream. In most cases, the majority of the space-savings comes from the compression, which means the majority of the space savings made in the file system flows through the backup as well.

When restoring files from a NDMP PAX backup, the files are restored in their compressed format.

Backing up a deduplication-enabled file system results in a shorter time window. However, the restore window required will increase because the system deduplicates the files being restored in real-time.

14.12.3 NDMP Volume Backups (NVB or VBB)

The VNX NDMP Volume Backup or Volume-Based Backup can backup deduplicated file systems (preserving deduplication and compression) and can restore them in full by using the full destructive restore method. NVB (or VBB) operates at the block level so it does not cause any reduplication of data during the backup or restore.

VNX does not support a single-file or file-by-file restore of files from NVB, so this method cannot be used unless the entire file system can be restored. Snapshots could be used to supplement the NVB/VBB backups for this purpose.

14.13 Checkpoints and Deduplication

The deduplication process releases space in the production file system immediately. However, blocks may be copied into the SavVol in the process. The exact impact of this affect is discussed in [SavVol Space Usage](#).

14.14 Replication and Deduplication

File systems with deduplication enabled can be replicated by using Replicator. Use of deduplication when replicating file systems will greatly reduce the amount of data sent over the network as a part of the initial baseline copy process. The impact of deduplication on replication after the initial baseline copy has completed, depends on the relative timing of replication updates and deduplication scans. Typically, replication updates (minutes or hours) are far more frequent than deduplication/compression scans of a file system (days or weekly). New and changed data in the file system are usually replicated in its non-deduplicated form first and any subsequent deduplication of that data prompts additional replication traffic due to the block changes within the file system. This is true of any post-process deduplication solution because replication updates more frequently than deduplication scans. The space savings realized by the production file system is reflected on the destination file system. *Deduplication can add up to 25MB/s of write activity during dedupe processing.*

14.15 File Level Retention (FLR) and Deduplication

You can enable deduplication on both the enterprise and compliance types of File Level Retention (FLR) file systems without compromising the protection on the files within the file system.

14.16 Alerting on DEDUPE events

It may be important to be notified when Deduplication fails on certain systems. For information on how to configure these alerts see [CAVA and DEDUPE Alerting \(Event Notifications\)](#).

15 Upgrading and Downgrading the VNX

15.1 Upgrades

There are three types of upgrades on VNX. They are:

- **In-Family Upgrades** – These are minor upgrades that are between two different revisions of the same major DART release. This will involve rebooting all of the datamovers and control station(s). Typically the datamovers are failed over to the standby datamover to keep the services available. Generally, replicated sites can have an in-family counterpart (source or target) without issue. Best practices dictate to upgrade the replication destination site first. For example going from 5.6.43.8 to 5.6.50.2
- **Patches** – These are similar to In-Family upgrades, but typically involve a smaller set of executables that have changed. This may or may not involve rebooting the control station. This will likely involve the rebooting of the data movers. Typically the datamovers are failed over to the standby datamover to keep the services available. Best practices dictate to upgrade the replication destination site first. For example 5.6.50.2 to 5.6.50.209.
- **Out-Of-Family Upgrades** – These upgrades are considered major and are disruptive. These upgrades require downtime for the entire array during the upgrade. They may be implemented as a series of smaller upgrades depending on the gap between the versions. Best practices dictate to upgrade the replication destination site first. This will likely include a FLARE microcode update to the storage (internal storage models). For example 5.6.50.2 to 6.0.40.8

15.2 Downgrades

EMC does not currently support the downgrades of VNX. EMC will rebuild a VNX to a specified revision of DART, but they will not guarantee preservation of data or configuration.

CATE lab testing has confirmed that a total loss of configuration and all user data (shares and contents) is required.

15.3 Upgrade Procedures

Upgrades are performed by EMC Staff only. EMC staff use a procedure generation system that defines the upgrade specifics as well as any required scripting to automate the upgrade. The high level steps are:

1. Suspend replication
2. Upgrade Control Station. (Local DM failover, GUI access, and alerting are suspended)
3. Restart Control Station
4. Upgrade Secondary Control Station (if applicable).
5. Upgrade DMs sequentially. (DMs will reboot at end of upgrade. The reboots can be manually delayed).
6. Resume replication.

Alternatively, you can avoid some interruption by failing over CIFS/NFS servers to remote site (depends on Dynamic DNS or 3DNS update). This may be more intrusive as a reboot of the DM.

16 Security Best Practices

In order to make the VNX environment as secure as possible and to compensate for lack of specific security features, the following security best practices must be implemented.

16.1 Functional IDs

These functional IDs must be managed by the storage operations teams and need to have passwords changed when appropriate. They should not be used by operational teams as a means of accessing the system. The exception would be if AD is inaccessible.

The table below provides information on the Functional IDs associated with OM according to new Functional Account technical requirements recently published by the CTI Information Security Administration (ISA) organization at:

http://www.citigroup.net/informationsecurityservices/doc/istg/TechnicalRequirementsforFunctionalAccounts_v1_2.pdf

User	Purpose	Role	Interactive	Privileged (Access to each VNX)	Password Expiration
Root	Administrative (super) user created with system	Admin	Yes	Yes	None (dual key control)
Nasadmin	Administrative user created with system	Admin	Yes	Yes	None (dual key control)
[domain]\srv_VNX_ecc	ECC connectivity for reporting and alerting	Read-only	Yes	Yes	None
[domain]\srv_VNX_ndmp	Backup processing	Read-only	Yes	Yes	None
[domain]\srv_VNX_filecensus	(Optional) Content auditing and reporting	Read-only	Yes	Yes	None
[domain]\srv_VNX_VirusScan	Account for proxy server to virus scan CIFS (windows) user data.	Read-only	No	Yes	None
[domain]\srv_emcsupport	Account for EMC Support	Admin (via sudo)	Yes	Yes	Per AD policy
[domain]\srv_netuitive	Account for Netuitive Reporting	Read-only*	No	Yes	None

**The Netuitive account (used only for VMWare) will be given elevated privileges for only server_stat command via sudo. All other commands are via read-only role.*

There are two functional IDs (root and nasadmin) that may be used under a dual key control methodology. That is, when a functional ID like root is needed (for upgrades, repair, etc.), no single user will retain the password. The password will be stored in a system such as CitiSafe such that two managers would need to request access to the password, with no single user holding both pieces at the same time. For example, if EMC support needed the root password for an upgrade or patch, then it would be obtained by two different managers, and then changed to a temporary value while support is using it, and then returned to the dual key value as before when support is finished. Alternatively, the EMC support account has specific root privileges for certain commands that are obtained via the sudo configuration.

16.2 Root Access for Storage Administrator and EMC support accounts (Sudo)

There are certain commands that need root privileges to run correctly. These commands are run with some frequency by EMC support staff and to a lesser degree by Citi Storage Admins (all member of nasadmin group). In order to prevent frequent access to the root password, root access can be granted to specific accounts via sudo. Any functional ids that need additional privileges should also be added to this file (i.e. Netuitive functional id).

This sudo privilege account access is fully logged (/var/log/secure). An example of a sudo entry is:

```
Jun  1 09:55:18 JVNT1C01 sudo:  gj06529 : TTY=pts/1 ; PWD=/tmp ; USER=root ;
COMMAND=/nas/sbin/nas_connecthome -modify -email_subject Call Home Alert c5-nsg8-cs01
```

This process is easily set up by the root account by modifying the /etc/sudoers file on the control station. It should be edited via the “visudo” command as it is very format sensitive. The following is approved content for this file. Additional command may be added at a later date. *It is also noted that some commands are based on the primary group of a user and may not function correctly despite the sudo configuration.*

```
# sudoers file- Citi Celerra NAS version : 1.4
#
## Sudoers allows particular users to run various commands as
## the root user, without needing the root password.
##
## Examples are provided at the bottom of the file for collections
## of related commands, which can then be delegated out to particular
## users or groups.
##
## This file must be edited with the 'visudo' command.

## Host Aliases
## Groups of machines. You may prefer to use hostnames (perhaps using
## wildcards for entire domains) or IP addresses instead.
# Host_Alias    FILESERVERS = fs1, fs2
# Host_Alias    MAILSERVERS = smtp, smtp2

## User Aliases
## These aren't often necessary, as you can use regular groups
## (ie, from files, LDAP, NIS, etc) in this file - just use %groupname
## rather than USERALIASES
# User_Alias ADMINS = jsmith, mikem

## Command Aliases
## These are groups of related commands...

## Networking
Cmdnd_Alias NETWORKING = /sbin/route, /sbin/ifconfig, /bin/ping, /sbin/dhclient, /usr/bin/net,
/sbin/iptables, /usr/bin/rfcomm, /usr/bin/wvdial, /sbin/iwconfig, /sbin/mii-tool

## Installation and management of software
Cmdnd_Alias SOFTWARE = /bin/rpm, /usr/bin/up2date, /usr/bin/yum

## Services
Cmdnd_Alias SERVICES = /sbin/service, /sbin/chkconfig

## Updating the locate database
Cmdnd_Alias LOCATE = /usr/sbin/updatedb

## Storage
Cmdnd_Alias STORAGE = /sbin/fdisk, /sbin/sfdisk, /sbin/parted, /sbin/partprobe, /bin/mount,
/bin/umount

## Delegating permissions
Cmdnd_Alias DELEGATING = /usr/sbin/visudo, /bin/chown, /bin/chmod, /bin/chgrp

## Processes
Cmdnd_Alias PROCESSES = /bin/nice, /bin/kill, /usr/bin/kill, /usr/bin/killall

## Drivers
Cmdnd_Alias DRIVERS = /sbin/modprobe

# Defaults specification
#
```

```
# Disable "ssh hostname sudo <cmd>", because it will show the password in clear.
#     You have to run "ssh -t hostname sudo <cmd>".
#
Defaults    requiretty

Defaults    env_reset
Defaults    env_keep = "COLORS DISPLAY HOSTNAME HISTSIZE INPUTRC KDEDIR \
                        LS_COLORS MAIL PS1 PS2 QTDIR USERNAME \
                        LANG LC_ADDRESS LC_CTYPE LC_COLLATE LC_IDENTIFICATION \
                        LC_MEASUREMENT LC_MESSAGES LC_MONETARY LC_NAME LC_NUMERIC \
                        LC_PAPER LC_TELEPHONE LC_TIME LC_ALL LANGUAGE LINGUAS \
                        _XKB_CHARSET XAUTHORITY NAS_DB"

## Next comes the main part: which users can run what software on
## which machines (the sudoers file can be shared between multiple
## systems).
## Syntax:
##
##     user    MACHINE=COMMANDS
##
## The COMMANDS section may have other options added to it.
##
## Allow root to run any commands anywhere
root    ALL=(ALL)    ALL

## Allows members of the 'sys' group to run networking, software,
## service management apps and more.
# %sys ALL = NETWORKING, SOFTWARE, SERVICES, STORAGE, DELEGATING, PROCESSES, LOCATE, DRIVERS

## Allows people in group wheel to run all commands
# %wheel    ALL=(ALL)    ALL

## Same thing without a password
# %wheel    ALL=(ALL)    NOPASSWD: ALL

## Allows members of the users group to mount and unmount the
## cdrom as root
# %users    ALL=/sbin/mount /mnt/cdrom, /sbin/umount /mnt/cdrom

## Allows members of the users group to shutdown this system
# %users    localhost=/sbin/shutdown -h now

User_Alias PAHC_USERS = %nasadmin, %fullnas, %root
User_Alias UPGRADERS = %fullnas, %root

User_Alias CLI_USERS =srv_netuitive

Defaults exempt_group = nasadmin
Defaults:PAHC_USERS !requiretty
Defaults:UPGRADERS !requiretty
Cmd_Alias PAHC_CMDS = /celerra/backendmonitor/bin/check_zone, /usr/bin/find, /sbin/sfdisk,
/sbin/fdisk, /usr/local/bin/nd-clnt, /usr/local/bin/nd-cfg,
/nas/sbin/check_workpart_and_log_slice_consistency, /nas/sbin/workpart, /nas/sbin/lvm_checkup,
/nas/sbin/setup_slot, /nas/sbin/lvm_checkup, /nas/tools/upgrade_check_tools/check_full_dump.pl,
/sbin/parted, /usr/sbin/pvdisplay, /var/tmp/.CUT/CUS*-*-/, /var/tmp/.CUT/CUS*-*-./dir/main,
/usr/sbin/smartctl
PAHC_USERS ALL = NOPASSWD: PAHC_CMDS
UPGRADERS ALL = NOPASSWD: /var/tmp/.CUT/CUS*-*-./dir/IMHC.dir/install_mgr
Cmd_Alias CLI_USERS_CMDS = /nas/bin/server_stats
CLI_USERS ALL =(ALL) NOPASSWD: CLI_USERS_CMDS

Cmd_Alias  EMCSUPPORT = /nas/sbin/cs_standby,/nas/sbin/nas_ca_certificate,\
                        /nas/sbin/nas_config,/nas/sbin/nas_connecthome,\
                        /nas/sbin/nas_halt,/nas/sbin/nas_mview,\
                        /nas/sbin/nas_rdf,/nas/sbin/server_user,\
                        /nas/sbin/rootfs_dhsm,/nas/sbin/rootfs_group,\
                        /nas/sbin/rootfs_rdf,/nas/sbin/rootfs_timefinder,\
                        /nas/sbin/rootnas_acl,/nas/sbin/rootnas_automountmap,\
                        /nas/sbin/rootnas_cel,/nas/sbin/rootnas_devicegroup,\
                        /nas/sbin/rootnas_disk,/nas/sbin/rootnas_diskmark,\
                        /nas/sbin/rootnas_fs,/nas/sbin/rootnas_fsck,\
```

```

/nas/sbin/rootnas_license,/nas/sbin/rootnas_pool,\
/nas/sbin/rootnas_quotas,/nas/sbin/rootnas_replicate,\
/nas/sbin/rootnas_server,/nas/sbin/rootnas_slice,\
/nas/sbin/rootnas_storage,/nas/sbin/rootnas_symm,\
/nas/sbin/rootnas_volume
emcadmin ALL = (ALL)      NOPASSWD: EMC SUPPORT
%celerra_admins ALL=(ALL) NOPASSWD: ALL

```


16.3 Control Station Security and Configuration

There are several items on the control station that need to be configured to make it secure and ready for deployment at Citi. These items include:


- Setup of DNS and NTP
- Access Control List (ACL)
- Authentication and Authorization via Active Directory

16.3.1 Control Station Setup of DNS and NTP


This configuration is controlled by the Storage Domain settings. You need to be logged in as a global domain administrator. When viewing all systems, select Domains. Under the Local Domain box on the right side, there will be a shortcut for DNS and NTP.



For DNS, enter the domain that is to be used on this array from a management perspective. Typically this is the NAM, APAC, LATAM, or EUR domains. We typically do not need the control station to communicate with other trusted domains (like REL) even though the datamover may have these domains configured. Use the Add button to add in a domain. If additional domains are needed, then the search list should also be configured giving priority to the most commonly used domain.



For NTP, enter up to 4x NTP (Time) servers in the following screen (reachable by the shortcut referenced above). The update cycle default is 30 minutes. This should be changed to at least 1 hour. There are no entries for *Key ID* or *Key Value*.



16.3.2 Access Control List (ACL) via IPTables

There must be an ACL implemented for limiting access to the VNX to known Citi Storage Administration systems. This will further prevent unauthorized access to the VNX.

EMC relies upon the internal firewall software called IPTables (formerly IPChains) to create an Access Control List that will limit access to the control station to known storage administration servers via HTTP, HTTPS, and SSH. This is an important compensating control to limit access to the VNX to known Storage Administration systems.

Warning: It is possible to configure the control station such that no IP traffic will get to the Control Station. If this happens and the configuration has not been saved, then a simple reboot of the CS should resolve the issue. It is very important to realize that saving the configuration will prevent this ability.

16.3.2.1 Environment Setup

To prepare the environment and to see what has been already configured use the following commands:

To view the current rules the following commands:

```
/sbin/iptables -L
/sbin/service iptables status
```

To turn on IPTables at boot time

```
chkconfig iptables on
```

16.3.2.2 Required Entries in ACL

ECC uses SSH, HTTPS, and HTTP to communicate with the VNX. Therefore, the ECC server running the Storage agent for NAS should be added to the ACL on the VNX control station.

ESRS also uses SSH to communicate with the VNX and should be added to the ACL. Note that there are typically 2 ESRS systems per site for redundancy.

Secondary control stations and control stations linked via replication should also be added to the ACL.

