



Global Distributed Storage Engineering
Citi Architecture & Technology Engineering

NetApp Data OnTap 7.1

GDSE-07003-SG

Solutions Guide

Version 0.2

1 Version Control

Version	Date	Author	Reviewed	Comments
0.1	16-FEB-2007	Richard Sopp	Stephen Hopkins	Peer Review
0.2	21-JUN-2007	Richard Sopp		RFC

2 Legal Notice

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

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

3 Contact

Please address any comments or questions regarding the information contained in this document via email to the “*GT Global Distributed Storage Engineering” distribution list or alternatively to gdse@citigroup.com.

4 Latest Versions

It is recommended that the latest version of the document be used from the GDSE website <http://engineering.ti.citigroup.net/gdse/>. Newer versions may include certification on additional platforms, or comments on issues discovered after deployment.

5 Table Of Contents

1	Version Control	2
2	Legal Notice	2
3	Contact.....	2
4	Latest Versions	2
5	Table Of Contents.....	3
5.1	Index of Tables.....	5
	Index of Figures.....	6
	Index of Figures.....	6
6	Introduction	7
7	Assumptions, Conventions, Definitions, Abbreviations and References.....	7
7.1	Assumptions.....	7
7.2	Issuing commands	7
7.3	Documentation conventions.....	7
7.4	Definitions and abbreviations	8
7.5	References.....	8
8	General Design	8
9	Best Practices	10
9.1	Installation	10
9.1.1	Platform requirments.....	10
9.1.2	Scalability.....	10
9.1.3	Upgrade prerequisites.....	10
9.1.4	Upgrade procedure	11
9.1.5	Unsupported configurations.....	11
9.2	ONTAP SOE configuration (options).....	11
9.3	Network Configuration.....	11
9.3.1	Gigabit Ethernet NIC Configurations and Naming	11
9.3.2	Quad Port Gigabit / Ethernet Configurations and Naming.....	12
9.3.3	Remote LAN Management (RLM) and naming	13
9.3.4	Remaining Hardware Conventions	14
9.3.5	Filer Naming Standards	14
9.3.6	Domain Name and Host Resolution	14
9.3.7	Time Synchronization and Time Zones	15
9.4	Storage Configuration	16
9.4.1	Data ONTAP Clustering.....	16
9.4.2	Aggregate Configuration and Sizing	16
9.4.3	Volume Configuration and Sizing	20
9.4.4	Spare Disks.....	23
9.4.5	Snapshots (Point-in-time copy).....	24
9.4.6	Snapmirror (Asynchronous Data Replication)	26
9.5	File System Configuration	27
9.5.1	CIFS.....	27
9.5.2	NFS.....	28
9.6	Monitoring and Alerting	30
9.6.1	AutoSupport.....	30
9.6.2	SNMP	31
9.7	VFiler (multistore).....	31
9.7.1	VFiler Overview.....	31
9.7.2	Hosting Filer / Default VFiler	32
9.7.3	VFiler Prerequisites.....	33
9.7.4	VFiler Limitations	34
9.7.5	Data ONTAP tasks allowed on VFilers	34
9.7.6	SOE VFiler Naming Standards	34
9.7.7	SOE VFiler Production Configuration	35
9.7.8	SOE VFiler DR/COB Configuration	37
9.8	Backup	38

9.8.1	NDMP Overview	38
9.8.2	Authorizing Access To the NDMP Host	39
9.8.3	Configuring the tape devices	39
9.8.4	Backup process	41
9.8.5	NDMP Operating Notes and Restrictions	41
10	Security Best Practices	42
10.1	Overview	42
10.2	Security Standards	42
10.3	Information Classification Certification	42
10.4	Anti-Virus	42
10.4.1	New features and enhancements in VirusScan for NetApp v7.1:	43
10.4.2	VirusScan for NetApp v7.1 ePO NetApp Policies configuration.....	43
10.4.3	Network Appliance Data ONTAP "vscan" Settings	47
10.4.4	NetApp Data ONTAP Anti-Virus Service Help	48
10.4.5	Resetting the flags in the filer's cache.....	49
10.4.6	VirusScan for NetApp - Virus Detection	49
10.5	Disk Sanitization.....	50
10.5.1	Disk Sanitization Limitations.....	50
10.5.2	Licensing Disk Sanitization.....	50
10.5.3	Sanitizing Disks	50
10.5.4	Reading Disk Sanitization Log Files.....	51
10.6	CITMP Information Security Policy Convergence	52
10.6.1	Information Ownership	52
10.6.2	Identification and Authentication	53
10.6.3	Authorization and Access Control	56
10.6.4	Compensating Control.....	60
10.6.5	Traceable Access Entitlements	61
10.6.6	Confidentiality and Integrity	62
10.6.7	Incident Detection and Response	63
10.6.8	Administration.....	64
11	Document version control	71
12	Appendix A: ePO & VirusScan v7.1 for NetApp Configuration.....	73
12.1	VirusScan for NetApp v7.1 ePO Default Policies configuration.....	73
12.2	VirusScan for NetApp v7.1 AutoUpdate and Weekly Scan Settings	80
13	Appendix B: NetApp SNMP MIB.....	83
14	Appendix C: Data ONTAP SOE Configuration Standards.....	92
15	Appendix D: ONTAP, Shelf and Disk Firmware installation procedures	97
15.1	Data ONTAP Install Procedure	97
15.2	Disk Firmware Upgrade Procedure.....	99
16	Appendix F: Data Schema for NetApp Inventory Agent	100
17	Appendix G: New disk test results	101
17.1	300GB Fiber Channel	101
17.2	500GB (S)ATA	101
18	Appendix H: FilerView process examples	102
18.1	Aggregate creation	102
18.2	Volume creation	106
19	Appendix I: Tape SAN connectivity	110
19.1	Observations from testing	111
19.1.1	HBA failure / Switch failure	111
19.1.2	On-lining HBA port during backup.....	111
19.1.3	World Wide Name behavioral considerations	111
19.1.4	ONTAP auto-aliasing of tape drives	112
19.1.5	Lab performance benchmarks.....	112
20	Appendix J: SOE VFile Creation Example	113
20.1	Basic Creation	113
20.1.1	Adding Resources	114
20.1.2	VFile Setup	115

20.1.3	CIFS Setup	116
20.1.4	Create /etc/exports File	119
20.1.5	Client NFS Mount Connectivity	119
20.1.6	Snapmirror Setup (Asynchronous Replication)	120
20.2	SOE Vfiler DR/COB Creation Example	122
20.2.1	Vfiler Inventory	122
20.2.2	DR/COB Filer Configuration	123
20.2.3	Special Setup – Volume FSID update.....	123
20.2.4	Recreate Production Vfiler on DR/COB	124
20.2.5	Vfiler Production Failback.....	127
21	Appendix J: Role Based Access Control capabilities	130
21.1	Overview	130
21.2	High-level Administrative User Concepts.....	131
21.2.1	Common Capabilities	131
21.2.2	System Administrator	132
21.2.3	Account Administrator	132

5.1 Index of Tables

Table 9.1 - System firmware prerequisites.....	10
Table 9.2 - Shelf firmware prerequisites.....	10
Table 9.3 - Default Name Resolution Order	14
Table 9.4 - NTP recommended settings.....	15
Table 9.5 - Recommended root volume size when implementing flexible volumes.....	17
Table 9.6 - Recommended RAID group configuration	18
Table 9.7 - Recommended RAID configuration for FAS270	18
Table 9.8 - Procedure to remove aggregate snapshot reserve.....	19
Table 9.9 - Suggested Snapshot settings	24
Table 9.10 - Snapshot standard schedule.....	27
Table 9.11 - Caching NIS group information	30
Table 9.12 - SNMP configuration settings.....	31
Table 9.13 - Citigroup recommended vFiler limits.....	32
Table 13.1 - SNMP trap list table definitions	83
Table 13.2 - SNMP MIB matrix.....	84
Table 19.1 – Tape SAN Infrastructure configuration steps	111
Table 19.2 – Tape SAN World Wide Node Name and Port Name considerations	111
Table 19.3 – Tape SAN Lab test performance benchmarks	112

Index of Figures

Figure 8.1 - NetApp Filer Architecture Design.....	8
Figure 9.1 - NetApp MultiStore (Vfiler)	32
Figure 9.2 – Backup Simple "Local" NDMP Configuration	38
Figure 9.3 – Backup Three-party NDMP Configuration.....	39
Figure 9.4 - FilerView GUI Telnet Session - Tape Devices Configuration	40
Figure 9.5 - FilerView GUI Telnet Session - Tape Devices Status	40
Figure 10.1 – VirusScan for NetApp v7.1 User File Access Data Flow.	43
Figure 10.2 - ePO Organizational container NetApp Policy configuration	44
Figure 10.3 - ePO Push Policy Changes to Server	45
Figure 10.4 - ePO Push Policy Changes to Server – Options	45
Figure 10.5 – VirusScan Server Network Associates Program Group.....	46
Figure 10.6 - VirusScan Server Network Appliance AV Server Console	46
Figure 10.7 - VirusScan Server NetApp Filers Serviced by AV Scanner Server	47
Figure 10.8 - NetApp Data ONTAP "useradmin" command.....	53
Figure 10.9 - NetApp Data ONTAP UserID password composition enforcement.....	54
Figure 12.1 - ePO General Policies - General tab Policy Settings.....	73
Figure 12.2 - ePO General Policies - Messages tab Policy Settings	74
Figure 12.3 - ePO General Policies - Reports tab Policy Settings	74
Figure 12.4 - ePO Default Processes Policies - Processes Policy Settings	75
Figure 12.5 - ePO Default Processes Policies - Detection Policy Settings.....	76
Figure 12.6 - ePO Default Processes Policies - Advanced Policy Settings	76
Figure 12.7 - ePO Default Processes Policies - Actions Policy Settings	77
Figure 12.8 - ePO User Interface Policies - Display Options Settings	78
Figure 12.9 - ePO User Interface Policies - Password Options Settings	78
Figure 12.10 - ePO Alert Manager Alert Policies – Alert Manager Alerts Settings	79
Figure 12.11 - ePO NetApp Policies - General settings.....	80
Figure 12.12 - ePO VirusScan for NetApp v7.1 Daily DAT Update Task Configuration.....	81
Figure 12.13 - ePO VirusScan for NetApp v7.1 Daily DAT Update Schedule Configuration	81
Figure 12.14 - ePO VirusScan for NetApp v7.1 Server Weekly Scan Task Configuration	82
Figure 12.15 - ePO VirusScan for NetApp v7.1 Server Weekly Scan Schedule Configuration....	82
Figure 15.1 - http://now.netapp.com Software Download Page.....	97
Figure 15.2 - Data ONTAP OS file selection	97
Figure 15.3 - Filer OS volume - mapping to local workstation	98
Figure 15.4 - Data ONTAP OS zip extract – screen 1.....	98
Figure 15.5 - Data ONTAP OS zip extract – screen 2.....	98
Figure 15.6 - Data ONTAP OS Filer Download.....	99
Figure 15.7 - Disk firmware update	99
Figure 16.1 - Citigroup NetApp Inventory Agent Schema	100
Figure 21.1 - RBAC entity concepts	130
Figure 21.2 - RBAC high level overview.....	131

6 Introduction

The operating system used by Network Appliance filers is Data ONTAP. This guide is a general release solution implementation for all Citi NetApp NAS Filer (with the exception of systems running the [LockVault*](#) compliance).

This guide describes how to set up and configure Network Appliance filers that run Data ONTAP 7.1 software. The document is intended for Systems and Storage Administrators responsible for the management of Citi SOE Network Appliance filers. The goal of this release is to provide supported and certified software for use on all Citi NetApp NAS platforms (with the exception of LockVault systems that follow a separate upgrade path).

*An ONTAP 7.1 system can be the source of a Lockvault relationship, but this version of the solutions guide not cover the use of ONTAP 7.1 on a Lockvault destination system. Lockvault (compliance support) will be included in a later version of this document.

7 Assumptions, Conventions, Definitions, Abbreviations and References

7.1 Assumptions

This document assumes the reader is familiar with the basics of NetApp NAS storage, including: hardware platforms, Data ONTAP, VFiles, quotas, snapmirror, snapshot and snapvault operations. Although experience of a previous version of ONTAP is not a requirement, familiarization with the product is strongly recommended before continuing to read this document.

7.2 Issuing commands

All commands listed in this document are to be executed on the Filer. There are two methods by which commands can be issued: the command line (CLI), and the web user interface (FilerView). Where necessary, the type of interface used will be indicated throughout this document.

7.3 Documentation conventions

The following conventions are employed within this document:

<u>Scenario</u>	<u>Example</u>
Command line output	Courier new 9 pt
Command line input	<i>Courier new 9 pt</i>
Input variable	< <i>Courier new 10 pt italic</i> >
Quotations	<i>Arial 10 pt italic</i>

From ONTAP 7.0, all GDSE documentation will refer to the versions of the ONTAP operating system by the family name, for example ONTAP 7.1. Storage administrators and managers should consult the GDSE NetApp product matrix at http://engineering.ti.citigroup.net/gdse/products/platforms/nasfiler/stdsdocs/supported_nas_config_matrix.pdf for the recommended minimum supported version on ONTAP in a given family. The current GDSE recommended and supported release of Data ONTAP 7.1 is 7.1.2.1P1. This is the only occasion where the exact ONTAP release will be mentioned in this document.

Unlike previous versions of GDSE NetApp Solutions Guide, this document also includes the information that was previously published in the GDSE NetApp Secure Configuration Guide. The concept of the SCG document is now defunct.

7.4 Definitions and abbreviations

Term or Acronym	Explanation or Definition
Cluster	A pair of Filers configured for High availability
EMS	Enterprise Management Services (snmp monitoring)
Filer	A single network attached storage device. The term is used in this document to distinguish it from the traditional Windows Server
ISA	Infrastructure Security Administrator

7.5 References

Item Name	Location / Value
GDSE Website	http://engineering.ti.citigroup.net/gdse/
GDSE NetApp Standards	http://engineering.ti.citigroup.net/gdse/products/platforms/nasfiler/
ONTAP 7.1 Project Scope	http://engineering.ti.citigroup.net/gdse/products/platforms/nasfiler/stdsdocs/07003-ps.pdf
ONTAP 7.1 Release Notes	http://engineering.ti.citigroup.net/gdse/products/platforms/nasfiler/proposed/07003-rn.pdf
ONTAP 7.1 Solutions Guide	http://engineering.ti.citigroup.net/gdse/products/platforms/nasfiler/proposed/07003-sg.pdf
DNS and NTP configuration standards	http://ipns.ny.smb.com/documents/Product-Development/CA-NTP/IGNAST/NTP_design_RFP.doc

8 General Design

This section explains the general design for the implementation of network-attached storage using filers from Network Appliance (NetApp). The diagram below provides different implementation examples using the FAS3000 series Filer for demonstration purposes.




Figure 8.1 - NetApp Filer Architecture Design

Filers and Clustering: NetApp filer clusters will be used in production with data being replicated to remote COB/Backup filer or filer cluster, where required by the business use case. Each cluster / filer will have locally attached storage. Data from multiple production clusters can be replicated to a single remote filer depending on capacity.

User Access: Users will directly access CIFS shares that are created on the clustered filers over the network. File access restrictions are configured using NTFS file permissions; these can be setup by system administrators as in the current Windows Network environment. Users can also mount NFS exports, where access restrictions are configured during the export process. These exports can be setup by system administrators as in the current Unix environment. NetApp also supports access by both NFS and CIFS users to the same data (see section 9.4.3.6).

Filer Fault Tolerance: Clustering is implemented so as to protect from the failure of one filer head (node). In the case of a node failure, the surviving cluster node takes over (fail-over) the resources of the failed/offline node. Each filer also has built in redundancy with dual power supplies, mirrored NVRAM (cluster only) and dual path connectivity to storage.

Networking and Fault Tolerance: Gigabit Ethernet will be used for better performance and throughput since more users will now be connecting to the filers. For fault tolerance, two gigabit Ethernet ports in each filer will be merged (trunked) in an active-standby role as well as connected to two separate network switches. This configuration will protect from the failure if either the gig-e NIC port, network switch/port or cable.

Data Protection: RAID-DP data parity protection, local Snapshots and remote SnapMirror all serve for optimal data availability and protection. Asynchronous data replication using SnapMirror enables tape backups to be executed at the remote location. This backup methodology provides the following benefits:

- Large amounts of data to be backed up with no impact to production users, and network performance
- Expands backup windows; allowing flexibility for increased storage capacities on production filers

Continuance of Business - COB: MultiStore (vfiler), SnapMirror and SnapVault software enable added flexibility for continuance-of-business and data migration. MultiStore allows creation of separate and completely private logical partitions in filer network and storage resources.

- Virtual filers with SnapMirror enable administrators to Move and migrate data quickly and easily with minimizing downtime.
- SnapMirror file-system asynchronous replication provides a point-in-time read-only copy of all replicated volumes/qtrees available at the remote location that can be brought on-line quickly in the event of a disaster.

Virus Scanning: Data ONTAP's "vscan" feature allows real-time on-access virus scanning of all incoming and outgoing files. GDSE and CTI Global Security Engineering has fully integrated and certified McAfee VirusScan for NetApp 7.1 On-Access scanning. Please refer to Antivirus Section later in this document.

9 Best Practices

This section of the document covers best practice configuration settings that relate to the use of ONTAP 7.1. For security related configuration information, refer to section 10.

9.1 Installation

The installation section describes the hardware platform requirements to support ONTAP 7.1, GDSE supported maximum raw storage configurations, upgrade prerequisites from earlier versions of ONTAP, and GDSE unsupported features or configurations.

9.1.1 Platform requirements

ONTAP 7.1 is not supported on the FAS700 platform. All other Citigroup certified hardware platforms will support ONTAP 7.1 (FAS200, F800, FAS900, FAS3000, R150, R200).

Note FAS3040, FAS3070 and FAS6000 series are not currently certified by GDSE and do not support ONTAP 7.1.

9.1.2 Scalability

As a general rule, GDSE supports NetApp Filer hardware configurations upto 75% of the maximum supported raw storage as listed in the product specifications on the <http://now.netapp.com> site. Implementations wishing to scale beyond this size will be reviewed on a case by case basis, and require engineering approval.

The 75% rule is also applied to the number of vFilers supported on a given hardware platform.

9.1.3 Upgrade prerequisites

Before upgrading your Filer, ensure the following pre-requisites are met:

9.1.3.1 System Firmware

Platform	Required system firmware version
R200	4.2.2 or above

Table 9.1 - System firmware prerequisites

There are no system firmware pre-requisites for the other NetApp NAS Filer hardware platforms.

9.1.3.2 Disk Shelf Firmware

Shelf type	Required disk-shelf firmware version
AT-FC and AT-FC2	02.08
AT-FCX	33.00
ESH	95.20
ESH2	18.50

Table 9.2 - Shelf firmware prerequisites

9.1.3.3 System State

- Ensure no flex volumes are in a WAFL inconsistent state prior to starting upgrade.
- For an up-to-date list of pre-requisites and limitations, please consult the [NetApp Release Notes page for ONTAP 7.1.2.1](#).

9.1.4 Upgrade procedure

For information on how to install or upgrade to Data ONTAP 7.1, refer to chapter 14. GDSE supports upgrading from ONTAP 6.5.4P4 or 7.0.4P9 to ONTAP 7.1. Refer to the Release Notes document for known issues with the upgrade process.

9.1.5 Unsupported configurations

GDSE has not certified, nor supports the use of the following configurations within ONTAP. This list is neither complete nor exhaustive but is designed to answer the regular support enquiries received for the NetApp Filer platform. If a certain configuration is not listed in this document (either as supported or unsupported) then please check with gdse@citigroup.com for confirmation.

- Database storage (excluding dump files)
- iSCSI protocol
- HTTP web serving protocol (except Filerview)
- FTP protocol
- SAN attached (except Tape SAN)
- User and group level quotas
- Synchronous SnapMirror (replication)
- Maintenance Center (due ONTAP 7.2)
- Mixed mode security style volumes (exception: vfiler pseudo root qtree)

9.2 ONTAP SOE configuration (options)

Refer to Appendix C: Data ONTAP SOE Configuration Standards for a complete list of the core options settings.

In addition to the values listed in Appendix C, ensure the registry value **options.admin.config.list** is defined as **/etc/registry.default,/etc/registry.local**. To check this value:

- set CLI to advanced privilege mode
- `registry get options.admin.config.list`

If the value is not defined as above, use the following command to configure this parameter:

```
registry set options.admin.config.list /etc/registry.default,
    /etc/registry.local
```

9.3 Network Configuration

9.3.1 Gigabit Ethernet NIC Configurations and Naming

- In a basic implementation 2 of the 4 Copper Gigabit Ethernet network interface ports will be used in each filer. More ports may be used. See section 9.3.2 for further information.
- The two gig-e ports should be configured (trunked) in an Active-Standby configuration for redundancy from both the filer and network switch perspective.
- The two connections from each filer **MUST** be connected to different network switches for redundancy in case of switch failure.
- Both nodes of the filer **MUST** be on the same network segment (subnet) since in a cluster fail-over scenario, one node has to takeover the name and IP address of the other node.
- When configuring the virtual interfaces, single-mode trunks should be used (active-passive).

Where Filers have onboard Ethernet ports (FAS3000), GDSE recommends single mode trunking ports e0a and e0c for Gigabit Ethernet connections in order to leverage the redundant buses on the system board.

- Interface e0, the management 10/100 Ethernet port is not used. By default it should be shut down, to prevent any possible future routing issues.

When creating the virtual interfaces (vif) on the filers, the following naming standard should be followed. In a cluster configuration, the vif's should have different names within the cluster pair. This is because in a fail-over, one of the nodes assumes the name for both nodes and thus different names have to be used for the fail-over to work. For all Filer models use the default naming convention for the Ethernet ports as shown in NetApps documentation.

All Gigabit Ethernet Virtual interfaces should be named gev#, where the first virtual interface is number one (gev1). Every additional virtual interface (cluster or not) should use the next sequential number. Vif's should be created in single-mode as active-standby configuration. In a cluster configuration all GigE virtual interfaces in the cluster pair **MUST** be unique.

Below is a sample command line for vif creation.

```
vif create single gev1 e0a e0c
```

9.3.2 Quad Port Gigabit / Ethernet Configurations and Naming

The following points should be noted when considering implementing multi-mode or aggregated IP network ports:

- NetApp Filers do not support Cisco's Port Aggregation Protocol
- NetApp Filers do not support Link Aggregation Protocol (dynamic 802.3ad)
- The multi-mode vif implementation in ONTAP is in compliance with IEEE 802.3ad (static)
- Operational owners should contact their local network support teams to determine what network configuration steps are required to set up a 'fixed' or static 'on-mode' EtherChannel on a server access switch.

The following standards should be followed when implementing multi-mode or aggregated IP network ports:

- All four Ethernet ports will be used in filers where Gigabit Ethernet is not available or where greater aggregate network bandwidth is required than offered by a single Gigabit Ethernet connection
- Each of the bundles from the filer **MUST** be connected to different network switches for redundancy in case of switch failure.
- Both nodes of the filer **MUST** be on the same network segment since in a cluster fail-over scenario; one node has to takeover the name and IP address of the other node.
- Filers that utilize the four Ethernet ports will have 3 virtual interfaces each. Two virtual interfaces will be created for each of the 2-port bundles using Multiple-mode (All active). All interfaces in a multimode VIF **MUST** be connected to the same network switch. Thereafter, another virtual interface will be created in Single-mode (Active-Passive) containing the previous 2 virtual interfaces. The result is that only one network interface (utilizing all 2 ports on it) will be active at one time.

Where Filers have onboard Ethernet ports (FAS3000), GDSE recommends multi-mode trunking ports e0a and e0b, and, e0c and e0d respectively; then single-mode trunking of the two multi-mode pairs, for quad Ethernet/Gigabit connections in order to leverage the redundant data paths on the system board.

- When increasing the number of active ports in a multi-mode interface, in order to match Citigroup Cisco network standards, the amount of new ports needs to be in the order of two. For example if there are two ports in the current multi-mode interface and aggregate bandwidth needs to be increased, then two further ports must be added (4 active ports). An odd number of active ports in a multi-mode interface are not supported.

When creating the virtual interfaces (vif) on the filers, the following naming standard should be followed. In a cluster configuration, the vif's should be named different within the cluster pair. This is because in a fail-over, one of the nodes assumes the name for both nodes and thus different names have to be used for the fail-over to work. On Filers with on-board Ethernet physical interfaces, the default port naming should be used (e0a, e0b, e0c and e0d). The letter at the end designates each port on system board. All 10/100Mb Ethernet Virtual interfaces should be named etv# and all Gigabit Ethernet Virtual interfaces should be named gev#. Every additional virtual interface (cluster or not) should use the next sequential number. In a cluster configuration all Ethernet virtual interfaces in the cluster pair MUST be unique.