The following is a summary of the required entries:

Entry	Purpose	Ports
ECC servers	Allow for reporting by ECC	SSH, HTTP, HTTPS
ESRS gateways	Allow for remote access by EMC	SSH, HTTP, HTTPS
Secondary control station	Allow for sessions between primary and secondary control stations. Any shared IP address should also be present in the list.	SSH, HTTP, HTTPS
Control stations from linked replication systems	Allow for sessions between local and remote systems	SSH, HTTP, HTTPS
NAS management systems	Allow for sessions initiated from NAS management servers / workstations	SSH, HTTP, HTTPS
Backup sever(s)	Allow for sessions initiated from NAS Backup servers	SSH, HTTP, HTTPS

16.3.2.3 Sample Script

The following is a sample script that should be used after modifying the IP addresses in variables and in the access configuration lines. It is provided as a sample because it is critical that you modify the content before implementing. This script is also provided in the VNX Utilities package (). It is located in c:\Program Files\VNX Utilities\IPTables\.

```
#!/bin/sh
# =====
# iptables_setup.sh
# Version 2.1
#
# Copyright 2009 by Citi. All rights reserved.
# ===Revision History=====
# 1.0 06/01/2009 written by Dave Celano, EMC and Gary Johnson, Citi
# 1.5 12/31/2010 Various updates by David Lee for added systems/workstations
# 2.0 05/13/2011 updated by Tad Taylor(EMC) and David Lee (Citi) to limit ESRS
#      jumping
# 2.1 modified list of servers for lab use
#
#
# =====
#
# This script will allow all outgoing connections, and
# block ssh/HTTP/HTTPS in from unauthorized IPAddresses. Pick your ports using
# the $SSHPort $HTTPPort $HTTPSPort variable below.
#
# You will need to modify the address of the control station.
# You will need to modify the address of each allowed management server.
# This script flushes the configuration so running it multiple times will
# only keep the last configuration. If the configuration fails or gets the
# control station to unreachable state, then reboot the cs and it will bring
# up the lastest saved configuration.
#
# ** Very important not to save until you are sure the configuration is correct
# and has been verified. **
#
# Save this script as iptables_setup.sh in your /root directory.
# and run it as root.
# just run: sh /root/iptables_setup.sh
#
# If you would like to have the firewall automatically
# come up at boot time, add the path to the script to
# the bottom of your /etc/rc.d/rc.local file. For instance
# /root/iptables_setup.sh
#
# Put the location of the IPtable binary here:
IPTABLES="/sbin/iptables"
#
# what SSH port to use? (Remember to configure /etc/ssh/sshd_config)
SSHPort="22"
#
# what HTTP port to use?
```

```

HTTPPort="80"

# what HTTPSPort port to use?
HTTPSPort="443"

# what Interface port to use? (default is eth3 on control station)
INT="eth3"

# what is Control Station Hostname?
#CSHOST="host.domain of the Control Station"
CSNAME=$(/bin/hostname)
DNSNAME=$(/bin/dnsdomainname)
CSHOST=$CSNAME.$DNSNAME

# set default policies
$IPTABLES -P INPUT ACCEPT
# flush the table
$IPTABLES -F INPUT

#
# ***** #
# List allowed management servers here for HTTP/HTTPS/SSH access
# You will not need to add entries for ECC or other systems unless they use
# HTTPS HTTP or SSH.
#
# Allow SSH HTTP HTTPS connections to this control station from authorized
# IPAddresses
# add additional lines for each IPAddress and change the -s IPAddress accordingly
# ***** #

#####
# NAS Admin Servers
#####

ALLOWED_ADMIN_SERVERS="ecctest01ms.nam.nsroot.net
                       gdsehds01.namdev.nsrootdev.net
                       gdsehds02.namdev.nsrootdev.net
                       gdsehds03.namdev.nsrootdev.net
                       nasmgmtgdse05.namdev.nsrootdev.net
                       CPNAWRNWV-107CX.nam.nsroot.net
                       nasmgmtgdse04dr.namdev.nsrootdev.net
                       10.40.42.211"

for admin in ${ALLOWED_ADMIN_SERVERS};do
    $IPTABLES -A INPUT -s ${admin} -i $INT -d $CSHOST -p tcp -m multiport --
dports $SSHPort,$HTTPPort,$HTTPSPort -j ACCEPT
done

#####
# NAS Admin Workstations
#####

ALLOWED_ADMIN_WORKSTATIONS="CTIREGLR8LBWTX.nam.nsroot.net
                             CPNAMLVWV-10MWN.nam.nsroot.net"

for admin_workstation in ${ALLOWED_ADMIN_WORKSTATIONS};do
    $IPTABLES -A INPUT -s ${admin_workstation} -i $INT -d $CSHOST -p tcp -m
multiport --dports $SSHPort,$HTTPPort,$HTTPSPort -j ACCEPT
done

#####
#ESRS Gateways
#####

ESRS_SERVERS="169.193.211.77
              169.193.211.78
              169.192.41.79"

for ESRS in ${ESRS_SERVERS};do

```

```

$IPTABLES -A INPUT -s ${ESRS} -i $INT -d $CSHOST -p tcp -m multiport --
dports $SSHPort,$HTTPPort,$HTTPSPort -j ACCEPT
done

#####
#Command Central Agent
#####

CCA_SERVERS="w2tsibdindia001.nam.nsroot.net"
for CCA in ${CCA_SERVERS};do
    $IPTABLES -A INPUT -s ${CCA} -i $INT -d $CSHOST -p tcp -m multiport --dports
    $SSHPort,$HTTPPort,$HTTPSPort -j ACCEPT
done

#####
# ECC Agents
#####

ECC_SERVERS="169.171.119.123
169.193.242.142"
for ECC in ${ECC_SERVERS};do
    $IPTABLES -A INPUT -s ${ECC} -i $INT -d $CSHOST -p tcp -m multiport --dports
    $SSHPort,$HTTPPort,$HTTPSPort -j ACCEPT
done

#####
# Symantec Data Insight
#####

DI_SERVERS=10.40.43.176
for DI in ${DI_SERVERS};do
    $IPTABLES -A INPUT -s ${DI} -i $INT -d $CSHOST -p tcp -m multiport --dports
    $HTTPSPort -j ACCEPT
done

#####
# Best Practice Script Host
#####

Allowed_SERVERS="165.203.254.65
153.40.196.139"
for bp_script in ${ALLOWED_SERVERS};do
    $IPTABLES -A INPUT -s ${bp_script} -i $INT -d $CSHOST -p tcp -m multiport
    --dports $SSHPort,$HTTPPort,$HTTPSPort -j ACCEPT
done

#####
##
## INSSH_ALLOWED is the set of peer Control Stations to which
## inbound SSH connections are explicitly
## allowed
#####

INSSH_ALLOWED=$(/nas/bin/nas_cel -list | awk '/^[0-9]/{print $4}')
for IN_addr in ${INSSH_ALLOWED}; do
    $IPTABLES -A INPUT -s ${IN_addr} -i $INT -d $CSHOST -p tcp -m multiport --
    dports $SSHPort -j ACCEPT
done

#####
# ***** #
#
# Block ALL OTHER IPAddresses from SSH HTTP and HTTPS
# Note: Host will see 'Address is not valid' or similar message

```

```
#####

$IPTABLES -A INPUT -s 0/0 -i $INT -p tcp --dport $SSHPort -j REJECT --reject-with
tcp-reset

$IPTABLES -A INPUT -s 0/0 -i $INT -p tcp --dport $HTTPPort -j REJECT --reject-with
tcp-reset

$IPTABLES -A INPUT -s 0/0 -i $INT -p tcp --dport $HTTPSPort -j REJECT --reject-
with tcp-reset

# Allow connections IN that were established FROM this control station
$IPTABLES -A INPUT -i $INT -p udp -m state --state ESTABLISHED -j ACCEPT
$IPTABLES -A INPUT -i $INT -p tcp -m state --state ESTABLISHED -j ACCEPT
$IPTABLES -A INPUT -i $INT -p icmp -m state --state ESTABLISHED,RELATED -j ACCEPT
# These lines will log all SSH/HTTP/HTTPS connections
$IPTABLES -A INPUT -i $INT -p tcp --dport $SSHPort -j LOG --log-prefix "***SSH
connection:"
$IPTABLES -A INPUT -i $INT -p tcp --dport $HTTPPort -j LOG --log-prefix "***HTTP
connection:"
$IPTABLES -A INPUT -i $INT -p tcp --dport $HTTPSPort -j LOG --log-prefix "***HTTPS
connection:"

#Allow loopback traffic. This is required for alot of local configuration sevice
s
$IPTABLES -A INPUT -i lo -j ACCEPT

#####
#####
##### IPTables OUTPUT Rules Below
#####
#####
#
#####
##
## SSH_ALLOWED is the set of peer Control Stations to which
## outbound SSH connections are explicitly
## allowed
#####
SSH_ALLOWED=$( /nas/bin/nas_cel -list | awk '/^[0-9]/{print $4}')

for addr in ${SSH_ALLOWED}; do
    $IPTABLES --append OUTPUT -o $INT -p tcp --dst ${addr} --dport ssh \
        --syn -m state --state NEW -j ACCEPT
done
#
#
#
#
#####
##
## TCP_OUT represents the set of outgoing TCP protocols
## to block from the $INT interface
## UDP_OUT represents the set of outgoing UDP protocols
## to block from the $INT interface
##
#####
TCP_OUT="ssh telnet"

# Disallow outgoing ${TCP_OUT} connections over interface $INT
for port in ${TCP_OUT}; do

    $IPTABLES --append OUTPUT -o $INT -p tcp --dport ${port} --syn -m state --state
NEW -j LOG \
        --log-level alert --log-prefix "REJECTED OUTBOUND"
    $IPTABLES --append OUTPUT -o $INT -p tcp --dport ${port} --syn -m state --state
NEW -j REJECT
done

# ***** #
echo The IPTables rules have been created.
echo Use 'service iptables save' to save rules to /etc/sysconfig/iptables
echo Use 'chkconfig iptables on' to enable them at bootup.
echo Save this ruleset after configuring for your system and run or
echo re-run this script each time you modify your rules. Do not forget to
echo run 'service iptables save' each time, or the rules will not persist.
```

```
echo You can also run 'iptables -L' or 'service iptables status' to
echo see the current state of your rules
#
# *END OF SCRIPT*
```

16.3.2.4 Saving the configuration

This step should not be taken lightly. As mentioned previously, saving the configuration removes the ability to reboot the non-responsive CS as a remedy to a mis-configured IPTables software. Verify that you can access the control station (by HTTPS and SSH) from all of the defined systems.

Warning: It is possible to configure the control station such that no IP traffic will get to the Control Station. If this happens and the configuration has not been saved, then a simple reboot of the CS should resolve the issue. It is very important to realize that saving the configuration will prevent this ability.

To view the current rules the following commands:

```
/sbin/iptables -L
/sbin/service.iptables status
```

To save the current IPTables configuration

```
(to /etc/sysconfig/iptables) use the following command:
/sbin/service iptables save*
```

*This Linux feature saves the iptables configuration file with a 600 rather than 644 permission set. This will result in an erroneous error on nas_checkup when running from a non-root user like the following:

```
-----Errors-----
Storage System : Check SPA SPB communication
Error HC_BE_14505017473: At least one iptables rule is missing.
Action : Run "spa_spb_comm -setup" to fix the problem.
```


16.3.3 Control Station Authentication & Authorization via Active Directory

To turn on external authentication and authorization via Active Directory, a LDAP connection from the control station(s) must be configured and tested before proceeding. This method of authentication is required to meet security standards.

The CLI only supports the single role of administrator and cannot use any other roles. Unisphere supports roles and does have the capability to run CLI based commands (including non-interactive scripting).

16.3.3.1 CA SSL Certificate

It is important to understand that authentication to the AD environment is authorized only via SSL. Therefore a valid CA SSL certificate is necessary. The VNX only accepts CA certificates in a PEM format (Base64). To obtain the certificate in the proper format, go to the Citi certificate webpage (<https://digitalcertificate7.citigroup.com>) and look for a display in PEM Encoding for the CA certificate.



Copy and paste the contents into a new file with a .crt extension (using an ASCII editor like notepad). You may have to fill out a request that is server name specific that has a certain expiration. The following is an example certificate in the correct format:


```

-----BEGIN CERTIFICATE-----
MIIDejCCAuOgAwIBAgIERLgHAzANBgkqhkiG9w0BAQUFADAfMQswCQYDVQQGEwJV
UzEQMA4GA1UEChMHV0FSIDE1OTAeFw0wNjA3MTQyMDM1MTBaFw0yNjA3MTQyMTA1
MTBaMB8xCzAJBgNVBAYTA1VTMRAdDgYDVQQKEwdXQVIGMTU5MIGfMA0GCSqGSIb3
DQEBAQUAA4GNADCBiQKBgQCTKcPCTpoDqN9dr0/MCQtsPG3awk0cTaKd28Sfcz5L
1eYzGN+E8mMOC+YM4/yq+hjwlnVqjO9QTYlRYLyG7zd0uYwHDDhg+V8yG+4ue4T
cP2egmVS3+o49SeF865kvjeKsOne44IldK3M50GAnph6HuLrkFaUVtLtZVXmo01b
GwIDAQABo4IBWtCCAb0wRwYDVR0gBEAwPjANBgtghkgBhvlvAWMBADANBgtghkgB
hvlvAWMBBDAOBgXghkgBhvlvAWMBAwEwDgYMYIZIAYb5bFjAQMCMBEgcWCGSAGG
+EIBAQQEAWIABzCBtwYDVR0fBIGvMIGsMDAgNKAypDAwLjELMAkGA1UEBhMCVVMx
EDA0BgNVBA0TBldBuiAxBNTkxDTALBgNVBAMTBENSTDEwcqBwoG6GQ2xkYXA6Ly8x
MC40MC40MS4xNjAvbz1XQVIlMjAxNTksYz1VUz9jZXJ0aWZpY2F0ZVJldm9jYXRp
b25MaXN0P2Jhc2WGJ2h0dHA6Ly8xMC40MC40MS4xNTkxQ29tYmluZWZWRDRFavQ1JM
LmNyYDARBgNVHRAEJDAigA8yMDA2MDcxNDIwMzUxMFQBDzIwMjYwNzE0MjEwNTEw
WjALBgNVHQ8EBAMCAQYwHwYDVR0jBBgwFoAUudow9atzRT4qUK4IU8FFCENyKMw
HQYDVR0OBBYEFiRnaMPWrc0U+KlCuCFPBRQhDcijMAwGA1UdEwQFMAMBAf8wHQYJ
KoZIHvZ9B0EABBAwDhsIVjcuMT00LjADAQSQMA0GCSqGSIb3DQEBBQUAA4GBAGIa
KpWX3vonpg80K0iSSi6eupduVR8hLAaizHM/Pon7ICiY2WnAVVHhlpqc/1GMSiKNO
OFVMrVW92xDrnZqR84H0J1VEXNvNI328YlmWf5BNi94TmtQjtDqizpuCW1jcQVg1
o1IiyUETNdHB0QsX6BH7dIiwuJy6Wt1/DLF9Sky3
-----END CERTIFICATE-----

```

16.3.3.2 Domain Configuration

To configure the domain settings you need to be logged in with global root privileges. You also need to configure and verify that DNS and NTP are working on the control station (independent of datamover configuration). Those settings can be configured on the control station properties screen. Settings → Network Settings.



The following fields should be used for Citi AD implementations:

Note: Current Citi policy is to use location specific dynamic list of domain controllers that is discoverable. However, VNX does not support this. Hard-coded DCs must be used. This may result in control stations unable to authorize admins when the DC list is updated and both DCs are not available.

- Domain Name – This is the fully qualified domain name
- Primary – This is the fully qualified name of the primary Domain Controller (DC) that was given to you.
- Backup – This is the fully qualified name of the backup DC.
- SSL Enabled – This box should be checked
- SSL Primary Certificate – If the certificate is in the right format it will display the expiration date here. If a path is displayed (even before a successful bind), then the certificate is not in the correct format. An incorrect format will never produce a successful bind to the primary DC.
- SSL Backup Certificate – If the certificate is in the right format it will display the expiration date here. If a path is displayed (even before a successful bind), then the certificate is not in the correct format. An incorrect format will never produce a successful bind to the backup DC.
- Upload New SSL Primary/Backup Certificate – This is where you specify the file location of the certificate to be saved on the control station.

- Port – This is where you specify the communication port. For SSL it is 636 (Non-SSL is 389). Only SSL is supported in production Citi environments.
- Directory Service Type – Use Custom Active Directory. This will open up several more fields.
- User ID Attribute – This is the account type to be used. Use the default value of sAMAccountName.
- Distinguished Name and Password – This is the account (in distinguished format) that will bind to the AD server.
- User Search Path – This is the directory location that includes all of the domain storage administrators
- User Name Attribute – Use cn
- Group Search Path – This is the directory location that includes the group that the domain storage administrators are members of.
- Group Name Attribute – Use cn
- Group Class/Member – Use group and member for Class and Member respectively.

The following are sample values from the legacy screen for your reference:

Domain Name: namdev.nsrootdev.net

Primary: namdev003.namdev.nsrootdev.net

Backup: namdev004.namdev.nsrootdev.net

SSL Enabled: ☒

SSL Primary Certificate: Jul 14 21:05:10 2026 GMT (Jul 14 21:05:10 2026 GMT)

Upload New SSL Primary Certificate: Browse...

SSL Backup Certificate: Jul 14 21:05:10 2026 GMT (Jul 14 21:05:10 2026 GMT)

Upload New SSL Backup Certificate: Browse...

Port: 636

Directory Service Type: ☐ Default Active Directory ☒ Custom Active Directory ☐ Other Directory Servers

User Id Attribute: sAMAccountName

Distinguished Name: cn=svr_emc_vnx_ad,ou=svc_accts,ou=StorageEng,dc=

Account Password:

User Search Path: ou=StorageEng,dc=namdev,dc=nsrootdev,dc=net

User Name Attribute: cn

Group Search Path: ou=dept_grps,ou=StorageEng,dc=namdev,dc=nsrootde

Group Name Attribute: cn

Group Class: group

Group Member: member

Show: Control Station Properties (for DNS and NTP settings)

OK Apply Test Cancel Help



This screen has been replaced by the following:

Select "Manage LDAP Domain" to configure AD binding.

Security Settings

[Manage Idle Timeout](#)
[Manage LDAP Domain](#)
[Change Password](#)

Select "Add" and configure a "Primary and Alternate domain".



Sample Configuration: *This is a sample only.. each region will be different.*

Primary: namdcnj051.nam.nsroot.net

Backup: namdcnj049.nam.nsroot.net

BindDN: CN=srv_nas_Celerra_Inst,OU=Svc_Accts,OU=NAS,OU=INFRA,DC=nam,DC=nsroot,DC=net

Domain Name: nam.nsroot.net



password : \$Da*****\$


User search path: OU=Users,OU=CPNA,OU=CS,DC=nam,DC=nsroot,DC=net

Group search path: OU=Misc_Grps,OU=CPNA,OU=CS,DC=nam,DC=nsroot,DC=net

Cert Path: G:\Where the cert is stored\CA.crt

Sample of configured output:





Once you have configured this screen, select apply and then the *Test* button at the bottom. If the test bind succeeds (Test Domains Settings: OK), then you may proceed to the next step. If it fails, then you should scrutinize the above parameters and re-test (also re-verify DNS and NTP are working). The test must work before logging in with domain accounts will work. You can see more detailed information in the `server_log`

16.3.3.3 Group Configuration

In order to map domain users to local roles, you have to configure an existing or new group to use a domain mapped group. Go to Settings → Security → User Customization for File. It is strongly recommended that you create a new group for domain access to keep the local and domain environments as separate as possible. Within group properties, change the selection from “Local only Group” to “Domain mapped Group”.

- Group Name – This is a descriptive name for the group.
- GID – This is the group ID field. Auto-select is recommended.
- Role – This is the role that will be assigned to this group. If a user is a member of multiple groups, then each of the roles will be assigned to the user.
- Group Type – LDAP Domain Mapped should be selected
- Domain Group Name – This is the name of the (permission) group that is used to identify users as VNX administrators. For example “per_gdse_VNX_admins”.
- Domain Name – This is the fully qualified domain name. For example “namdev.nsrootdev.net”.

For more information about creating groups based on roles, see the section below regarding group creation.

16.3.3.4 Logging Into CS with Domain Accounts (Unisphere)

To log into Unisphere with a domain account, use only the username along with the domain password in the password field. **Use of the domain\username will fail.**

You will notice that local accounts are created as they log into the CS. These accounts are not viable to log into the system with. They cannot be used with or without AD. It is harmless to delete these accounts for cleanup purposes, but they will be recreated when that domain user logs back in.

16.3.3.5 Logging Into CS with Domain Accounts (SSH CLI)

In order to access the CLI (SSH to control station) using domain accounts, a domain user must first log into Unisphere as described above. A local user account is created on the VNX for that domain user.

To login into the CS via SSH, use an approved SSH client and connect to the CS IP address using the domain\username along with the domain password for that user. **Use of only the username will fail.**

16.3.4 Default Password Expiration

To avoid having the local user password expire (because it is a local Unix account mapped to a domain account), you will need to make the following change to the /etc/login.defs file:

```
Default value:
PASS_MAX_DAYS    120
New Value:
PASS_MAX_DAYS    99999
```

The password expiration on the domain account will still be enforced. This change will prevent the Linux OS from incorrectly thinking the local user password has expired. Password expirations for domain accounts are managed by active directory. The local user is only a placeholder and not a real user (its password is not used to authenticate the user).

If you already have accounts that have been created (already logged into GUI), you can update their password expiration by using the chage command:

```
chage -l user.domain (to list current value)
chage -M 99999 user.domain
```

Alternatively, the local account could be deleted (from Unisphere) and re-added by logging into the GUI again.

16.3.5 Idle Session Timeout

The idle session timeout should be set to 60 minutes (or less) to meet minimum security standards. To set this value on Unisphere (as root) go to Security → Unisphere

For SSH users, the Linux OS must be modified.

As root, modify the /etc/profile file:

```
Default value:
TMOUT=300
New Value:
TMOUT=3600
```

This will change future sessions to timeout after 60 minutes (3600 seconds).

16.4 Role-Based Access Control

An administrative user account is always associated with a primary group and each group is assigned a role. A role defines the privileges (that is, the operations) the administrative user can perform on a particular VNX object.

16.4.1 Groups and Roles

An administrative user account is always associated with a primary group. It can also be associated with other groups, to a maximum of 17. Like user accounts, groups can be either a local group or a local group mapped to a domain group.

VNX offers both predefined and custom groups. Predefined groups cannot be modified or deleted. Each group is assigned a role. A group can be assigned only one role at a time, although a role can be associated with many groups. If a group is not assigned a role, the group is given the default role of operator, which has read privileges only.

The default system groups are:

Group Name	Associated Role
administrator	Administrator

dataprotection	Data Protection
datarecovery	Data Recovery
filemover	FileMover Application
fullnas	NAS Administrator
localdataprotection	Local Data Protection
nasadmin	Operator
network	Network Administrator
opadmin	Operator
root	Administrator
sanadmin	SAN Administrator
security	Security Administrator
storage	Storage Administrator

A role defines the privileges (read, modify, or full control) you can perform on a particular Unisphere object. VNX offers both predefined and custom roles. An administrative user account is always associated with a primary group and each group is assigned a role. A role defines the privileges (read, modify, or full control) the administrative user can perform on a particular VNX object. By default, all administrative users have read privileges, allowing them to view objects. Modify privileges allow an administrative user to make changes to an existing object. Full control privileges allow an administrative user to create and delete objects and make significant changes.

The default roles are:

Role	Description	Objects with Read Privileges	Objects with Modify Privileges	Objects with Full Control Privileges*
Administrator	System administrator role	All	None	All
Data Protection	Similar to Operator	All	None	Tasks
Data Recovery	Checkpoints and Replication	All	None	Tasks, Checkpoints, Replication
FileMover Application	FileMover application	All	None	Tasks, Filemover
Local Data Protection	Checkpoints	All	None	Tasks, Checkpoints
NAS Adminsitrator	Many NAS functions	All	Licenses, Network Services, Log Collection and Transfer	Tasks, Datamovers, Alerts, Events, Public Key Certificates, Connect Home
Operator	Read-only role	All	None	Tasks
SAN Administrator	Similar to Operator	All	None	Tasks
Security Administrator	Security Functions	All	Licenses	Tasks, Administrative Access, Public Key Certificates
Storage Administrator	Similar to Administrator	All	Licenses, Network Services	Tasks, Data Movers, Alerts, Events, All Storage, All Protocols, All Data Protection, Public Key Certificates

A user can have multiple roles by belonging to multiple groups. In this case, a union of these roles determines what operations the administrative user can perform.

16.5 VNX Audit Log Files

There are several audit files that are used to trace user and system based events. These audit files include:

/nas/log/cmd_log

This file displays a list of most successfully-executed VNX-specific commands. Commands that only display information or commands that reside in the /nas/sbin directory are not logged. Log entries look like:

```
2007-02-21 16:35:55.532 server_2:0:21096:S: /nas/bin/server_mount server_2 -p root_fs_2 /
2007-02-21 16:35:56.526 server_2:0:21096:E: /nas/bin/server_mount server_2 -p root_fs_2 /
```

The fields contained in these log entries and their values in the example above are:

Data & Time: 2007-02-21 16:35:55.532

Source: server_2

User ID: 0

PID: 21096

Start or End Marker: S or E

The Actual Command: /nas/bin/server_mount server_2 -p root_fs_2 /

/nas/log/sys_log

This log file displays a cumulative list of system event and log messages from the most recent Control Station reboot.

/nas/log/osmlog (aka /var/log/messages)

This is the standard Linux “syslog.” Many of the non-VNX specific commands will record security-relevant messages to this file.

/nas/http/logs/access_log

This file records activity via Unisphere and includes the authenticated user identity of the issuer of HTTP requests. The fields most likely to be of interest are:

Originating IP address

User Identity (as a name)

Date and Timestamp

URL accessed

HTTP Return Status (for example, 200 to indicate success)

16.6 Terminal Server (Uplogix)

Citi uses its own terminal server (Uplogix) as a backup means of communicating with the VNX in case of network loss/failure. The Uplogix terminal server will need to be connected to the service port of the VNX. If the CE needs to use this port for access, then the terminal server will need to be temporarily disconnected. The terminal server will authenticate users to specific ports. Once connected, a user will then need to authenticate to the control station as with any normal SSH session. This connection will not be defined in the firewall for the control station.

Uplogix uses a RJ-45 connection and will need to be converted to the DB9 connection that the VNX uses. This is a standard pin conversion (no special pinout).

Port settings:

- Baud Rate: 19200
- Data Bits: 8
- Parity Bit: N (none)

- Stop Bit: 1
- Serial Flow Control: N (none)
- Null modem: auto

Uplink commands:

- port 1/3 (select port "1/3")
- term (go to interactive terminal mode)
- config serial (adjust baud rate and other settings)
- ~ <enter> (logout)
- ~l (lock the port)

16.7 Standard Build Security Checklist

A summary of the required specific security parameters and settings are listed in the [Appendix Standard Build Security Checklist](#). To automate, the checking of these and related parameters, see the Collect Logs scripting in [Appendix Additional Scripting - Check Configuration](#).

16.8 Required Scripting

In order to compensate for the audit log product deficiencies, there is some scripting that will need to be implemented to copy the logs to a remote NFS share. The scripting can be found in [Appendix Additional Scripting - Collect Logs](#).

16.9 Disabling the FTP service on the Data Movers

The following commands are used to control the FTP service. Changing the entries in netd (/nas/server/slot_x/netd) will have no effect on the ftp service. In order to meet minimum CISS security standards, every datamover must have its ftp service disabled.

To check the status of the ftp service, run the following:
server_ftp ALL -service -status

To stop the ftp service, run the following:
server_ftp ALL -service -stop

To start the ftp service, run the following (not recommended):
server_ftp ALL -service -start

* Note: This will only disable the Data Mover FTP service. Currently sFTP outbound from the control station is not restricted for troubleshooting and file transfers.

16.10 Apache Web Server Security

There are two settings that are needed in order for Unisphere to meet minimum CISS security standards.

```
SSLProtocol all -SSLv2
SSLCipherSuite HIGH:MEDIUM
```

These configuration changes are for SSL version and Cryptographic length. Both of these settings are managed by the same Apache Web Server configuration file httpd.conf (/nas/http/conf/httpd.conf). As root, add the following lines (search for Cipher):

```
# SSL Cipher Suite:
# List the ciphers that the client is permitted to negotiate.
# See the mod_ssl documentation for a complete list.
```

```
#SSLCipherSuite ALL:!ADH:!EXPORT56:RC4+RSA:+HIGH:+MEDIUM:+LOW:+SSLv2:+EXP:+eNULL
#
# * the following line added to refuse SSLv2 (will force to SSLv3)
SSLProtocol all -SSLv2
#
# * the following line added to only allow >= 128bit
SSLCipherSuite HIGH:MEDIUM
#
```

16.11 Altering the Local User Path (ECC SSH account)

In some instances, it may be necessary to alter the path parameter on a CLI user (SSH access0. This is when it is difficult or impossible to add a full path when referring to commands (for example “nas_version” instead of “/nas/bin/nas_version”).

To modify the path, the .bash_profile file may need to change ownership to the local ecc account. Then the file can be modified and saved.

As root, you may need to perform the following if root is the owner of the profile:

```
[root@c5-ns480-cs0 eccuser]# chown eccuser /home/eccuser/.bash_profile
```

As the local account, you can then modify the file /home/<useraccount>/.bash_profile where useraccount is the name of the user account on the VNX control station that you are modifying. Add or modify the path variable definition the export of the variable PATH.

```
PATH=$PATH:$HOME/bin:/nas/bin
export PATH
alias useradd='echo Please use Unisphere for User/Group management.'
alias userdel='echo Please use Unisphere for User/Group management.'
alias usermod='echo Please use Unisphere for User/Group management.'
alias groupadd='echo Please use Unisphere for User/Group management.'
alias groupdel='echo Please use Unisphere for User/Group management.'
alias groupmod='echo Please use Unisphere for User/Group management.'
```

16.12 Security Standards

Citi Information Security Standards (“CISS”) have been used as a guideline to produce this document. See appendix A for a summary of the CISS standards. Please refer to the Information Security Services (“ISS”) website at <http://www.citigroup.net/informationsecurityservices/> for more details.

16.13 Information Classification Certification

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 Appendix A.1 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.

17 Unisphere Navigation

17.1 Unisphere Navigation

Unisphere is the GUI interface for the VNX platform. It can be accessed from the control station or from the SPs of the array. If the VNX is part of a domain, it can be accessed by the IP address of any node in the domain (SPs, Control Stations). In comparison to previous releases of Unisphere (i.e. DART 6), the core organization is similar, but the menus have been rearranged. The following is a screenshot of the interface with references to the names of each element:



The following is list of the details under each subject:

After first logging into Unisphere, you are brought to a screen that is at the domain level. The following subjects apply:

17.1.1 HOME

This is a link back to dashboard

17.1.2 Systems List

This is link to any System

17.1.3 Dashboard

Initial view when logging into Unisphere.

17.1.4 SYSTEM LIST

- Shows all systems manageable through Unisphere Domains
- Select Master Domain
- Configure NTP
- Change Unisphere Login Password
- Alerts Displays a table of all systems alerts

After a system, has been selected, the following subjects apply:

17.1.5 SYSTEM

- General System information (File/Block)
- Graphical View of Hardware
- Manage SP Cache, Fast Cache
- Manage SPA, SPB
- View all components and their status
- View and configure IP networking
- Create Event Notifications
- Monitor Performance
- Monitor overall health of Array
- View Task Status
- Reports Wizard, generate system views

17.1.6 STORAGE

- View, Modify & Create File systems
- Manage CIFS connectivity
- Manage NFS connectivity
- Manage ISCSI connectivity
- Manage Storage Groups
- View, Create Virtual Tape Library for NDMP
- View, Manage & Create Storage Pools
- View NAS Volumes
- View & Manage Raid Groups, & LUNs
- Array Storage Summary
- Manage FAST
- Create & Manage File System Migrations
- Create & Manage SAN Copy sessions

17.1.7 HOSTS

- Hosts Summary View
- Manage Host connectivity
- View Virtual Host Infrastructure
- Data Protection
- Manage & Create Local Replicas
- LUN Snapshots
- Clones
- File System checkpoints
- Manage Remote Replicas
- Mirrors
- Consistency Groups
- File System Replication
- View storage systems Reserved LUN Pool

17.1.8 SETTINGS

- Create Users to Manage Systems
- Manage LDAP Domain
- Manage Data mover Parameters
- Enable Disable UI and CLI

17.1.9 SUPPORT

- Access System specific documents & procedures
- Access Online Interactive Resources
- Download Software for system updates
- Access Best Practice, White Papers & live chat
- Manage service requests

Many of the screens under Unisphere will seem familiar to the DART 6 and DART 5 users. Whereas the overall menu structure has changed, the individual screens are in many cases the same screen.

17.2 Customizing the Dashboard

The initial screen that comes up for Unisphere can be customized.

The screenshot displays the EMC Unisphere web interface. The top navigation bar includes a search box, an 'Advanced Search' button, and a 'Customize' button (circled in red). The main content area is divided into three sections: 'System Alerts (10 of 113)', 'System Information', and 'Storage Capacity Summary'.

Severity	Created	Message	Event Code
Critical	Feb 15, 201...	Enclosure 1 ... 0x1260830...	
Critical	Feb 15, 201...	Enclosure 1 ... 0x1260830...	
Error	Apr 5, 2013...	A Celerra d... 0x1360890...	
Error	Apr 5, 2013...	A Celerra d... 0x1360890...	

System Information:


- Status: ✖ Critical (2)
- Name: c5-vnx5700-cs0
- Model: VNX5700 (Unified)
- SP A IP Address: 10.40.85.25
- SP B IP Address: 10.40.85.26
- Control Station IP Address: 10.40.85.23

Storage Capacity Summary:

- Used (33019.99 GB)
- Free Raw Disk (10258.59 GB)
- Free Storage Pool (6462.67 GB)
- Free Space for File (11505.93 GB)

Alerts: 113 ✖ 2 Critical Certificates: 3 ⚠ 1 Unaccepted User: gj06529

To change the contents of the dashboard, select the customize link on the upper right of the screen.



Drag the option to an open area or remove any undesirable option by selecting the X. Select *Done* when all changes have been made.

The current dashboard options available are:

- System Information
- System Alerts
- Storage Capacity Information
- Storage Pools for File
- Storage Pool Capacity
- RAID Group Capacity
- Storage Status
- Storage Resources
- Tiering Summary
- Host Allocation
- Operating Systems
- Hosts
- Connectivity Issues
- Virtualization
- Replication Summary
- Data Protection Resources
- Reserved LUN Pool Usage
- Mirror Connections

18 Unisphere Storage Domains

18.1.1 DART and Unisphere

The VNX Management GUI interface is Unisphere for both File and Block. This updated interface is a step in the unification of the VNX and CLARiiON interfaces. In addition, the VNX has adopted a security concept known as the storage domain from the CLARiiON environment. These domain concepts are independent of the Active Directory or LDAP Domains discussed elsewhere in this document.

18.1.2 Legacy and Multi Domains

Legacy domains are CLARiiON systems (FLARE 23 and above) and Celerra systems (DART 6.0). VNX Domains are comprised of VNX systems. Legacy domains support different user roles and subsequently have a different domain database. This limits VNX systems from joining legacy domains and legacy systems from joining VNX domains. However, there is a multi domain which is comprised of both legacy and VNX domains. **This may not contain all security features of the VNX domain. Therefore the recommendation is to use separate storage domains for VNX and the previous CX series.**

18.1.3 Domain Attributes

18.1.3.1 Nodes and Master Nodes

The SPs and Control Stations in a domain are all nodes in the domain. The Unisphere GUI can be accessed from all nodes. When you configure a domain, you create a master node. The master node stores the master copy of the domain database and distributes updates to the nodes of the domain. The nodes are able to operate without a working master, but the domain database cannot be updated until a new master is assigned.

18.1.3.2 Global Users

Global users are users with domain-wide access. Global users are administered on a domain basis. There must be at least one global administrator in each domain.

18.1.3.3 Gateways

Gateways are multi-domain connections from the local domain to other domains. The connections are unidirectional.

18.1.4 Storage Domain Design

The following domains are not nested and have no hierarchical relationship (at this time). They are only logical groupings of VNXs. A VNX can be the member of more than one domain at a time. There is a recommended limit of approximately 200 members for a given domain. Each VNX (DART 6.0 and later) should be a member of the following domains. VNX arrays should be the master nodes in these domains (when a VNX is available):

Recommended Domains

1. Regional Storage Domain
This domain is the primary management domain and allows the administrative user to access all VNXs in the same region.

Optional Domains

2. Global Storage Domain
The primary purpose of this domain is reporting on an entire region.
3. Business Storage Domain (optional)
The primary purpose of this domain is reporting purposes. It is not required and should be used only when needed for reporting purposes. For example a very large business such as CARDS may have a need for a business domain.

18.1.5 Reporting

Reporting features of the VNX (under development) will be based on the storage domains. Therefore, use of these domains as described above will help to organize the VNX reporting.

18.1.6 Storage Domain Masters

Storage Domains should be controlled by a single master and single back up master VNX.

18.1.7 Local vs. Global

When logging into a VNX, you can select Local or Global. Global will log you into the storage domain and potentially all of the VNXs in that domain. Local limits the login to that specific VNX.

19 Unified Storage Considerations

Unified storage is a term to describe a VNX array that is configured for both SAN and NAS. This configuration will typically be deployed in smaller regions where there is not enough capacity requirements to justify the purchase of separate arrays. All of the current best practices (including security) apply to this configuration.

The following VNX documents are the primary source of configuration and best practice information for the NAS and SAN teams:

- [VNX for SAN Solutions Guide](#)
- [VNX for NAS Solutions Guide](#) (this document)

19.1 Unified Storage Support Ownership

There should be a defined operational owner of all of the VNX Unified arrays. This will be the primary point of contact for EMC will reach out to for specific issues and customer alerting. This ownership needs to be agreed to by the NAS and SAN OPS teams. This ownership could change in the future, but needs to be consistent for all Unified systems.

The Unified VNX Array will have the same look and feel as a standard VNX NAS system. The only physical difference may be additional fiber ports added to the SPs. Since the NAS side ultimately controls access to the system via the IPTables ACL configuration on the control stations, the NAS team will need to configure the ACL upfront, regardless of support ownership. The SAN team would not need to configure their ACL as it is secondary to the IPTables ACL. Support issues that are team specific should be routed by the owning OPS team.

19.2 Unified Storage Notifications

Internally configured notifications and alerts (i.e. SNMP traps, email) should still go to the separate teams. From a support perspective, the one OPS team will be the primary owner, but will transition any specific issues to the partner OPS team.