Below is a sample command line for vif creation.

```
vif create multiple etv1 e0a e0b
vif create multiple etv2 e0c e0d
vif create single etv3 etv1 etv2
```

9.3.3 Remote LAN Management (RLM) and naming

The Remote LAN Management port is a management interface introduced with new Filer models, starting with the FAS3000 series. GDSE evaluated this component and determined the best practice implementation would be as a “management interface of last resort”. The RLM exhibits the following functionality:

- SSH is the only supported management protocol
- Only local administrators on the Filer who are members of the local Administrators group are allowed to logon via RLM.

Once RBAC is fully implemented, the only account that will have the rights to login to the RLM will be root (see below for special case in using root).

- Where an operations management by-pass process requires the use of the **root** account, the device local **naroot** proxy credentials must be used for SSH authentication to the RLM otherwise authentication will fail.
- RLM does not support the system options of `autologout.console.enable` nor `autologout.telnet.enable`, instead upon successful authentication an administrator can choose to logout an existing active session or logoff without completing the new logon attempt to the RLM interface.
- RLM supports connection restrictions by leveraging the settings in the `ssh.access` option (see section 9.1 for `ssh.access` configuration standard). Therefore SSH access to the RLM card will be restricted to the IP addresses or host names defined within the `ssh.access` system option.

GDSE recommends the following best practices when implementing and using the RLM interface:

- Where the local network infrastructure permits, the RLM interface should be on a separate network segment to the console and data network interfaces in order to provide access to the Filer where a server farm switch has failed and management is not possible via any other interface.
- The DNS name for the interface must be the physical filer host name followed by the characters ‘-rlm’. Therefore a Filer called **nasny01** would have RLM interface named **nasny01-rlm**.

- The RLM port should not be used for day-to-day management processes, but as an interface of last resort when regular SSH, FilerView and Console access cannot be achieved to the Filer. This is especially useful when the Filer is in a remote location, the console has hung and the system requires a reboot.
- Performing a power cycle from the **RLM**

Where a Filer is still responsive to console commands, and a power cycle is required to clear an error condition, GDSE recommends halting the Filer before performing a *system power off* or *system power cycle* command as these RLM commands can result in a 'dirty' shutdown of the system. A graceful termination of processes is always desirable so, where possible, the halt command should be used first.

9.3.4 Remaining Hardware Conventions

For all other cabling requirements and slot assignments on all Citigroup supported Filer models, storage administrators should follow NetApp's hardware configuration guides, available from <http://now.netapp.com>

9.3.5 Filer Naming Standards

9.3.5.1 Primary (Production) Site

NAS<2 or 3 letter location code><2 digit sequential number>

Example: NASNY01 or NASLDN01.

NAS stands for Network Attaches Storage. In a cluster pair, the nodes should be numbered on sequence i.e. 1 & 2, 3 & 4 etc. would be two cluster pairs.

9.3.5.2 Backup / Disaster Recovery Site

To distinguish primary filers from disaster recovery filers, "DR" will be appended at the end of the above naming convention. Therefore, the names will be:

NAS<2 or 3 letter location code><2 digit sequential number>DR

Example: NASRUT01DR.

9.3.6 Domain Name and Host Resolution

Like UNIX and Windows servers, Data ONTAP resolves host names by searching maps (or databases) for services to use. The Filer tries name resolution services in default order or in the order specified in the `/etc/nsswitch.conf` file. For example, by default:

- 1) Initially, the filer attempts to resolve host names locally by searching the `/etc/hosts` file (designated by "files" in the table below and in the `/etc/nsswitch.conf` file)
- 2) If name is not resolved, the filer tries NIS (if NIS is enabled)
- 3) If name is still not resolved, the filer requests services from a DNS server (if DNS is enabled)

Data ONTAP Default Name Resolution Order	
Map	Services in each Map
hosts	files NIS DNS
Passwd	files NIS
netgroups	files NIS
group	files NIS
shadow	files NIS

Table 9.3 - Default Name Resolution Order

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

- During the initial filer setup the operator will be asked to enter the Domain and DNS server information, and NIS domain and server information. **Filers must be configured to use the global DNS settings as described in [IP Name Services Engineering NTP and DNS standards document](#). 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 Filer 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.**
- During CIFS setup the operator will be asked to enter WINS servers for netbios name resolution, as a normal Windows NT server would.

9.3.7 Time Synchronization and Time Zones

9.3.7.1 NTP

Time synchronization needs to be configured for the filers 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 Filers 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.

If the difference in time between the filer and the time source is:

- Greater than the max_skew, time has to be set manually by the filer administrator using the “date” command.
- Less than the max_skew, but more than the min_skew, time is gradually converged to the correct time at every scheduled update.
- Less than the min_skew, the filer will not attempt to modify the time.

The following commands need to be run on all filers.

```
options timed.servers 192.193.215.65,192.193.215.69,192.193.215.73
options timed.enable on
```

Refer to Appendix C of [IP Name Services Engineering NTP and DNS standards document](#) for the order in which these IP addresses should be configured at each site.

If the site at which the Filer is installed is not listed in the aforementioned document, then use the default order for NTP servers as shown in the above example.

All regions must conform to the time source servers identified above; noting that changes to the order of IP addresses may be required depending on the installation site.

The following time related settings should also be used.

Command	Description
<code>options timed.enable on</code>	Enable NTP daemon
<code>options timed.log on</code>	NTP events recorded in message log
<code>options timed.max_skew 30m</code>	NTP correct clock up to 30 mins off
<code>options timed.min_skew 50</code>	NTP only correct clock if more than 50ms
<code>options timed.proto ntp</code>	Use the NTP protocol to manage time
<code>options timed.sched 1h</code>	Check system time every 1 hour

Table 9.4 - NTP recommended settings

9.3.7.2 TimeZones

GDSE recommends that the local timezone be used for the geographic location at which the Filer is installed. For example, Filers in California should be configured PST8PDT. If the timezone in which the Filer is located supports Daylight Savings Time (DST), then the Filer should be configured with the timezone information that supports DST. Contact gdse@citigroup.com if you are unsure which timezone value to use for your Filer.

To check your current timezone configuration, type `timezone` at the command line.

To update your timezone configuration, type `timezone zone_name` at the command line (where `zone_name` is the name of your timezone. For example PST8PDT is the timezone name for California).

9.4 Storage Configuration

9.4.1 Data ONTAP Clustering

A cluster consists of a pair of NetApp filers that are connected and configured in a special way to provide fault tolerance.

- NetApp Data ONTAP clustering is Active/Active from the perspective of the client accessing the system as both nodes are simultaneously serving data. When a node is triggered offline or fails, the partner node initiates a take-over the failed node's services and disks, and continues providing all file and block access protocols to clients and applications. However, due to the nature of this failover process and the absence of a secondary active path to the partner's disks, NetApp Data ONTAP clustering is actually considered an Active/Passive solution, where a particular CIFS share or NFS export is only active on one filer node at a time.
- The clustering technology used allows the nodes use each other's NVRAM (Operational Memory), where all system and operational information is kept as a backup space, in case of a failure. Using the NVRAM information, upon failure the active node creates a virtual filer. The virtual filer appears to users and administrators like a functional version of the failed filer.
- Once the failed node is operational again, the "*cf giveback*", command on the active node fails-back all services and disks back to the restored node.
- All cluster takeover and fail-back normally occurs with no user impact.

Clustering is enabled by executing the following command on one of the cluster nodes:

```
cf enable
```

9.4.2 Aggregate Configuration and Sizing

9.4.2.1 Aggregates and Flex-volumes versus Traditional Volumes

One of the key new features in ONTAP 7.0 is the introduction of flexible volume (flex-vol) and aggregate technology. Using the original NetApp volume type, now referred to as a traditional volume (trad-vol), a group of disks were brought together as a volume and assigned a RAID level. There was a direct relationship between a disk and the volume in which it was a member. With the introduction of aggregates and flex-vols a layer of abstraction is introduced between the physical disks and the logical volume. As such, disks are assigned to an aggregate, and the volumes are carved from the total space available in the aggregate. The RAID level is configured at the aggregate level. The volume itself is striped across all the disks in the aggregate.

GDSE strongly recommends implementing flex-vols to increase storage utilization across NetApp NAS platforms, and to migrate existing systems from traditional volumes to flexible volumes. The following tools can assist with the migration from traditional to flexible volumes: NetApp's snapvault and ndmcopy utilities and Microsoft's Robocopy utility (CIFS only).

9.4.2.1.1 When to choose traditional volumes over flexible volumes

Traditional volumes can still be considered for root volumes, as essentially root volumes are a two disk mirrored pair, containing the system volume. Traditional volumes should no longer be considered for storing business data, as they do not offer an efficient store for the data. GDSE recommends using flex-vols for all business data volumes, and suggests that flex-vols would also be a good fit for the root volume. Refer to the following section for considerations when choosing traditional volumes over flexible volumes for the root volume.

9.4.2.1.2 Considerations for using an aggregate to house the root volume

In deciding whether an aggregate, and thereby a flex-vol, should be used to house the root volume, the following points should be considered regarding additional uses of the root aggregate (the aggregate containing the root volume):

- The root aggregate must be separate from other aggregates on the Filer that contain business data. The exception to this rule is the FAS270C that only has a small number of drives available. Therefore the FAS270C can have both business data and the root volume on the same aggregate.
- The flexible root volume must have the space guarantee set to the volume value.
- The flexible root volume cannot be shrunk to a value less than the pre-determined minimum volume size for the given Filer platform. Any attempt to shrink a volume smaller than this value will generate a warning message and abort the root volume resizes. GDSE recommends setting the root volume size at triple the minimum size valued provided by NetApp to allow adequate space for core dump files, and other verbose log files.

Filer Model	NetApp Minimum Root Volume Size	GDSE Recommended Root Volume Size
FAS270	10GB	30GB
FAS920	12GB	36GB
FAS940	14GB	42GB
FAS960	19GB	57GB
FAS3020	12GB	36GB
FAS3050	16GB	48GB
F820/825	10GB	30GB
F840	13GB	39GB
F880	13GB	39GB
R200	19GB	57GB

Table 9.5 - Recommended root volume size when implementing flexible volumes

- Additional space on the root aggregate could be used to create a second flexible volume for uses such as a restore volume, or a separate meta-data volume for each MultiStore (Vfiler) instance. In all instances, care should be taken to ensure there is enough space on the root aggregate for these additional requirements.

9.4.2.1.3 Common misconceptions regarding aggregates

The following misconceptions are often made regarding aggregates:

- *Aggregates can be shared using CIFS or exported using NFS.* Only the volume level or below contains a file system. The aggregate is a physical collection of storage resources (disks) as opposed to a logical file system such as a volume.
- *Snapmirror replication occurs at the aggregate level.* In ONTAP 7.0, snapmirror replication continues to be supported at the volume or qtree level only.

- *Traditional volumes can be converted to flexible volumes.* The data must be moved off the existing traditional volume, the volume destroyed, a new aggregate and subsequent flex-vol created, then the data migrated back in order to achieve the migration from traditional to flexible volumes.

9.4.2.2 Aggregate Naming Standards

The **root aggregate**, that is the aggregate containing the root volume, should be named **aggr0**. All subsequent aggregates should be named sequentially, aggr1, aggr2 and so on. For implementations where a traditional volume has been used as a root volume, such that a root aggregate does not exist, then the first data aggregate will be named aggr1.

Aggregate names should not include business codes or any other type of descriptive naming convention because business ownership of the data will be at the volume level and not the aggregate level.

9.4.2.3 Aggregate Sizing Standards

For implementations where the root volume is to be housed on an aggregate, then the aggregate should comprise of the first two disks in the first shelf in the system, configured as a RAID4 group of two disks. See the notes in section 9.4.2.1.2 with regards to the size of the root volume.

For aggregates which will contain business data, the sizing standards are based on the underlying RAID group configuration to be used in the aggregate. Use the following table to determine the RAID group size and the recommended number of aggregates. This is based upon the model of Filer and the type of drives (fiber channel or SATA) you are using.

Drive Type	Filer Model	Recommended RAID Type	Max # of disks per RAID Group	Max # of RAID Groups per Aggregate
Fiber Channel	F800	RAID-DP	18	3
	FAS270C	See note below		
	FAS900	RAID-DP	18	3
	FAS3000	RAID-DP	18	3
(S)ATA	R200	RAID-DP	14	3
	FAS3000	RAID-DP	14	3

Table 9.6 - Recommended RAID group configuration

The exception to the above table is the FAS270C, due to the limited number of disks available within the system from the small form factor design, and the ability to expand upon the initial small level of storage with additional shelves. As the FAS270C contains both nodes of the cluster in the single shelf, there are only seven disks available to each node, with one of these disks reserved as a hot spare to each node. Therefore the standardizing model for FAS270C depends on the number of shelves attached to the system:

# Shelves	RAID Type	# Disks per RAID group
1	RAID-DP	6
2	RAID-DP	13
3	RAID-DP	20

Table 9.7 - Recommended RAID configuration for FAS270

The figure for the maximum number of RAID groups per aggregate in the above table has been determined by GDSE, and is not a limitation of the ONTAP operating system or a

recommendation from NetApp. The GDSE recommended number of RAID groups in an aggregate protects the business data on the system by limiting the amount of storage that could be impacted by an aggregate related issue; while balancing the I/O load for a given volume over all the disks in the aggregate. Thereby providing more spindles to serve data and hence increasing back end I/O performance.

To prevent unnecessary fragmentation of the file systems, GDSE does not recommend allocating aggregates at over 90% utilization rates.

9.4.2.3.1 Aggregate Snapshots / Snap Reserve

By default an aggregate retains 5% of available space as an aggregate level snapshot reserve. If the aggregate is 1TB in size, then the snapshot reserve will be 50GB, leaving 0.95TB as usable. As a default rule NetApp recommends maintaining this 5% snap reserve level, but after extended consultation GDSE were able to determine the instances when an administrator may wish to leverage the features of the aggregate snap reserve.

As such GDSE will not standardize on the value of the snap reserve (except to say that the value should not be increased above 5%), but instead provide operational owners with a list of situations where aggregate snap reserves may be beneficial, to allow for an informed decision to be reached. Obviously reducing the aggregate snap reserve to a value of 0% has the immediate benefit of making more usable space available. However the following instances are situations where the snap reserve will be useful:

- RAID SyncMirror and MetroCluster (both of these components are not certified by GDSE for use in Citigroup)
- Where WAFL needs to perform a previous consistency pointer check
- Where the Filer is rebooted to roll back to a previous aggregate consistency pointer
- To provide an additional layer of reassurance when conducting potential data destructive commands at the aggregate level by providing an aggregate level snap.

When considering the aforementioned situations, operational owners also need to factor in what actions would be taken in the event of invoking contingency, and whether the ability to perform a full/partial aggregate restore would improve the recovery time objective to return to production as normal status. After assessing the situations where the aggregate snap reserve may be utilized, if operational owners deem that the reserve space is not required in their environment, then the reserve can be removed by issuing the following command.

Command	Explanation	Output
<code>df -A</code>	Show aggregate space	<pre> Kbytes used avail capacity aggr0 119052776 94411920 24640856 79% aggr0/.snapshot 6265932 326252 5939680 5% aggr1 2500108212 157289036 2342819176 6% aggr1/.snapshot 131584640 1752 131582888 0% </pre>
<code>snap reserve -A aggr1 0</code>	Reduce aggregate snap reserve of aggr1 to zero	<pre> Deleting snapshot 'hourly.0' in aggregate 'aggr1' to recover storage Deleting snapshot 'hourly.1' in aggregate 'aggr1' to recover storage Deleting snapshot 'hourly.2' in aggregate 'aggr1' to recover storage Deleting snapshot 'nightly.0' in aggregate 'aggr1' to recover storage </pre>
<code>df -A</code>	Show aggregate space	<pre> Kbytes used avail capacity aggr0 119052776 94411920 24640856 79% aggr0/.snapshot 6265932 326252 5939680 5% aggr1 2631692852 157290788 2474402064 6% aggr1/.snapshot 0 1752 0 ---% </pre>

Table 9.8 - Procedure to remove aggregate snapshot reserve

9.4.2.3.2 Comparing Aggregate Snapshots and Volume Snapshots

Aggregate snapshots can be used where implementation requirements meet criteria outlined in the previous section. Aggregate snapshots are only accessible by storage administrators from the Filer console, as compared to volume snapshots, which are utilized at the business aligned support or user level by mounting to the snapshot directory on a volume, in order to restore a file or group of files from the file system snapshot.

The aggregate snapshot and volume snapshot are separate disk-based point in time copies, representing different levels of file system objects within the Filer.

Aggregate snapshot space reserves can be reinstated, or reverted to the default snapshot reserve in the case of volume snapshots, if there is enough space on the respective aggregate or volume to create/increase the reserve.

9.4.2.4 Aggregate Creation

The example in section 18.1 shows how to create an aggregate using the NA_Admin FilerView GUI. Please note the example is using 206GB SATA drives on a FAS3000.

Even though disks can be chosen manually for aggregate creation, it is best to use the *recommended* disks in the contiguous order in which they are listed. This is best for management of disk errors. The figures below show the foolproof way to create volumes using the “na_admin” GUI utility.

9.4.2.5 Aggregate Expansion

Data ONTAP aggregates can be expanded while on-line without any user access impact, and minimal performance impact. When disks are added to an existing aggregate, the filer will attempt to keep the amount of data striped across each disk about equal. For example, if you have 14 data disks in a RAID group that is part of an aggregate, you can expand the size of the RAID group using the ‘aggr add’ command.

Storage administrators wishing to expand the capacity of an aggregate should consult the NetApp Data ONTAP documentation, available as part of the ONTAP operating system.

9.4.3 Volume Configuration and Sizing

9.4.3.1 Naming convention

9.4.3.1.1 Primary Filer

Since we anticipate that multiple departments will probably be sharing filers, we are recommending the naming of volumes by department. We are also recommending that the volume name be kept generic and not to include the purpose of the volume (like apps or data etc.). We have also taken into consideration the fact that one department could have multiple volumes on the same filer.

<Abbreviated Department Name>_V<One digit sequential number>

Examples of the above are EQTY_V1, EQTY_V2, FIR_V1 etc.

9.4.3.1.2 Backup / Disaster Recovery Filer

Since data replication is being done in either a many to one or one to one fashion it is important to be able to distinguish the source filer (and volume) by looking at the destination (snapmirror

copy) volume name on the DR filer. This becomes even more important because backups are also being performed on the mirrored copy of the primary volume.

<Source Filer Name>_<Source Volume Name>DR

As an example, if NASNY01 is the primary filer and EQTY_V1 is the primary volume, the Snapmirror volume name on the disaster recovery filer (NASRUT01DR in our example) will be NASNY01_EQTY_V1DR.

9.4.3.2 Volume Creation

A volume is a logical file system on a Filer striped across all the disks in the aggregate. The example in section 18.2 shows how to create a volume using the NA_Admin FilerView GUI

- A single filer can support up to 100 flex-vols, and a cluster can support up to 200 flex-vols. During a cluster failover situation a single node in a cluster can support up to 200 flex-vols.
- A single filer can support up to 100 aggregates. Note that where traditional volumes are used these will count as both an aggregate and a volume.
- Unless being used for NFS data, volumes should be configured with the NTFS security style using the following command: `qtree security /vol/<vol_name> ntfs` NTFS should be used as the default security style on vol0 of the physical filer.
- Volumes should be configured to the appropriate language standard. For example, Filers in North America should have each volume configured with the language type en_US. This is implemented via the following command: `vol lang <vol_name> en_US`. The volume language value will vary from region to region and depends upon local requirements.

- Volumes over 90% full will begin to suffer noticeable performance degradation. As such GDSE recommends restricting the utilization of volumes to a value less than 90%
- When creating volumes, including the root volume (vol 0), ensure the create_ucose and convert_ucose are options are enabled on each volume. This achieved with the following commands:
`vol options <vol_name> convert_ucose on`
`vol options <vol_name> create_ucose on`

9.4.3.3 Volume Sizing

Filers today employ 72GB, 144GB and 300GB Fiber Channel disk drives, all certified for use.

For CIFS/NTFS and NFS file system protocols the maximum volume size is limited to the maximum file size the OS can write. In both cases the maximum file size is 2TB.

When creating a new volume, one must take into account the requirements for backup and COB in determining the optimal size for a volume.

9.4.3.4 Volume Resizing

Data ONTAP flex-vols can be resized while on-line without any user access impact, and minimal performance impact. Flex-vols can either be grown or shrunk according to the size of the parent aggregate and the quantity and type of space reservations (guarantees) made by other flex-vols in the same aggregate. You can change the size of the flex-vol using the 'vol size' command.

Storage administrators wishing to change the size of a volume should consult the NetApp Data ONTAP documentation, available as part of the ONTAP operating system.

9.4.3.5 Volume Cloning

Data ONTAP flex-vols can be cloned while on-line without any user access impact, and minimal performance impact. Cloning involves creating a set of pointers to the original volume and only saving updated blocks within the clone volume itself. This method of only saving the changed blocks between the original source volume and the clone makes for an extremely efficient way to perform multiple update scenarios on the same piece of data, and as such will probably be very beneficial to programmers and developers. You can create a clone of a flex-vol using the 'vol clone' command.

In addition a clone can be split such that the original image blocks are copied into the clone, and merged with the changed blocks to form a new and independent volume.

Clones exhibit the following behavior:

- Are writeable copies
- Cannot be refreshed with changes from the source volume
- Cannot be reestablished with the original source volume once split

Storage administrators wishing to clone a volume should consult the NetApp Data ONTAP documentation, available as part of the ONTAP operating system.

9.4.3.6 Multi-protocol Volumes

Data ONTAP supports a mixed mode on volumes, whereby either NFS or CIFS Access Control Lists are implemented, and files can moved between either type of ACL. However, both NetApp and GDSE do not recommend implementing mixed mode volumes, as the ACL of the last user to access a file will overwrite the existing ACL, if that ACL was created or modified using the opposing file access protocol. As there is no direct functionality mapping between an NFS ACL and a CIFS ACL, this can lead to a denial of access to the file object when an ACL is overwritten. Therefore GDSE does not support mixed mode volumes on a NetApp Filer.

Instead it is possible to allow access to a single file system from both Unix and Windows clients, by simultaneously exporting and sharing the file system. Data ONTAP will natively perform the UID to CIFS or NT Name to UNIX translation.

The prerequisites for this implementation are listed below:

- CIFS license for Filer
- NFS license for Filer
- Option `wafn.nt_admin_priv_map_to_root` **on**
- NIS+ server set to NIS compatibility mode (ONTAP does not support NIS+)
- Add filer (or VFile) owning the volume to Active Directory domain
- Add filer (or VFile) owning the volume to NIS domain
- CIFS User ID must equal Unix ID otherwise explicit translation required

The process for translating a NFS request for a file with NTFS security is shown below:

```
Lookup numeric UID in /etc/passwd or NIS
  yielding UNAME
Lookup UNAME in /etc/usermap.cfg
  yielding NTDOMAIN and NTNAME
If NTNAME is null ("")
  access is denied
If UNAME is not matched in usermap.cfg
  set NTNAME equal to UNAME
Lookup NTNAME in NTDOMAIN yielding
  NT_SID + GROUP_MEMBERSHIP_SIDS
If NTNAME not found set NTNAME to wafn.default_nt_user
```

The process for translating a CIFS request when a user logs in is shown below:

```
Lookup NTDOMAIN+NTNAME in usermap.cfg
yielding UNAME
If UNAME is null ("")
    access is denied
If NTDOMAIN+NTNAME is not matched in usermap.cfg
    set UNAME to lowercased NTNAME
Lookup UNAME in UNIX password database
yielding UID + GROUP_IDS
if UNAME not found
    set UID to wafl.default_unix_user
If wafl.nt_priv_map_to_root option is set AND
    NT_GROUPS includes BUILTIN\Administrator
    set UID to 0 (root)
```

GDSE recommends the following settings when implementing this type of volume:

- Update the last entry in the /etc/usermap.cfg file to disallow NT logins in the name of root from all domains
`*\root => ""`
- If NFS user's equivalent CIFS User ID is in a different Windows domain than the Filer (or Vfiler), ensure that the relevant domain appears in the **cifs.search_domains** option
- If granting application access to the file system, whereby a specific functional ID is used to access data and that functional ID is different in the Windows and Unix domains, perform the user translation by adding the appropriate entry to the /etc/usermap.cfg file
- For NTFS implementations, ensure that the Unix mount path is not restricted by an NTFS ACL to avoid a failed mount condition

The business requirements in each case will determine which type of file system is best suited. This will depend on factors such as whether a Windows or Unix support group is administering the data; and the predominant access vector to the data (for example is the majority of the user population in the Windows or Unix environment).

As stated above, the majority user environment should be considered in each case to avoid the Filer having to perform any unnecessary user account translation (thereby reducing the requirement for system resources). There are no benchmarks available at this time to help in determining which file system type to use when balancing conflicts between the support organization and the majority user access method requirements (for example a conflict where the SA group is a Unix team, but 75% of the user access is from a Windows environment).

9.4.4 Spare Disks

Spare disks are global. Spare disks can be used for both the creation of new volumes as well as replacement of failed drives. Always ensure that there is at least one spare disk available on each filer at all times to guard against disk failure.

- GDSE recommend allocating at least 1 spare disk (for failure protection) for each node in a cluster.
- Hot spare drives are not disk shelf dependant. Upon disk failure, spares disks can be allocated from any shelf that is "owned" by the filer.
- When a spare disk needs to be allocated for a failed one, the disk that is on the top of the "Volumes | Report | Show Spare Disks" list is used.

- Any spare disk in the shelf can perform the function of a hot-spare disk and does not have to be pre-allocated. Spare disk configuration is performed on the fly and no restarts are required.
- Ensure the spare disks on the system are zeroed. Use `vol status -s` to list the spare disks on the system. Look for the term “(not zeroed)” at the end of each entry. If a spare disk is listed as not zeroed, GDSE recommends using the `disk zero spares` command to prepare the spare disk for use prior to a disk failure event on the Filer; otherwise the spare disk will zero itself once another disk fails, and thereby delay the rebuild and spin up time for the disk.

9.4.5 Snapshots (Point-in-time copy)

A snapshot is a point in time read-only image of the entire Data ONTAP file system that reflects the state of the file system at the time the snapshot is created.

Snapshot files carry the same permissions as the original file. A user who has permission to read a file in the active file system can still read that file in the snapshot. A user without read permission to the same file cannot read that file in a snapshot.

Write permission does not apply because snapshots are read-only. Any user with the appropriate permissions can access snapshots.

Every directory in the filer's active file system contains a directory named `~snapshot`, through which users can access old versions of files in that directory. This allows the recovery of files without having to go to tape backup except for archival purposes.

9.4.5.1 Snapshot configuration

The frequency of snapshots and the number of snapshots to keep before overwriting can be set when the volume is created. As of Data ONTAP v6.4 *one volume can have up to a maximum of 255 snapshots*, however **GDSE does not recommend exceeding 250 scheduled snapshots**. The following are the suggested settings, and can be modified as needed to provide increased storage utilization.

Snapshot Type	Schedule	Retention	Comments
Hourly Snapshots	Every day at 12:00 Noon and 5:00PM	2	These hourly snapshots provide the ability to recover files during the current working day. The snapshots will be labeled <code>hourly.0</code> and <code>hourly.1</code> . These snapshots will be recycled every day.
Nightly Snapshots	Monday through Saturday at 12:00AM	22	This retention in combination with the weekly snapshots, will provide 25 days worth of data to be available on-line for any data recovery. The snapshots will be labeled <code>nightly.0</code> to <code>nightly.21</code> . After the 22 nd snapshot, it will start deleting the oldest snapshot when new ones are created. NOTE: When new snapshots are created, the value of <code>n</code> in <code>nightly.n</code> is adjusted such that the higher the value of <code>n</code> , the OLDER the snapshot. The daily snapshots run at a pre-determined time, which cannot be changed.
Weekly Snapshots	Every Week on Sunday 12:00AM	3	The weekly snapshots run at a pre-determined time, which cannot be changed.

Table 9.9 - Suggested Snapshot settings

The snapshot in ~snapshot\hourly.0 on the DR filer will be used for backup (incremental or full). As mentioned above, hourly.0 is always the latest hourly snapshot since it is numbered the lowest. See backup section (9.8.4) for details.

9.4.5.2 Snapshot Reserve

You can specify the percentage of disk space that snapshots can occupy. The default setting for this amount is **20%** of the total (both used and unused) space on the disk. However, for installations that consistently fail to utilize this 20% reserve, the 'snap reserve <vol-name> %' command can be used to reduce the snapshot reserve on a per volume basis, therefore providing improved storage utilization. The snapshot reserve value should be decreased in fixed amounts of 5%, to a value no less than 5% (for example 20 to 15, 15 to 10).

9.4.5.3 Snapshot Directory

All snapshots are stored in the ~snapshot directory on the root of every volume that has snapshots enabled. The ~snapshot directories on the filer volumes should be configured as *hidden but accessible*.

When the snapshot directory is not visible, users or administrators can still access snapshot data manually either from:

- DOS command line (cd ~snapshot)
- Connecting to the directory via [\\servername\volume_share_name\~snapshot](#)
- Sharing out the ~snapshot directory and connecting directly to the share \\servername\snapshot_share_name.

When the snapshot directory is made visible, it can be viewed using explorer by enabling the 'show hidden files and folders' option.

Filer configuration for enabling/disabling the snapshot directory visibility using Filerview is under CIFS | Configure | General Setup | Show Snapshot. The command line enabling/disabling the visibility of the snapshot directory is

`options cifs.show_snapshot on`

The snapshot directory can be made accessible or not accessible volume-by-volume, on an as needed basis, using the following console command on the filer:

`vol options <volume name> nosnapdir off`

However, when the snapshot directory is made inaccessible using the above command, the snapshot directory will not be visible using the previous specified option (cifs.show_snapshot).

9.4.5.4 Snapshot Permissions

When making the snapshot directory visible to users, consideration must be taken for all snapshots for the entire volume will be displayed, though file access controls still apply. Snapshot files carry the *same* permissions as the original file. A user who has permission to read a file in the active file system can read that file in a snapshot. A user without read permission to the same file cannot read that file in a snapshot.

Write permission does not apply because snapshots are read-only.

9.4.5.5 Snapshot 'file-folding'

File folding is the process of checking the data in the most recent snapshot, and, if it is identical to the snapshot currently being created, just referencing the previous snapshot instead of taking up disk space writing the same data in the new snapshot.

Disk space is saved by sharing unchanged file blocks between the active version of the file and the version of the file in the latest snapshot, if any. The filer must compare block contents when folding a file, so there is a tradeoff between performance and space utilization.

- Performance trade-off will be minimal compared to disk space saved.
- This option is not available to NFS users in this release.

"snapshot_file_folding" is enabled with the options command.

`options cifs.snapshot_file_folding.enable on`

9.4.6 Snapmirror (Asynchronous Data Replication)

SnapMirror replicates data from one volume (the source) to another volume (the mirror) and periodically updates the mirror to reflect incremental changes on the source. The result of this process is an on-line, *read-only* volume on the disaster recovery site that contains the same data as the source volume at the time of the most recent update.

The first time a snapmirror pair is created; a full volume copy is initiated. Only the initial process requires the destination (DR) volume to be offline. Once the initial copy is done, the destination volume automatically comes back online. The incremental copy works with both volumes online. The source and destination volume naming standards are as noted in the volume naming standards section above.

For replication to work it must be enabled on both the source and destination Filers. This is achieved by using the command `snapmirror on`.

9.4.6.1 Snapmirror Replication Frequency

Snapmirror is an asynchronous copy of data, not real-time. Below is the recommended replication schedule. The recommendations have taken into consideration the maximum number of concurrent replications, the latency required for the data transfer to be completed and avoidance of 'stacking' more than 8 replications at the same time.

NOTE: The maximum number of snapmirror entries supported by ONTAP in the `/etc/snapmirror.conf` file is 600. The model and load of the Filer in question determines the maximum number of concurrent snapmirror operations. Operational owners should consult NetApp's documentation to confirm the number of concurrent snapmirror operations supported by each device.

Snapmirror volumes on the filers should be divided into groups of 6 volumes each. Snapmirror replication frequency for volumes within each group will be **EVERY 10 MINUTES**.

The groups will be staggered to run at 2-minute intervals.

Volumes should be evenly distributed across the groups i.e. 1st to 5th volumes would be one each in groups 1 to 5, 6th volume would be in group 1, 7th in group 2 etc. Careful management, tracking and documentation of volume replication will be required.

When configuring Snapmirror entries, a **custom schedule** has to be used to achieve the above goals. The appropriate custom schedules to be used during configuration are provided in the table below as well as a sample Filerview screen shot. Configuration changes made in the

Filerview interface are reflected in the `/etc/snapmirror.conf` file on the filer. Verification that these settings have been written to the file is recommended. Note that snapmirror schedule configurations are always made at the *destination* filer.

Note, to avoid conflicts between the snapmirror and snapshot operations, GDSE recommends that the first replication group is offset by one minute so that the first snapmirror event does not happen at the same time of a snapshot event.

Replication Group	Custom Schedule	Minute Replication will Run
Group 1	1-59/10 * * *	1, 11, 21, 31, 41, 51
Group 2	3-59/10 * * *	3, 13, 23, 33, 43, 53
Group 3	5-59/10 * * *	5, 15, 25, 35, 45, 55
Group 4	7-59/10 * * *	7, 17, 27, 37, 47, 57
Group 5	9-59/10 * * *	9, 19, 29, 39, 49, 59

Table 9.10 - Snapshot standard schedule

The Snapmirror command also supports a throttling argument in order to limit the amount of network bandwidth used by the snapmirror process. This is particularly beneficial when production and contingency sites are separated by a narrow link. This throttling is achieved by using the `-k` option on the snapmirror command to specify the maximum number of kilobytes per second that can be utilized by the replication process. Note that setting this throttle does not guarantee that the snapmirror replication will run at this maximum rate, it merely implements a cap on the network utilization of the process.

9.5 File System Configuration

WAFL (Write Anywhere File Layout) is the file system used in all NetApp file servers. WAFL was designed to meet four primary requirements:

- It provides fast file access (NFS, CIFS) services.
- It supports large file systems (tens of GB) that grow dynamically as disks are added.
- It provides high performance while supporting RAID (Redundant Array of Independent Disks).
- It restarts quickly, even after an unclean shutdown due to power failure or system crash.

9.5.1 CIFS

9.5.1.1 Share Naming Recommendation

Since clustering is being used on the production filers, as in our current MSCS configurations, and multiple shares of the same type (`data_grp`, `data_tmp`, `data_pub`, `apps` etc.) can reside on the same cluster, 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.

NOTE: During the pilot phase of the original NetApp implementation in the year 2000, some concerns had been raised regarding the use of application shares and running packaged (wrapped) applications from the filers. GDSE are continually reviewing these issue(s).

9.5.1.2 File Permissions

Directory and file level permissions are managed using Global Groups. This is required for both cluster fail-over as well as the ability for disaster recovery to another filer. File level permissions get replicated during SnapMirror. If the primary and backup/DR filers 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 filer operations and user access controls.

9.5.2 NFS

9.5.2.1 NFS Overview

NFS is one of the file level access protocols supported by NetApp's Data ONTAP operating system (both NFSv3 and NFSv4, although NFSv4 has not yet been certified by GDSE).

The features of WAFL make NFS on NetApp Filers ideal as the new standard for Citigroup NFS File Server implementations going forward.

9.5.2.2 NFS Configuration

Like VFILER, NFS is an optional NetApp feature that requires a license purchase. The current certified version of Data ONTAP v7.0 supports both NFS v3 and NFSv4 (although NFSv4 has not yet been certified by GDSE), and like SUN Solaris NFS server supports most of the NFS standards export options. During certification it was determined that one could use NIS+ netgroups with NetApps NFS with some limitations:

- File systems can be exported to a netgroup either from local file on the filer (/vol/vol0/etc/netgroup) or NIS+ map.
- **Using NIS+, we CANNOT export with 'root' or 'rw' to a group, but we CAN export a file system as 'ro' to a group.**
- You can specify up to 255 names with the rw option. The names can be host names or subnets; you cannot specify netgroups.
- A single NIS+ entry (Solaris 8) cannot be more than 8192 characters.

NFS clients can mount only those resources that have been exported. You can export resources in two ways:

- Specify resources in the /etc/exports file and execute the `export fs` command either at the command line or from the /etc/rc file to export the resources according to the specifications in the /etc/exports file.
 - This method is for exporting resources that you want to export after each filer reboot. NOTE: This is the proposed method for SOE NetApps NFS support.
- Use the `export fs` command to export resources that are not specified in the /etc/exports file or to export resources specified in the /etc/exports file with different access restrictions.
 - This method is for exporting resources temporarily because the `export fs` command is not persistent across filer reboots.

Example of /etc/rc and /etc/exports files with NFS exports:

```
filer1> rdfile /vol/vol0/etc/exports
#Auto-generated by setup Thu Dec 5 09:11:38 EST 2002
/vol/vol0 -access=canonball.ny.ssmb.com,root=canonball.ny.ssmb.com
#/vol/philtest -ro,access=roothost:FIADM
/vol/philtest -access=roothost:FIADM
filer1>

filer2> rdfile /vol/vol0/etc/rc
#Regenerated by registry Fri Dec 27 17:17:52 GMT 2002
#Auto-generated by setup Thu Dec 5 17:08:02 EST 2002
hostname filer2
vif create single vif1 e4 e5
ifconfig vif1 `hostname`-vif1 partner vif0
route add default 150.110.93.3 1
routed on
options dns.domainname ny.ssmb.com
options dns.enable on
options nis.domainname ue.ssmb.com
options nis.servers 162.124.248.195
options nis.enable on
savecore
exportfs -a
nfs on
filer2>
```

9.5.2.3 Licensing

To enable the NFS license, complete the following steps.

```
filer1>license nfs=<code>
```

To disable NFS complete the following steps:

```
filer1>nfs off
filer1>license delete NFS
```

9.5.2.4 Using NIS / NIS+ for user information

Every vfiler must be configured to use one or more DNS servers for name resolution, and one or more NIS servers for user/host authentication to NFS shares, during the “setup” procedure. NIS+ is the management database of choice for CTI SOE.

IMPORTANT: NetApp Data ONTAP supports NIS, not NIS+. In order to authenticate to NIS+, the NIS+ manager must **enable NIS compatibility** for NetApp NFS shares. In addition, a Filer can only be in a single NIS domain, therefore vFilers (see section 9.7) should be used to export data to clients in different NIS domains.

Additionally, you should set up caching of group information. Using cached group information enables Data ONTAP to quickly determine NFS users' secondary GIDs when users try to connect to the filer. Without the cached information, Data ONTAP has to access the NIS server for group information each time an NFS user requests a connection.

To cache NIS group information, complete the following steps.

1. Enter the following command	<code>Options nis.group_update_schedule <hours_list></code> <i>hours_list</i> can be a comma-separated list of numbers that range from 1 to 24, representing the hours of the day, or now, which means that the cache is updated immediately.
2. To make Data ONTAP use the cached group information, enter the following command:	<code>options nis.group_update.enable on</code> The default is not to use cached group information.

Table 9.11 - Caching NIS group information

9.5.2.5 User Name Search Order

Like Solaris, Data ONTAP uses the `/etc/nsswitch.conf` file to determine the order of name resolution search. If the `nsswitch.conf` file does not already exist in the `/etc` directory, Data ONTAP creates a default `nsswitch.conf` file when NIS is enabled on your filer. The contents of the default `nsswitch.conf` file are as follows:

```
hosts: files nis dns
passwd: files nis
netgroup: files nis
group: files nis
shadow: files nis
```

The default `nsswitch.conf` file searches for user information in the following order:

- Configuration files in the `/etc` directory of the filer's default volume
- NIS

To change the order in which Data ONTAP searches for user information, map a drive to the administrative volume of the filer `/vol/vol0` and edit the `/etc/nsswitch.conf` file.

9.6 Monitoring and Alerting

9.6.1 AutoSupport

The autosupport process monitors the filer's operations and sends automatic email messages to Network Appliance Technical Support to alert them to potential filer problems. If necessary, Network Appliance Technical Support contacts you at the email address that you specify to help resolve a potential system problem.

Autosupport should be enabled (default) and configured with the correct mailhost (will vary by region) so that e-mail messages can be sent externally. An example, for NY the mail host would be `mailhub-nyc3.ny.smb.com`.

Autosupport should also be set to notify the filer administrator(s) by adding additional e-mail addresses in the "To:" boxes. Due to the security classification of content of the autosupport message and the clear text nature of SMTP messages, only Citigroup internal addresses are permitted in "To:" boxes, or in the `autosupport.to` option. Even when the transport protocol is set to HTTPS, emails are still sent to the system administrators defined in `autosupport.to` option. The target URL, when using HTTPS, is not user definable but is included here for reference: `support.netapp.com/asupprod/post/1.0/postAsup`.

A sample FilerView configuration screen is shown below and can be accessed by “Filer – Configure AutoSupport”. Advanced options are available from the command line using “**options autosupport**” at the console prompt, including the ability to configure a web proxy server and select the autosupport transport protocol.

9.6.2 SNMP

The recommended way to monitor filers is by configuring them to send SNMP traps to a management console. The Nerve Center / Netcool or Tivoli monitoring solutions is capable of receiving and interpreting these traps. These traps are shown in chapter 13 of this document. Certification of NetApp Data ONTAP 7.1 includes EMS certification of NetApp MIB 1.18.1, which adds one new trap to the traps available in ONTAP 7.0.

The filer needs to be configured with the destination name of the SNMP trap target server and SNMP traps have to be initialized on the filer via the console. The following example demonstrates how to configure SNMP using the CLI.

Command	Description
<code>snmp authtrap 1</code>	Enable authentication traps
<code>snmp community delete ro admin</code>	Delete default SNMP community on new installs
<code>snmp community delete ro public</code>	
<code>snmp community add ro fe03h20</code>	Add Citi default SNMP community string (or the community name provided by your local EMS group)
<code>snmp contact "Name, Email, Phone"</code>	Contact info used in SNMP trap. For example: use support group name, email and group phone line in the quotes.
<code>snmp traphost add server1.nam.nsroot.net</code>	Add trap recipient server name provided by your local EMS group (FQDN recommended)
<code>snmp traphost add server2.nam.nsroot.net</code>	
<code>snmp inti 1</code>	Initialize SNMP daemon

Table 9.12 - SNMP configuration settings

9.7 Vfiler (multistore)

9.7.1 Vfiler Overview

NetApp MultiStore software lets you easily create separate and completely private logical partitions of a filer's, network, and storage resources. Each virtual filer (VFILER) partition maintains absolute separation from every other virtual filer partition. The result is that business departments sharing the same filer cannot access or in any way find resources allocated to vfilers other than their own.

VFILER enables a single filer to do the work of multiple servers. NFS exports and CIFS shares can be set up independently for each partition. In addition, you can also consolidate multiple disparate NIS and/or Windows domains on one filer.

VFILER yields the following benefits:

- Maintains domain infrastructure while providing multi-domain storage consolidation - Multi NIS domain support.

- Win/UNIX storage consolidation without reconfiguration – allowing varied CIFS and NFS file services on single filer.
- Enables dynamic assignment/reassignment of storage and network resources - Seamless non-intrusive disaster recovery features COB.
- Enables repurposing of the file servers to application servers. – Reallocation of current HA Solaris NFS file servers.




Figure 9.1 - NetApp MultiStore (Vfiler)

9.7.2 Hosting Filers / Default Vfiler

The filer where you create vfilers is called the hosting filer. The storage and network resources used by the vfilers exist on the hosting filer. Licensing MultiStore automatically causes Data ONTAP to create a default vfiler on the hosting filer.

The name of the default vfiler is vfiler0. Initially, vfiler0 owns all the resources on the filer. After you create vfilers, vfiler0 owns the resources that are not owned by the vfilers you create. The default vfiler exists as long as the vfiler license is enabled. On a filer with the vfiler license enabled, you cannot destroy vfiler0.

The storage resource assigned to a vfiler can be one or more qtrees or volumes. The network resource assigned can be one or more base IP addresses or IP aliases associated with network interfaces. You can add or remove resources at any time.

NetApp MultiStore (VFILER) was certified for the primary purpose of segregating CIFS/Windows and NFS/UNIX file systems. VFILERS were limited to a total of 10 usable VFILERS per cluster. This limitation allowed for only 5 VFILERS on any one-filer node, allowing for cluster failover and COB capacity requirements.

As of DataONTAP v6.5, you can configure up to 32 usable VFILERS on a single filer node (up to 64 in a cluster), depending on the memory capacity of the hosting filer. The maximum limits for number of vFilers are shown in the table below, and highlights the maximum value supported by NetApp and the value certified by GDSE.

Filer Memory Size	NetApp	Citigroup
less than 1 GB	10	7
between 1 GB and 2GB	25	18
over 2 GB	32	24

Table 9.13 - Citigroup recommended vFiler limits

9.7.3 VFiler Prerequisites

The following list describes the prerequisites for using VFILER:

- Because VFILER is an optional software product, you must enable the software license for it. The name of the license is `vfiler`.
- There must be Ethernet interfaces on the filer.

9.7.3.1 VFiler IP Address Procurement

Please refer to your System Administrator for IP segment management. System Administrators usually have access to manage IP request for his/her IP segment, using the Corporate IP Management System –QIP. Please contact your regional DNS administration team for the current web address of the QIP system.

9.7.3.2 VFiler Planning Worksheet

The VFILER Planning Worksheet is a table that aids the System Administrator with the planning of `vfiler` creation on Netapp filers. System Administrators are highly encouraged to treat this worksheet as a work-in-progress document that can also serve to aid in the maintenance of NetApp filers.

Please reference the link for a copy of the VFILER planning Worksheet. A sample of the worksheet can be found in the Sample Worksheet tab.

<http://engineering.ti.citigroup.net/gdse/products/platforms/nasfiler/opsbestpractices/>

9.7.3.3 3DNS Name Services

3DNS name services allow a single hostname to be serviced from multiple locations and or IP addresses. This feature, in combination with the `vfiler` system, `qtree`, and `Snapmirror` facilitate the faultless failover of `vfiler` DNS host names and NFS mounts, in the event of disaster in the production environment.

3DNS requires an IP port to be monitored for every VFILER. As a standard best practice, all VFILERS should be configured for 3DNS name services, with monitoring on TCP port 445. Requests for 3DNS name services requests require the following form, http://engineering.ti.citigroup.net/gdse/products/platforms/nasfiler/opsbestpractices/3-DNS_request_form.xls. The request process for 3DNS requests can be found on the Whitepages website (<https://ip03rut.nam.nsroot.net/webreq/>), and the 3DNS website (<http://lbsupport.ny.ssmb.com/>). For implementations outside the North America region that are not supported via the Whitepages team, check with you local directory and name services support group.

9.7.3.4 Licensing

Enabling the `vfiler` license has the following effects on the filer:

You can use the `vfiler` and `ipspace` commands. Data ONTAP starts logging `vfiler` status information and sends `vfiler` status information using the `autosupport` feature. The `routed` daemon is disabled. The `ip.match_any_ifaddr` option is set to `Off`.

To enable the `vfiler` license, complete the following step.

`License vfiler=<code>`

To disable the `vfiler` license, complete the following step.

`License vfiler=disable`

You can disable the vfiler license only when no vfilers other than vfiler0 are on the filer. If there are vfilers on the filer, you must destroy them before disabling the vfiler license.

Disabling the vfiler license has the following effects:

MultiStore becomes unavailable immediately. You can no longer use the `vfiler` and `ipSPACE` commands. The routed daemon is enabled after the next reboot.

9.7.4 VFile Limitations

The following list describes the limitations of vFiler:

- VFile management is limited to individual filer nodes. There is no centralized management utility.
- VFilers can only be managed from the command line, not from the Filerview GUI
- The “`vfiler limit`” command requires a reboot to increase or decrease the number of vfilers supported.
- Filer cache/memory resources are reserved for the number of VFilers configured by the “`vfiler limit`” command, irrespective of usage. Each potential vFiler (whether configured or not) is allocated 400K of memory. For example if you have only created five VFilers but you have a vFiler limit value of sixteen, then 16x400K memory will be reserved for VFilers.
- Not all Data ONTAP commands are supported on vfilers.
- Only Ethernet interfaces are supported for VFilers.
- Only the CIFS, NFS, Internet SCSI (iSCSI), HTTP, and FTP protocols are supported for VFilers. However it should be noted that while the Filer supports iSCSI, HTTP and FTP protocols, these protocols are not certified for use within Citigroup on the NetApp NAS Filer platform.

9.7.5 Data ONTAP tasks allowed on VFilers

Data ONTAP allows a limited command set for management of VFilers. In addition there is no “`na_admin`” GUI for management of VFilers. The most common tasks administrators will employ are as follows:

- Managing local vFiler users: “`useradmin`”
- Managing files access protocols: “`setup`” for NIS/NFS configuration and “`cifs setup`” for CIFS configuration
- Managing quotas. “`quota`”
- Preparing vfilers for virus protection: “`vscan`”

9.7.6 SOE VFile Naming Standards

9.7.6.1 Naming Guidelines

NetApp recommends the following guidelines:

- The hostname can contain up to 31 alphanumeric ASCII characters.
 - o NOTE: Citigroup AD host names cannot be longer than 15 characters. Hence, all filer hostnames should be limited to 8 characters or less, to allow for future allocation of vfiler hostnames, which add 7 characters to the filer hostname.
- The name can contain the dash (-) and the underscore (_) characters, but the dash character must not be the first character.
 - o NOTE: Citigroup DNS naming standards do not allow the use of the underscore (_) in its host names.
- The name is case-insensitive.
- The name `vfiler0` is reserved.