19.3 Unified Storage Audit Logging

Any audit logging reconciliation will continue as normal as the Unified VNX will be configured to send its NAS and SAN audit logs separately. The SEMS environment will parse these logs separately and will notify the appropriate OPS teams for any exceptions.

19.4 Unified Storage Activities and Roles

The following is a chart showing the roles for Unified, File, and Block Configurations.

Product	Activity	Role in RBAC	Support Team	Notes
ALL	Create accounts and grant access level	Admin or Security Admin	ISA	ISA role is not currently in production
ALL	User Account Entitlement Review	Admin Security Admin	ISA	ISA role is not currently in production
ALL	Basic Config, Hardware Assembly, Hardware Replacement, Upgrades, Firmware, Support	N/A	EMC	
ALL	Storage Domain configuration	Admin	NAS or SAN	Domain will include LDAP for Block/File in next release
Block, Unified	VNX OE for Block (FLARE) Management (config, network settings, ...)	Admin or SAN Admin	SAN	
Block, Unified	Create/manage RAID groups and Storage Pools in FLARE	Admin or SAN Admin	SAN	

Block, Unified	SAN hardware monitoring (FLARE generated SNMP alerts)	Admin or SAN Admin	SAN	
Block, Unified	SAN (SP, FC, ...) performance monitoring/troubleshooting	Admin or SAN Admin	SAN	
Block, Unified	SAN capacity utilization monitoring (at RAID Group or Storage Pool level in FLARE)	Admin or SAN Admin	SAN	
Block, Unified	Create/manage LUN's	Admin or SAN Admin	SAN	
Block, Unified	LUN snapshot creation/maintenance	Admin or SAN Admin	SAN	
Block, Unified	SAN zoning	N/A	SAN	
Block	Configure ACL for SP access	Admin	SAN	
Block, Unified	Audit log review (generated from FLARE) and response	N/A	SEMS and SAN	
File, Unified	VNX OE for File (DART) Management (config, network settings, ...)	Admin or NAS Admin	NAS	
File, Unified	Configure ACL for array access	Admin or NAS Admin	NAS	
File, Unified	Create/manage Storage Pools in DART	Admin or NAS Admin	NAS	Block storage pools instead of LUNs may be used in next release
File, Unified	Create/manage virtual data movers	Admin or NAS Admin	NAS	
File, Unified	NAS hardware monitoring (Data Mover, Control Station, DART generated SNMP alerts)	Admin or NAS Admin	NAS	DART alerting includes block system issues
File, Unified	NAS (Data Mover, Network, ...) performance monitoring and troubleshooting	Admin or NAS Admin	NAS	
File, Unified	NAS space utilization monitoring (at DART Storage Pool level)	Admin or NAS Admin	NAS	
File, Unified	NAS Capacity Planning (completed with CIFS and NFS admins)	Admin or NAS Admin	NAS	
File, Unified	Filesystem creation/replication/snapshot creation/maintenance	Admin or NAS Admin	NAS	
File, Unified	Filesystem replication/snapshot monitoring	Admin or NAS Admin	NAS	
File, Unified	NDMP backup configuration and monitoring	Admin or NAS Admin	NAS	
File, Unified	CAVA config/maintenance	Admin or NAS Admin	NAS and Wintel SA	
File, Unified	Audit log review (generated from DART)	N/A	SEMS and NAS	
File, Unified	CIFS directory/file structure and quota management	N/A	Wintel SA	
File, Unified	CIFS capacity utilization monitoring (on CIFS File System level)	N/A	Wintel SA	
File, Unified	CIFS file restores	N/A	Wintel SA	
File, Unified	CIFS Data Management (migrations, restructures, quotas, qtree layout, security layout, ...)	N/A	Wintel SA	
File, Unified	Create/manage CIFS file shares	N/A	Wintel SA	
File, Unified	NFS directory/file structure and quota management	N/A	UNIX SA	
File, Unified	NFS capacity utilization monitoring (on NFS File System level)	N/A	UNIX SA	
File, Unified	NFS file restores	N/A	UNIX SA	
File, Unified	NFS Data Management (migrations, restructures, quotas, security layout, ...)	N/A	UNIX SA	
File, Unified	Create/manage NFS file shares	N/A	UNIX SA	
Unified	Redirect support issues to partner team	N/A	NAS or SAN	
Unified	Coordinate upgrade schedules	N/A	NAS or SAN	

19.5 Unified Storage NAS / SAN Team Coordination

The control station will be the physical network interface to access the SPs and will also have an ACL (in the form of IPTables) that supersede the ACL defined on the SPs. As such, the NAS and SAN teams will have to work cooperatively on the following:

- The SAN and NAS teams should identify and establish the correct notification and escalation procedures for each team. This may be specific individuals, but is more likely a distribution group.
- Addition / remove of SAN Administrative Client (NAS OPS will add/remove SAN Admin Client to IPTables)
- Any Update of the OE for File (DART) will likely require a similar update for the OE for Block (FLARE) as the releases are tied together. For current recommendations on DART see [EMC NAS CATE Supported Versions](#) and on FLARE see [VNX SAN and CLARiiON Supported Config Matrix](#).
- Transfer of SAN specific issues from NAS OPS to SAN OPS (i.e. non-shared SP components, OE for Block issues, or hardware components supporting only the designated SAN storage). The NAS Datamovers are logically a SAN client of the block side and will likely be affected by only significant block hardware issues.
- Change Control of the array should be managed separately by the SAN and NAS teams. There should be two different records for the SAN array versus the NAS array. In the case of overlaps (i.e. microcode changes), the change record would need additional approvals from the partner team (i.e. NAS OPS would want SAN OPS as an additional approver for that change).

19.6 Unified Storage Security

All of the current security best practices for NAS and SAN should continue to be followed. There are separate configurations for LDAP (Active Directory) on the block and file configuration pages. The AD configuration should point to the same location for both teams. This may require the search path to start at a higher level in the AD forest. Future releases of code will only have a single LDAP/AD configuration page.

Like with other NAS systems, the local accounts for root and nasadmin should be configured as domain Administrators in the storage domain for the array. This will allow these users to access the block configurations when needed and will eliminate login errors after initial configuration.

The IPTables ACL on the control stations will administer access to the array via HTTPS, SSH, and navisecli. Therefore, only the NAS team will have the ability to update the allowed client list for both the SAN and NAS teams.

Once a formal ISA team is setting up account configurations on the system, the NAS Administrator and SAN Administrator roles can be used instead of the Administrator. This will isolate the NAS, SAN, and Security functions.

19.7 Unified Storage Drive Configuration

Storage used for NAS and SAN should be isolated as much as possible.

- RAID groups should not be shared.
- Control LUNs for NAS may be placed on SAN vault drives.
- LUNs allocated to NAS should be on dedicated RAID groups.
- Any block storage pools that SAN and NAS rely upon should not be shared.
- DAEs and spare drives can be shared, but avoided if possible.

All other best practices for NAS and SAN should be followed.

20 APPENDICIES

20.1 APPENDIX A: SNMP MIB

```
-----
-- Copyright (c) 1998-2002, All rights reserved by
--
-- EMC Corporation
-- 171 South Street
-- Hopkinton, MA 01748 USA
--
-- This software is furnished under a license and may be used and copied
-- inclusion of the above copyright notice. This software or any other
-- only in accordance with the terms of such license and with the
-- copies thereof may not be provided or otherwise made available to any
-- other person. No title to and ownership of the software is hereby
-- transferred.
--
-- The information in this software is subject to change without notice
-- and should not be construed as a commitment by EMC Corporation.
--
-- EMC Corporation assumes no responsibility for the use or reliability
-- of this software.
--
-----
```

```
--
--
-- EMC Celerra Management Information Base description
--
-----
```

```
--
-- Author Identification
--
-- gma      Gary Ma, EMC
-- WZ      Weining Zhu, EMC
-- alc      Angela Cocks, EMC
--
-----
```

```
--
-- 003 alc 07-19-02 Added Web GUI event
-- 002 WZ 07-02-02 Added Jserver Alert trap type
-- 001 WZ 04-18-01 Added Control Station started trap type
-- 000 gma 11-05-98 Created Original Module
--
-----
```

```
EMC-CELERRA DEFINITIONS ::= BEGIN
```

```
IMPORTS
    enterprises, Opaque
        FROM RFC1155-SMI
    DisplayString
        FROM RFC1213-MIB
    OBJECT-TYPE
        FROM RFC-1212
    TRAP-TYPE
        FROM RFC-1215;
--
-- emc
--     FROM EMC-MIB;

    emc OBJECT IDENTIFIER ::= { enterprises 1139 }
    emcCelerra OBJECT IDENTIFIER ::= { emc 2 }
```

```
-----
--
-- Celerra Event Group
```

```
--
-----

celEventTable OBJECT-TYPE
    SYNTAX SEQUENCE OF CelEvent
    ACCESS not-accessible
    STATUS mandatory
    DESCRIPTION
        "A table containing information
         about an particular event."
    ::= { emcCelerra 1}

celEvent OBJECT-TYPE
    SYNTAX CelEvent
    ACCESS not-accessible
    STATUS mandatory
    DESCRIPTION
        "Information about an particular event."
    INDEX {celEventFacility,
           celEventID}
    ::= { celEventTable 1}

CelEvent ::= SEQUENCE {
    celEventFacility INTEGER,
    celEventID INTEGER,
    celEventSeverity INTEGER,
    celEventDescr DisplayString}

celEventFacility OBJECT-TYPE
    SYNTAX INTEGER
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The number representing the facility which generates
         the event. For now, the value could be:
         first 128 - default Dart facility
         129 - master control
         130 - event log
         131 - box monitor
         132 - video service
         133 - callhome application
         134 - AAF backup control station."
    ::= { celEvent 1 }

celEventID OBJECT-TYPE
    SYNTAX INTEGER
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The event id which identifies the particular event
         within a facility."
    ::= { celEvent 2 }

celEventSeverity OBJECT-TYPE
    SYNTAX INTEGER
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "Severity of the event."
    ::= { celEvent 3 }

celEventDescr OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "text description of the event."
    ::= { celEvent 4 }

-----
--
```

```
--      Trap definitions
--
-----

celReboot TRAP-TYPE
ENTERPRISE emcCelerra
VARIABLES { celEvent }
DESCRIPTION
    "Trap message will be sent in the event of a Data Mover
    Reboots. The trap message includes the facility, event id,
    severity and a text description of the event."
    ::= 1

celMasterCtlFault TRAP-TYPE
ENTERPRISE emcCelerra
VARIABLES { celEvent }
DESCRIPTION
    "Trap message will be sent in the event of Master Control
    serious fault.
    unexpected daemon exit or control station heartbeat
    missing."
    ::= 2

celHWFailure TRAP-TYPE
ENTERPRISE emcCelerra
VARIABLES { celEvent }
DESCRIPTION
    "Trap message will be sent in the event of celerra hardware
    failure."
    ::= 3

celSlotStale TRAP-TYPE
ENTERPRISE emcCelerra
VARIABLES { celEvent }
DESCRIPTION
    "Trap message will be sent in the event of celerra stale
    reason code."
    ::= 4

celSlotPanicked TRAP-TYPE
ENTERPRISE emcCelerra
VARIABLES { celEvent }
DESCRIPTION
    "Trap message will be sent in the event of panicked
    slot."
    ::= 5

celIntfFailure TRAP-TYPE
ENTERPRISE emcCelerra
VARIABLES { celEvent }
DESCRIPTION
    "Trap message will be sent in the event of both interfaces
    failure."
    ::= 6

celAAF TRAP-TYPE
ENTERPRISE emcCelerra
VARIABLES { celEvent }
DESCRIPTION
    "Trap message will be sent in the event of AAF."
    ::= 7

celCSStart TRAP-TYPE
ENTERPRISE emcCelerra
VARIABLES { celEvent }
DESCRIPTION
    "Trap message will be sent in the event of Control Station started."
    ::= 8

celJServer TRAP-TYPE
ENTERPRISE emcCelerra
```

```

VARIABLES { celEvent }
DESCRIPTION
    "Trap message will be sent in the event of a JServer alert."
::= 9

celWebGUI TRAP-TYPE
ENTERPRISE emcCelerra
VARIABLES { celEvent }
DESCRIPTION
    "Trap message will be sent in the event of a Web GUI notification request."
::= 10

END

```

20.2 APPENDIX B: Operational Commands

The following command line functions are demonstrated here:

- Check Versions of Software
- Check System Status
- Gather Support Materials
- Restart Control Stations
- Restart Datamovers
- Backup VNX Database
- Detailed Storage Configuration

20.2.1 Check Versions of Software

To check the version of DART running on the control station, run the *nas_version* command.

```
[nasadmin@c5-vnx5500-cs0 ~]$ nas_version -l
Name       : emcnas                      Relocations: /nas
Version    : 7.0.14                      Vendor: EMC
Release    : 0                          Build Date: Wed 27 Apr 2011 07:56:49 AM EDT
Size       : 629880885                   License: EMC Copyright
Signature  : (none)
Packager   : EMC Corporation
URL        : http://www.emc.com
Summary    : EMC nfs base install
Description:
EMC nfs base install
```

To check the version of DART running on the datamovers, run the *server_version* command.

```
[nasadmin@c5-vnx5500-cs0 ~]$ server_version ALL
server_2 : Product: EMC Celerra File Server Version: T7.0.14.0
server_3 : Product: EMC Celerra File Server Version: T7.0.14.0
server_4 : Product: EMC Celerra File Server Version: T7.0.14.0
```

20.2.2 Check System Status

To run a self diagnostic series of scripts and to gain a detailed status of the system components, run the *nas_summary*, *getreason*, *nas_checkup*, and *server_checkup* commands. In addition to the status of components, this command also produces a section on issues (warnings and errors) and specific recommendations on how to make changes.

```
[nasadmin@c5-vnx5500-cs0 ~]$ /nas/tools/nas_summary
nas_summary v2.1
Collecting Replication info...Please Wait
#####
System type: VNX5500      (Gateway)
Version: 7.0.14-0
#####
Number of data movers: 3      VDMs: 1
server_2: VNX5500      7.0.14.0      State:5      Primary
server_3: VNX5500      7.0.14.0      State:5      Primary
server_4: VNX5500      7.0.14.0      State:5      (Standby)
#####
Number of arrays: 1
APM00111601868: Clariion VNX5500
SPA:  Flare 31  Read Cache: 200  Write Cache: 3689  OPERATIONAL
SPB:  Flare 31  Read Cache: 200  Write Cache: 3689  OPERATIONAL
#####
Filesystems:      UXFS:39 Checkpoint:39      Total:95
Number of disk vols (LUNs):      12
#####
Replication mode: RepV2      # Replications: 1      # Copies: 0
      Session types:      Filesystem: 1      VDM: 0      iSCSI: 0
      States:      OK: 1 (100%)      Inactive: 0
      Failed Over: 0      Switched Over: 0
# of Interconnects: 3      Remote Celerras: 1
#####
```

```

Features in use:      CIFS: *True*      iSCSI: False      NFS: *True*
                    Rep: *True*      SRDF: False      DHSM: False
#####
DBCHK State: No errors found

[nasadmin@c5-vnx5500-cs0 ~]$ /nas/sbin/getreason
10 - slot_0 primary control station
11 - slot_1 secondary control station
 5 - slot_2 contacted
 5 - slot_3 contacted
 5 - slot_4 contacted

[nasadmin@c5-vnx5500-cs0 ~]$ nas_checkup
Check Version: 7.0.14.0
Check Command: /nas/bin/nas_checkup
Check Log      : /nas/log/checkup-run.110914-085400.log

-----Checks-----
Control Station: Checking statistics groups database..... Pass
Control Station: Checking if file system usage is under limit..... Pass
Control Station: Checking if NAS Storage API is installed correctly..... Pass
Control Station: Checking if NAS Storage APIs match..... Pass
Control Station: Checking if NBS clients are started..... Pass
Control Station: Checking if NBS configuration exists..... Pass
Control Station: Checking if NBS devices are accessible..... Pass
Control Station: Checking if NBS service is started..... Pass
Control Station: Checking if PXE service is stopped..... Pass
Control Station: Checking if standby is up..... Pass
Control Station: Checking integrity of NASDB..... Pass
Control Station: Checking if primary is active..... Pass
Control Station: Checking all callhome files delivered..... Pass
Control Station: Checking resolv conf..... Pass
Control Station: Checking if NAS partitions are mounted..... Pass
Control Station: Checking ipmi connection..... Pass
Control Station: Checking nas site eventlog configuration..... Pass
Control Station: Checking nas sys mcd configuration..... Pass
Control Station: Checking nas sys eventlog configuration..... Pass
Control Station: Checking logical volume status..... Pass
Control Station: Checking valid nasdb backup files..... Pass
Control Station: Checking root disk reserved region..... Pass
Control Station: Checking if RDF configuration is valid..... N/A
Control Station: Checking if fstab contains duplicate entries..... Pass
Control Station: Checking if sufficient swap memory available..... Pass
Control Station: Checking for IP and subnet configuration..... Pass
Control Station: Checking auto transfer status..... Warn
Control Station: Checking for invalid entries in etc hosts..... Pass
Control Station: Checking the hard drive in the control station..... Pass
Control Station: Checking if Symapi data is present..... Pass
Control Station: Checking if Symapi is synced with Storage System..... Pass
Blades       : Checking boot files..... Pass
Blades       : Checking if primary is active..... Pass
Blades       : Checking if root filesystem is too large..... Pass
Blades       : Checking if root filesystem has enough free space..... Pass
Blades       : Checking if using standard DART image..... Pass
Blades       : Checking network connectivity..... Pass
Blades       : Checking status..... Pass
Blades       : Checking dart release compatibility..... Pass
Blades       : Checking dart version compatibility..... Pass
Blades       : Checking server name..... Pass
Blades       : Checking unique id..... Pass
Blades       : Checking CIFS file server configuration..... Warn
Blades       : Checking domain controller connectivity and configuration..... Pass
Blades       : Checking DNS connectivity and configuration..... Warn
Blades       : Checking connectivity to WINS servers..... Pass
Blades       : Checking I18N mode and unicode translation tables..... Pass
Blades       : Checking connectivity to NTP servers..... Warn
Blades       : Checking connectivity to NIS servers..... Pass
Blades       : Checking virus checker server configuration..... Warn
Blades       : Checking if workpart is OK..... Pass
Blades       : Checking if free full dump is available..... Pass
Blades       : Checking if each primary Blade has standby..... Pass