9.7.6.2 SOE Production VFilers

In cooperation with Storage Integration, GDSE determined the Citigroup naming standards for VFILERS are as follows:

filenamev##xxx

Where:

filename: name of hosting filer,
v specifies host is a virtual filer,
and viler number of that particular filer,
xxx: is a three-letter abbreviation of the business unit the VFILER services,

Examples of the vfiler naming rule:

nasrut110v01int.nam.nsroot.net
nasrut110v02int.nam.nsroot.net

9.7.7 SOE VFile Production Configuration

9.7.7.1 SOE Prerequisites

The following prerequisites must be met before you can create a vfiler:

- Creating a vfiler requires storage. The units of storage can be qtrees or volumes. You must create the qtrees or volumes before creating the vfilers.
- The storage unit for vfiler configuration information must be writable. It must not be a read-only file system, such as the destination volume or qtree in a SnapMirror relationship.
- The IP address assigned by the vfiler must be un-configured or unassigned to the network interface when creating the vfiler.

9.7.7.2 SOE Storage Guidelines

When deciding which storage units to assign to vfilers, remember the following guidelines:

- The first storage unit assigned to the vfiler contains the vfiler configuration information (i.e. "/etc" directory). It is called the primary storage unit. Although you can remove storage units from a vfiler at any time after the vfiler is created, the primary storage unit must remain for as long as the vfiler exists.
- If the qtree or volume to be used as the vfiler system storage unit contains an /etc directory, the data in the directory is lost after you add the qtree or volume to the vfiler. All other data is preserved.
- A volume assigned to a vfiler must not be the host filer's root volume. You can, however, assign qtrees in the root volume to a vfiler.
- If the vfiler administrator needs to create qtrees on the vfiler, assign volumes instead of qtrees to the vfiler. This is because qtrees can be created only at the root of a volume.
- When managing NFS exports, CIFS shares, quotas, and options, vfiler administrators need to enter the complete path names of the storage resources used by the vfilers in commands and configuration files. Therefore, choose volume and qtree names appropriately, so that you can share with the vfiler administrators the complete path names beginning with *filer_name:/vol/vol_name*.

9.7.7.3 SOE VFile Create Guidelines

For the purpose consistency of NetApp filer administration, it was determine that all configuration/system files for the host filer and vfiler(s) be kept on /vol/vol0.

NOTE: Because we cannot create vfiler(s) off the root volume (vol/vol0), for it will overwrite the /etc directory, a qtree will be created off the root volume of /vol/vol0 to accommodate the vfiler(s) system files, which in that case becomes the root volume of the vfiler(s).

The following guidelines must be adhered when creating vfilers:

1. All primary storage (vfiler configuration information, i.e. /etc) assigned to a vfiler will be a qtree created from the host filer's root configuration volume (/vol/vol0).
2. Create IP alias on hosting filer's VIF NIC for the vfiler.
 - **NOTE: Create VFILER with IP alias before assigning IP alias to the host filer VIF. This allows us to create vfilers without having to temporarily disable the VIF.**
3. One qtree per VFILER.
4. qtree naming rules are as follows:

```
filename:/vol/vol0/ nasrut110v01int
filename:/vol/vol0/ nasrut110v02int
filename:/vol/vol0/ nasrut110v03int
```

Where:

filename:	name of hosting filer,
/vol/vol0/	is the hosting filer configuration or system volume,
filename:	name of hosting filer,
v	specifies host is a virtual filer,
##	and viler number of that particular filer,
xxx:	is a three-letter abbreviation of the business unit the VFILER services,

Examples of the qtree naming rule:

```
nasrut110:/vol/vol0/nasrut110v01int
nasrut110:/vol/vol0/nasrut110v02int
nasrut110:/vol/vol0/nasrut110v03int
```

9.7.7.4 SOE VFile Create Example

1. create qtree for vfiler configuration files
`qtree create /vol/vol0/nasrut110v01tffi`
2. create vfiler using the qtree created in the above step, as the primary volume, and the IP address assigned to the VFILER
`vfiler create nasrut110v01tffi -i 150.110.93.104 /vol/vol0/ nasrut110v01tffi`
3. assign designated data volume(s) to vfiler
`vfiler add nasrut110v01tffi /nasrut110v01tffi _v1`
4. create IP address alias on host filer VIF using IP assigned to VFILER
`ifconfig gev1 alias 150.110.93.104 netmask 255.255.255.0`
5. edit "/vol/vol0/etc/rc" file on the host filer; add the alias just created to initialize upon filer reboot.
`ifconfig gev1 alias 150.110.93.104 netmask 255.255.255.0 up`

9.7.7.5 SOE Cluster Considerations

- Number of Citigroup supported vFilers is defined in Table 9.13.
- A member of a cluster that does not have a vfiler license can take over its partner with a vfiler license. However, each node must have a vfiler license to support vfilers.
- The vfilers hosted by the filers of the cluster are created and configured independently. That is, each filer can host a different number of vfilers, and the vfiler configuration on the filers can be different from each other.
- In takeover mode, the functioning filer takes over all vfilers created on the failed filer. These vfilers include the vfilers you create and vfiler0. Therefore, for vfilers on the failed filer to work correctly after the takeover, **each network interface used by a vfiler in a cluster must have a partner interface.**

9.7.8 SOE Vfiler DR/COB Configuration

9.7.8.1 Installation and Setup

As part of the installation and setup of Production vfilers COB/DR vfilers will be configured to failover using the same Production VFILER names.

9.7.8.2 Replicating Data

All COB/DR filers supporting production vfilers must have the vfiler license enabled. Data replication for vfilers will employ the current SOE processes for NetApp filer volume replication. However, with the introduction of qtrees for the vfiler configuration files, an additional replication rule will be introduced for qtree snapmirror.

VFILER introduces a more efficient, faster process for filer disaster recovery (COB/DR). Using snapmirror, once the vfiler qtree and data volumes have been copied to the COB/DR filer, you can re-create the vfiler on the destination filer.

Re-creating the vfiler from the qtree snapmirror allows for seamless recovery, whereby the vfiler is recovered with all its resources as if it were the source filer. In other words, the vfiler name (hostname), data volumes, and all other configuration options are available to the end user without any end-user modifications to their operating environment.

Where legacy filers are configured for recovery using different names for its data volumes on the COB/DR filer, and 3DNS host name failover is not available, special considerations (such as, remapping of user drives) must be taken to ensure an effective failover.

All vfiler COB/DR configurations should be created and tested at the time the production vfiler is installed.

9.8 Backup

The back-up section of this solutions guide has been considerably revised since the previous Citigroup certified release of Data ONTAP. Due to the different backup products used globally within Citigroup, 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.

9.8.1 NDMP Overview

Network Data Management Protocol (NDMP) controls backups and restores on the Network Appliance filer. 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.

In the simplest NDMP configuration, the tape device or library is directly attached to the NAS system to be backed up. This approach works well for network systems with large amounts of data because it keeps data from flowing across the network, preventing network congestion.




Figure 9.2 – Backup Simple "Local" NDMP Configuration

Simple "Local" NDMP Configuration - All data flows directly from the NAS device to the directly attached library. Only metadata and NDMP control information traverse the network between the backup server and the NAS device.

More advanced NDMP configurations allow data to flow from one NDMP-compliant system over the network to a system hosting a tape storage device. This configuration is often referred to as *three-party* or *three-way*.




Figure 9.3 – Backup Three-party NDMP Configuration

Three-party NDMP Configuration - Data flows from NAS Device #2 over the network to NAS device #1 and then to tape. Metadata is sent from NAS device #2 to the backup server. NDMP control information flows between the backup server and NAS device #1.

9.8.2 Authorizing Access To the NDMP Host

The server controlling the backup job will need to be authorized for NDMP backups of the target Filer. Refer to the backup software vendor's documentation for this step. ***As of Data ONTAP v6.5, the NDMP Api allows for any filer account to be used for backup. Hence the backup account no longer need be "root".***

All regional operators/support are required to create a backup functional ID for NetApp NDMP backup, for example: filename\NASbackup.

The process for generating the password for the NDMP backup account is:

```
ndmpd password <account_name>
```

Where <account_name> is the name of the account to be used for NDMP backups. This command will return a 16 character hash that is used by the backup software to authenticate the NDMP session.

Follow the backup software vendor's documentation to verify that the NDMP password has been configured correctly.

9.8.3 Configuring the tape devices

Log on to the NetApps file server.

Execute **sysconfig -t** and verify that the system recognizes the drives. For example:




Figure 9.4 - FilerView GUI Telnet Session - Tape Devices Configuration

Execute an **mt status** command to verify that the system can access the drives. For example:




Figure 9.5 - FilerView GUI Telnet Session - Tape Devices Status

Start the NDMP server daemon by executing the following command:

```
ndmpd on
```

Check the NDMP server daemon status by executing the **ndmpd status** command. For example:

```
filer1> ndmpd status
ndmpd ON.
No ndmpd sessions active.
```

9.8.4 Backup process

Regular backups will be made from the snapmirror on DR filer. The backup schedule will typically include a monthly full backup, weekly cumulative incremental, and daily differential incremental backups. All backups are kept for 3 years, mandated by SEC requirements.

In order to back up data from a mirror, you must back up an existing snapshot, not the active file system, on the mirror. This is because the filer must create a snapshot in the volume containing the data before NDMP writes data from the active file system to tape. The filer cannot create snapshots on the mirror because it is read only.

The snapshot in ~snapshot\hourly.0 on the DR filer will be used for backup (incremental or full). Hourly.0 will always contain the latest hourly snapshot since it is numbered the lowest.

If a volume is taken offline (such as to do a snapmirror), then backups cannot occur. NDMP has to see the volume in an on-line state in order to do the backup.

NOTE: During the NDMP backup, on the Backup/DR filer, all data replication for the volume in question is suspended until the backup is completed.

9.8.5 NDMP Operating Notes and Restrictions

- 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.

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 Security Best Practices

This section only applies to all Citigroup certified NetApp Filers (with the exception of the LockVault compliance release) running ONTAP 7.1, and supercedes the previous separate Secure Configuration Guide document found at http://gdse.ny.ssmb.com/GDSE_docs/GDSE-03032-S.pdf and the security configuration for ONTAP 7.0.

10.1 Overview

This section's main focus is to design the security, network topology, and configuration of Network Appliance NAS appliances for use as SOE file servers in order to meet the requirements of Citigroup Information Security Standards ("CISS").

This section, together with the Information Security ("IS") EP, ICQ and RCSA processes, provide the foundation for the assessment of operational risk associated with production NAS environments.

10.2 Security Standards

Citigroup 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.

10.3 Information Classification Certification

The NAS infrastructure processes information classified as "public", "internal" and "confidential". This classification by definition makes it a CITMP Information Security "Medium Risk" system. Information classified as "restricted" cannot be processed on the NAS infrastructure unless adequate controls, such as encryption, are built into infrastructure.

10.4 Anti-Virus

Three options exist today for virus scanning on the NT/CIFS space for Network Appliance filers:

1. Client-side virus scanning. All client workstations, as is the SOE standard, run local virus scanning engines and thus any files that are accessed from filers will get scanned for viruses
2. Scheduled virus scanning from Wintel Servers.
3. On-access virus scanning using 3rd-party scan-engines acting as gateways/proxies.

Both option #1 and #3 are proactive scanning. Option #2 is a passive approach to virus scanning. Previous versions of the NetApp Standards certified the combination of option #1 and option #2.

CTI employs Network Associates Inc., McAfee AV Suite of products as its SOE standard for virus scanning. With the availability of VirusScan for NetApp on-access scanning, the new certified scanning method is the combination of options #1 and #3. These options create the most robust virus protection scheme for real-time file access. In other words, the scanner examines files when they are read from disk, written to disk, and renamed. Additionally, the management of scanning batch schedules and the need for user interaction are eliminated.

VirusScan for NetApp v7.1 protects NetApp filers from viruses by running the anti-virus program on a separate scanner-server. When a client system attempts to access or store a file on the filer, the filer's operating system (Data ONTAP) virus scan API determines (in real-time) if the file should be scanned. If the file needs to be scanned, the filer sends a request to the scanner-server. The VirusScan for NetApp scanner-server scans the file and takes appropriate action, if necessary.

10.4.1 New features and enhancements in VirusScan for NetApp v7.1:

Refer to Figure 8.1 - NetApp Filer Architecture Design on page 8 for the Anti-Virus for NetApp infrastructure design.

- Multiple-to-multiple configurations. Ability to register multiple scanner-servers with one filer, and also configure one scanner-server to manage multiple filers.
- Automatically reset the filer clean file cache, after each dat or engine update.
- ePolicy Orchestrator manageability:
 - o Configure policies and manage VirusScan for NetApp tasks from the ePolicy Orchestrator console.
 - o View ePolicy Orchestrator tasks in the local VirusScan server VirusScan Console.




Figure 10.1 – VirusScan for NetApp v7.1 User File Access Data Flow.

10.4.2 VirusScan for NetApp v7.1 ePO NetApp Policies configuration

The policies related to NetApp filers, e.g. what servers will manage what filers, should be configured at the ePO organization container level where the VirusScan for NetApp v7.1 servers are managed.

NOTE: The following steps are to be completed by the local ePO Administrator, **and only**, after the ePolicy Orchestrator NAP file for VirusScan for NetApp v7.1 has been installed and configured. ePO Administrators refer to [http://virus-defense.nam.nsroot.net/Pub/MCAFEe/ePO3.5/VirusScan for NetApp 7.1/Documentation/VS7netapps.doc](http://virus-defense.nam.nsroot.net/Pub/MCAFEe/ePO3.5/VirusScan%20for%20NetApp%207.1/Documentation/VS7netapps.doc), or contact [Distributed Security Engineering](#) for NAP file installation procedures. Please refer to [Appendix A](#) for VirusScan for NetApp v7.1 ePO configuration settings.

1. Select the Repository folder.
 - a. Scroll down and open your organization container
 - b. Create a sub-container for VirusScan for NetApp v7.1 servers. This is recommended in order to quickly identify the VirusScan servers.
 - c. Within your organizational container find your VirusScan for NetApp v7.1 servers (check Lost and Found, and other Servers containers), and move or cut and paste to the new container created in Step b.
 - d. Select the VirusScan for NetApp v7.1 Servers container or name created in Step b.
 - e. In the right pane "Policies", select the VirusScan for NetApp 7.1
 - f. Deselect "Inherit"
 - g. Select "NetApp Policies".
 - h. Click Add button, add name of NetApp Filer serviced by your organization.
 - i. Repeat above step for multiple NetApp Filers.
 - j. Select the additional properties shown in the figure below.




Figure 10.2 - ePO Organizational container NetApp Policy configuration

2. From your organization container
 - a. Open the VirusScan for NetApp v7.1 servers container or name created in Step 1b
 - b. Select any of your VirusScan for NetApp v7.1 servers
 - c. From the right pane, select the VirusScan for NetApp 7.1 software
 - d. Select "NetApp Policies"
 - e. With the "Inherit" radio button selected, all the servers configured in the above policies (Step 1) are filtered down.
 - f. Deselect "Inherit"
 - g. Remove filers not to be scanned by the selected server.

NOTE: The recommended ratio of NetApp Filers to VirusScan for NetApp servers is 4:1. This ratio is considered optimal for Filer/VirusScan server I/O performance.

3. Once the granular NetApp policies have been completed. It is necessary to push the changes to the server.
 - a. Right mouse-click on the server
 - b. Select “Agent Wakeup Call” – Figure 10.3
 - c. Select default options – Figure 10.4




Figure 10.3 - ePO Push Policy Changes to Server




Figure 10.4 - ePO Push Policy Changes to Server – Options

4. Verify ePO configuration options are inherited by VirusScan Server
 - a. From VirusScan for NetApp v7.1 Server
 - b. Run Programs/Network Associates/Network Appliance AV Server Console – Figure 10.5
 - c. Select the appropriate server
 - d. Right-mouse click, and select properties – Figure 10.6
 - e. The “General” tab should display the NetApp Filers configured in Step 3 – Figure 10.7.




Figure 10.5 – VirusScan Server Network Associates Program Group




Figure 10.6 - VirusScan Server Network Appliance AV Server Console




Figure 10.7 - VirusScan Server NetApp Filers Serviced by AV Scanner Server

10.4.3 Network Appliance Data ONTAP “vscan” Settings

The NetApp Filer like the VirusScan and ePO Management servers requires certain configuration settings enabled to allow for On-Access scanning.

NOTE: The following steps are to be completed **only** by the local NetApp Filer administrator (e.g. in North America, Storage Integration). These procedures are not dependent on the VirusScan or ePO Manager server configurations, and can be completed at anytime.

1. Enable scanning on the NetApp filer. To do so, enter the following command at the filer's console:

```
filer1> vscan on
```

```
Mon Feb 10 10:12:20 EST [filer1: rc:notice]: CIFS virus scanning is enabled.
Virus scanning is enabled
```

```
filer1>
```

2. Enable vscan options

- a. Set vscan options mandatory_scan off

The default setting is “On”. When this option is set to “On”, file access is denied if a successful scan cannot be performed. When this option is set to “Off”, access is allowed even if no successful scan can be performed.

b. Set vscan options `client_msgbox off`

Displays the current setting for the `client_msgbox` option. If set to "on", the filer will attempt to send a pop-up MsgBox to the opener of an infected file.

3. Run "vscan" command

Displays vscan settings and provides summary information about scan requests.

```
filer1> vscan
Virus scanning is enabled.
Virus scanners(IP and Name)  Connect time (dd:hh:mm)  Reqs    Fails
-----
146.128.134.142  \\NAMGDSEAV1      00:02:54      0       0
List of extensions to scan:
??_, ARJ, ASP, BAT, BIN, CAB, CDR, CL?, COM, CSC, DL?, DOC, DOT, DRV, EML, EXE, GMS, GZ?, HLP,
HT?, IM?, INI, JS?, LZH, MD?, MPP, MPT, MSG, MSO, OCX, OFT, OLE
, PP?, RAR, RTF, SCR, SHS, SMM, SYS, VBS, VS?, VXD, WBK, WPD, XL?, XML
Number of files scanned:  7
Number of scan failures:  0
filer1>
```

10.4.3.1 Controlling File Types to be scanned

To control what files are scanned, you can add to, subtract from or replace the default extensions by "vscan extensions" command.

- Displaying the file types to be scanned:
`vscan extensions`
- Adding file types to be scanned:
`vscan extensions add ext[, ext...]` , where "ext" is the extension you want to add.
Example: `vscan extensions add txt`
NOTE: Up to 255 extensions can exist in the file extensions list.
- Removing file types from the extension list:
`vscan extensions remove ext[, ext...]` , where "ext" is the extension you are adding.
- Replacing the file types in the extensions list:
`vscan extensions set ext[, ext...]"` , where "ext" is the extension you want to set.

10.4.4 NetApp Data ONTAP Anti-Virus Service Help

To display list and brief description of all vscan commands use the "vscan help" command.

`vscan help`

10.4.5 Resetting the flags in the filer's cache.

Resetting the vscan cache assures that files previously flagged as virus-free will be rescanned. This is especially necessary when .DAT files are updated. The latest version of VirusScan for NetApp v7.1 has a setting that automatically resets the vscan cache post .DAT file updates.

To reset the flags in the filer's cache, access the filer's console and issue the command **vscan reset**.

```
filer1> vscan reset
vscan reset completed successfully.
filer1>
```

10.4.6 VirusScan for NetApp - Virus Detection

When an infected file is detected by NetShield for NetApp the following set of events take place:

1. The filer console logs an rpc0_warning message from the CIFS protocol about the AV Scanner server that includes all relevant attributes about the infected file;

```
Mon Mar  1 14:37:16 EST [filer1: CIFSvScan:warning]: Cannot connect
to server. Session setup error: Unknown error code
Mon Mar  1 14:37:16 EST [filer1: rpc_0:warning]: CIFS: Virus scanner
146.128.134.142 completed a scan on
VSCAN_ADMIN$\HOME\EICAR.COM for client 146.128.1.37
(NAMCPSRNYCW158) as user eg55388 but returned the following
status: [0x5] and status message: Error 5
...Continuing
```

- a. To view the status of vscan after it has detected a virus type the command "vscan":

```
filer1> vscan
Virus scanning is enabled.
Virus scanners(IP and Name)  Connect time (dd:hh:mm)  Reqs    Fails
-----
146.128.134.142  \\NAMGDSEAV1          00:00:13          71     55
```

List of extensions to scan:

```
??_, ARJ, ASP, BAT, BIN, CAB, CDR, CL?, COM, CSC, DL?, DOC, DOT, DRV, EML, EXE, GMS,
GZ?, HLP, HT?, IM?, INI, JS?, LZH, MD?, MPP, MPT, MSG, MSO, OCX, OFT
, OLE, OV?, PIF, POT, PP?, RAR, RTF, SCR, SHS, SMM, SYS, TAR, VBS, VS?, VXD, WBK, WPD
, XL?, XML, ZIP
```

```
Number of files scanned:  71
Number of scan failures:  55
filer1>
```

NOTE: The "Number of files scanned" shows how many files were intercepted by the scanner server. The "Number of scan failures" shows the number of files that were viral infected and denied access by the scanner server.

10.5 Disk Sanitization

Disk sanitization is the process of physically obliterating data by means of overwrites patterns on filer or NearStore appliance disks in a manner that precludes recovery of that data by any known recovery methods. The Data ONTAP “*disk sanitize*” is an advanced “*disk*” command feature that enables you to overwrite the contents of a disk. It accomplishes this by using three successive byte overwrite patterns per cycle and a default of six cycles per operation, in compliance with United States Department of Defense and Department of Energy security requirements.

Disk sanitization is a requirement mandated by CITMP IS Policy. As such, all failed disks within every CTI data center is subject to stringent sanitization or degaussing procedures.

NOTE: Current CITMP limitations restricting movement of electronic media warrant that all disk sanitization cycles occur within the production system. The concern here is that an administrator could possibly erase production data. Fortunately, with the NetApp disk subsystem, Data ONTAP will only allow a disk in “spare” mode (spare pool) to be selected for sanitization. Other limitations are outlined below.

10.5.1 Disk Sanitization Limitations

The disk sanitization operations described in this section are subject to the following limitations:

- The operations are not supported on V-Series appliances (note GDSE has not certified V-Series filers).
- The operations are not supported in takeover mode on clustered systems or on disabled cluster systems.
- Disk sanitization cannot be carried out on disks that were failed due to readability or write-ability problems.
- Disk sanitization cannot be carried out on disks that have ever belonged to a SnapLock Enterprise (SEC 17a-4-compliant) volume.
- The formatting phase of the disk sanitization process is skipped on Advanced Technology Attachment (ATA) drives.

10.5.2 Licensing Disk Sanitization

The disk sanitization feature for Data ONTAP 7.0 requires installation of a license. Once installed, the license enables the disk sanitize command and limited SnapMirror replication abilities that facilitate selective sanitization operations.

Caution: Once installed on a NetApp appliance, the license for disk sanitization is permanent.

To install the disk sanitization license, complete the following step:

Enter the following command: `license add sanitize_code`

Note: `sanitize_code` is the disk sanitization license code that Network Appliance provides.

10.5.3 Sanitizing Disks

You can sanitize any disk that has “spare” status. This includes disks that exist on the filer or NearStore appliance as spare disks after the volume that they belong to has been destroyed and disks that were removed from the spare disk pool by the disk remove command but have been returned to spare status after an appliance reboot.

To sanitize a disk or a set of disks on a filer or NearStore appliance, complete the following steps:

1. From the Data ONTAP CLI, change privilege to advanced mode.
 - a. *"priv set advanced"*
2. Print list of disks and disk status. Disks that you want to sanitize cannot belong to any existing RAID group in any existing volume.
 - a. *"sysconfig -r"*
3. Enter the following command to physically sanitize the specified disk or disks of all existing data.
 - a. *"disk sanitize start" [-p pattern1 [-p pattern2] [-p pattern3]] [-c cycle_count] disk_list*
 - i. *-p* pattern specifies a cycle of one to three user-defined hex byte overwrite patterns that can be applied in succession to the disks being sanitized. The default hex pattern specification is *-p 0x55 -p 0xAA -p 0x36*.
 - ii. *-c cycle_count* specifies the number of cycles for applying the specified overwrite patterns. The default value is 6 cycles. The maximum value is 7 cycles.
 - iii. *disk_list* specifies a space-separated list of spare disks to be sanitized.

Example: The following command applies the default disk sanitization three overwrite patterns for the default six cycles (for a total of 18 overwrites) to the specified disks, 7.6, 7.7, and 7.8.

disk sanitize start 7.6 7.7 7.8

Result: The specified disks are sanitized, put into the pool of broken disks, and marked as "sanitized." A list of all the sanitized disks is stored in the *vol/vol0/etc* directory.

Verification: To check the status of the disk sanitization process, enter the following command:

disk sanitize status [disk_list]

Note If you need to abort the sanitization operation, enter *disk sanitize abort [disk_list]*

Caution: Do not turn off the filer, disrupt the disk loop, or remove target disks during the sanitization process. If the sanitization process is disrupted, the target disks that are in the formatting stage of disk sanitization will require reformatting before their sanitization can be completed. If the sanitization process is interrupted by power failure, filer panic, or a user-invoked disk sanitize abort command, the disk sanitize command must be re-invoked and the process repeated from the beginning.

The time to complete the sanitization process for each drive varies with the size of the drive, the number of patterns specified, and the number of cycles specified.

For example, the following command invokes 1 format overwrite pass and 18 pattern overwrite passes of disk 7.3.

disk sanitize start -p 0x55 -p 0xAA -p 0x37 -c 6 7.3

- If disk 7.3 is 73 GB and each formatting or pattern overwrite pass on it takes 30 minutes, then total sanitization time is 19 passes times 30 minutes, or 570 minutes or 9.5 hours.

To release sanitized disks from the pool of broken disks for reuse as spare disks, enter the following command:

disk sanitize release disk_list

10.5.4 Reading Disk Sanitization Log Files

The disk sanitization process outputs two types of log files. One file, `sanitization.log`, lists all the drives that have been sanitized. For each disk being sanitized, a file is created where the progress information will be written.

Storage administrators who are familiar with previous versions of Data ONTAP will note that the name of the log file has changed from `sanitized_disks` to `sanitization.log`.

Listing the sanitized disks: The `/etc/sanitization.log` file contains the serial numbers of all drives that have been successfully sanitized. For every invocation of the `disk sanitize start` command, the serial numbers of the newly sanitized disks are appended to the file.

A read of the `/etc/sanitization.log` file shows output similar to the following:

```
filer1> rdfile /etc/sanitization.log
Tue Jun 24 02:54:11 Disk 8a.44 [S/N 3FP0RFAZ00002218446B] sanitized.
Tue Jun 24 02:54:15 Disk 8a.43 [S/N 3FP20XX400007313LSA8] sanitized.
Tue Jun 24 02:54:20 Disk 8a.45 [S/N 3FP0RJMR0000221844GP] sanitized.
Tue Jun 24 03:22:41 Disk 8a.32 [S/N 43208987] sanitized.
```

Reviewing the disk sanitization progress: A progress file is created for each drive sanitized and the results are consolidated to the `/etc/sanitization.log` file every 15 minutes during the sanitization operation. Entries in the log resemble the following:

```
Tue Jun 24 02:40:10 Disk sanitization initiated on drive 8a.43 [S/N 3FP20XX400007313LSA8]
Tue Jun 24 02:40:10 Disk sanitization initiated on drive 8a.44 [S/N 3FP0RFAZ00002218446B]
Tue Jun 24 02:40:10 Disk sanitization initiated on drive 8a.45 [S/N 3FP0RJMR0000221844GP]
Tue Jun 24 02:53:55 Disk 8a.44 [S/N 3FP0RFAZ00002218446B] format completed in 00:13:45.
Tue Jun 24 02:53:59 Disk 8a.43 [S/N 3FP20XX400007313LSA8] format completed in 00:13:49.
Tue Jun 24 02:54:04 Disk 8a.45 [S/N 3FP0RJMR0000221844GP] format completed in 00:13:54.
Tue Jun 24 02:54:11 Disk 8a.44 [S/N 3FP0RFAZ00002218446B] cycle 1 pattern write of 0x47
    completed in 00:00:16.
Tue Jun 24 02:54:11 Disk sanitization on drive 8a.44 [S/N 3FP0RFAZ00002218446B]
    completed.
Tue Jun 24 02:54:15 Disk 8a.43 [S/N 3FP20XX400007313LSA8] cycle 1 pattern write of 0x47
    completed in 00:00:16.
Tue Jun 24 02:54:15 Disk sanitization on drive 8a.43 [S/N 3FP20XX400007313LSA8]
    completed.
Tue Jun 24 02:54:20 Disk 8a.45 [S/N 3FP0RJMR0000221844GP] cycle 1 pattern write of 0x47
    completed in 00:00:16.
Tue Jun 24 02:54:20 Disk sanitization on drive 8a.45 [S/N 3FP0RJMR0000221844GP]
    completed.
Tue Jun 24 02:58:42 Disk sanitization initiated on drive 8a.43 [S/N 3FP20XX400007313LSA8]
Tue Jun 24 03:00:09 Disk sanitization initiated on drive 8a.32 [S/N 43208987]
Tue Jun 24 03:11:25 Disk 8a.32 [S/N 43208987] cycle 1 pattern write of 0x47 completed in
    00:11:16.
Tue Jun 24 03:12:32 Disk 8a.43 [S/N 3FP20XX400007313LSA8] sanitization aborted by user.
Tue Jun 24 03:22:41 Disk 8a.32 [S/N 43208987] cycle 2 pattern write of 0x47 completed in
    00:11:16.
Tue Jun 24 03:22:41 Disk sanitization on drive 8a.32 [S/N 43208987] completed.
```

10.6 CITMP Information Security Policy Convergence

10.6.1 Information Ownership

The head of Distributed Engineering has information ownership of NAS products and all related engineering documents. NAS products are available in the Citigroup Technology Catalogue. All security and technology best practices for NAS products are the responsibility of Global Distributed Storage Engineering and can be referenced from the following website: <http://engineering.ti.citigroup.net/gdse/products/platforms/nasfiler/>

10.6.2 Identification and Authentication

10.6.2.1 Verify Session Authentication

NetApp Data ONTAP requires user authentication for every console or FilerView GUI system administration session. Users are authenticated by their respective userid. Every user requires a static password for system authentication. Passwords are encrypted using a variant on DES. The encrypted passwords are kept in the registry.

10.6.2.2 Support for Unique UserID

The following SOE NetApp best practices outline the procedures required for the creation of a unique userids:

By default the built-in root account has full administration capability of the filer. The root account can be used when accessing the filer via a telnet session (command line interface), serial connected console or the FilerView web based user interface. All three utilities can be utilized for making configuration changes. As of NetApp Data ONTAP 7.1, Role Based Access Controls allows segregation of duties for administration of filer security and account administration.

Local user accounts are created and managed by ISAs with the useradmin command from the console command line interface (CLI).

```
useradmin user add <user name> -c "<comment>" -n "<user full name>" -g
<group_name>
```




Figure 10.8 - NetApp Data ONTAP "useradmin" command

With ONTAP 7.1, NetApp supports segregation of duties for security and administration of the filer, where each account can inherit roles and capabilities based on group membership. In order to comply with CITMP Information Security policies every Citigroup NetApp support region will move to a segregated administration model after deploying ONTAP 7.1. The migration to a segregated model is not required upon immediate upgrade, and project management aspect of this infrastructure change will be handled via the CATE organization. Once segregation of duty has been implemented the following principles of account management still apply:

- Post Filer installation; create new system administrator userId's. These userId's must be created using the user's Citigroup SOE GEID (e.g. rs01146). This procedure ensures all ID's are unique and auditable to a known individual.
- "root" and "administrator" accounts (system accounts) should be handed over to an Information Security Officer/Administrator to ensure that these accounts are not used for filer administration.
- Should there be a need to use a system account, the requester should use an auditable and traceable Change Management system to request the use of the account.

- NetApp Filer administration should be performed only by authorized local accounts that are also members of the Domain\NAS_Admns group. Storage administrators should not create any local groups on the filer.
- Ensure that only the groups and roles defined in Appendix J: Role Based Access Control capabilities are defined on the Filer (GDSE is developing an automated process using Operations Manager scripts to centralize this reporting requirement for account management that will be available early in Q3 2007).
- Changes to add or delete local users must use an auditable and traceable Request Management system.
 - o This procedure ensures that there is an auditable trail for all administrative activity done on the filer.

10.6.2.3 Password Composition

Password composition will be enforced using the certain options within ONTAP. These options, which are prefixed with the option name `security.passwd`, can be found in Appendix C: Data ONTAP SOE Configuration Standards.

User account passwords and user names have the following restrictions:

- The password must contain at least six characters, including at least two alphabetical characters and one non-alphabetical character.
- Passwords must contain no more than 14 characters.
- The login name is case-sensitive.
- The login name cannot contain any of the following characters:
* \ : , = / > [< | + ?] ; space




Figure 10.9 - NetApp Data ONTAP UserID password composition enforcement

10.6.2.4 Automatic disabling of accounts

NetApp Data ONTAP currently has the capability to disable accounts based on invalid login attempts. The system log also captures invalid or failed login attempts, when a user has 3 consecutive invalid failed logins.

Starting with ONTAP 7.1, account disabling is a native Data ONTAP feature that is introduced via the following system security option:

`options security.passwd.lockout.numtries`

Refer to section Appendix C: Data ONTAP SOE Configuration Standards for the appropriate value for this option.

When an administrator attempts to ssh to a Filer when their account is locked out they will receive a permission denied message in the SSH client that is similar to the response when an incorrect password is entered.

At the point where an administrator account is locked out, an entry similar to the following example will be recorded in the system message log:

```
Fri Jun 15 09:32:16 EDT [filer1: useradmin.lockedout.user:warning]: User
'nn12345' is locked out of the appliance for failing authentication '6' times.
```

To determine which accounts have been locked out, an account administrator (ISA) who has the useradmin command privilege must run the 'useradmin user list' command against each user account. For example:

```
useradmin user list nn12345
```

Accounts that have been locked out will show a status of disabled in response to the user list query. GDSE is developing an automated process using Operations Manager scripts to centralize this reporting requirement for account management that will be available early in Q3 2007.

In order to reset the password for the locked out administrator, an account administrator (ISA) with the 'security-passwd-change-others' privilege reset the password on the account to change the account status to enabled. At next login the administrator will then be prompted to change their password due to the '**security.passwd.firstlogin.enable**' option. For more information on the security.passwd.firstlogin.enable option consult Appendix C: Data ONTAP SOE Configuration Standards.

- Root cannot be locked out through repeated invalid password attempts
- Unlike Windows, the CIFS administrator account can be locked out via invalid password login attempt. As a privileged functional account CIFS administrator should not be used for day-to-day or automated/scripted device management

10.6.2.5 Password Aging

Automated password expiration should be enforced after 63 days (similar to the account management policies implemented in Active Directory), and account inactivity expiration should be enforced after 120 days. These account policy features are not available in the current version of NetApp Data ONTAP. Password ageing policies are implemented on a per account basis and are not available at the system level in order to provide policy enforcement granularity on different accounts.

At the time of writing this document, functional IDs are currently exempt from ONTAP password aging policies until guidance can be provided by Citigroup's controls team.

All storage administrator accounts must use the password aging process. Password aging is implemented by modifying an administrator's account as follows:

```
useradmin user modify nn12345 -m 1 -M 63
```

The -m argument is the minimum age of the password (in days) before an administrator is allowed to change **their own** password, and the -M argument is the maximum age of the password (in days) before the account is considered expired. Expired accounts are listed with a status of expired when performing a 'useradmin user list nn12345' query. GDSE is developing an automated process using Operations Manager scripts to centralize this reporting requirement for account management that will be available early in Q3 2007. The minimum password age value does not apply if an Account Administrator (ISA) is changing another administrator's password, such that if the previous password is less than 24 hours old then the ISA can still change the password.

There is no ability within ONTAP to determine the current age of a users password.

Refer to the [Release Notes](#) for details of a password aging related bug found when upgrading to ONTAP 7.1 from a previous version.

10.6.2.6 Functional IDs

Data ONTAP today uses two local and one Win2K domain SOE functional ID's for specific processes:

root - Overall system management and Backup.

The current Veritas NetBackup NDMP backup client requires the NetApp Data ONTAP "root" account authentication for its backup services. This limitation will be resolved in the next certified SOE version Data ONTAP, where a Backup functional ID will be used.

Administrator - Win2K domain and CIFS management.

Created when a file is added to Windows Domain. This account is used to manage a CIFS shares created for Windows domain access.

Domain\SRV_filerVirusScan or equivalent - VirusScan server process.

This account is a Windows domain functional id required for the McAfee process used for File virus scanning. This account is added to local administrator group (\filename\administrators) on the filers, and has full access to the CIFS file system.

Creation and use of functional id's should follow the best practices described in section 10.6.2.2. All Functional Ids audit events are captured in the "/etc/log/auditlog" file in the system volume (/vol/vol0 or \filename\c\$).

NetApp Data ONTAP recycles the "/etc/log/auditlog" file every week for six weeks. In order to comply with CITMP IS policies for review of functional Id audit events dating back three months (twelve weeks), Operational Owners must capture a minimum of seven weekly Snapshots of /vol/vol0.

10.6.3 Authorization and Access Control

10.6.3.1 User Access Controls

Access controls for NetApp filers must be considered from two differing perspectives:

- System Administration access controls
- File system access controls

10.6.3.2 System Administrator Access Controls

System Administration access controls are supported in NetApp Data ONTAP. The RBAC model is described in Appendix J: Role Based Access Control capability.

10.6.3.3 File System Access Controls

CIFS Access Controls are managed from the Windows Domain Infrastructure. As such, respective Business' Windows system administrator manage CIFS/NTFS file share access.

10.6.3.4 Windows NT Domain – Active Directory Support

From a Windows domain perspective, each filer looks and behaves like a Windows Server including NetBIOS name resolution. The filers must therefore be members of resource domains. If the primary and backup/DR filers are in a different domain, it should be verified that the same

trust(s) to the master user account domain(s) exist. The computer account for the filer must be created before it is joined into the domain.

Note: Current SOE certified versions of Data ONTAP allow filers to join and participate in mixed-mode and native-mode active directory domains which utilize Windows NT 4.0, Windows 2000 Server and Windows 2003 Server domain controllers for directory lookups and authentication.

During the setup process and when the CIFS protocol restarts, the filer will attempt to sense what type of domain exists on the network. It accomplishes this by identifying what types of domain controllers are available.

Filers will first search for an active directory domain controller by querying DNS for the SRV record of an active directory domain controller. If the filer cannot locate an active directory domain controller it will switch to "NT4 mode" and search for a Windows NT 4.0

10.6.3.5 File Share Access Administration (CIFS)

In addition to the root account mentioned above, filers also have default built-in local groups such as Administrators, Power Users, and Users. These groups can be used to assign different levels of CIFS security administration.

There are certain CIFS limitations to modifying local users and groups:

- Local users on filers can only be created and deleted by using the useradmin command. You cannot use User Manager in Windows NT 4.0 (or Windows 2000) to manage local user accounts.
- Local groups on filers can be created and modified using User Manager on NT 4.0 or Computer Manager on Windows 2000, or the useradmin group command introduced in ONTAP 7.0.

10.6.3.6 System Volume Security (/vol/vol0)

The root volume, e.g. "/vol/vol0" is the system volume where all configuration settings and logs are kept. This volume is shared as **c\$** by default, the same naming convention as Windows Servers. It is important to lock it down, but at the same time, not to impede the security requirements for adequate access and system auditing.

- Principle:
 1. Network Appliance's filers are pre-configured with c\$, etc\$ and home shares on the root volume. The etc\$ and home shares are never used and are covered by the c\$ share. All shares on /vol/vol0 other than the c\$ share MUST be removed.
 2. Define a three-level hierarchy for security management - Users who need access to volume (/vol/vol0) may have very different administrative roles. You should consider three global group types with different rights and permissions, ranging from full-control, limited-access, and read-only.
- Preparation:
 1. **Create two NT Global Groups:** Within every Citigroup Windows domain infrastructure two global groups must be created for c\$ access control. One group will have full control to c\$, and the other will have read-only access. Create "Domain\NetApp_Admins", and "Domain\NetApp_Users".

2. **Add or Verify account, “administrator”, on the filer:** By default, this account is a member of local “Administrators” group. It gives you full access right to CIFS-related setup. You can verify if this account exists with command “*useradmin user list*”.
3. **Test account, “administrator”, on the filer:** You should map a drive to c\$ on the filer with this account to make sure you'll have full access to c\$. You should also test full-control right using NT User Manager on the filer, and verify that you can add/remove/change members to the local groups.

Note: Steps #2, and #3 above ensure that at least one account that has full access to the filer. This account is critical for initial security setup. As stated in section 10.6.2.2, *the “administrator” is a system account, its use should be managed with a Change Management system, and should be the only account with full access control to c\$.*

- Access Control:

1. **Local Administrator Group:** Verify these settings for accounts that are allowed for CIFS full access to the filer.
 - a. Open user manager to remove “[DOMAIN]/administrators” from the local “Administrators” group.
 - b. Add global group “[Domain]\NetApp_Admins” to filer local group “Administrators”.

Note: To securely manage the filer, the domain administrators do not and shall not assume the role of administrators of the filers. This step is necessary because Data ONTAP will add “domain administrators” group once it joins a NT domain.

2. **Share-Level Access Control:** Follow these two steps to give share-level permission to two groups.
 - a. Verify local group “Administrators” has full-control access to c\$ share. This is default setting.
 - b. Grant global group “[Domain]\NetApp_Users” full-control access to c\$ share.

Note: For c\$, this should be done through FilerView.




Figure 10.4 - NetApp Administrator CIFS Shares Management

3. File-Level Access Control: Assign the permission of all files, and directories as follows:
 - a. Map \\filename\c\$ to local host using administrator account.
 - b. Give "[File]\Administrators" group the full-control right to all files and directories.
 - c. Give "[Domain]\NetApp_Users" group the read-only right to all files and directories.




Figure 10.5 - NetApp_Users and NetApps_Administrators Groups C\$ Security Permissions

- Ongoing management:

From time to time, you need to assign permission for c\$ share, you could follow these rules that are based on the design principle stated above.

1. To give full access to share c\$ of all filers, assign the intended account to global group, e.g. "[Domain]\NetApp_Admin".
2. To give full access to share c\$ of a specific filer, assign the intended account to local group "administrators" of the target filer.
3. To give read-only access to share c\$ of all filer, assign the intended account to global group "[Domain]\NetApp_Users".
4. To grant full-access to specific files under share c\$, assign the intended account to global group "[Domain]\NetApp_Users". Then, assign file permission individually.

In figure 10.6 you can view an example. In that example, users in global group "NYC\NetApp_Inventory" will have full access to a configuration file "inventory.src".




Figure 10.6 - NetApp_Inventory Group "inventory.src" File Permissions

NFS Access Controls are managed by NIS Domain access controls. As such, respective Business' NIS administrator manages NFS file systems access.

10.6.3.7 System Administration Entitlements

System administration entitlements are supported in NetApp Data ONTAP. The RBAC model is described in Appendix J: Role Based Access Control capability.

10.6.3.8 File System Access Entitlements

CIFS and NFS file access entitlements are captured through domain authentication, and NIS authentication, respectively. CIFS/NTFS entitlements are managed and controlled by the Business' Windows Infrastructure system administrator.

10.6.4 Compensating Control

10.6.4.1 System Administration Access Controls

System administration access controls must be implemented whereby Filer management is restricted to five specific management stations. GDSE mandates that these systems will consist of:

- 2 SOE terminal servers per support group
 - o One installed at the main production site
 - o One installed at the main contingency site
 - o They have the SSH SOE package installed
 - o These terminal servers should be specifically used for NAS management and should be configured in Application Terminal Services mode.
- 2 DFM servers per support group
 - o One installed at the main production site
 - o One installed at the main contingency site
- 1 Global inventory WebAdmin server (ISOENY12 [162.124.198.180])

10.6.5 Traceable Access Entitlements

10.6.5.1 System Administration Access

The “/etc/log/auditlog” captures every command input and output entered by a user, making all actions traceable back to that specific user who is granted a specific role. A sample of the “/etc/log/auditlog” looks as follows:

```

Tue Apr 13 09:39:21 EDT [netapps02: rc:auditlog]: eg55388:0TT:console shell:netapps02">
Tue Apr 13 09:39:21 EDT [netapps02: rc:auditlog]: eg55388:IN:console shell:
Tue Apr 13 09:39:28 EDT [netapps02: rc:auditlog]: eg55388:0TT:console shell:netapps02">
Tue Apr 13 09:39:28 EDT [netapps02: rc:auditlog]: eg55388:IN:console shell:vscan scanner server
Tue Apr 13 09:39:28 EDT [netapps02: rc:auditlog]: eg55388:0TT:vscan:The following commands are available: for more information
Tue Apr 13 09:39:28 EDT [netapps02: rc:auditlog]: eg55388:0TT:vscan:type "vscan help <command>"
Tue Apr 13 09:39:28 EDT [netapps02: rc:auditlog]: eg55388:0TT:vscan:extensions off options scansec
Tue Apr 13 09:39:28 EDT [netapps02: rc:auditlog]: eg55388:0TT:vscan:help on reset
Tue Apr 13 09:39:33 EDT [netapps02: rc:auditlog]: eg55388:0TT:console shell:netapps02">
Tue Apr 13 09:39:33 EDT [netapps02: rc:auditlog]: eg55388:IN:console shell:vscan scanners
Tue Apr 13 09:39:33 EDT [netapps02: rc:auditlog]: eg55388:0TT:vscan:No virus scanners are registered with the filer.
Wed Apr 14 09:46:51 EDT [netapps02: rc:auditlog]: eg55388:0TT:console shell:netapps02">
Wed Apr 14 09:46:51 EDT [netapps02: rc:auditlog]: eg55388:IN:console shell:
Wed Apr 14 09:47:03 EDT [netapps02: rc:auditlog]: eg55388:0TT:console shell:netapps02">
Wed Apr 14 09:47:03 EDT [netapps02: rc:auditlog]: eg55388:IN:console shell:priv set
Wed Apr 14 09:47:26 EDT [netapps02: rc:auditlog]: eg55388:0TT:console shell:netapps02">
Wed Apr 14 09:47:26 EDT [netapps02: rc:auditlog]: eg55388:IN:console shell:useradd eg55388
Wed Apr 14 09:47:52 EDT [netapps02: rc:auditlog]: eg55388:0TT:useradmin:New password:
Wed Apr 14 09:47:52 EDT [netapps02: rc:auditlog]: eg55388:0TT:useradmin:Password must contain at least 2 alphabetic characters and
Wed Apr 14 09:47:52 EDT [netapps02: rc:auditlog]: eg55388:0TT:useradmin:at least 1 numeric or special character
Wed Apr 14 09:48:33 EDT [netapps02: rc:auditlog]: eg55388:0TT:useradmin:New password:
Wed Apr 14 09:48:33 EDT [netapps02: rc:auditlog]: eg55388:0TT:useradmin:Password is too short - must be atleast 6 characters
Wed Apr 14 09:48:36 EDT [netapps02: rc:auditlog]: eg55388:0TT:useradmin:New password:
Wed Apr 14 09:48:36 EDT [netapps02: rc:auditlog]: eg55388:0TT:useradmin:Password is too short - must be atleast 6 characters
Wed Apr 14 09:48:36 EDT [netapps02: rc:auditlog]: eg55388:0TT:useradmin:Too many failures - try later
Wed Apr 14 09:48:44 EDT [netapps02: rc:auditlog]: eg55388:0TT:console shell:netapps02">
Wed Apr 14 09:48:44 EDT [netapps02: rc:auditlog]: eg55388:IN:console shell:
Wed Apr 14 09:48:48 EDT [netapps02: rc:auditlog]: eg55388:0TT:console shell:netapps02">
Wed Apr 14 09:48:48 EDT [netapps02: rc:auditlog]: eg55388:IN:console shell:useradd eg55382
Wed Apr 14 09:48:56 EDT [netapps02: rc:auditlog]: eg55388:0TT:useradmin:New password:
Wed Apr 14 09:49:05 EDT [netapps02: rc:auditlog]: eg55388:0TT:useradmin:Retype new password:
Wed Apr 14 09:49:05 EDT [netapps02: rc:auditlog]: eg55388:0TT:useradmin:User <eg55382> added.
Wed Apr 14 09:49:08 EDT [netapps02: rc:auditlog]: eg55388:0TT:console shell:netapps02">
Wed Apr 14 09:49:08 EDT [netapps02: rc:auditlog]: eg55388:IN:console shell:

```

Figure 10.7 - NetApp Data ONTAP “/etc/log/auditlog”

10.6.5.2 File System Access

CIFS file access entitlements are captured through domain authentication. Following the inventory collection procedures enables TI Distributed Engineering to capture all CIFS/NTFS file access entitlements. These can be referenced from the following:

<https://webadmin.nam.nsroot.net/main.asp> (Login required). Then select:

1. DPC container, then select;
2. North America container, then select;
3. NetApp Filers container.

Note: The address for WebAdmin will be different on a region to region basis, and the container in step 2 will be different for each region.




Figure 10.8 - NetApp CIFS/NTFS File Access Entitlements

10.6.6 Confidentiality and Integrity

10.6.6.1 Approved Anti-Virus

NetApp Filers configured for SOE CIFS/NTFS file system access, must conform to CITMP IS policies requiring approved anti-virus software be installed and maintained to protect the file system against virus attacks. Citigroup SOE certifies McAfee VirusScan as its standard anti-virus software. As of this writing, TI Distributed Engineering has certified McAfee VirusScan for NetApp v7.1 as the standard anti-virus software for NetApp filers.

Refer to section 10.4 for more information on Anti-Virus configuration.

10.6.6.2 Support by Approved and Certified Vendor

Network Appliance (NetApp) is a certified vendor in the current Citigroup Technology Catalog. TI Distributed Engineering-GDSE owns and maintains all current SOE certification and best practices documents including a product matrix of current and future certified products and operating system versions. Please refer to the following for all certification information: <http://engineering.ti.citigroup.net/gdse/products/platforms/nasfiler/>. The NetApp Data ONTAP operating system is also in the Citigroup Vulnerability Threat Matrix (VTM) inventory, and is scrutinized for UNIX and Windows vulnerabilities.