```

```

Blades      : Checking if Blade parameters use EMC default values..... Info
Blades      : Checking VDM root filesystem space usage..... Pass
Blades      : Checking if file system usage is under limit..... Pass
Blades      : Checking slic signature..... Pass
Storage System : Checking disk emulation type..... Pass
Storage System : Checking disk high availability access..... Pass
Storage System : Checking disks read cache enabled..... Pass
Storage System : Checking disks and storage processors write cache enabled. Pass
Storage System : Checking if FLARE is committed..... Pass
Storage System : Checking if FLARE is supported..... Pass
Storage System : Checking array model..... Pass
Storage System : Checking if microcode is supported..... N/A
Storage System : Checking no disks or storage processors are failed over... Warn
Storage System : Checking that no disks or storage processors are faulted.. Pass
Storage System : Checking that no hot spares are in use..... Pass
Storage System : Checking that no hot spares are rebuilding..... Pass
Storage System : Checking minimum control lun size..... Pass
Storage System : Checking maximum control lun size..... N/A
Storage System : Checking maximum lun address limit..... Pass
Storage System : Checking system lun configuration..... Pass
Storage System : Checking if storage processors are read cache enabled.... Pass
Storage System : Checking if auto assign are disabled for all luns..... N/A
Storage System : Checking if auto trespass are disabled for all luns..... N/A
Storage System : Checking storage processor connectivity..... Pass
Storage System : Checking control lun ownership..... N/A
Storage System : Checking if Fibre Channel zone checker is set up..... N/A
Storage System : Checking if Fibre Channel zoning is OK..... N/A
Storage System : Checking if proxy arp is setup..... Pass
Storage System : Checking if Product Serial Number is Correct..... Pass
Storage System : Checking SPA SPB communication..... Fail
Storage System : Checking if secure communications is enabled..... Pass
Storage System : Checking if backend has mixed disk types..... Pass
Storage System : Checking for file and block enabler..... Pass
Storage System : Checking if nas storage command generates discrepancies... Pass
Storage System : Checking if Repset and CG configuration are consistent.... Pass
-----

```

```
[nasadmin@cs5-vnx5500-cs0 ~]$ server_checkup ALL
```

```
server_2 :
```

```
-----Checks-----
```

```
Component REPV2 :
```

```
F_RDE_CHEC: Checking the F-RDE compatibility of Repv2 sessions..... Pass
```

```
Component HTTPS :
```

```
HTTP      : Checking the configuration of HTTP applications..... Pass
```

```
SSL       : Checking the configuration of SSL applications..... Pass
```

```
Component CIFS :
```

```
ACL       : Checking the number of ACLs per file system.....*Pass
```

```
Connection: Checking the load of CIFS TCP connections..... Pass
```

```
Credential: Checking the validity of credentials..... Pass
```

```
DC        : Checking the connectivity and configuration of Domain Controller*Pass
```

```
DFS       : Checking the DFS configuration files and DFS registry..... Pass
```

```
DNS       : Checking the DNS configuration and connectivity to DNS servers. Fail
```

```
EventLog  : Checking the configuration of Windows Event Logs..... Pass
```

```
FS_Type   : Checking if all file systems are in the DIR3 format..... Pass
```

```
GPO       : Checking the GPO configuration..... Pass
```

```
HomeDir   : Checking the configuration of home directory shares..... Pass
```

```
I18N      : Checking the I18N mode and the Unicode/UTF8 translation tables. Pass
```

```
Kerberos  : Checking password updates for Kerberos..... Pass
```

```
LDAP      : Checking the LDAP configuration..... Fail
```

```
LocalGrp  : Checking the database configuration of local groups..... Fail
```

```
NIS       : Checking the connectivity to the NIS servers..... Pass
```

```
NS        : Checking the naming services configuration..... Pass
```

```
NTP       : Checking the connectivity to the NTP servers..... Pass
```

```

Ntxmap    : Checking the ntxmap configuration file..... Pass
Secmap    : Checking the SECMAP database..... Pass
Security  : Checking the CIFS security settings..... Pass
Server    : Checking the CIFS file servers configuration..... Fail
Share     : Checking the network shares database..... Fail
SmbList   : Checking the range availability of SMB IDs.....*Pass
Threads   : Checking for CIFS blocked threads..... Pass
UM_Client : Checking the connectivity to usermapper servers..... Fail
UM_Server : Checking the usermapper server database.....*Pass
UnsupOS   : Checking for unsupported client network operating systems..... Pass
UnsupProto: Checking for unsupported client network protocols..... Pass
VC         : Checking the configuration of Virus Checker servers..... Fail
WINS      : Checking the connectivity to WINS servers..... Pass

Component FTPDS :

FS_Type    : Checking if all file systems are in the DIR3 format..... Pass
FTPD       : Checking the configuration of FTPD..... Fail
NIS        : Checking the connectivity to the NIS servers..... Pass
NS         : Checking the naming services configuration..... Pass
NTP        : Checking the connectivity to the NTP servers..... Pass
SSL        : Checking the configuration of SSL applications..... Pass

NB: a result with a '*' means that some tests were not executed. use -full to run them

-----
-----CIFS : DNS Warnings-----

Warning 17456037931: server_2 : The DNS domain 'namdev.nsrootdev.net' is defined with
only one DNS server. High availability is compromised.
--> Add another DNS server using the 'server_dns' command. The recommendation is to
configure DNS with DNS servers accessible from different subnets.

-----CIFS : LDAP Warnings-----

Warning 17453416470: server_2 : The current LDAP configuration for the domain
'infra.wlab.net' contains only one server. LDAP's high availability is compromised. Your
service provider recommends defining at least two servers.
--> Use the server_ldap command to add another server (openLDAP) or define an alternate
server in the client profile preferredServerList (iPlanet).

-----CIFS : LocalGrp Warnings-----

Warning 17451974718: server_2 : The CIFS server 'NASEMCCA02' is no longer valid but is
still present in the local group database. This server might have been temporarily
removed.
--> This is not a critical issue. There is no user tool yet available. If there is a real
need, contact your Customer Service to clean up the local group database and remove the
obsolete servers.

Warning 17451974725: server_2 : emcnasvdm1: The local administrators password for the
CIFS server NASEMCTEST01 must be changed. This situation occurs if a StandAlone server is
created or the local users have been activated on a regular server, but the administrators
password set by the control station, is not yet changed. The connections to this server
using this account are currently being refused, until the password gets changed.
--> Using MMC, change the password of the local administrator. The easiest way is to use
Ctrl-Alt-Sup on a regular Windows machine, and click on the 'change password' button.

-----CIFS : Server Warnings-----

Warning 17456037929: server_2 : The DNS server is not up-to-date. It contains an invalid
IP address for CIFS server 'nasemctest02.namdev.nsrootdev.net'. '10.40.42.203' is
currently not valid. The CIFS clients may not be able to access the Data Mover. By
default, the Data Mover registers its own IP address in DNS. This certainly means static
address have been added manually.
--> Verify if automatic update is allowed (check the dns.updateMode parameter). Remove
manually the incorrect IP address from the DNS server.

```

```

Warning 17456037929: server_2 : The DNS server is not up-to-date. It contains an invalid
IP address for CIFS server 'nasebccava04.namdev.nsrootdev.net'. '10.40.42.217' is
currently not valid. The CIFS clients may not be able to access the Data Mover. By
default, the Data Mover registers its own IP address in DNS. This certainly means static
address have been added manually.
--> Verify if automatic update is allowed (check the dns.updateMode parameter). Remove
manually the incorrect IP address from the DNS server.

-----CIFS : Share Warnings-----

Warning 17451974859: server_2 : The share database still contains shares for
'NASEMCCA02' which seems to be an obsolete server. This probably corresponds to shares
associated to servers that have not been unexported after corresponding servers have been
removed.
--> If the server has been definitively removed, the shares should be unexported using
server_export server_x -unexport -name shareName -o netbios=serverName. If some server
information remains in the share database afterward, contact your Customer Service to
completely clean up it.

-----CIFS : UM_Client Warnings-----

Warning 17454399560: server_2 : Both Usermapper and LDAP are currently enabled. This may
lead to a risk of access issues if the LDAP server cannot be accessed intermittently.
--> It is not recommended to run the LDAP and Usermapper services simultaneously. Use the
server_usermapper command to manage the Usermapper service. Use the server_ldap command to
manage LDAP.

-----CIFS : VC Warnings-----

Warning 17453613256: server_2 : The configuration of the virus checker seems correct, but
this service is currently stopped. It is recommended to start it, in order to block and
remove viruses from the data disk partitions.
--> Start the virus checker. To start the virus checker use the CLI command "server_setup
'server_X' -Protocol viruschk -option start".

-----

-----FTPDS : FTPD Errors-----

Error 13156024423: server_2 : An error occurred when trying to load the FTPD server
configuration.
--> Check the Data Mover root file system state before retrying to start the FTPD server
or to modify the FTPD server configuration. If the FTPD configuration file exists with bad
parameters, you can force the default configuration using "server_ftp <servername> -
modify -reset"

Error 13156876314: server_2 : The persona 'default' contains nor certificate neither
private keys sets. So, this persona can not be used by a SSL feature on the Data Mover.
--> Run the server_certificate command to generate a new key set and certificate for this
persona.

-----

Total : 2 errors, 9 warnings

server_3 :

-----Checks-----

Component REPV2 :

F_RDE_CHEC: Checking the F-RDE compatibility of Repv2 sessions..... Pass

Component HTTPS :

HTTP      : Checking the configuration of HTTP applications..... Pass
SSL       : Checking the configuration of SSL applications..... Pass

Component CIFS :

ACL       : Checking the number of ACLs per file system.....*Pass

```

```

Connection: Checking the load of CIFS TCP connections..... Pass
Credential: Checking the validity of credentials..... Fail
DC       : Checking the connectivity and configuration of Domain Controller Fail
DFS      : Checking the DFS configuration files and DFS registry..... Fail
DNS      : Checking the DNS configuration and connectivity to DNS servers. Fail
EventLog : Checking the configuration of Windows Event Logs..... Fail
FS_Type  : Checking if all file systems are in the DIR3 format..... Pass
GPO      : Checking the GPO configuration..... Pass
HomeDir  : Checking the configuration of home directory shares..... Fail
I18N     : Checking the I18N mode and the Unicode/UTF8 translation tables. Fail
Kerberos : Checking password updates for Kerberos..... Fail
LDAP     : Checking the LDAP configuration..... Pass
LocalGrp : Checking the database configuration of local groups..... Fail
NIS      : Checking the connectivity to the NIS servers..... Pass
NS       : Checking the naming services configuration..... Fail
NTP      : Checking the connectivity to the NTP servers..... Fail
Ntxmap   : Checking the ntxmap configuration file..... Pass
Secmap   : Checking the SECMAP database..... Fail
Security : Checking the CIFS security settings..... Pass
Server   : Checking the CIFS file servers configuration..... Fail
Share    : Checking the network shares database..... Fail
SmbList  : Checking the range availability of SMB IDs.....*Pass
Threads  : Checking for CIFS blocked threads..... Pass
UM_Client: Checking the connectivity to usermapper servers..... Pass
UM_Server: Checking the usermapper server database..... Pass
UnsupOS  : Checking for unsupported client network operating systems..... Pass
UnsupProto: Checking for unsupported client network protocols..... Pass
VC       : Checking the configuration of Virus Checker servers..... Pass
WINS     : Checking the connectivity to WINS servers..... Fail

Component FTPDS :

FS_Type  : Checking if all file systems are in the DIR3 format..... Pass
FTPD     : Checking the configuration of FTPD..... Fail
NIS      : Checking the connectivity to the NIS servers..... Pass
NS       : Checking the naming services configuration..... Fail
NTP      : Checking the connectivity to the NTP servers..... Fail
SSL      : Checking the configuration of SSL applications..... Pass

NB: a result with a '*' means that some tests were not executed. use -full to run them

-----
server_4 :
Error 4010: server_4 : tests can't be run, server_4 is a standby server

```

20.2.3 Collect Support Materials and SP Collect

To collect the file based support materials (typically under the direction of EMC support), run the *collect_support_materials* command. Depending on the user shell, you may need to add the bash shell command to the *collect_support_materials* command (i.e. *bash bash /nas/tools/collect_support_materials*). The output folder is */nas/var/emcsupport/*.

```

[nasadmin@cs5-vnx5500-cs0 ~]$ /nas/tools/collect_support_materials
collect_support_materials[17131]: The collection script revision 4.3.14 has started.
PPID is 17131
Monitor_collect_script called
  Discovering Clariion SPs
  Collecting system times
  Collecting output from server_log
  Collecting output from internal commands
  Collecting event log configuration files
  Collecting files from .etc dir of each DM
  Collecting /nas/http/domain
  Collecting Mirrorview DR logs
  Collecting Backend Monitor logs
  Collecting /var logs
  Collecting upgrade logs
  Collecting files from /etc
  Collecting files from /root
  Collecting files from /proc
  Collecting /http/logs and /tomcat/logs

```

```

Collecting Celerra Manager tasks
Collecting cron files
Collecting Control Station process information and versions
Collecting /nas/jsrver/debug_of_core* files
Collecting /nas/bdb files
Collecting C3 log files
Collecting output from other CS commands
Collecting Backend audit logs
Collecting files from /usr
Collecting files from /tmp
Collecting Boot logs
Collecting ConnectEMC xml files from /opt/connectemc
Now running material collection script for longer running commands.
Collecting complete nas dir listing
Collecting output from nas commands
Collecting Replicator information
Collecting RDF information
Collecting F-RDE information
Collecting DHSM information
Collecting Symmetrix information
Collecting Clariion information
Collecting other files from /nas, /nas/site, /nas/sys,
    /nas/rdf, and /nas/dos
Collecting complete replication v2 namedb listing

Collecting /nas/log/*, /nas/log/webui/*, /nas/ConnectHome/*,
    /nas/jsrver/logs, /nas/log/connectemc/* and /nas_standby/log/*
Material Collection File:
/nas/var/emcsupport/support_materials_APM00111601868.110914_1004.zip has been generated.

*****
Please include file /nas/var/emcsupport/support_materials_APM00111601868.110914_1004.zip
with materials submitted to EMC for problem investigation.
*****

collect_support_materials[17131]: The collection script has finished successfully.

```

To collect the block-side support materials, run the *get_spcollect* command. The output folder is */nas/var/log*. The output (SPCOLLECT.ZIP) will be overwritten every time this command is run.

```

[nasadmin@c5-vnx5500-cs0 ~]$ /nas/tools/.get_spcollect

Generating spcollect zip file for Clariion(s)

Creating spcollect zip file for the Service Processor A_APM00111601868. Please wait...

spcollect started to pull out log files(it will take several minutes)...
Wait until new _data.zip file size becomes final(it will take several minutes)
Retrieving new _data.zip file...
Wait until new _data.zip file size becomes final(it will take several minutes)
Retrieving new _data.zip file...
Wait until new _data.zip file size becomes final(it will take several minutes)
Retrieving new _data.zip file...

spcollect zip file APM00111601868_SPA_2011-09-14_14-10-18_2a8004_data.zip for the Service
Processor A_APM00111601868 was created
Creating spcollect zip file for the Service Processor B_APM00111601868. Please wait...

spcollect started to pull out log files(it will take several minutes)...
Wait until new _data.zip file size becomes final(it will take several minutes)
Retrieving new _data.zip file...
Wait until new _data.zip file size becomes final(it will take several minutes)
Retrieving new _data.zip file...
Wait until new _data.zip file size becomes final(it will take several minutes)
Retrieving new _data.zip file...

spcollect zip file APM00111601868_SPB_2011-09-14_14-15-46_2ad030_data.zip for the Service
Processor B_APM00111601868 was created

Deleting old SPCOLLECT.zip file from /nas/var/log directory...
Old SPCOLLECT.zip deleted

```

```

Zipping all spcollect zip files in one SPCOLLECT.zip file and putting it in the
/nas/var/log directory...
  adding: APM00111601868_SPA_2011-09-14_14-10-18_2a8004_data.zip (stored 0%)
  adding: APM00111601868_SPB_2011-09-14_14-15-46_2ad030_data.zip (stored 0%)

```

20.2.4 Backing Up the VNX Database

The VNX automatically backs up the entire database and saves it to a file named /celerra/backup/_nasbkup.01.tar.gz. The twelve most recent copies are located here as well. To manually backup the database, use the `nasdb backup` command.

```

[nasadmin@c5-vnx5500-cs0 /]$ /nas/sbin/nasdb_backup /nas /celerra/backup 110914
NAS_DB backup in progress .....
DBMS backup: done
NAS_DB checkpoint in progress .....NAS_DB checkpoint done
dbchk: done
NAS DB Backup done

[nasadmin@c5-vnx5500-cs0 /]$ ls -al /celerra/backup
total 50216
drwxrwxr-x  3 nasadmin nasadmin   4096 Sep 14 10:33 .
drwxr-xr-x 11 root      root      4096 Jun  6 23:03 ..
-rw-r--r--  1 nasadmin nasadmin  11584 Sep 14 10:33 _dbms_backup.01.tar.gz
-rw-r--r--  1 nasadmin nasadmin  11575 Sep 14 10:01 _dbms_backup.02.tar.gz
-rw-r--r--  1 nasadmin nasadmin  11584 Sep 14 10:33 _dbms_backup.0K.tar.gz
drwx-----  2 root      root      16384 Jun  6 19:02 lost+found
-rw-r--r--  1 nasadmin nasadmin 2550888 Sep 14 01:01 _nasbkup.01.tar.gz
-rw-r--r--  1 nasadmin nasadmin 2551166 Sep 14 02:01 _nasbkup.02.tar.gz
-rw-r--r--  1 nasadmin nasadmin 2551220 Sep 14 03:01 _nasbkup.03.tar.gz
-rw-r--r--  1 nasadmin nasadmin 2551452 Sep 14 04:01 _nasbkup.04.tar.gz
-rw-r--r--  1 nasadmin nasadmin 2551633 Sep 14 05:01 _nasbkup.05.tar.gz
-rw-r--r--  1 nasadmin nasadmin 2551707 Sep 14 06:01 _nasbkup.06.tar.gz
-rw-r--r--  1 nasadmin nasadmin 2551897 Sep 14 07:01 _nasbkup.07.tar.gz
-rw-r--r--  1 nasadmin nasadmin 2552075 Sep 14 08:01 _nasbkup.08.tar.gz
-rw-r--r--  1 nasadmin nasadmin 2552192 Sep 14 09:01 _nasbkup.09.tar.gz
-rw-r--r--  1 nasadmin nasadmin 2552479 Sep 14 10:01 _nasbkup.10.tar.gz
-rw-r--r--  1 nasadmin nasadmin 2552787 Sep 14 10:33 _nasbkup.11.tar.gz
-rw-r--r--  1 nasadmin nasadmin 2550671 Sep 14 00:01 _nasbkup.12.tar.gz
-rw-r--r--  1 nasadmin nasadmin 8074773 Sep 14 10:33 nasdb_backup.110914.tar.gz
-rw-r--r--  1 nasadmin nasadmin 8047732 Sep 14 10:01 nasdb_backup.1.tar.gz
-rw-r--r--  1 nasadmin nasadmin 4488543 Aug 25 13:40 nasdb_backup.b.tar.gz
-rw-r--r--  1 nasadmin nasadmin 12470 Sep 14 10:33 stat_groups_backup.tar.gz

```

The manually created backup file `nasdb_backup.XXXXXX.tar.gz` should be ftp'd (or copied via an NFS share) to a remote system. Loss or corruption of the `nasdb` can result in a failover of all of the datamovers.