10.6.6.3 Disposal of failed disks

NetApp filers use RAID-4 and RAID-DP parity data protection for all volumes. In the event that any disk drive fails, the volume enters a degraded state, while reconstructing the data "on the fly" from a spare disk, with no interruption in file service.

The Filer uses industry standard media (disks) to store Citigroup "Medium Risk" Confidential information. Currently, all failed media should be processed in accordance Global Data Center Disk Media Destruction policies.

10.6.7 Incident Detection and Response

10.6.7.1 Login Violations

As of version 7.1, ONTAP includes a policy threshold for login violations; whereby the system message log captures when an account is locked out. For more information on this feature refer to section 10.6.2.4.

10.6.7.2 Audit Trails

NetApp Data ONTAP captures audit events for both:

- System administration, and
- File system access.

10.6.7.3 System Administration Audit Trails

For system administration all events are captured to the "/etc/log/auditlog" file. The /etc/log/auditlog file has two significant limitations:

- a. As mentioned in section 10.6.3 the "/etc/log/auditlog" is recycled after six weeks.
- b. In addition, the filer administrator has the capability to manipulate the file's content.

Current audit trail logging implementations rely on the Syslog process to extract the device log.

To mitigate these risks shown above, TI Distributed Engineering will certify a regulatory compliance version of NetApp Data ONTAP that will allow for storage volumes to be created as Write-Once-Read-Many (WORM) volumes. Snapshots of the filer system volume will be replicated to the WORM volume with a retention period, whereby the replicated content cannot be altered until the retention period has expired.

Operational Owners will be responsible for producing Level 3 PCMs that describe the process, number of snapshots and retention periods used to assure CITMP IS policy controls.

10.6.7.4 File System Access Audit Trails

NetApp Data ONTAP has the capability to capture every file system access, logon, and logoff events. This capability requires the "cifs.audit.enable" option **on**. The output of these events are captured in the "/etc/log/auditlog" and "/etc/log/adtlog.evt" files. The "/etc/log/adtlog.evt" is viewed using any Windows server or workstation "Event Viewer" utility. This option default is off. To capture these events the following options must also be enabled on:

1. For file access events to be generated, the option "cifs.audit.file_access_events.enable" must also be on. This option will only capture access events if system access control list exist. This option default is on.
2. For logon and logoff events to be generated, the option cifs.audit.logon_events.enable must also be on. Logon and logoff events reflect CIFS session connects and disconnects, respectively. This option default is on.

Enabling these options generates an overflow of information to be captured to both the system log and "adtllog.evt" files, and does not relate to any operational concerns of the system. TI Distributed Engineering-GDSE SOE policy recommends the "cifs.audit.enable" option remain in its default setting **off**. Operational Owners should enable this option only upon request of a Business Owner to monitor or audit file(s), if foul play is suspected.

The CIFS audit options can be found in Appendix C: Data ONTAP SOE Configuration Standards, where the options begin cifs.audit.




Figure 10.10 - NetApp Data ONTAP "adtllog.evt" file Event View

10.6.7.5 Audit Trails Review

NetApp Data ONTAP captures all audit events. However, Operational Owners must produce a Level 3 PCM that outlines procedures for periodic review of the "/etc/log/auditlog" file. Review should encompass audit events of possible incident detection and authorized usage.

10.6.8 Administration

10.6.8.1 Information Security Administration (ISA)

The ISA role is introduced in ONTAP 7.1. Please refer to Appendix J: Role Based Access Control capability for details as to ISA capabilities. The migration to a segregated model is not required upon immediate upgrade to ONTAP 7.1, and the project management aspect of this infrastructure change will be handled via the CATE organization.

10.6.8.2 Security Standards

TI Distributed Engineering-GDSE in cooperation with TI Storage Integration has produced a standard set of Data ONTAP security configuration options that must be adhered in order to comply with CITMP IS policy. These options are defined in Appendix C: Data ONTAP SOE Configuration Standards as are highlighted as the entries with the red background.

10.6.8.3 SOE Security Options

The SOE security options can be found in Appendix C: Data ONTAP SOE Configuration Standards. The relevant entries are highlighted in the cells with a red background.

10.6.8.4 Inventory Agent for Network Appliance Filers

The Inventory Agent for Network Appliance is an in-house tool developed by Wintel Engineering. The agent runs a list of Data ONTAP system commands from a central inventory server to pull system and volume information off the filers. The output is then parsed and saved to the SOE Database. The inventory data is archived for 30 days. The data schema is shown in Appendix F: Data Schema for NetApp Inventory Agent.

Operational Owners must include their filers in the inventory process. To facilitate filer inventory collection the following steps must be taken:

- a. The agent needs read-only rights to root volume (/vol/vol0), and
- b. Full-control right to 'hosts.equiv' and 'inventory.src' files under '/etc/' directory.

10.6.8.5 Configuration

Ensure the NAM\NetApp_Inventory group is trusted in your domains.

1. System Inventory

- a. Share Permission:
 - i. Grant full-control right of c\$ share to group NAM\NetApp_Inventory.
- b. File Permission:
 - i. Grant group "NAM\NetApp_Inventory" the read-only right to all files/directories on c\$.
 - ii. Grant group "NAM\NetApp_Inventory" the full-control right to 'hosts.equiv' under '/etc' directory on c\$. the file
 - iii. Create a blank file "inventory.src" under '/etc' directory, and grant group "NAM\NetApp_Inventory" the full-control right of this file.
- c. The inventory collection initiation is a manual process. Contact the North America Lan Ops team to start the inventory process for your filer.

Post-Configuration Examples:

Once the process is successfully completed, the 'hosts.equiv' file will have the entries similar defined below:

```
isoeny12.ny.ssmb.com    srv_sqlserver_soedb
isoedl01.nam.nsroot.net srv_sqlserver_soedb
```

In this example, the agent runs as 'srv_sqlserver_soedb' account on server 'isoeny12.ny.ssmb.com'.

The following are the list of commands in the 'inventory.src' file that is executed by the inventory agent:

df	options	snmp
df -i	snap sched	storage show disk
ifconfig -a	snap list	sysconfig -a
nbtstat	sysconfig -r	vif status
		vol status

Table 10.4 – Citigroup SOE NetApp Inventory Collection - Data ONTAP commands




Figure 10.11 - Citigroup SOE WebAdmin NetApp Filer Inventory Report

2. Entitlement Review Inventory

In addition to system options and configuration inventory the agent also collects Entitlement or Access Controls inventory for audit review. The collection process parses this data and stores it to the Wintel Engineering SOE database where the information can be viewed from the DPC tree in WebAdmin.

For inventory agent to collect the permission assignments of the shared folder, the inventory agent must have proper access. Granting this access is usually performed by SA group that assigned the shares ACLs. The following configuration is required to collect file system user entitlements data:

- Make NetApp_Inventory a member of your local "*filename*\PowerUsers" group.
- grant 'read attribute' and 'read permission' access. Assign 'read-only' rights of the shared folder to group "*filename*\Power Users".

NOTE: Obviously, SA should grant the access only to the shared folder. Collecting shared folder permission is not mandatory. However, without granting the proper rights, the shared folder permissions data will not be available in WebAdmin.

The inventory data is published in the DPC\Region\NetApp_Filers container in WebAdmin.




Figure 10.12 - Citigroup SOE WebAdmin NetApp Filer Entitlement Review Report

10.6.8.6 Vulnerability Assessment

NetApp Data ONTAP is a UNIX based kernel susceptible to some of the same vulnerabilities that affect UNIX systems. As such, NetApp Data ONTAP is tracked by Citigroup's VTM database as a UNIX system. Please refer to the VTM website for the latest vulnerability resolutions:
<https://stcprod001.nj.ssmb.com/eii/main/program/index.aspx?programid=231>

10.6.8.7 Cryptography and PKI

CITMP IS policy requires all default access protocols are disabled for remote system management. The use of an approved encryption protocol must be used for remote administration of system resources. NetApp provides optional software that enables support for SSH and SSL access to the filer system console.

Unlike pre 7.0 releases of Data ONTAP, SecureAdmin is installed by default on systems shipped directly from the vendor. There is no requirement to have to separately install SecureAdmin, however SSH and SSL will still require configuration during the system installation phase.

10.6.8.8 Secure Filer Access (SecureAdmin)

SecureAdmin provides data encryption for transporting encrypted passwords, commands, arguments, and results to and from administrative clients through a secure authenticated channel between a management host and Data ONTAP using either of the following protocols:

- Secure Shell (SSH) protocol—Provides an encrypted remote shell and interactive network session.
Note: SecureAdmin v3.0 supports SSH v1 and SSH v2 clients, but SSH v1 is disabled as part of the Filers best practice security configuration.
- Secure Sockets Layer (SSL) protocol—Provides encrypted Web access to NetApp FilerView web interface.

SecureAdmin supports the following encryption algorithms:

- RSA/DSA 2048 bit
- 3DES in CBC mode
- HMAC-SHA1
- HMAC-MD5

NOTE: The system inventory process described in section 10.6.8.4 requires SecureAdmin v3.0 on all filers and uses OpenSSH v3.7.1.1 for encrypted communications between the inventory host agent and the filers.

For detailed SecureAdmin information please refer to the ONTAP System Administration Guide.

10.6.8.9 SecureAdmin Configuration - SSH

1. Start web interface 'FilerView'.
2. Select Menu item 'SecureAdmin->SSH->Configure'.
3. Enter host key strength – 2048 (default).
4. Enter server key strength – 1536 (default).
5. Click 'Generate Keys'.
6. Select Menu item 'SecureAdmin->SSH->Enable/Disable'.
7. Enable SSH.

(Note: The generated host key is stored under /etc/sshd.)



Figure 10.13 - NetApp FilerView - SecureAdmin SSH Configuration Screen

10.6.8.10 Citigroup SOE SSH Clients

Citigroup SOE standard SSH client is OpenSSH v 3.7.1.2. OpenSSH client is available from GDSGOLD as SPAC ID "ossh3712".

The install path for the package requires a valid disk based D: drive on the local system. The complete install path is hard-coded into the package. Refer to the package readme file for details.

10.6.8.11 SecureAdmin Configuration – SSL

1. Start web interface 'FilerView'
2. Select Menu item 'SecureAdmin->SSL->Configure'.
3. Select 'Generate Certificates'.
4. Enter self-signed certificate information. (See figure below)
5. Click 'Generate Certificates'.
6. Select Menu item 'SecureAdmin->SSL->Enable/Disable'.
7. Enable SSL.

(Note: The generated certificate is stored under /etc/ssld.)

As of ONTAP 7.0, when creating SSL keys via GUI, the default key length used is **2048** bits. When creating SSL certificates, set the 'Days to Live' value of the certificate at **730 days**.

Figure 10.15 - NetApp FilerView - SecureAdmin SSL Configuration Screen

10.6.8.12 Citigroup SOE SSL Web Browser Support

SSL is built into all major browsers; therefore, installing a digital certificate on the filer enables the SSL capabilities between filer and browser. SOE web browsers, Netscape 4.7 and Internet Explorer 5.5 or above, are compatible with SecureAdmin v3.0 SSL implementation.

Unlike using FilerView to send the filer password in plain text, using SSL and Secure FilerView improves security by encrypting the filer password and all administrative communication when you manage your filer from a browser.

11 Document version control

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

Version	Changes
0.1	Changes from ONTAP 7.0 solutions guide: • Remove Certification Scope (formerly Chapter 8) as this is repetition of the contents of the Release Notes document • Change template default font size to 10pt • Move Anti Virus section from Best Practice chapter to Security chapter • Remove Appendic D, approved Citigroup Cryptographic Standards as this is repetition of information found in CISS policy doc • Add Assumptions, Definitions, Conventions, Abbreviations, References section • Move core ONTAP options settings section to start of Best Practice chapter • Added installation section to start of Best Practice chapter; detailing platform requirements, upgrade prerequisites, system scalability and upgrade procedure • Move Disk Sanitization section from Best Practice chapter to Security chapter • Remove old references chapter from end of document, replace with references sub-section at start of document • Move Domain Name and Host Resolution; Filer Naming Standards and Time Synchronization sections under Network Configuration section • Create Storage Configuration chapter and move Aggregate, Volume, Spare Disks, Clustering, Snapshot and Snapmirror sections into new chapter • Create Monitoring and Alerting chapter and move AutoSupport and SNMP into new chapter • Remove redundant Filerview screenshot from SNMP section and Networking (VIF creation) section • Update NTP and DNS sections with new IPNS VIP standards • Add TimeZone best practices to NTP section • Added Unsupported Configurations section to Best Practices chapter in response to FAQs • Add ONTAP version number, recommended alert severities for Netcool and Tivoli platforms, and minimum Netcool / Tivoli rule engine requirements to SNMP MIB appendix • Remove redundant Alert Group and Category columns from SNMP MIB appendix • Rename original 'severity' column to default (vendor) severity column in SNMP MIB appendix • Remove redundant FilerView screenshot from Snapshot section • Remove redundant cli screenshot from Snapshot File-folding section • Remove redundant FilerView screenshot from Snapmirror section • Remove redundant FilerView screenshot from Autosupport section • Move SOE vFiler creation example to Appendix J • Consolidated list of all options settings in Appendix C. Removed list of ONTAP options settings from Best Practice section SOE security options section and referenced Appendix C • Remove redundant Filerview screenshots from Secureadmin section • Add GDSE recommended vFiler limits to vFiler section • Remove redundant CLI screenshot from COFS audit security section • Add 'supported from' column to ONTAP options matrix

Version	Changes
0.2	- Remove local group table for CIFS/NTFS privileges, superceded by RBAC appendices - Update 'Support for Unique User ID', 'Password Composition', 'Automatic Disabling of Account', and 'Password Aging' chapters in the Security Best Practices section to reflect RBAC

12 Appendix A: ePO & VirusScan v7.1 for NetApp Configuration

12.1 VirusScan for NetApp v7.1 ePO Default Policies configuration

As a result of the .nap file installation, the new VirusScan for NetApp Software option appears in the Directory Tree of the ePO console. This is the level where all VirusScan default configuration settings will be managed for all VirusScan for NetApp v7.1 servers managed by ePO.

1. Select the Directory globe.
 - a. In the Policies tab,
 - b. Open the VirusScan for NetApp 7.1 software.
 - c. select "General Policies".
 - d. In the "General" tab
 - e. Deselect "Inherit"
 - f. Select the properties shown in the figure below.




Figure 12.1 - ePO General Policies - General tab Policy Settings

2. Select the "Messages" tab
 - a. Deselect "Inherit".
 - b. Select all the properties shown in the figure below.




Figure 12.2 - ePO General Policies - Messages tab Policy Settings

3. Select "Reports" tab
 - a. Deselect "Inherit"
 - b. Select all the properties shown in the figure below.




Figure 12.3 - ePO General Policies - Reports tab Policy Settings

4. Select the “Default Processes Policies” option
 - a. In the “Processes” tab
 - b. Deselect “Inherit”
 - c. Select all the properties shown in the figure below.




Figure 12.4 - ePO Default Processes Policies - Processes Policy Settings

5. Select the “Detection” tab
 - a. Deselect “Inherit”
 - b. Select all the properties shown in the figure below.
 - c. NOTE: There are no file exclusions configured. Be aware that any file exclusions set at the server will be overwritten by this policy.




Figure 12.5 - ePO Default Processes Policies - Detection Policy Settings

6. Select the "Advanced" tab
 - a. Deselect "Inherit"
 - b. Select all the properties shown in the figure below.




Figure 12.6 - ePO Default Processes Policies - Advanced Policy Settings

7. Select the “Actions” tab
 - a. Deselect “Inherit”
 - b. Select all the properties shown in the figure below.




Figure 12.7 - ePO Default Processes Policies - Actions Policy Settings

Skip Low-Risk and High-Risk Process Policies. These are inherited from the “Default Processes Policies”. See Step 4.

8. Select the “User Interface Policies” option
 - a. In the “Display Options” tab
 - b. Deselect “Inherit”
 - c. Select all the properties shown in the figure below.




Figure 12.8 - ePO User Interface Policies - Display Options Settings

9. Select the "Password Options" tab
 - a. Deselect "Inherit"
 - b. Select all the properties shown in the figure below.




Figure 12.9 - ePO User Interface Policies - Password Options Settings

10. Select the “Alert Manager Alert Policies” option
 - a. In the “Alert Manager Alerts” tab
 - b. Deselect “Inherit”
 - c. Select all the properties shown in the figure below




Figure 12.10 - ePO Alert Manager Alert Policies – Alert Manager Alerts Settings

11. Select the “NetApp Policies” option
 - a. In the “General” tab
 - b. Deselect “Inherit”
 - c. Select all the properties shown in the figure below.

NOTE: No NetApp filers will be added at this level of policy because adding them here will cause every VirusScan for NetApp server managed by ePO to possibly scan every NetApp filer configured here. The policy configuration of NetApp Filers will be managed from the organizational containers in the Repository tree.




Figure 12.11 - ePO NetApp Policies - General settings

This completes the ePO Directory level policies configuration for the VirusScan for NetApp v7.1 software. The following section will address the NetApp Policies form an organizational and server level.

12.2 VirusScan for NetApp v7.1 AutoUpdate and Weekly Scan Settings

In addition to the previous policy configuration settings. The ePO administrator must also schedule the AutoUpdate and Weekly Scan job for the VirusScan for NetApp v7.1 server.

1. From the Global Directory
 - a) In the right pane, Select the Tasks tab.
 - b) Right-mouse click, select Schedule Tasks
 - c) In the Task tab
 - d) Name – VirusScan for NetApp v7.1 7AM Daily DAT Update
 - e) Select all other default properties
 - f) Select the Schedule tab
 - g) Use the property setting s shown in the figure below




Figure 12.12 - ePO VirusScan for NetApp v7.1 Daily DAT Update Task Configuration




Figure 12.13 - ePO VirusScan for NetApp v7.1 Daily DAT Update Schedule Configuration

2. Create a Weekly Scan of the VirusScan for NetApp v7.1 server
 - a) Right-mouse click, Create New Task
 - b) In the Task tab
 - c) Name – VirusScan for NetApp v7.1 Weekly Scan on Sunday 8AM
 - d) Select all other default properties
 - e) Select the Schedule tab
 - f) Use the property settings s shown in the figure below




Figure 12.14 - ePO VirusScan for NetApp v7.1 Server Weekly Scan Task Configuration




Figure 12.15 - ePO VirusScan for NetApp v7.1 Server Weekly Scan Schedule Configuration

13 Appendix B: NetApp SNMP MIB

New traps added to this version of ONTAP are highlighted with a green background  .

For puposes of certification a trap is considered new to ONTAP if it hase been added to the MIB since the previous Citigroup certified release of ONTAP (7.0). Table 13.1 (below) contains the column definitions used in the SNMP MIB matrix (Table 13.2)

Table 13.1 - SNMP trap list table definitions

ONTAP Version	The earliest Citigroup release of ONTAP to include this SNMP trap
MIB Version	The MIB version number corresponding to the ONTAP version from the previous column
OID / Trap #	The Object Identifier / trap number An OID may relate to one component such as RAID disk, but each OID may contain more than one message string meaning it is possible to receive multiple differing alerts with the same OID.
Trap Name	The name of the trap
Description	Trap description
# Different Message Strings (sub-types)	The number of different message string permutations assigned to a single OID. This information is useful when trying to generate a trap with a specific message string. If the number in this field is a hyperlink, click on the link to view the ONTAP 'ems invoke' command used to test the trap. If the number is greater than one then there will be more that one 'ems invoke' command per OID. Click on the entry in this field to jump to the relevant section of this SNMP appendix for the procedure on how to test each trap. An entry of 'n/a' means that sub-types are not supported with that particular trap type and therefore cannot be tested with the 'ems invoke' command.
Default Severity	The severity assigned to the trap by default (by NetApp) using the trap severity rules (see trap definitions at end of /etc/mib/netapp.mib)
Tivoli Severity	Citigroup recommended severity level for use with Tivoli monitoring infrastructure. Tivoli is currently in the process of being deployed globally as the standard monitoring tool for distributed systems. Check with your regional monitoring team to confirm whether Tivoli or Netcool is available at your location. Citigroup command center policy should be to invoke support callout for <u>Fatal</u> and <u>Critical</u> severity events. Check with your regional command center team that the Tivoli rule set used in their environment passes Fatal and Critical messages to a command center console. The minimum Tivoli rule engine required to support ONTAP 7.1 is <u><insert Tivoli rule engine version number here></u>
Netcool Severity	Citigroup recommended severity level for use with Netcool monitoring infrastructure. Check with your regional monitoring team to confirm whether Tivoli or Netcool is available at your location. <u><insert statement regarding severity levels that will invoke callout></u> The minimum Netcool rule engine required to support ONTAP 7.1 is <u><insert Netcool rule engine version number here></u>
Auto Clear	This event supports automatic event clearing / cancellation
Clear OID	The OID / trap number that clears this trap / alert condition. N/A means this alert does not support auto-clearing functionality

Table 13.2 - SNMP MIB matrix

Introduced		OID / Trap #	Trap Name	Description	# Different Message Strings (sub-types)	Default Severity	Citigroup Tivoli Severity	Citigroup Netcool Severity	Auto Clear	Clear OID
ONTAP Version	MIB #									
6.4.x	1.13	2	userDefined	A polling-style trap built using the 'snmp traps' command on the filer.	n/a	Alert	Harmless		Yes	2
6.4.x	1.13	6	dhmNotice DegradedIO	Disk Health Monitor - Reported a Disk Degraded-I/O Event	5	Notification	Warning		No	N/A
6.4.x	1.13	7	dhmNotice PFAEvent	Disk Health Monitor - Reported a Disk Predictive-Failure Event	n/a	Informational	Warning		No	N/A
6.4.x	1.13	11	emergency Trap	A user-defined trap indicating an extremely urgent situation, usually indicating that the system has failed and is shutting down.	n/a	Emergency	Fatal		No	N/A
6.4.x	1.13	12	alertTrap	A user-defined trap indicating a condition that should be corrected immediately, such as a failed disk.	n/a	Alert	Critical		No	N/A
6.4.x	1.13	13	criticalTrap	A user-defined trap indicating a critical condition, such as a hard device error.	n/a	Critical	Critical		No	N/A
6.4.x	1.13	14	errorTrap	A user-defined trap indicating an error condition, such as a mistake in a configuration file.	n/a	Error	Critical		No	N/A
6.4.x	1.13	15	warningTrap	A user-defined trap indicating a condition which is not an error, but which may require special handling.	n/a	Warning	Critical		No	N/A
6.4.x	1.13	16	notification Trap	A user-defined trap meant to provide information, such as an hourly uptime message.	n/a	Notification	Critical		No	N/A
7.0.3 / 7.0.4	1.17.3	17	informational Trap	A user-defined trap used for informational purposes	n/a	Informational	Critical		No	N/A
7.0.3 / 7.0.4	1.17.3	18	dbgTrap	A user-defined trap used for debugging purposes. This trap was previously mis-documented as being trap number 17, with the name debugTrap. It has always, actually, been trap 18. The name change is to avoid issues with environments that already have a record for debugTrap	n/a	Debug	Warning		No	N/A
6.4.x	1.13	21	diskFailed Shutdown	System is shutting down because the system has been running in degraded mode for 24 hours. The trap includes a string describing the failed disk.	1	Emergency	Fatal		Yes	26
6.4.x	1.13	22	diskFailed	One or more disks failed. The trap includes a string describing the failed disk(s).	34	Alert	Critical		Yes	26
6.4.x	1.13	26	diskRepaired	The failed disks have been repaired.	n/a	Notification	Harmless		Yes	26
6.4.x	1.13	31	fanFailure Shutdown	Critical chassis or cpu fans have failed and the system is shutting down.	1	Emergency	Fatal		Yes	36
6.4.x	1.13	33	fanFailed	One or more chassis fans failed. The trap includes a string describing the failed fan(s)	1	Critical	Critical		Yes	36
6.4.x	1.13	35	fanWarning	One or more chassis or cpu fans are in warning state. The trap includes a string describing the fan(s) in the warning state.	1	Warning	Warning		Yes	36
6.4.x	1.13	36	fanRepaired	All fans are repaired.	1	Notification	Harmless		Yes	36

Introduced		OID / Trap #	Trap Name	Description	# Different Message Strings (sub-types)	Default Severity	Citigroup Tivoli Severity	Citigroup Netcool Severity	Auto Clear	Clear OID
ONTAP Version	MIB #									
6.4.x	1.13	41	powerSupply Failure Shutdown	Critical power supplies or power rails failed and the system is shutting down.	1	Emergency	Fatal		Yes	46
6.4.x	1.13	43	powerSupply Failed	One or more redundant power supplies failed. Include in the trap a string describing the failed power supply(ies).	1	Critical	Critical		Yes	46
6.4.x	1.13	45	powerSupply Warning	One or more power supplies or power rails in warning state. Include in the trap a string describing the power supply(ies) or the power rail(s) in the warning state.	1	Warning	Warning		Yes	46
6.4.x	1.13	46	powerSupply Repaired	Previously-failed power supplies or power rails have been repaired.	1	Notification	Harmless		Yes	46
6.4.x	1.13	55	cpuTooBusy	CPU utilization exceeds 90%. This trap is not enabled by default. To enable this trap set the registry entry options.monitor.cpu.enable to on. Note that as the threshold for this trap is checked once a minute it is possible to receive multiple instances of this trap in a short time.	1	Warning	Warning		Yes	56
6.4.x	1.13	56	cpuOk	CPU utilization has dropped back below 90%. This trap is a placeholder - it is not currently sent by ONTAP.	1	Notification	Harmless		No	N/A
6.4.x	1.13	62	nvrAmBattery Discharged	The NVRAM battery is fully discharged.	2	Alert	Critical		No	N/A
6.4.x	1.13	63	nvrAmBattery Low	The charge in the NVRAM battery is low.	8	Critical	Fatal		No	N/A
6.4.x	1.13	72	clusterNode Failed	A node in a Cluster FailOver configuration failed. Its partner will assume service for the failed node.	1	Alert	Fatal		Yes	76
6.4.x	1.13	75	clusterNode TakenOver	The partner has taken over for a failed cluster node.	1	Warning	Fatal		Yes	76
6.4.x	1.13	76	clusterNode Repaired	A cluster node has resumed operation.	1	Notification	Harmless		Yes	76
6.4.x	1.13	82	volumeFull	At least one volume is more than 98% full. The string sent with the trap gives the name of the volume or volumes which exceed the threshold.	1	Alert	Critical		Yes	86
6.4.x	1.13	85	volumeNearly Full	At least one volume is more than 95% full. The string sent with the trap gives the name of the volume or volumes which exceed the threshold.	1	Warning	Critical		Yes	86
6.4.x	1.13	86	volume Repaired	All volumes are now under 95% full.	1	Notification	Harmless		Yes	86
6.4.x	1.13	91	overTemp Shutdown	System temperature is too high to continue operating. The system is shutting down.	3	Emergency	Fatal		Yes	96
6.4.x	1.13	95	overTemp	System temperature is too high and in the warning level.	1	Warning	Critical		Yes	96
6.4.x	1.13	96	overTemp Repaired	System temperature has returned to an acceptable value.	1	Notification	Harmless		Yes	96
6.4.x	1.13	103	shelfFault	A disk storage shelf reported a fault, probably due to a problem with drive placement, fans, power, or temperature.	2	Critical	Critical		Yes	106

Introduced		OID / Trap #	Trap Name	Description	# Different Message Strings (sub-types)	Default Severity	Citigroup Tivoli Severity	Citigroup Netcool Severity	Auto Clear	Clear OID
ONTAP Version	MIB #									
6.4.x	1.13	106	shelfRepaired	A previously-reported shelf fault is now corrected.	2	Notification	Harmless		Yes	106
6.4.x	1.13	111	GlobalStatus Non Recoverable	The appliance's overall status changed to 'nonRecoverable', indicating a problem so severe that the appliance is shutting down.	1	Emergency	Fatal		Yes	116
6.4.x	1.13	113	globalStatus Critical	The appliance's overall status changed to 'critical', indicating a problem that needs immediate attention.	1	Critical	Critical		Yes	116
6.4.x	1.13	115	globalStatus NonCritical	The appliance's overall status changed to 'nonCritical', indicating a problem that needs attention.	1	Warning	Warning		Yes	116
6.4.x	1.13	116	globalStatus Ok	The appliance's overall status returned to normal.	1	Notification	Harmless		Yes	116
6.4.x	1.13	126	softQuota Exceeded	A user has exceeded their soft quota limit.	1	Notification	Warning		Yes	127
6.4.x	1.13	127	softQuota Normal	A user is safely back under their soft quota limit.	1	Informational	Harmless		Yes	127
6.4.x	1.13	134	autosupport SendError	Unable to send autosupport. The trap includes a string describing the reason for the failure.	25	Error	Warning		Yes	136
6.4.x	1.13	135	autosupport Configuration Error	Autosupport may be configured incorrectly. The trap includes a string describing the mis-configuration.	7	Warning	Critical		Yes	136
6.4.x	1.13	136	autosupport Sent	Autosupport was sent successfully.	6	Notification	Harmless		Yes	136
7.0.3 / 7.0.4	1.17.3	142	upsLinePower Off	UPS: Input line power has failed and UPS is now on battery	1	Alert	Harmless		Yes	146
7.0.3 / 7.0.4	1.17.3	143	upsBattery Critical	UPS: Battery is nearly exhausted, starting graceful shutdown	1	Critical	Harmless		Yes	146
7.0.3 / 7.0.4	1.17.3	144	upsShutting Down	UPS: Shutting down now: Time left on battery is exhausted	1	Error	Harmless		Yes	146
7.0.3 / 7.0.4	1.17.3	145	upsBattery Warning	UPS: Warning: Time left on battery is getting critical	1	Warning	Harmless		Yes	146
7.0.3 / 7.0.4	1.17.3	146	upsLinePower Restored	UPS: Input line power has been restored and UPS is now off battery	1	Notification	Harmless		Yes	146
7.0.3 / 7.0.4	1.17.3	151	app Emergency	The application encountered an extremely urgent situation and requires an immediate response	1	Emergency	Harmless		No	N/A
7.0.3 / 7.0.4	1.17.3	152	appAlert	The application is in a condition that should be corrected immediately	1	Alert	Harmless		No	N/A
7.0.3 / 7.0.4	1.17.3	153	appCritical	The application encountered a critical condition	1	Critical	Harmless		No	N/A
7.0.3 / 7.0.4	1.17.3	154	appError	The application encountered an error condition	1	Error	Harmless		No	N/A

Introduced		OID / Trap #	Trap Name	Description	# Different Message Strings (sub-types)	Default Severity	Citigroup Tivoli Severity	Citigroup Netcool Severity	Auto Clear	Clear OID
ONTAP Version	MIB #									
7.0.3 / 7.0.4	1.17.3	155	appWarning	The application is in a condition which is not an error, but which may require special handling	1	Warning	Harmless		No	N/A
7.0.3 / 7.0.4	1.17.3	156	appNotice	The application is notifying of a certain event	1	Notification	Harmless		No	N/A
7.0.3 / 7.0.4	1.17.3	157	appInfo	The application's message is meant for informational purposes	1	Informational	Harmless		No	N/A
7.0.3 / 7.0.4	1.17.3	158	appTrap	The application requires to be debugged	1	Debug	Harmless		No	N/A
6.4.x	1.13	162	alfFilewrap	The internal audit file has wrapped. You are currently losing event records.Warning the user	1	Alert	Warning		No	N/A
6.4.x	1.13	166	alfFileSaved	The internal audit log has been autosaved to an external evt file. Notifying the user.	1	Notification	Harmless		No	N/A
6.4.x	1.13	167	alfFileNearly Full	The internal audit log file is nearly full. The file is going to wrap. Notifying the user	1	Informational	Warning		No	N/A
6.4.x	1.13	176	Quota Exceeded	One of the quota limits has been exceeded. The trap includes a string describing the specific event which triggered the trap.	1	Notification	Critical		Yes	177
6.4.x	1.13	177	quotaNormal	One of the quota limits has gone back down to a normal level. The trap includes a string describing the specific event which triggered the trap.	1	Informational	Harmless		Yes	177
6.4.x	1.13	187	wafDirFull	The directory has been filled to its limit	1	Informational	Critical		No	N/A
6.4.x	1.13	192	eccSummary	Memory ECC: number of new correctable ECC errors	1	Alert	Warning		No	N/A
6.4.x	1.13	195	eccMasked	Memory ECC: High frequency of ECC errors	1	Warning	Critical		No	N/A
7.0.3 / 7.0.4	1.17.3	204	ftpdError	Ftp daemon: service stopped	2	Error	Warning		No	N/A
7.0.3 / 7.0.4	1.17.3	206	ftpdMaxConn Notice	Ftp daemon: number of connections hits maximum number	2	Notification	Warning		No	N/A
7.0.3 / 7.0.4	1.17.3	216	ftpdMaxConn Threshold Notice	Ftp daemon: number of connections nearly hits maximum number	2	Notification	Warning		No	N/A
7.1.2P1	1.18.1	222	scsitgtFCP LinkBreak	SCSI Target: Link Break on FCP adapter	2	Alert	Critical		No	N/A
7.0.3 / 7.0.4	1.17.3	224	scsitgtPartner Path Misconfigured	SCSI Target: FCP Partner Path Misconfigured	1	Error	Critical		No	N/A
7.0.3 / 7.0.4	1.17.3	226	ScsitgtThrottle Notice	SCSI Target: Throttle limit event notification	2	Notification	Warning		No	N/A
6.5.x	1.16	237	vifPrimaryLink Failed	Vif: The primary Interface on a Single mode vif has failed	1	Informational	Critical		No	N/A
6.5.x	1.16	238	vifAllLinks Failed	Vif: All the links of the vif have failed	1	Debug	Critical		No	N/A

Introduced		OID / Trap #	Trap Name	Description	# Different Message Strings (sub-types)	Default Severity	Citigroup Tivoli Severity	Citigroup Netcool Severity	Auto Clear	Clear OID
ONTAP Version	MIB #									
6.5.x	1.16	245	vfStopped	A vFiler has stopped.	1	Warning	Critical		No	246
6.5.x	1.16	246	vfStarted	A vFiler has started.	1	Notification	Harmless		Yes	246
6.5.x	1.16	254	vscanVirus DetectedError	VScan: Vscan scanner has detected a virus on the Filer	2	Error	Warning		No	N/A
6.5.x	1.16	255	vscanDis Connection	VScan: There has been a connection to the vscan servers dropped	2	Warning	Critical		No	N/A
6.5.x	1.16	256	vscan Configuration Change	VScan: There has been a change to the vscan configuration	11	Notification	Warning		No	N/A
7.0.3 / 7.0.4	1.17.3	257	vscan Connection	VScan: A virus scanner has made a connection to the Filer	1	Informational	Warning		No	N/A
6.5.x	1.16	266	vscanServerU pgrade	VScan: The Vscan server has been upgraded	1	Notification	Harmless		No	N/A
6.5.x	1.16	272	volume Restricted ByMirrorBiglo	A volume that experienced a medium error during reconstruction is restricted and marked waff-inconsistent, but starting waffiron has failed. This trap is issued to alert operator that a volume is not accessible and waffiron must be started to allow access to it.	2	Alert	Critical		No	N/A
6.5.x	1.16	274	volume Inconsistent Umount	This trap is issued when we unmount a volume due to an inconsistency.	1	Error	Critical		No	N/A
6.5.x	1.16	275	volumeState Changed	Volume is being taken offline or being restricted. The string sent with trap specifies name of affected volume and its state.	3	Warning	Warning		Yes	275
6.5.x	1.16	276	volumeOnline	Volume is online now. The string sent with trap specifies name of volume which is online now.	2	Notification	Harmless		Yes	276
6.5.x	1.16	283	rmcCard Needs Replacement	Remote Management Controller card needs replacement. The trap includes a string specifying the reason for replacement	1	Critical	Harmless		No	N/A
7.0.3 / 7.0.4	1.17.3	284	rmcCard Missing Cables	Remote Management Controller card is missing its internal cable, or LAN cable or power supply cable. The trap includes a string specifying the missing component	1	Error	Harmless		No	N/A
7.0.3 / 7.0.4	1.17.3	294	volume Remote Unreachable	Local volume encountered an error while communicating to remote volume	1	Error	Warning		Yes	296
7.0.3 / 7.0.4	1.17.3	296	volume RemoteOK	The communication between remote volume and local volume returned to normal	1	Notification	Warning		Yes	296
7.0.3 / 7.0.4	1.17.3	297	volume Remote Restored	The data on remote volume has been fully restored to local volume	1	Informational	Warning		No	N/A
7.0.3 / 7.0.4	1.17.3	298	volume Remote RestoreBegin	The data on remote volume has started being restored to local volume by Restore-on-Demand	1	Debug	Warning		No	N/A

Introduced		OID / Trap #	Trap Name	Description	# Different Message Strings (sub-types)	Default Severity	Citigroup Tivoli Severity	Citigroup Netcool Severity	Auto Clear	Clear OID
ONTAP Version	MIB #									
6.5.x	1.16	304	volume Restricted RootConflict	Volume is restricted due to a root volume conflict. The string sent with trap specifies name of conflicting volume which is being restricted	1	Error	Critical		No	N/A
6.5.x	1.16	314	volumeOffline TooBig	Volume can not be brought online because its raw size is larger than maximum allowed size. The string sent with trap specifies name of affected volume and its raw size, and maximum allowed size.	1	Error	Critical		No	N/A
6.5.x	1.16	324	volumeOffline	Volume is being taken offline. The string sent with trap specifies name of affected volume and reason for being taken offline.	2	Error	Warning		No	N/A
6.5.x	1.16	334	volume Restricted	Volume is being restricted. The string sent with trap specifies name of affected volume and reason for being restricted.	1	Error	Warning		No	N/A
6.5.x	1.16	344	volume Degraded Dirty	Volume is degraded and has dirty parity. WAFL_check must be run on this volume before it can be brought online. The string sent with trap specifies name of affected volume.	1	Error	Critical		No	N/A
6.5.x	1.16	354	volumeError	This trap is issued when a volume can not be brought online due to an error. The string sent with trap specifies name of affected volume and error description.	5	Error	Critical		No	N/A
6.5.x	1.16	364	snapmirror SyncFailed	Synchronous SnapMirror failed and went into asynchronous mode.	1	Error	Harmless		Yes	366
6.5.x	1.16	366	snapmirror SyncOk	Synchronous SnapMirror went into synchronous mode.	1	Notification	Harmless		Yes	366
7.0.3 / 7.0.4	1.17.3	371	chassis Temperature Shutdown	The chassis temperature is extreme. The appliance has initiated a shutdown to protect itself. The operating environment should be monitored and corrected before restarting the appliance	2	Emergency	Fatal		Yes	376
7.0.3 / 7.0.4	1.17.3	372	chassis Temperature Warning	The chassis temperature is either too high or too low. The temperature should be monitored and, if possible, corrected	2	Alert	Critical		Yes	376
7.0.3 / 7.0.4	1.17.3	375	chassis Temperature Unknown	The chassis temperature is unknown, because reading can't be obtained from the chassis temperature sensor	1	Warning	Critical		Yes	376
7.0.3 / 7.0.4	1.17.3	376	chassis Temperature OK	The chassis temperature is OK	1	Notification	Harmless		Yes	376
7.0.3 / 7.0.4	1.17.3	381	chassisCPU FanStopped	One or more CPU fans have stopped. The appliance has initiated a shutdown to protect itself. A new motherboard may be required to correct the fan. Contact Network Appliance for direction	1	Emergency	Fatal		Yes	386
7.0.3 / 7.0.4	1.17.3	383	chassisCPU FanSlow	A CPU fan is spinning too slowly. A new motherboard may be required to correct the fan. Contact Network Appliance for direction	1	Critical	Critical		Yes	386
7.0.3 / 7.0.4	1.17.3	386	chassisCPU FanOK	All CPU fan(s) are properly functioning	1	Notification	Harmless		Yes	386

Introduced		OID / Trap #	Trap Name	Description	# Different Message Strings (sub-types)	Default Severity	Citigroup Tivoli Severity	Citigroup Netcool Severity	Auto Clear	Clear OID
ONTAP Version	MIB #									
7.0.3 / 7.0.4	1.17.3	391	chassisPower Supplies Failed	Multiple chassis power supplies failed	1	Emergency	Fatal		Yes	396
7.0.3 / 7.0.4	1.17.3	392	chassisPower Supply Degraded	One or more chassis power supplies are degraded. A description of the degraded state has been logged to the console and message log file	1	Alert	Critical		Yes	396
7.0.3 / 7.0.4	1.17.3	393	chassisPower Supply Failed	One chassis power supply failed	1	Critical	Critical		Yes	397
7.0.3 / 7.0.4	1.17.3	394	chassisPower Supply Removed	One or more chassis power supplies are removed	1	Error	Critical		Yes	396
7.0.3 / 7.0.4	1.17.3	395	chassisPower Supply Off	One or more chassis power supplies are off	2	Warning	Critical		Yes	396
7.0.3 / 7.0.4	1.17.3	396	chassisPower SuppliesOK	The chassis power supplies are all functioning properly	1	Notification	Harmless		Yes	396
7.0.3 / 7.0.4	1.17.3	397	chassisPower SupplyOK	This chassis power supply is functioning properly	1	Informational	Harmless		Yes	397
7.0.3 / 7.0.4	1.17.3	403	chassisPower Degraded	The power within the chassis is degraded. A description of the degraded state has been logged to the console and message log file	1	Critical	Critical		Yes	406
7.0.3 / 7.0.4	1.17.3	406	chassisPower OK	The power within the chassis is functioning properly	1	Notification	Harmless		Yes	406
7.0.3 / 7.0.4	1.17.3	412	chassisFan Degraded	A chassis fan has been degraded	1	Alert	Critical		Yes	416
7.0.3 / 7.0.4	1.17.3	413	chassisFan Removed	A chassis fan FRU has been removed	1	Critical	Critical		Yes	416
7.0.3 / 7.0.4	1.17.3	414	chassisFan Stopped	One or more chassis fans have stopped	1	Error	Critical		Yes	416
7.0.3 / 7.0.4	1.17.3	415	chassisFan Warning	One or more chassis fans is spinning slowly or too fast	1	Warning	Critical		Yes	416
7.0.3 / 7.0.4	1.17.3	416	chassisFan OK	All chassis fans are functioning properly	1	Notification	Harmless		Yes	416
7.0.3 / 7.0.4	1.17.3	424	write Verification Failed	A write has failed a verification test on a SnapValidator enabled volume	1	Error	Harmless		No	N/A
7.0.3 / 7.0.4	1.17.3	435	domain Controller Disconnect	A CIFS domain controller connection to the filer has failed	1	Warning	Warning		No	N/A

Introduced		OID / Trap #	Trap Name	Description	# Different Message Strings (sub-types)	Default Severity	Citigroup Tivoli Severity	Citigroup Netcool Severity	Auto Clear	Clear OID
ONTAP Version	MIB #									
7.0.3 / 7.0.4	1.17.3	444	plexFailed	This trap is issued when one plex of a mirrored traditional volume or aggregate is failed. The string sent with this trap specifies name of affected plex or mirrored traditional volume or aggregate	2	Error	Warning		No	N/A
7.0.3 / 7.0.4	1.17.3	454	plexOffline	This trap is issued when a plex becomes offline. The string sent with this trap specifies name of affected plex or mirrored traditional volume or aggregate	2	Error	Warning		No	N/A
7.0.3 / 7.0.4	1.17.3	482	maxDirSize Alert	A directory has reached it's maxdirsize limit. Either increase the maxdirsize or clean up the directory	1	Alert	Critical		No	N/A
7.0.3 / 7.0.4	1.17.3	485	maxDirSize Warning	A directory is getting close to it's maxdirsize limit. Either increase the maxdirsize or clean up the directory	1	Warning	Critical		No	N/A

14 Appendix C: Data ONTAP SOE Configuration Standards

Items listed with a light red background  form the Citigroup Standard Build for NetApp NAS Filer.

Items listed with a light green background  are configured according to regional requirements, where these regional settings are part of Level III PCMs.

The commands list below follow the form options <option_name> <value>, so to enable the audit log, for example, the command would be options auditlog.enable on. Most columns are self explanatory, except 'Supported from' that details the release of ONTAP at which the given option was supported or updated. The contents of some cells in this table have been wrapped for presentation purposes.

Supported from	ONTAP Options Configuration	SOE Security Setting	Description
6.4	auditlog.enable	on	Enables/disables the audit logging of commands executed at the console/telnet shell or by using ssh or GUI. The default is on . The data is logged to the file /etc/log/auditlog for a filer. Assuming they do not get full, auditlog files are saved for a total of six weeks.
6.4	autologout.console.enable	on	Enables/disables the autologout 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 autologout.console.timeout value. Any change to this option is effective after a command is entered.
6.4	autologout.console.timeout	5	The number of minutes the console is idle after which console connection is disconnected.
6.4	autologout.telnet.enable	on	Enables/disables the autologout of telnet connections. The default is on which causes telnet connections to be disconnected after the number of minutes specified by the autologout.telnet.timeout value. Any change to this option requires a logout before it takes effect.
6.4	autologout.telnet.timeout	5	The number of minutes after which telnet connections are disconnected if autologout.telnet.enable is on. The default is 60 minutes. Any change to this option requires a logout before it takes effect.
6.4	autosupport.content	complete	Dictates the type of content that the autosupport notification should contain if autosupport.enable and/or autosupport.support.enable are set to on. The default for this complete . Any change to this option will take effect the next time an autosupport message is generated.
6.4	autosupport.enable	on	Dictates whether recipients defined in autosupport.to will receive autosupport emails.
6.4	autosupport.from	<first.last@citi.com>	Filer Administrator recipient email address, in the standard Citigroup email address form. Alternatively this email address can be a distribution list. Whichever value you use, it must be a valid entry in the Citigroup Global Address List (GAL).
6.4	autosupport.mailhost	mailhub-nyc3.ny.ssmb.com	The address of the SMTP server used to send the autosupport email to the recipients defined in the autosupport.to option.

Supported from	ONTAP Options Configuration	SOE Security Setting	Description
6.4	autosupport.to	<first.last@citi.com>	Filer Administrator recipient email address, in the standard Citigroup email address form. Alternatively this email address can be a distribution list. Whichever type of address is used, the email address must be internal to the Citigroup email system. Due to the clear text nature of autosupport messages, internet email addresses are forbidden in this field
6.4	autosupport.support.enable	on	Dictates whether NetApp receives the autosupport message. The default value is on .
7.0	autosupport.support.proxy	webproxy.wlb.nsroot.net:8080	Dictates the address of the Citigroup web proxy service used to send the Autosupport posting to NetApp. For North America and EMEA this address is resolved via 3DNS. For all other regions this address is resolved via DNS round robin. The webproxy.wlb.nsroot.net service is supported by Ken Tunstall's organization that can be called out via North America DSO / TP21.
7.0	autosupport.support.transport	https	Dictates the type of transport protocol used to send the autosupport message. This option default is SMTP , however SMTP is a clear text protocol and autosupport messages contain sensitive information. Therefore secure HTTP is required.
6.5	cifs.audit.autosave.onsize.enable	on	When CIFS audit log reaches the value defined cifs.audit.logsize the audit log will be automatically saved and a new log started
6.4	cifs.audit.enable	off	When this option is on , CIFS audit events may be generated for file access or for logon and logoff. For file access events to be generated, the option ifs.audit.file_access_events.enable must also be on. For logon and logoff events to be generated, the option cifs.audit.logon_events.enable must also be on. This option default is off .
6.4	cifs.audit.file_access_events.enable	on	When both this option and the cifs.audit.enable option are on , file access events will be audited when a file is accessed by an account for an operation and the file has a System Access Control List (SACL) entry that matches the access. If no SACL entry matches the access, then no event will be generated. The default is on .
6.4	cifs.audit.logon_events.enable	on	When both this option and the cifs.audit.enable option are on , logon and logoff events will be generated. Logon and logoff events reflect CIFS session connects and disconnects, respectively. The default is on .
6.4	cifs.bypass_traverse_checking	off	When on (the default), directories in the path to a file are not required to have the 'X' (traverse) permission. This option does not apply in UNIX qtrees.
6.4	cifs.oplocks.enable	on	
6.5	cifs.weekly_w2k_password.change	on	
6.4	ftpd.enable	off	If on the appliance accepts FTP connections on port 21. If off connections on port 21 are refused. The default is off.
6.4	httpd.enable	off	Enables HTTP access to the filer. Valid values for this option are on or off . The default value for this option is off .

Supported from	ONTAP Options Configuration	SOE Security Setting	Description
6.4	nfs.webnfs.enable	off	nfs.webnfs.enable When enabled, the NFS server supports WebNFS lookups. Valid values for this option are on (enabled) or off (disabled). The default value for this option is off .
6.4	pcnfsd.enable	off	Enables/disables the PCNFS (PC)NFS authentication request server (see na_pcnfsd(8)). Valid values for this option are on or off . The default value for this option is off .
6.5	raid.reconstruct.wafliron.enable	on	Enables starting wafliron (see na_vol(1)) when reconstruction encounters a medium error. Valid values for this option are on and off . The default value for this option is on . When a medium error is encountered during reconstruction, access to the volume is temporarily restricted and reconstruction proceeds ypassing media errors. If this option is enabled, wafliron is started automatically, thus bringing the volume back online. If this option is disabled, the volume stays restricted.
6.5	raid.scrub.duration	600	
6.4	rsh.enable	off	Enables the RSH server on the filer. Valid values for this option are on or off . The default value for this option is on .
7.1	security.passwd.firstlogin.enable	on	Dictates whether an administrator is required to change their password the first time they log onto the Filer after their password has been changed. Valid values for this option are on or off . The default value for this option is on .
7.1	security.passwd.lockout.numtries	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 0 . Configuring this option to '6' meets the CITMP policy requirements.
6.5	security.passwd.rules.enable	on	This option controls whether a check for minimumlength and password composition is performed when new passwords are specified. A value of on means that the check will be made, and the password rejected if it doesn't pass the check. A value of off means that the check won't be made. The default value for this option is on . This option does not apply to the users "root" or "Administrator" (the NT Administrator account).
7.0	security.passwd.rules.everyone	on	This option controls whether a check for password composition is performed for all users, including "root" and "Administrator". A value of off means that the checks do not apply to "root" or "Administrator" (but still may apply to all other users). The default value for this option is off . security.passwd.rules.enable (above) must have the value on or this option is ignored.
7.1	security.passwd.rules.history	6	Defines the number of previous passwords that ONTAP will remember for each administrator account. The default value for this option is 0 .
7.0	security.passwd.rules.maximum	14	This option controls the maximum number of characters a password can have. Though there is no default value for this option, only the first 16 characters are saved. Users with passwords greater than 14 characters will not be able to log in via the Windows interfaces. security.passwd.rules.enable must have the value on or this option is ignored.