20.2.5 Restart Control Stations

There are several conditions that may require a reboot of the control station. As the VNX series all have dual control stations, it may be desirable to turn off the secondard control station before rebooting the primary control station to prevent a failover.

```

[nasadmin@c5-vnx5500-cs0 /]$ /nas/sbin/t2reset pwroff -s 1
[nasadmin@c5-vnx5500-cs0 /]$ reboot

(new session)
[nasadmin@c5-vnx5500-cs0 /]$ /nas/sbin/t2reset pwron -s 1

```

20.2.6 Restart Datamovers

To restart or reboot the datamovers, the `server_cpu` command should be used. If you do not specify warm or cold as a parameter, then the default warm reboot is performed. This is much faster than a cold reboot (`server_cpu <movername> -reboot cold -monitor time`).

```

server_cpu server_2 -reboot -monitor now
server_2 : reboot in progress 3.3.3.3.3.3.3.3.3.4.4.done

```

20.2.7 Detailed Storage Configuration

To gather detailed storage information about the system, the `nas_storage` and `setup_clariion2` commands can be used. Use the output of the first command as a parameter on the second command.

```
[nasadmin@cs5-vnx5500-cs0 ~]$ nas_storage -list
id  acl  name  serial_number
1   0    APM00111601868  APM00111601868

[nasadmin@cs5-vnx5500-cs0 ~]$ /nas/sbin/setup_backend/setup_clariion2 list config
APM00111601868
System 10.40.42.207 is up
System 10.40.42.208 is up

Detecting Data Movers...2 3 4 Done.

Clariion Array: APM00111601868  Model: VNX5500  Memory: 12288

Enclosure(s) 0_0,1_0,0_1,1_1,0_2,1_2 are installed in the system.

Enclosure info:

-----
      0  1  2  3  4  5  6  7  8  9  10  11  12  13  14
-----
1_2: 600 600 600 600 600 600 600 600 600 600 600 600 600 600 600
SAS   9   9   9   9   9   9   9   9   9   9   9   9   9   9  10  R6
-----

      15 16 17 18 19 20 21 22 23 24
-----
1_2: 600 600 600 600 600 600 600 600 600 600
SAS  10  10  10  10  10  10  10  10  HS  HS  R6
-----

      0  1  2  3  4  5  6  7  8  9  10  11  12  13  14
-----
0_2: 600 600 600 600 600 600 600 600 600 600 600 2000 100 100 100
INV  SP  SP  SP  SP  SP  SP  SP  10  10  10  10  8  SP  SP  SP  R6
-----

      0  1  2  3  4  5  6  7  8  9  10  11  12  13  14
-----
1_1: 2000 2000 2000 2000 2000 2000 2000 2000 2000 2000 2000 2000 2000 2000
NL
SAS  SP   8   8   8   8   8   8   8   8   8   8   8   8  HS  R6
-----

      0  1  2  3  4  5  6  7  8  9  10  11  12  13  14
-----
0_1: 600 600 600 600 600 600 600 100 100 100 100 100 100 100 100
INV  SP  SP  SP  SP  SP  SP  SP  SP  SP  SP  SP  SP  SP  SP  UB
-----

      0  1  2  3  4  5  6  7  8  9  10  11  12  13  14
-----
1_0: 2000 2000 2000 2000 2000 2000 2000 2000 2000 2000 2000 2000 2000 2000
NL
SAS  SP  SP  SP  SP  SP  SP  SP  SP  SP  SP  SP  SP  SP  SP  UB
-----

      0  1  2  3  4  5  6  7  8  9  10  11  12  13  14
-----
0_0: 300 300 300 300 600 600 600 100 100 100 100 100 100 100 100
INV   0   0   0   0  SP  SP  SP  HS  SP  SP  SP  SP  SP  SP  R5
-----
```

Disk group info:

Disk Group ID: 0 r5 Disks: 0_0_3,0_0_2,0_0_1,0_0_0
Disk Group ID: 8 r6 Disks:
0_2_11,1_1_1,1_1_2,1_1_3,1_1_4,1_1_5,1_1_6,1_1_7,1_1_8,1_1_9,1_1_10,1_1_11,1_1_12,1_1_13
Disk Group ID: 9 r6 Disks:
1_2_0,1_2_1,1_2_2,1_2_3,1_2_4,1_2_5,1_2_6,1_2_7,1_2_8,1_2_9,1_2_10,1_2_11,1_2_12,1_2_13
Disk Group ID: 10 r6 Disks:
0_2_6,0_2_7,0_2_8,0_2_9,0_2_10,1_2_14,1_2_15,1_2_16,1_2_17,1_2_18,1_2_19,1_2_20,1_2_21,1_2_22

Lun info:

Lun ID: 13 RG ID: 0 State: Bound ??
Lun ID: 16 RG ID: 8 State: Bound d7
Lun ID: 17 RG ID: 8 State: Bound d10
Lun ID: 18 RG ID: 9 State: Bound d8
Lun ID: 19 RG ID: 9 State: Bound d11
Lun ID: 20 RG ID: 10 State: Bound d9
Lun ID: 21 RG ID: 10 State: Bound d12

Spare info:

Spare ID: 200 Disk: 1_2_23
Spare ID: 8083 Disk: 0_0_7
Spare ID: 8120 Disk: 1_1_14
Spare ID: 8121 Disk: 1_2_24

Size	Type	Disks	Spares
100	EFD	19	1
2000	NL SAS	31	1
600	SAS	46	2
300	SAS	4	0

20.3 APPENDIX C: Standard Build Security Checklist

To meet the basic security requirements, the following items need to be configured:

Supported from	Dart Options Configuration	Required Security Setting	Description
5.6.43	/nas/sbin/nas_config - sessiontimeout	60	Enables/disables the auto logout of console connections. The default is on which causes console connections to be disconnected after the console has been idle for the number of minutes specified by the Sessiontimeout value. Any change to this option is effective after a command is entered. The number of minutes the console is idle after which console connection is disconnected.
5.6.43	/nas/sbin/nas_connecthome -modify -email_from	<first.last@citi.com>	Filer Administrator recipient email address, in the standard Citi email address form. This must be a valid entry in the Citi Global Address List (GAL).
5.6.43	/nas/sbin/nas_connecthome -modify -email_server	mailhub-3dns.nam.nsroot.net	The address of the SMTP server used to send the autosupport email to EMC. This needs to be a Citi email server. There is also an optional email_server_2 field that can point to another Citi email server. Both email servers should be defined. The primary email server is required. Each region may have its own email server list or dns lists to use.
5.6.43	/nas/sbin/nas_connecthome -modify -email_to	emailalert@emc.com	This is the standard emc alerting email address. It is the default value.
5.6.43	Use of Collect Audit Logs Scripting	Run daily	See Collect Logs Script
5.6.43	HTTPS Control station	off	The "off" notations in the required settings column are in reference to the IPTables Script that is implemented to restrict access to approved host with everything denied as default. These are not a parameter settings as these point to an ACL list where these are controlled. See IPTables Script
5.6.43	HTTP Control Station	off	The "off" notations in the required settings column are in reference to the IPTables Script that is implemented to restrict access to approved host with everything denied as default. These are not a parameter settings as these point to an ACL list where these are controlled. See IPTables Script
5.6.43	SSH Control Station	off	The "off" notations in the required settings column are in reference to the IPTables Script that is implemented to restrict access to approved host with everything denied as default. These are not a parameter settings as these point to an ACL list where these are controlled. See IPTables Script
5.6.43	/nas/sbin/nas_config - password -retries	6	This option dictates how many failed password login attempts are permitted before a storage administrators account is locked out. The default value for this setting is 3. Configuring this option to '6' meets the CITMP policy requirements. <i>Must have root privileges for this command.</i>
5.6.43	/nas/sbin/nas_config - password -min	6	<6..15> = minimum length of the new password. The default length is 8 characters. The length has to be a value between 6 and 15 characters. <i>Must have root privileges for this command.</i>
5.6.43	/nas/sbin/nas_config -	2	This option controls the minimum number of new

	password -newchars		alphabetic characters a password must have. (IE: A password cannot be just digits and symbols.) These are capital and lowercase letters from a to z. The default value for this option is 3. <i>Must have root privileges for this command.</i>
5.6.43	/nas/sbin/nas_config – password -digits	1	This option controls the minimum number of digit characters a password must have. These are numbers from 0 to 9. The default value for this option is 1. <i>Must have root privileges for this command.</i>
5.6.43	/nas/sbin/nas_config – password –spechars	0	This option controls the minimum number of symbol characters a password must have. These are whitespace and punctuation characters. The default value for this option is 0. <i>Must have root privileges for this command.</i>
5.6.43	server_ftp server_<n> - service –status	stopped	<n>=Datamover server number (slot number). This option controls the FTP service for the datamovers. FTP must be disabled. See Disabling the FTP service on the Data Movers .

20.4 APPENDIX D: Additional VNX Scripting

To further automate the management functions of the VNX, several scripts have been co-developed by Citi and EMC.

20.4.1 Check Configuration (Recommended)

This *optional* script (check_config_1.9.pl or later version) will be run from a windows client machine that will connect to all of the Celerra control stations to verify the security configuration. This script is managed by an excel spreadsheet. It produces a report that is also an excel spreadsheet.

20.4.1.1 Prerequisites

In order to run the script to report on the status of each Celerra, you will need the following:


- Create directory C:\Celerra_Uilities\Check_Config\Scripts
- Create directory C:\Celerra_Uilities\Check_Config\Scripts\Configurations
- Create directory C:\Celerra_Uilities\Check_Config\Scripts\Reports
- Save “check_config_1.x.pl” into “Scripts” sub-directory, x is the current version of script
- Save “nas_info.xls” into “Configurations” sub-directory
- After running the “check_config_1.x.pl” script, the results will be saved under “Reports” sub-directory
- Active Perl v5.3.809 (spac perls580) installed, or latest version
- Tectia SSH Client v4.4.10.8 (spac tccl4430) installed, or latest version
- Additional Perl libraries (see below).
- SSH Public and private keys on the windows server and the corresponding public key on the control stations, See Section 18.4.1.3 SSH Public and Private Keys

20.4.1.2 Perl Libraries

The script makes use of several perl libraries that need to be stored in the library directory structure on the host running the script. The required libraries are in addition to the base Active Perl v5.3.809 Citi package (perls580).

The following libraries (directories, sub directories, and contents) can be found in:

- Tie (RegexpHash) copy to \app_srv\ActivePerl\Perl\site\lib\
- OLE(Storage_Lite) copy to \app_srv\ActivePerl\Perl\site\lib\
- Spreadsheet (Parse and Write Excel) copy to \app_srv\ActivePerl\Perl\site\lib\



Note: These libraries will have to be manually copied to the specified directories. The package does not copy them.

20.4.1.3 SSH Public and Private Keys

The script makes use of the ssh protocol to automate the communication with the Celerras. To avoid repeated requests for username and passwords for sessions, public and private keys are established. It is assumed that the windows host has the Tectia SSH Client v4.4.10.8 Citi package (tccl4430), or later version is installed. The Tectia package updates the path, but is not effective until the next logon. To configure the Celerras for key-based communication perform the following steps:

1. On the windows host execute as the current user (after a logoff):


```
ssh-keygen2 -P
Generating 2048-bit dsa key pair
55 Oo.oOo.oOo.o
Key generated.
2048-bit dsa, gj06529@NASMGMTGDSE05, Wed Oct 21 2009 19:10:20
Private key saved to C:/Documents and Settings/gj06529/Application Data/SSH/User
Keys/id_dsa_2048_a
Public key saved to C:/Documents and Settings/gj06529/Application Data/SSH/UserK
eys/id_dsa_2048_a.pub
```

This will generate the public (.pub) and private keys (no ext) in the ssh2 format. It places the keys into an Application Data subdirectory for the current user.
2. On the Celerra Control Station (as root) execute:


```
mkdir /root/.ssh
```
3. On the windows host execute (where "celerra#1" is the IP of the first control station)
 1. `cd C:\Documents and Settings\<your username>\Application Data\SSH\UserKeys`
 2. `scp2 id_dsa_2048_a.pub root@celerra#1:/root/.ssh`

```
Host key digest saved to
d:\data\home\Administrator\.ssh2\hostkeys\keys_d996c8ea
96ad262733ea5bc6a4dac3b18dcc2d4d
warning: Invalid permissions for directory
'd:\data\home\Administrator\.ssh2',user keys not read
root@10.40.42.211's password:
id_dsa_2048_c.pub | 1.2kB | 1.2kB/s | TOC: 00:00:01 | 100%
```

You will be prompted for the root password in order to copy the file. You will likely get a warning message about permissions (ignore this). This will copy the public key file to the Celerra control station. If you are managing files from multiple administration servers, you can rename the key files before copying them to the Control Station.

Note: If you get a "destination is unreachable" error message, then you will need to add this windows server to the IPTables ACL for SSH access. Retry the scp2 command once this has been completed.

4. On the Celerra Control Station (as root) execute:
 1. `cd /root/.ssh`
 2. `ssh-keygen -i -f id_dsa_2048_a.pub >> authorized_keys`

This will convert the public key in ssh2 format to the openssh format and will append it to the authorized key file (that may contain multiple keys).

5. After the first Celerra Control Station has been configured, just copy the `authorized_keys` and `id_dsa_2048_a.pub` to the `/root/.ssh` folder from the first Celerra to the remaining Celerras.

```
scp authorized_keys root@celerra#2:/root/.ssh
```

```
authorized_keys 100% 2202 2.2KB/s 00:00
```

This will establish automated communication from the windows host to each of the Celerra Control Stations. If another windows host is desired, then all of the steps must be repeated.

20.4.1.4 Running the Script

Input

This script reads an excel spreadsheet that contains the Citi security configuration. For any new Celerra server add a new line to the spreadsheet following the example format:

Hostname	SessionTimeou	Number_Digits_On_Password	Password_Minimum	Number_New_Characters_On_Password	Number_Allowed_Retries	Number_Of_Special_Characters	Email_Sender	Email_Recipient	Primary_Mail_Server	Secondary_Mail_Server
10.40.42.201	60	1	6	2	6	0	connectemc@emc.com	connectemc@citi.com	168.159.48.78	168.159.48.79
10.40.42.211	5	1	6	2	6	0	Jorge.Martinez@citi.com	emcalert@emc.com	168.159.48.78	168.159.48.79
10.5.22.96	5	1	6	2	6	0	Jorge.Martinez@citi.com	emcalert@emc.com	168.159.48.78	168.159.48.79

There should be a line for every Celerra with all of the values filled out (i.e. copy previous line).

Output

The script will produce an excel spreadsheet (C:\Celerra_Utility\Check_Config\Reports\citi_report_mmddyy.xls) with the corresponding results (Yes = complies with defined value, No = does not comply with defined value) defined by the input excel spreadsheet.

Hostname	Session_Timeout_Time	Password_Digits	Password_Min	Password_New_chars	Password_Retries	Password_Spechars	FTP_Disable	Email_From	Email_To	Primary_Email_Server	Secondary_Email_Server	SSH_IPTables	HTTP_IPTables	Auditd_Disable
10.40.42.201	Yes	Yes	Yes	Yes	No	Yes	No	No	No	No	Yes	No	No	No
10.40.42.211	Yes	Yes	Yes	Yes	No	Yes	Yes	No	No	No	No	No	No	No
Host 10.5.22.96 did not respond														

20.4.1.1 Frequency

This script is intended to be run on a regular basis (at least annually) in order to verify deployments of recommended security configurations.

20.4.2 Collect Logs (Required)

This *required* script (`collect_logs_1.7.pl` or later version) will be run from each VNX control station. The script will be called from a shell script (`run_collection.sh`) that gets executed as a cron job. The scripts (perl and shell) will be stored on an NFS share that will also contain the log files. The script also maps the userids from `/etc/passwd` to the `cmd_log` file so that it is more readable. The following log files are copied by this script:

- `/nas/log/osmlog`
- `/nas/log/cmd_log`
- `/nas/log/cmd_log.err`
- `/nas/http/logs/ssl_request_log`
- `/nas/http/logs/access_log`
- `/var/log/messages`

This script is required for corrective action due to gaps in the audit log. Until all logs are delivered via Syslog, this script will be required.

20.4.2.1 Create the NFS Shares

Establish an NFS Share that will contain the audit logs from every VNX (organized by hostname_serialnumber). Establish another NFS share to contain the scripts required to collect the logs. There should be a subdirectory (rather than using the root) so that security can be set for the directory. The subdirectory (for example "audit") should have permissions set to 755 and the owner set to nasadmin.

1. Create nfs share and give root access to a remote host
2. Configure the subdirectory. As root:
 - Attach to the nfs share from that remote host (mount xxx.xxx.xxx.xxx/nfs_share /mnt/my_dir
 - Create a subdirectory (mkdir /mnt/my_dir/audit)
 - Change the permission to 755 (chmod 755 /mnt/my_dir/audit)
 - Make nasadmin the owner (chown nasadmin /mnt/my_dir/audit)
 - Disconnect the mount (umount /mnt/my_dir)
3. Delete the nfs share.
4. Recreate the nfs share putting the subdirectory in the path of the share, listing the VNXs only as read/write hosts and access hosts.
5. Repeat all steps for the second nfs share.

The amount of storage required is based on the number of days logs kept on the NFS share and the number of VNXs. The number of extra day's logs is specified as a parameter to the perl script.

20.4.2.2 Mount the NFS Shares

Each NFS Share should be mounted to a local directory (/mnt/NAS_LOGS and /mnt/NAS_SCRIPTS) that is consistent among all of the VNXs and the scripting. There are two mounts as there may be retention and security differences between the script area and auditlog area. The mounts should also be added to the /etc/fstab file for persistence. Unfortunately, the local directories listed below are hard coded into the scripts.

Sample /etc/fstab entry:

```
10.40.42.213:/fs88/audit      /mnt/NAS_LOGS      nfs  noauto,rw  0 0
10.40.42.213:/fs89/audit      /mnt/NAS_SCRIPTS    nfs  noauto,rw  0 0
```

20.4.2.3 Copy the scripts to the NFS Share

The scripts will be located in a local directory on the windows host with the installed VNX Utilities package (such as c:\Program Files\VNX_Utilities\Collect_Logs). Use SCP to move the files to the VNX.

```
scp2 c:\Program Files...collect_logs_1.7.pl nasadmin@xx.xx.xx.x:/mnts/NAS_SCRIPTS
```

SCP may not put the executable bit on these scripts, so a "chmod 744" may need to be run on them.

20.4.2.4 Configure Cron

Cron needs to be configured on every VNX server. To configure cron do the following:
From the control station (as nasadmin)

1. cd /nas/site/cron.d
2. vi nas_user
3. Add a new line as follows (or at another time):


```
15 7 * * * nasadmin /mnt/NAS_SCRIPTS/run_collection.sh >/dev/null 2>&1
```
4. Save the file

5. Restart the crond service (/sbin/service crond restart) to make sure changes are effective immediately

20.4.2.5 Run Collection

The cron entry should call a shell script call run_collection.sh (provided) that runs the perl script. The purpose of this shell script is to make sure that all of the VNXs call the perl script in the same way. This is also an easy way to update the name of the script or the parameter without having to modify it on each CS. The run_collection.sh shell script should look like:

```
#!/bin/bash
perl /mnt/NAS_SCRIPTS/collect_logs_1.7.pl 367
```

20.4.2.6 Create Logs Subdirectory

For additional organization and to test the security on the NAS_LOGS nfs share, create a subdirectory called “Logs”. Unfortunately, this directory is hard-coded into the scripting.

20.4.2.7 Frequency and Cleanup

The cron definition determines the time of day that the script will be run. The script should be run once per day. The run_collection.sh shell script specifies the number of days worth of logs to keep. A value of 0 (zero) will keep no previous days. Using 1 (one) will keep yesterday’s logs and today’s. **The recommended value is 367 (three hundred sixty seven) days.** This will provide an entire year of logs available at any time. This is a records and retention requirement.

20.4.2.8 WORM Audit Repository (Optional)

The use of WORM based storage for the audit log repository is an optional configuration. It is important to know that the script will try to clean up old configuration files based on the passed parameter. Therefore, the WORM environment should be set to a value that is less than the script value. WORM environments should allow the user(s) to remove files after their expiration has passed. Otherwise, the script will try to delete a file that is not allowed for deletion.

Based on the 1 year log retention requirement:

- Set the WORM for 366 days expiration
- Set the script for 367 days

20.5 APPENDIX E: Sample Configuration of VNX Models

The following are expected common configurations of the VNX for NAS deployment.

20.5.1 VNX 5300 Sample Configuration

2-blade, 70x 600GB SAS, 8x TB NL SAS, 4x 300GB SAS Vault

VNX Components with EMC Rack Rail Kit		
Part	Description	Qty
VNX5300CS	VNX5300 CONTROL STATION - EMC RACK	1
VNX5300CS2	VNX5300 CONTROL STATION - 2ND CS - EMC RACK	1
VNX5300DME	VNX5300 DME: 1 DATA MOVER+FC SLIC-EMC RACK	1
VNX5300DM	VNX5300 ADD ON DATA MOVER+FC SLIC-EMC RACK	1
VNX53D153015	VNX5300 DPE;15X3.5 DRIVES EMC RACK 8X300GB	1
V31-DAE-R-15	3U DAE WITH 15X3.5 INCH DRIVE SLOTS WITH RACK	2
VNXRACK-40US	VNX 40U RACK WITH CONSOLE	1

VNX Components with Customer Rack Rail Kit		
VNX5300CSM	VNX5300 CONTROL STATION -MINIRACK	1
VNX5300CS2M	VNX5300 CONTROL STATION - 2ND CS - Mini RACK	1
VNX5300DMEM	VNX 5300 DME: 1 DATA MOVER+FC SLIC-MINI RACK	1
VNX5300DMM	VNX5300 ADD ON DATA MOVER+FC SLIC-MINI RACK	1
VNX53D153015M	VNX5300 DPE; 15X3.5 DRIVES - MINI RACK 8X300G	1
V31-DAE-N-15	3U DAE WITH 15X3.5 INCH DRIVE SLOTS	2

VNX Network and Disk Options		
Datamover Network Interface Options		
VDMM1GCUAS	1GBE DM MODULE 4 PORT FOR VNX5300	2
VDMMXG2TWAS	10GBE 2 OP MODULE (TWINAX) FOR VNX5300	2
Disk Options		
V3-VS15-600	600GB 15K SAS DISK DRIVE	72
V3-VS07-020	2TB 7200RPM 6GB SAS DISK DRIVE	10

Standard items for VNX 5300		
PW40U60-USS	RACK-40U-60 PWR CORD US	1
VNXFCSFPS	ADDITIONAL 8 G FC SFP FOR VNX 51/53	1
VNXSPSAS	2ND OPTIONAL SPS FOR VNX 51/53	1
Software License Options		
456-103-679	CC / PROSPHERE MGMT PKG VNX 5300	1
457-100-732	CC / PROSPHERE VNX MGMT PKG BASE	1
ADVP-V53	ADV FILE LICENSE (NFS; MPFS AND PNFS) FOR VNX5300	1
BASEP-V53	BASE FILE LICENSE (CIFS AND FTP) FOR VNX5300	1
LPS-V53	LOCAL PROTECTION SUITE FOR VNX5300	1
RPS-V53	REMOTE PROTECTION SUITE FOR VNX5300	1
SCS-V53	SECURITY & COMPLIANCE SUITE FOR VNX5300	1
UNIF-V53	UNISPHERE FILE & VNX OE VNX5300	1
Support and Documentation		
M-PRESW-001	PREMIUM SOFTWARE SUPPORT	1
WU-PREHW-001	PREMIUM HARDWARE SUPPORT - WARR UPG	1
VNX53-KIT	DOCUMENTATION KIT FOR VNX5300	1

20.5.2 VNX 5500 Sample Configuration

3-blade, 140x 600GB SAS, 14x TB NL SAS, 4x 300GB SAS Vault

Part	Description	Qty
PW40U-60-USV	CAB QUAD POWER CORD US TWISTLOCK	1
V-VX-VS1530	3.5 IN 300GB 15K VAULT PACK FOR 6GSDAE-15/15 drive DPE	1

VDMM1GCUA	VNX 1GBASE-T DM MODULE 4 PORT	3
VDMMXG2TWA	VNX 10GBE 2 OP MODULE (TWINAX)	3
VNX5500DM	VNX5500 ADD ON DATA MOVER+FC SLIC-EMC RACK	2
VNX5500DME	VNX5500 DME: 1 DATA MOVER+FC SLIC-EMC RACK	1
VNX5500DP15	VNX5500 DPE; 15X3.5 DRIVE SLOTS-EMC RACK	1
VNX55DMEX	EMPTY VNX55 DM ENCLOSURE - EMC RACK	1
VNX6GSDAE15	VNX 15X3.5 IN 6GB SAS EXP DAE -EMC RACK	10
VNXCS	VNX CONTROL STATION-EMC RACK	1
VNXCS2	VNX 2ND CONTROL STATION-EMC RACK	1
VNXFCSFP	ADDITIONAL 8 G SFP FOR VNX5500	1
VNXRACK-40U	VNX 40U RACK WITH CONSOLE	1
VX-VS07-020	2TB 7200RPM 6GB SAS 3.5 CARRIER	16
VX-VS15-600	600GB 15K 520BPS 6GB SAS 3.5 CARRIER	143
456-103-680	CC / PROSPHERE MGMT PKG VNX 5500	1
457-100-732	CC / PROSPHERE VNX MGMT PKG BASE	1
ADVP-V55	ADV FILE LICENSE (NFS; MPFS AND PNFS) FOR VNX5500	1
BASEP-V55	BASE FILE LICENSE (CIFS AND FTP) FOR VNX5500	1
LPS-V55	LOCAL PROTECTION SUITE FOR VNX5500	1
RPS-V55	REMOTE PROTECTION SUITE FOR VNX5500	1
SCS-V55	SECURITY & COMPLIANCE SUITE FOR VNX5500	1
UNIF-V55	UNISPHERE FOR FILE FOR A VNX5500	1
VNXOE-55	VNX OE LICENSE MODEL FOR VNX5500	1
VNXOECAPT	VNX OE PER TB HIGH CAPACITY FOR VNX5500;5700;7500	32
VNXOEPERFTB	VNX OE PER TB PERFORMANCE FOR VNX5500;5700;7500	87
M-PRESW-001	PREMIUM SOFTWARE SUPPORT	1
WU-PREHW-001	PREMIUM HARDWARE SUPPORT - WARR UPG	1
VNX55-KIT	DOCUMENTATION KIT FOR VNX5500	1

20.5.3 VNX 5700 Sample Configuration

VNX 5700: 4-blade, 210x 600GB SAS, 28x TB NL SAS, 4x 300GB SAS Vault

Part	Description	Qty
MSAS-MSAS-10	One pair of Mini-SAS to Mini-SAS 10M cables	1
MSASHD-MSAS10	ONE PR 8M MINI-SAS HD TO MINI-SAS CABLES	2
PW40U-60-USV	CAB QUAD POWER CORD US TWISTLOCK	2
V-VX-VS1530	3.5 IN 300GB 15K VAULT PACK FOR 6GSDAE-15/15 drive DPE	1
VDMM1GCUA	VNX 1GBASE-T DM MODULE 4 PORT	4
VDMMXG2TWA	VNX 10GBE 2 OP MODULE (TWINAX)	4
VNX5700DM	VNX5700 ADD ON DATA MOVER+FC SLIC-EMC RACK	2
VNX5700DME	VNX5700 DME: 2 DATA MOVER+FC SLIC-EMC RACK	1
VNX5700SPE	VNX5700 SPE 4X6G SAS BE-EMC RACK	1
VNX57DMEX	EMPTY VNX57 DM ENCLOSURE - EMC RACK	1
VNX6GSDAE15	VNX 15X3.5 IN 6GB SAS EXP DAE -EMC RACK	16
VNX6GSDAE15P	VNX57/75 15X3.5 IN 6GB SAS PRIMARY DAE-EMC RACK	1
VNXCS	VNX CONTROL STATION-EMC RACK	1
VNXCS2	VNX 2ND CONTROL STATION-EMC RACK	1
VNXRACK-40U	VNX 40U RACK WITH CONSOLE	2
VNXSPS1KW	VNX57/75 1.2KW SPS 15/25 DRV VAULT DAE-EMC RACK	1
VSPM8GFFEA	VNX 4 PORT 8G FC IO MODULE PAIR	1
VX-VS07-020	2TB 7200RPM 6GB SAS 3.5 CARRIER	30
VX-VS15-600	600GB 15K 520BPS 6GB SAS 3.5 CARRIER	214
MSAS-MSAS-10	One pair of Mini-SAS to Mini-SAS 10M cables	1
MSASHD-MSAS10	ONE PR 8M MINI-SAS HD TO MINI-SAS CABLES	2
PW40U-60-USV	CAB QUAD POWER CORD US TWISTLOCK	2
V-VX-VS1530	3.5 IN 300GB 15K VAULT PACK FOR 6GSDAE-15/15 drive DPE	1
VDMM1GCUA	VNX 1GBASE-T DM MODULE 4 PORT	4
VDMMXG2TWA	VNX 10GBE 2 OP MODULE (TWINAX)	4
VNX5700DM	VNX5700 ADD ON DATA MOVER+FC SLIC-EMC RACK	2
VNX5700DME	VNX5700 DME: 2 DATA MOVER+FC SLIC-EMC RACK	1
VNX5700SPE	VNX5700 SPE 4X6G SAS BE-EMC RACK	1
VNX57DMEX	EMPTY VNX57 DM ENCLOSURE - EMC RACK	1
VNX6GSDAE15	VNX 15X3.5 IN 6GB SAS EXP DAE -EMC RACK	16
M-PRESW-001	PREMIUM SOFTWARE SUPPORT	1
WU-PREHW-001	PREMIUM HARDWARE SUPPORT - WARR UPG	1
VNX57-KIT	DOCUMENTATION KIT FOR VNX5700	1

20.5.4 VNX 7500 Sample Configuration

VNX 7500: 8-blade, 490x 600GB SAS, 56x TB NL SAS, 4x 300GB SAS Vault

Part	Description	Qty
MSAS-MSAS-10	One pair of Mini-SAS to Mini-SAS 10M cables	3
MSASHD-MSAS10	ONE PR 8M MINI-SAS HD TO MINI-SAS CABLES	6
PW40U-60-USV	CAB QUAD POWER CORD US TWISTLOCK	4
V-VX-VS1530	3.5 IN 300GB 15K VAULT PACK FOR 6GSDAE-15/15 drive DPE	1
VDMM1GCUA	VNX 1GBASE-T DM MODULE 4 PORT	8
VDMMXG2TWA	VNX 10GBE 2 OP MODULE (TWINAX)	8
VNX6GSDAE15	VNX 15X3.5 IN 6GB SAS EXP DAE -EMC RACK	37
VNX6GSDAE15P	VNX57/75 15X3.5 IN 6GB SAS PRIMARY DAE-EMC RACK	1
VNX7500DM	VNX7500 ADD ON DATA MOVER+FC SLIC-EMC RACK	6
VNX7500DME	VNX7500 DME: 2 DATA MOVER+FC SLIC-EMC RACK	1
VNX7500SPE	VNX7500 SPE 4X6G SAS BE-EMC RACK	1
VNX75DMEX	EMPTY VNX75 DM ENCLOSURE - EMC RACK	3
VNXCS	VNX CONTROL STATION-EMC RACK	1
VNXCS2	VNX 2ND CONTROL STATION-EMC RACK	1
VNXRACK-40U	VNX 40U RACK WITH CONSOLE	4
VNXSPS1KW	VNX57/75 1.2KW SPS 15/25 DRV VAULT DAE-EMC RACK	1
VSPEXPSASBEA	VNX7500 SAS BE EXPANSION TO 8 BUSES ADD-ON	1
VSPM8GFFEA	VNX 4 PORT 8G FC IO MODULE PAIR	2
VX-VS07-020	2TB 7200RPM 6GB SAS 3.5 CARRIER	58
VX-VS15-600	600GB 15K 520BPS 6GB SAS 3.5 CARRIER	499
MSAS-MSAS-10	One pair of Mini-SAS to Mini-SAS 10M cables	3
MSASHD-MSAS10	ONE PR 8M MINI-SAS HD TO MINI-SAS CABLES	6
PW40U-60-USV	CAB QUAD POWER CORD US TWISTLOCK	4
V-VX-VS1530	3.5 IN 300GB 15K VAULT PACK FOR 6GSDAE-15/15 drive DPE	1
VDMM1GCUA	VNX 1GBASE-T DM MODULE 4 PORT	8
VDMMXG2TWA	VNX 10GBE 2 OP MODULE (TWINAX)	8
VNX6GSDAE15	VNX 15X3.5 IN 6GB SAS EXP DAE -EMC RACK	37
VNX6GSDAE15P	VNX57/75 15X3.5 IN 6GB SAS PRIMARY DAE-EMC RACK	1
VNX7500DM	VNX7500 ADD ON DATA MOVER+FC SLIC-EMC RACK	6
VNX7500DME	VNX7500 DME: 2 DATA MOVER+FC SLIC-EMC RACK	1
VNX7500SPE	VNX7500 SPE 4X6G SAS BE-EMC RACK	1
VNX75DMEX	EMPTY VNX75 DM ENCLOSURE - EMC RACK	3
M-PRESW-001	PREMIUM SOFTWARE SUPPORT	1
WU-PREHW-001	PREMIUM HARDWARE SUPPORT - WARR UPG	1
VNX75-KIT	DOCUMENTATION KIT FOR VNX7500	1

20.5.5 VNX 5200 Sample Configuration

VNX 5200: 2-blade, 13x 300GB SAS 10k

Qty	Part	Description
1	VNX52DP62510	VNX5200 DPE 25X2.5" DR-25X600G10K-EMC
1	VNXBRACK-40U	VNXB 40U RACK WITH FRONT PANEL
1	VNXB6GSDAE25	VNXB 25X2.5 6G SAS EXP DAE-EMC RACK
1	VNXBCS2	VNXB 2ND CONTROL STATION-EMC RACK
1	VNXBCS	VNXB CONTROL STATION-EMC RACK
1	VNXB52DM	VNX5200 ADD ON DM+FC SLIC-EMC RACK
1	VNXB52DME	VNX5200 DME: 1 DM+FC SLIC-EMC RACK
13	V4-2S15-300	VNX 300GB 15K SAS 25X2.5 DPE/DAE
1	VBPW40U-US	CAB QUAD POWER CORD US TWISTLOCK
1	VNX52-KIT	VNX5200 Documentation Kit=IC
2	VDMBM1GCUA	VNXB 1GBASE-T DM MODULE 4 PORT
2	VDMBMXG2OPA	VNXB 10GBE 2 OP MODULE (2 SFP+)
1	VSPBM8GFFEA	VNXB 4 PORT 8G FC IO MODULE PAIR
1	ESRS-GW-200	EMC SECURE REMOTE SUPPORT GATEWAY CLIENT
1	RP-LS	RECOVERPOINT LICENSE SOLUTION
1	456-104-619	RP/SE REM FOR RPS V51 V52=IC
1	456-104-614	RP/SE LOC FOR LPS V51 V52=IC
4	VNXBOEPERFTB	VNXB OE PER TB PERFORMANCE
1	VNXOE-5200	VNX5200 Operating Environment
1	SCS-VNX5200	VNX5200 Security & Compliance Suite=IC
1	RPS-VNX5200	VNX5200 Remote Protection Suite=IC
1	UNISF-VNX5200	VNX5200 Unisphere File Suite=IC
1	LPS-VNX5200	VNX5200 Local Protection Suite=IC
1	M-PRESWE-001	PREMIUM SW SUPPORT Includes months 1-36 @ \$ 284/mo.
1	WU-PREHWE-01	PREMIUM HW SUPPORT-WARR UPG

20.5.6 VNX 5400 Sample Configuration

VNX 5400: 2-blade, 115x 600GB SAS 10k

Qty	Part	Description
1	VNXB54DP25	VNX5400 DPE 25X2.5" DRIVE SLOTS-EMC RA
1	VNXBRACK-40U	VNXB 40U RACK WITH FRONT PANEL
4	VNXB6GSDAE25	VNXB 25X2.5 6G SAS EXP DAE-EMC RACK
1	VNXBCS2	VNXB 2ND CONTROL STATION-EMC RACK
1	VNXBCS	VNXB CONTROL STATION-EMC RACK
1	VNXB54DME	VNX5400 DME: 1 DM+FC SLIC-EMC RACK
1	VNXB54DM	VNX5400 ADD ON DM+FC SLIC-EMC RACK
115	V4-2S10-600	VNX 600GB 10K SAS 25X2.5 DPE/DAE
1	V-V4-260010	VNX 600GB 10K VAULT 25X2.5 DPE/DAE
1	VBPW40U-US	CAB QUAD POWER CORD US TWISTLOCK
1	VNX54-KIT	VNX5400 Documentation Kit=IC
2	VDMBM1GCUA	VNXB 1GBASE-T DM MODULE 4 PORT
2	VDMBMXG2OPA	VNXB 10GBE 2 OP MODULE (2 SFP+)
1	VSPBM8GFFEA	VNXB 4 PORT 8G FC IO MODULE PAIR
1	ESRS-GW-200	EMC SECURE REMOTE SUPPORT GATEWAY CLIENT
1	RP-LS	RECOVERPOINT LICENSE SOLUTION
1	456-104-620	RP/SE REM FOR RPS V53 V54=IC
1	456-104-615	RP/SE LOC FOR LPS V53 V54= IC
71	VNXBOEPERFTB	VNXB OE PER TB PERFORMANCE
1	VNXOE-5400	VNX5400 Operating Environment
1	UNISF-VNX5400	VNX5400 Unisphere File Suite=IC

1	SCS-VNX5400	VNX5400 Security & Compliance Suite=IC
1	LPS-VNX5400	VNX5400 Local Protection Suite=IC
1	RPS-VNX5400	VNX5400 Remote Protection Suite=IC
1	M-PRESWE-001	PREMIUM SW SUPPORT Includes months 1-36 @ \$ 452/mo.
1	WU-PREHWE-01	PREMIUM HW SUPPORT-WARR UPG

20.5.7 VNX 5600 Sample Configuration

VNX 5600: 2-blade, 115x 600GB SAS 10k

Note: This model supports 3x blades with RPQ for Citigroup

Qty	Part	Description
1	VNXB56DP25	VNX5600 DPE 25X2.5" DRIVE SLOTS-EMC RA
1	VNXBRACK-40U	VNXB 40U RACK WITH FRONT PANEL
8	VNXB6GSDAE15	VNXB 15X3.5 6G SAS EXP DAE-EMC RACK
1	VNXBCS2	VNXB 2ND CONTROL STATION-EMC RACK
1	VNXBCS	VNXB CONTROL STATION-EMC RACK
1	VNXB56DM	VNX5600 ADD ON DM+FC SLIC-EMC RACK
1	VNXB56DME	VNX5600 DME: 1 DM+FC SLIC-EMC RACK
1	V-V4-230015	VNX 300GB 15K VAULT 25X2.5 DPE/DAE
115	V4-VS15-600	VNX 600GB 15K SAS 15X3.5 DAE
1	VBPW40U-US	CAB QUAD POWER CORD US TWISTLOCK
1	VNX56-KIT	VNX5600 Documentation Kit=IC
2	VDMBM1GCUA	VNXB 1GBASE-T DM MODULE 4 PORT
2	VDMBMXG2OPA	VNXB 10GBE 2 OP MODULE (2 SFP+)
1	VSPBM8GFFEA	VNXB 4 PORT 8G FC IO MODULE PAIR
1	ESRS-GW-200	EMC SECURE REMOTE SUPPORT GATEWAY CLIENT
1	RP-LS	RECOVERPOINT LICENSE SOLUTION
1	456-104-616	RP/SE LOC FOR LPS V55 V56=IC
1	456-104-621	RP/SE REM FOR RPS V55 V56=IC
69	VNXBOEPERFTB	VNXB OE PER TB PERFORMANCE
1	SCS-VNX5600	VNX5600 Security & Compliance Suite=IC
1	VNXOE-5600	VNX5600 Operating Environment
1	RPS-VNX5600	VNX5600 Remote Protection Suite=IC
1	LPS-VNX5600	VNX5600 Local Protection Suite=IC
1	UNISF-VNX5600	VNX5600 Unisphere File Suite=IC
1	M-PRESWE-001	PREMIUM SW SUPPORT Includes months 1-36 @ \$ 715/mo.
1	WU-PREHWE-01	PREMIUM HW SUPPORT-WARR UPG

20.5.8 VNX 5800 Sample Configuration

VNX 5800: 3-blade, 229x 600GB SAS 10k

Qty	Part	Description
1	VNXB58DP25	VNX5800 DPE 25X2.5" DRIVE SLOTS-EMC RA
1	VNXBRACK-40U	VNXB 40U RACK WITH FRONT PANEL
9	VNXB6GSDAE25	VNXB 25X2.5 6G SAS EXP DAE-EMC RACK
1	VNXBCS2	VNXB 2ND CONTROL STATION-EMC RACK
1	VNXBCS	VNXB CONTROL STATION-EMC RACK
2	VNXB58DM	VNX5800 ADD ON DM+FC SLIC-EMC RACK
1	VNXB58DME	VNX5800 DME: 1 DM+FC SLIC-EMC RACK
1	VNXB58DMEX	EMPTY VNX58 DM ENCLOSURE-EMC RACK
229	V4-2S10-600	VNX 600GB 10K SAS 25X2.5 DPE/DAE
1	V-V4-260010	VNX 600GB 10K VAULT 25X2.5 DPE/DAE
1	VBPW40U-US	CAB QUAD POWER CORD US TWISTLOCK

1	VNX58-KIT	VNX5800 Documentation Kit=IC
6	VDMBMXG2OPA	VNXB 10GBE 2 OP MODULE (2 SFP+)
1	VSPBM8GFFEA	VNXB 4 PORT 8G FC IO MODULE PAIR
1	ESRS-GW-200	EMC SECURE REMOTE SUPPORT GATEWAY CLIENT
1	RP-LS	RECOVERPOINT LICENSE SOLUTION
1	456-104-617	RP/SE LOC FOR LPS V57 V58=IC
1	456-104-622	RP/SE REM FOR RPS V57 V58=IC
139	VNXBOEPERFTB	VNXB OE PER TB PERFORMANCE
1	SCS-VNX5800	VNX5800 Security & Compliance Suite=IC
1	UNISF-VNX5800	VNX5800 Unisphere File Suite=IC
1	LPS-VNX5800	VNX5800 Local Protection Suite=IC
1	RPS-VNX5800	VNX5800 Remote Protection Suite=IC
1	VNXOE-5800	VNX5800 Operating Environment
1	M-PRESWE-001	PREMIUM SW SUPPORT Includes months 1-36 @ \$ 1,261/mo.
1	WU-PREHWE-01	PREMIUM HW SUPPORT-WARR UPG

20.5.9 VNX 7600 Sample Configuration

VNX 7600: 4-blade, 343x 600GB SAS 10k

Qty	Part	Description
1	VNXB76DP25	VNX7600 DPE 25X2.5" DRIVE SLOTS-EMC RA
1	VNXBRACK-40U	VNXB 40U RACK WITH FRONT PANEL
13	VNXB6GSDAE25	VNXB 25X2.5 6G SAS EXP DAE-EMC RACK
1	VNXBCS2	VNXB 2ND CONTROL STATION-EMC RACK
1	VNXBCS	VNXB CONTROL STATION-EMC RACK
2	VNXB76DM	VNX7600 ADD ON DM+FC SLIC-EMC RACK
1	VNXB76DME	VNX7600 DME: 2 DM+FC SLIC-EMC RACK
1	VNXB76DMEX	EMPTY VNX76 DM ENCLOSURE-EMC RACK
343	V4-2S10-600	VNX 600GB 10K SAS 25X2.5 DPE/DAE
1	V-V4-260010	VNX 600GB 10K VAULT 25X2.5 DPE/DAE
1	VBPW40U-US	CAB QUAD POWER CORD US TWISTLOCK
1	VNX76-KIT	VNX7600 Documentation Kit=IC
4	VDMBM1GCUA	VNXB 1GBASE-T DM MODULE 4 PORT
8	VDMBMXG2OPA	VNXB 10GBE 2 OP MODULE (2 SFP+)
1	VSPBM8GFFEA	VNXB 4 PORT 8G FC IO MODULE PAIR
1	ESRS-GW-200	EMC SECURE REMOTE SUPPORT GATEWAY CLIENT
1	RP-LS	RECOVERPOINT LICENSE SOLUTION
1	456-104-623	RP/SE REM FOR RPS V75 V76=IC
1	456-104-618	RP/SE LOC FOR LPS V75 V76=IC
207	VNXBOEPERFTB	VNXB OE PER TB PERFORMANCE
1	LPS-VNX7600	VNX7600 Local Protection Suite=IC
1	UNISF-VNX7600	VNX7600 Unisphere File Suite=IC
1	SCS-VNX7600	VNX7600 Security & Compliance Suite=IC
1	VNXOE-7600	VNX7600 Operating Environment
1	RPS-VNX7600	VNX7600 Remote Protection Suite=IC
1	M-PRESWE-001	PREMIUM SW SUPPORT Includes months 1-36 @ \$ 1,959/mo.
1	WU-PREHWE-01	PREMIUM HW SUPPORT-WARR UPG

20.5.10 VNX 8000 Sample Configuration


VNX 8000: 8-blade, 685x 600GB SAS 15k

Qty	Part	Description
1	VNXB80SPE	VNX8000 SPE-EMC RACK
4	VNXBRACK-40U	VNXB 40U RACK WITH FRONT PANEL
45	VNXB6GSDAE15	VNXB 15X3.5 6G SAS EXP DAE-EMC RACK
1	VNXB6GSDAE15P	VNXB 15X3.5 6G SAS PRI DAE-EMC RACK
1	VNXB80DME	VNX8000 DME: 2 DM+FC SLIC-EMC RACK
1	VNXBCS2	VNXB 2ND CONTROL STATION-EMC RACK
1	VNXBCS	VNXB CONTROL STATION-EMC RACK
6	VNXB80DM	VNX8000 ADD ON DM+FC SLIC-EMC RACK
3	VNXB80DMEX	EMPTY VNX80 DM ENCLOSURE-EMC RACK
1	V-V4-V30015	VNX 300GB 15K VAULT 15X3.5 DAE
685	V4-VS15-600	VNX 600GB 15K SAS 15X3.5 DAE
6	VBMSASHDMSAS8	ONE PR 8M MINI SAS HD TO MINI SAS CBLS
3	VBMSAS-MSAS10	ONE PR OF 10M MINI SAS TO MINI SAS CBLS
4	VBPW40U-US	CAB QUAD POWER CORD US TWISTLOCK
1	VNX80-KIT	VNX8000 Documentation Kit=IC
16	VDMBMXG2OPA	VNXB 10GBE 2 OP MODULE (2 SFP+)
2	VSPBM8GFFEA	VNXB 4 PORT 8G FC IO MODULE PAIR
1	ESRS-GW-200	EMC SECURE REMOTE SUPPORT GATEWAY CLIENT
1	RP-LS	RECOVERPOINT LICENSE SOLUTION
1	456-104-626	RP/SE LOC FOR LPS V80=IC
1	456-104-627	RP/SE REM FOR RPS V80=IC
411	VNXBOEPERFTB	VNXB OE PER TB PERFORMANCE
1	RPS-VNX8000	VNX8000 Remote Protection Suite=IC
1	LPS-VNX8000	VNX8000 Local Protection Suite=IC
1	UNISF-VNX8000	VNX8000 Unisphere File Suite=IC
1	VNXOE-8000	VNX8000 Operating Environment
1	SCS-VNX8000	VNX8000 Security & Compliance Suite=IC
1	M-PRESWE-001	PREMIUM SW SUPPORT Includes months 1-36 @ \$ 2,953/mo.
1	WU-PREHWE-01	PREMIUM HW SUPPORT-WARR UPG

20.6 APPENDIX F: Software Enablers of VNX Models

Required Packages (Software Enablers)

To view the current enablers (packages), go to System → System Properties → Software Tab



Note: The above example has every enabler Active and should not be used as a guide.

The following are required packages and should be Active on all VNX NAS (File or Unified) Arrays:

Software Enabler / Package	Reboot?	Notes
UnisphereFile	No	
UnisphereBlock	No	
UnisphereAnalyzer	No	Part of FAST Suite (FAST, FAST Cache, QoS Manager, Analyzer)
Unisphere	No	
ThinProvisioning	Yes	Part of VNX Installation Toolbox (SAN Copy, Thin Provisioning, Compression)
FASTCache	No	Part of FAST Suite (FAST, FAST Cache, QoS Manager, Analyzer)
Deduplication		
Compression	Yes	Part of VNX Installation Toolbox (SAN Copy, Thin Provisioning, Compression)

Other enablers are not required, but may be installed with certain packages (i.e. QoS Manager and FAST come with FAST Suite).**Enablers should be installed by an EMC CE.

20.7 APPENDIX G: CAVA 4.6.8 CAVA Support and Dual CPU Recommended Ratio

20.7.1 CAVA System Requirements

The following table details the minimum system requirements recommended for each AV server. The CAVA server may be a VMWare host (if using VSE 8.7i or later).

Subject	Minimum Requirement
Hardware	600 MHz (or faster) Pentium III compatible processor 512 MB RAM 80 MB of free disk space VMWare or physical server
Operating System	Windows 2003 Server or Windows 2003 x64 Server or Windows 2008 x64 Server
AD and Network	The Windows network must contain a domain controller with Active Directory and DNS enabled.

Note: This server must be dedicated and is not to be used for any other function (i.e. scripting or otherwise).

20.7.2 VNX to AV scanner ratio

This table is only a guideline based upon best practices. This guideline provides the flexibility to add or remove CAVA servers to maximize scanner utilization as needed. CAVA server performance should be closely monitored via BMC agents to ensure VNX responsiveness to end-user requests. AV farm members must be similar in performance characteristics and hardware (cpu, memory, network) since round-robin load balancing does not have insight into individual server utilization. VMware is a candidate to host AV scanners, provided vendor support is available. McAfee supports VMware for version 8.7i and later.

For reference, please consider the follow resources:

VNX Active Datamovers	CAVA Scanners
1	2+
2	3+
3-4	4+
5-7	5+

21 Revision History

0.2 Initial Version partially based on v1.6.3 of Celerra SG with contributions from Tim Armstrong, Jackson Tan, and Sean Tai.

1.0 Production Release

- Updated sudoers example configuration file (v 1.4) [16.2]
- Updated Architecture section and added model specific configuration [1.1]

1.1 Unified Release

- Increased DM Limit from 32TB to 100TB when using 10GigE [3.11, 3.12]
- Updated Usermapper backup procedures [3.14.7.5]
- Added Unisphere and SSH Configuration/Login (moved 2.1 and below to 2.2) [2.1, 2]
- Added link to VNX for SAN solutions guide [1]
- Added FLR (WORM) [9]
- Added Robocopy Configuration [2.16]
- Updated audit scripting [6.1.2]
- Added 3DNS support, updated VDM configuration information [8]
- Added CAVA sizing tool and updated CAVA setup instructions [13]
- Added Unified Considerations (moved appendices to 19, revision history to 20) [19]

1.2 Minor update

- Updated NTP example for typo [2.11.1]
- Updated Time zone example [2.11.2]
- Updated Unified Security to describe account specifics and AD recommendations [19]
- Removed redundant statement about FLARE/DART releases [19]
- Added Activities & Roles Chart to Unified [19]
- Removed ownership recommendation [19]

1.3 Storage Pools Update

- Added Microsoft reference for setspn command [8.6.1]
- Added CIFS server rename [3.14.7]
- Added quota restriction statement [2.16]
- Updated NFS for VDM [3.15]
- Updated CS Storage Domain based DNS and NTP [16.3.1, updated 16.3.2+, 2.10, 2.9]
- Removed redundant blank VDM naming standards heading [2.9 -]
- Updated Instructions for AD configuration for Storage Domain [16.3.3]
- Added SMB 3.0 info to SMB Protocol [3.14.6, updated 3.14.7+]
- Added Block Storage Pool and modified "AVM" & File Pool [3, 3.1, 3.2, 3.3, Incremented 3.3+, 3.4, removed 3.13]
- Updated FASTCache Recommendations [3.8]
- Updated FSN Planning [2.5]
- Added more CLI output for performance statistics [7.4.2]
- Added Customizing the Dashboard [17.2]
- Modified Sudo setup file for Netuitive and functional id list [16.1,16.2]

1.4 VNX NG Update

- Added VNX NG Models to supported models and model details [1.1,1.2]
- Updated Spare Drives [3.5.1]
- Updated Supported Drives [3.5.3]
- Updated NDMP Datastreams [10.7]
- Updated Software Package List [1.3]

-
- Updated FAST Cache [1.1.2, 3.9]
- 1.4.1 VNX NG Additional Items
- Added model info for VNX NG [1.1.1]
 - Added sample configs for VNX NG [20.5]
 - Added File Extension Filtering [3.14.8]
 - Updated comments for IPTables “off” for HTTPS, HTTP, SSH [20.3] Updated Spares (CLI Commands for Spare Policy) [3.5.1]
- 1.5 VNX Update
- Updated McAfee Agent and MVSE versions [13.3]
 - Added Citi Information Security Standards [16.12]
 - Added Information Classifications Certifications [16.13]
 - Added HTTP, iSCSI, and HTTP are not support protocols [1.1.1.4]
 - Added CS sFTP ability is not restricted with the FTP service disabled [16.9]
 - Updated Vault Disk design [3.5.2]
 - Added Software Enablers section [20.6]
- 1.5.1 CAVA 6.5.0 details update
- Release for Testing
 - Edited the CAVA Version, Platform and Package Name [1.3]
 - Removed Special Notes/ Known Issues for Packages [1.3.1]
 - Replaced Celerra wording with VNX [13.1]
 - Removed Windows 2003 Server support and added Windows 2012 x64 Server support [13.2]
 - Screen output for checking CAVA Status replaced [13.21]
 - Removed screenshot for Log Files output [13.23.2]
 - Updated VNX to AV Scanner Ratio [13.4]
 - Added dedicated FLR infrastructure requirement [9]
 - Added Appendix G for previous CAVA requirements [20.7]
- 1.7 NFSv4 certified support for NFS only use [3.15.5]
- NFSv4 detail [3.15.5.1]
 - NFSv4 ACL [3.15.5.2]
 - NFSv4 and NFSv3 coexistence [3.15.5.3]
 - Useful NFSv4 Host information [3.15.5.4]
 - Data mover configuration to enable NFSv4 [3.15.5.5]