Supported from	ONTAP Options Configuration	SOE Security Setting	Description
7.0	security.passwd.rules.minimum	6	This option controls the minimum number of characters a password must have. The default value for this option is 8 . Configuring this option to '6' meets the CITMP policy requirements. security.passwd.rules.enable must have the value on or this option is ignored.
7.0	security.passwd.rules.minimum.alphabetic	2	This option controls the minimum number of alphabetic characters a password must have. (IE: A password cannot be just digits and symbols.) These are capitol and lowercase letters from a to z. The default value for this option is 2 . security.passwd.rules.enable must have the value on or this option is ignored.
7.0	security.passwd.rules.minimum.digit	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 . security.passwd.rules.enable must have the value on or this option is ignored.
7.0	security.passwd.rules.minimum.symbol	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 . security.passwd.rules.enable must have the value on or this option is ignored.
6.4	snmp.enable	'on'	Enables the SNMP server on the filer. Valid values for this option are on or off . The default value for this option is on .
7.0	ssh.access	See notes	Restricts the SSH access to the Filer. The default value for this option is '*', meaning that access from any host is allowed. The secure configuration requires the following settings to limit SSH access to each Filer. The SOE servers and DFM servers consist of the IP addresses for the designated management stations for each support group. 'host=<soe_server1_IP>,<soe_server2_IP>,<dfm_server1_IP>,<dfm_server2_IP>,<soe_inventory_server>'
6.4	ssh.enable	'on'	Enables or disables the SSH protocol on the filer. Valid values for this option are on or off. The default value for this option is off.
7.0	ssh1.enable	'off'	Enables or disables the SSH version 1 protocol on the filer. Valid values for this option are on or off. The default value for this option is on.
7.0	ssh2.enable	'on'	Enables or disables the SSH version 2 protocol on the filer. Valid values for this option are on or off. The default value for this option is on.
6.4	ssl.enable	on	
6.4	telnet.enable	'off'	Enables the Telnet server on the filer. Valid values for this option are on or off. The default value for this option is on.
6.4	ftpd.enable	'off'	Enables the ftp (Trivial File Transfer Protocol) server on the filer. Valid values for this option are on or off . The default value for this option is off . When enabled, the filer's ftp server allows <i>get</i> requests, but does not allow <i>put</i> requests.

Supported from	ONTAP Options Configuration	SOE Security Setting	Description
7.0	trusted.hosts	see notes	Restricts the FilerView access to the Filer. The default value for this option is '*', meaning that access from any host is allowed. The secure configuration requires the following settings to limit HTTPS access to each Filer. The SOE servers and DFM servers consist of the IP addresses for the designated management stations for each support group. <code>'<soe_server1_IP>, <soe_server2_IP>, <dfm_server1_IP>, <dfm_server2_IP>'</code>
6.4	wafl.default_security_style	'ntfs'	Specifies the default security style assigned to a new volume. All qtrees created on the volume get this as their security style. Legal values for this option are 'unix', 'ntfs', or 'mixed'. The default value for this option is 'unix', unless the filer is an NTFSonly filer, in which case the default is 'ntfs'.
6.4	wafl.nt_admin_priv_map_to_root	'on'	When turned on, an NT administrator is mapped to the Unix root account. The default value for this option is 'on'.

15 Appendix D: ONTAP, Shelf and Disk Firmware installation procedures

15.1 Data ONTAP Install Procedure

1. Login to NetApp Support Website:
 - a. <http://now.netapp.com/NOW/cgi-bin/software>
 - b. Goto end of Software list "To access specific version"
 - c. Select Data ONTAP and type version number (7.1.2.1P1)
 - d. Select Go
 - e. Click the download button at the bottom of the screen
 - f. Review the material on the must read notification screen, click the accept radio button and then click continue
 - g. Accept the license agreement
 - h. Scroll down to the software images until you find the 7121P1_setup.exe files for deployment from the Windows platform
 - 7121P1_setup_i.exe required for Intel Platform (800, 900 and R200 series)
 - 7121P1_setup_m.exe required for MIPS Platform (200 series)
 - 7121P1_setup_e.exe required for PC Elf Platform (3000 series)
 - i. Select the appropriate version for your hardware platform according to the previous step
 - j. Right-click on the appropriate 712P1_setup.exe file link and select "Save as" to download the software

Figure 15.1 - <http://now.netapp.com> Software Download Page






Figure 15.2 - Data ONTAP OS file selection



2. From your desktop, map \\filename\c\$ (vol/vol0) with an account that has “write” or administrator privileges.

Figure 15.3 - Filer OS volume - mapping to local workstation




3. From your download directory:
 - a. Run “7121P1_setup_e.exe”, “7121P1_setup_i.exe” or “7121P1_setup_m.exe” file (depending on your Filer hardware platform)
 - b. Accept the instructions by clicking “OK”

Figure 15.4 - Data ONTAP OS zip extract – screen 1



- c. Using the drive letter defined in step 0,
 - enter the target folder information
 - ensure “overwrite files without prompting” tick box is checked
 - ensure “when done unzipping open: inst_perm.bat” is checked
 - click the “unzip” button

Figure 15.5 - Data ONTAP OS zip extract – screen 2



4. SSH or Remote Console to Filer CLI.
 - a. If filer is a cluster node; type “cf disable” command to disable cluster failover to partner node.
 - b. Type “download” command. Once completed, your filer will have the downloaded DOT in place.
 - c. Type “reboot” command to complete upgrade.
 - d. Repeat steps 2 – 4 to upgrade partner node.
 - e. Upon completion of both nodes, type “cf enable” from any node to re-enable cluster failover.

Figure 15.6 - Data ONTAP OS Filer Download

```

$ ssh -l root nas3020gdse1
nas3020gdse1> download
download: Reminder: upgrade both the nodes in the Cluster
download: You can cancel this operation by hitting Ctrl-C in the next 6 seconds.
download: Depending on system load, it may take many minutes
download: to complete this operation. Until it finishes, you will
download: not be able to use the console.
Tue Feb 14 13:09:22 EST [nas3020gdse1: download.request:notice]: Operator requested download initiated
download: Downloading boot device
Version 1 ELF86 kernel detected.

```

15.2 Disk Firmware Upgrade Procedure

1. Download latest disk firmware pack from NOW site
2. Expand zip file to \\filename\c\$\etc\disk_fw
3. Run “disk_fw_update” at the filer console prompt.

```

$ ssh -l root nas3020gdse1
nas3020gdse1>
nas3020gdse1>
nas3020gdse1>
nas3020gdse1>
nas3020gdse1>
nas3020gdse1>
nas3020gdse1>
nas3020gdse1>
nas3020gdse1>
nas3020gdse1>
nas3020gdse1>
nas3020gdse1>
nas3020gdse1>
nas3020gdse1>
nas3020gdse1>
nas3020gdse1> disk_fw_update
Downloading disk firmware can be disruptive to the system.
It involves spinning down and then spinning up all
affected disks, and suspending disk I/O to disks being
downloaded to. This may take up to 5 minutes to complete.
This delay may cause clients to hang, and may cause
CIFS clients to drop sessions.

NOTE: Updating firmware on some Fibre Channel disks
may create an open loop condition which can only be cleared
by power cycling the affected shelf or shelves.

In a cluster, you should not disable the cluster or perform
any takeover/giveback actions until the firmware update
is complete. Disk firmware update should only be attempted
on one cluster partner at a time.

Downloading disk firmware can only be performed when the cluster is
enabled, and both nodes are able to communicate with each other.
Firmware download cannot be performed when in takeover mode.
Refer to the man page for disk_fw_update for complete details.

Are you sure you want to continue with firmware update? y
<You can ignore disk error messages for disks being updated until this
update completes>
nas3020gdse1> Disk Firmware: No disks in primary FC loop eligible for download
Tue Feb 14 13:14:16 EST [nas3020gdse1: dfu.firmwareUpToDate:info]: Firmware is up-to-date on all disk drives
nas3020gdse1>

```

Figure 15.7 - Disk firmware update

16 Appendix F: Data Schema for NetApp Inventory Agent




Figure 16.1 - Citigroup NetApp Inventory Agent Schema

17 Appendix G: New disk test results

This appendix includes testing and performance details on the new 300GB Fiber Channel and 500GB (S)ATA drives.

17.1 300GB Fiber Channel

Test environment:	FAS3050 w/ 16 x 300GB Fiber Channel drives RAID group of 16 disks from two shelves on a single loop
Load generator:	12 sets of nfs_hammer and 1 filer hammer driving both CPU and Disk Utilization in the range of 70-80%
Zeroing time:	1 hour, 15 minutes to 1 hour, 24 minutes depending upon drive model (Hitachi vs. Seagate)
Rebuild time:	19 hours, 31 minutes and 22.75 seconds (for full reconstruction) 4 hours, 16 minutes and 09.78 seconds (using disk to disk copy - rapid disk recovery)
Sanitization time:	5 hours 40 minutes

17.2 500GB (S)ATA

Test environment:	FAS3050 w/ 14 x 500GB (S)ATA drives RAID group of 14 disks from two shelves on a single loop
Load generator:	12 sets of nfs_hammer and 1 filer hammer driving both CPU and Disk Utilization in the range of 70-80%
Zeroing time:	2 hours, 55 minutes to 3 hours, 3 minutes depending upon drive model
Rebuild time:	41 hours, 02 minutes and 44.11 seconds (for full reconstruction) 12 hours, 5 minutes and 40.29 seconds (using disk to disk copy - rapid disk recovery)
Sanitization time:	8 hours 49 minutes

18 Appendix H: FilerView process examples

This appendix includes detailed screen shot examples of certain administrator activities within FilerView.

18.1 Aggregate creation

Please note this example is using 206GB SATA drives on a FAS3000.



From the 'FilerView' navigation menu, select 'Aggregates', and then 'Add'



Click the 'Next' button

Aggregate Wizard - Aggregate Parameters

Aggregate Name:
Enter a name for the new aggregate.

Double Parity:
Select to enable double parity on this aggregate. Enabling this option requires an extra disk per RAID group. ☒ Double Parity

< Back Cancel Next >


Enter a name for the aggregate following the naming standard discussed in section 9.4.2.2, select RAID-DP (except for root aggregates) and click the 'Next' button

Aggregate Wizard - RAID Parameters


RAID Group Size:
Enter the number of disks per RAID group on this aggregate. Disks will be organized into RAID groups of this size.

< Back Cancel Next >


Choose the RAID group size according to the standards discussed in section 9.4.2.3, and then click the 'Next' button



Select 'Automatic' and click the 'Next' button




Select the appropriate size from the available disks, and click the 'Next' button




The screenshot shows a web browser window titled 'nas3020gdse02: Aggregate Wizard - Internet Explorer 6'. The main heading is 'Aggregate Wizard - Number of Disks'. Below the heading, the text reads: 'Number of Disks: Select the number of disks of size 206G to add to the aggregate. There are 25 spares of that size available.' A dropdown menu is set to '24'. At the bottom, there are three buttons: '< Back', 'Cancel', and 'Next >'.

Select the number of disks using the RAID group size standards discussed in section 9.4.2.3, and then click the 'Next' button




The screenshot shows a web browser window titled 'nas3020gdse02: Aggregate Wizard - Internet Explorer 6'. The main heading is 'Aggregate Wizard - Commit'. Below the heading, the text reads: 'Below is a summary of your changes.' A text box contains the following summary: 'Create New Aggregate', 'Volume Name: aggr1', 'RAID Group Size: 16', 'Number of Disks: 24', 'Disk Size: 206G', and 'Double Parity: yes'. At the bottom, there are three buttons: '< Back', 'Cancel', and 'Commit'.

Review the settings at the confirmation screen, and, if acceptable, click 'Commit'




Upon receipt of the success message, click the 'Close' button


18.2 Volume creation



From the 'FilerView' navigation menu, select 'Volumes', and then 'Add'



Click the 'Next' button to start the wizard




Select the 'Flexible' volume type and click the 'Next' button

The screenshot shows a web browser window titled "nas3020gdse01: Volume Wizard - Internet Explorer 6". The main content area has a blue header bar with the text "Volume Wizard - Volume Parameters". Below the header, there are two sections. The first section is labeled "Volume Name:" and contains a text input field with the value "vol1" and a help icon. The second section is labeled "Language:" and contains a dropdown menu with "English (US)" selected and a help icon. At the bottom of the dialog, there are three buttons: "< Back", "Cancel", and "Next >".

Enter a name for the volume following the naming standard discussed in section 9.4.3, define the language type and click the 'Next' button

The screenshot shows a web browser window titled "nas3020gdse01: Volume Wizard - Internet Explorer 6". The main content area has a blue header bar with the text "Volume Wizard - Flexible Volume Parameters". Below the header, there are three sections. The first section is labeled "Containing Aggregate" and contains a dropdown menu with "aggr1 (2.3 TB, raid_dp)" selected and a help icon. The second section is labeled "Total Volume Size:" and contains a text input field with the value "100" and a unit dropdown menu with "GB" selected and a help icon. The third section is labeled "Space Guarantee" and contains a dropdown menu with "volume" selected and a help icon. At the bottom of the dialog, there are three buttons: "< Back", "Cancel", and "Next >".

Select the aggregate on which you want to create the volume, the size of the volume (note the volume size will include the volume snapshot reserve), and the space guarantee type for the volume (volume recommended). Once completed click the 'Next' button



Review the settings at the confirmation screen, and, if acceptable, click 'Commit'



Upon receipt of the success message, click the 'Close' button

19 Appendix I: Tape SAN connectivity

This appendix was created as a result of Tape SAN connectivity testing conducted for the North America storage group. It contains recommendations and considerations that apply globally when integrating the NetApp NAS Filer with a Tape SAN.

It is important to note that NDMP does not have the capacity to manage a path failover. In this solution the NetBackup server will initially fail all streams that were utilizing the broken path and freeze the tapes. Then through the job retry parameter, recovery is achieved by opening the stream on the remaining active path(s).

The table below describes the steps required in order to prepare and configure the infrastructure for Tape SAN backup of NetApp NAS Filers. The implementation steps shown in this table are listed in a specific order. This order must be followed to avoid issues when initially connecting the NAS Filers to the SAN Fabric.

Step	Component	Instruction
1.	NAS Filer	The steps in this section are to be completed by the NAS administrator This step must be completed <u>prior</u> to physically attaching the Filer to the Tape SAN Fabric: Disable OSMS support on the Filer - Halt the Filer to the OK, CFE> or LOADER> prompt (prompt type is dependant on the model of the Filer you are using) - At the prompt type 'setenv fc-no-mgmt? true' (without the ' ' marks) - To verify the variable type 'printenv' (without the ' ' marks) - Once confirmed, boot the Filer as normal - The QLogic HBAs, with ONTAP, will detect the type of protocol to used based on the connection topology. Therefore if the HBA is connected to a switch, PTP mode will be selected, otherwise FCAL when connected to a disk shelf
2.	Tape SAN	The steps in this section are to be completed by the SAN administrator - Attach the Tape drives to Tape SAN switch (under the assumption that the Tape Drives/NDMP do not support load-balancing and failover, all the drives have a single path to the Tape SAN while testing in the lab) - The Switch will discover the Tape WWNN (Node Name) and WWPN (Port Name). They may be same or different depending upon which Tape drive port is attached to the switch (assuming there are 2 ports per Tape drive) - Create a new Zone for the filers following standard naming convention. The Zoning will be based on WWPN of the Tape drives. (The filer will see the WWNN and report the WWNN when any queries are made from the Filer console, and not the WWPN) * See note 3 in "Observations" section - Enable Port Binding - Add Zone to the Zone set - Activate the Zone set - Unlike with Solaris hosts, no configuration information needs to be passed back to the NAS administrators once zoning is activated

Step	Component	Instruction
3.	NAS Filer	The steps in this section are to be completed by the NAS administrator a) Using the 'storage show tape' command confirm that the tapes are now visible to the Filer b) Run 'storage unalias -t' to remove the initial auto-assigned aliases for the tape drives c) Alias each tape drive manually based on tape drive configuration supplied by backup integration / operations e.g. 'storage alias st0 WWN[5:001:04f000:5952d0]' d) The desired outcome is to have the tape drives numbered on the Filer in the same order as the NetBackup infrastructure to improve the storage and backup administrators ability to diagnose issues with the backup infrastructure

Table 19.1 – Tape SAN Infrastructure configuration steps

19.1 Observations from testing

This section annotates some behavioral observations recorded as part of the testing. Reviewing this section is worthwhile as it indicates expected behavior given certain predefined failure conditions.

19.1.1 HBA failure / Switch failure

Through simulations by blocking the switch port, the scenarios of HBA / Switch failure were tested. In this instance at least one of the available tape drives experienced a media write error (error 84). The tape will then be frozen by the NetBackup application.

In order to clear this condition, intervention is required on the Filer host in terms of off-lining the drive. For example to offline the first tape drive issue the following command from the Filer: 'mt -t rst0a offline'. Once off-lined, the backup administrator will need to move the tape out of the drive and request a new tape load.

19.1.2 On-lining HBA port during backup

If an off-lined host port on the Filer is brought on-line via administrator intervention during an active backup session, media write errors will occur and the Netbackup server will be required to restart the job. No manual intervention is required on either the Filer by the NAS administrator or by the backup server administrator to recover from this condition.

19.1.3 World Wide Name behavioral considerations

Assuming the tape drives are connected to the SAN using the first fiber port on each drive then the following WWN addressing resolution convention is applied when implementing this configuration:

WWN type	World Wide Node Name	World Wide Port Name
Implemented	NAS Filer	Switch
Experienced	Executing command 'Storage Show Tape'	When zoning
Example	WWN[5:001:04f000:5952 cd]	WWN[5:001:04f000:5952 ce]
Difference	WWPN minus one	WWNN plus one

Table 19.2 – Tape SAN World Wide Node Name and Port Name considerations

Note if the fiber is connected to the second port of the tape drive then the difference between the WWNN and the WWPN will be two.

19.1.4 ONTAP auto-aliasing of tape drives

At stage 3a) in the infrastructure configuration list (Table 1), the NAS Filer will auto-alias the tape drives using the “st” naming convention (refer to NAS Filer vendor documentation for full details). For operational implementation purposes the automatic nature of this naming convention is not convenient as it lacks a consistency with how the NetBackup infrastructure refers to each tape device.

The Filer will auto-assign aliases based on the World Wide Node Name of each device into alpha-numerical order, such that a tape drive with the WWNN of 50:01:04:F0:00:59:52:BE will be assigned a lower device alias than a drive with 50:01:04:F0:00:59:52:CD.

The purpose of destroying and creating new aliases in steps 3b) and 3c) is to give the tape devices the same number in both the NetApp interface and the NetBackup interface, such that rst0a on the filer will be the first drive on the NetBackup console, and rst1a will be the second (and so on).

19.1.5 Lab performance benchmarks

The following performance statistics were recorded as part of the lab test:

No. of Tape Drives	No. of Volumes	Vol Size (GB)	Data Size (GB)	Speed (GB/hour)	NetBackup Config Optimization ([Y]es/[N]o)	Transaction Type ([B]ackup/[R]estore)
1	1	50	30	112.74	N	B
2	2	50	30	192.87	N	B
3	6	50	30	255.56	N	B
3	3	50	30	289.09	Y	B
4	8	50	30	402.53	Y	B
1	1	50	30	48.02	Y	R
2	2	50	30	96.09	Y	R

Table 19.3 – Tape SAN Lab test performance benchmarks

20 Appendix J: SOE Vfiler Creation Example

20.1 Basic Creation

Using the naming conventions in the NetApp VFILER and NFS Best Practices guide, VFILER creation should be completed using the following steps:

Action	Example	Results
Create VFILER system qtree	<code>qtree create /vol/vol0/vfxxxp #</code>	<code>filer1> qtree create /vol/vol0/vfguen1</code>
Verify qtree creation	<code>qtree</code>	<pre> Volume Tree Style Oplocks Status ----- ---- dse_v1 ntfs enabled normal we_v1 ntfs enabled normal vol0 ntfs enabled normal vol0 vfguen1 mixed enabled normal gue_v1 unix enabled normal </pre>
For NFS; change qtree security style to mixed (this allows both Unix and Wintel management access to system vol/qtree). This mixed mode qtree security setting is the one exception to the mixed mode not recommended rule).	<code>qtree security /vol/vol0/vfxxxp # mixed</code>	<code>filer1> qtree security /vol/vol0/vfguen1 mixed</code> Mon Jun 2 10:36:37 EDT [filer1: rc:notice]: security style for /vol/vol0/vfguen1 changed from unix to mixed
Check vfiler status, ensure vfiler license is active	<code>vfiler status</code>	<code>filer1> vfiler status</code> vfiler0 running
Create vfiler using qtree created in step 1 as the system path and an assigned IP address	<code>vfiler create vfilername-xxxp# -i ip address /vol/vol0/qtree</code>	<code>filer1> vfiler create vfiler1-guen1 -i 150.110.93.104 /vol/vol0/vfguen1</code> Mon Jun 2 10:55:54 EDT [vfiler1-guen1@filer1: rc:notice]: Vfiler vfiler1-guen1 initialized. The IP addresses of vfiler vfiler1-guen1 are still unconfigured. To use these addresses, you must configure them using ifconfig, and make those settings persistent by adding the ifconfig command lines to /etc/rc filer1>
Check vfiler status, ensure newly created vfiler is active	<code>vfiler status</code>	<code>filer1> vfiler status</code> vfiler0 running vfiler1-guen1 running filer1>
Edit "/etc/rc" of hosting filer, add ifconfig interface up setting for persistence across reboots.	Map to <code>filernane\CS\$</code> directory, edit /etc/rc file, add the following line: ifconfig interface alias xxx.xxx.xxx.yyy netmask 255.255.xxx.yyy up	<code>ifconfig vif0 alias 150.110.93.104 netmask 255.255.255.0 up</code>
Add ip alias to selected vfiler interface	<code>ifconfig interface alias ip address netmask</code>	<code>filer1> ifconfig vif0 alias 150.110.93.104 netmask 255.255.255.0</code> filer1>

Action	Example	Results
Check interface status. In addition you should ping the IP address from a host other than the filer to test connectivity and name resolution.	<pre>ifconfig interface</pre>	<pre>filer1> ifconfig vif0 vif0: flags=1948043<UP,BROADCAST,RUNNING,MULTICAST,TC PCKSUM> mtu 1500 inet 150.110.93.101 netmask 0xffffffff00 broadcast 150.110.93.255 inet 150.110.93.104 netmask 0xffff0000 broadcast 150.110.255.255 partner vif1 (not in use) ether 02:a0:98:00:fb:90 (Enabled virtual interface) filer1></pre>
Check vfiler status using -r switch to see assigned resources	<pre>vfiler status -r</pre>	<pre>filer1> vfiler status -r vfiler0 running ipspace: default-ipspace IP address: 150.110.93.101 Path: / UUID: 00000000-0000-0000-0000-000000000000 vfiler1-guen1 running ipspace: default-ipspace IP address: 150.110.93.104 Path: /vol/vol0/vfguen1 UUID: f72b65ca-91d3-11d7-9936-00a09800fb90 filer1></pre>

20.1.1 Adding Resources

The VFILER is created and running, however, the only resources assigned thus far are the system qtree and its IP address. Employing current NetApp Standards, the intent is to keep data volumes separate from system files (vfiler system qtree). As such, the NetApp administrator must create data volumes and add these as additional resources to the vfiler. When a volume is added to the vfiler as a resource, the hosting filer no longer has management control of the volume.

To add volume resources to the VFILER, complete the following steps:

Action	Example	Results
Add resource to vfiler	<pre>vfiler add vfiler path</pre>	<pre>filer1> vfiler add vfiler1-guen1 /vol/gue_v1 WARNING: reassigning storage to another vfiler does not change the security information on that storage. If the security domains are not identical, unwanted access may be permitted, and wanted access may be denied. Mon Jun 2 12:01:16 EDT [filer1: rc:notice]: Path /vol/gue_v1 moved to vfiler "vfiler1- guen1" filer1></pre>
Check vfiler status using -r switch to see assigned resources	<pre>vfiler status -r</pre>	<pre>filer1> vfiler status -r vfiler0 running ipspace: default-ipspace IP address: 150.110.93.101 Path: / UUID: 00000000-0000-0000-0000-000000000000 vfiler1-guen1 running ipspace: default-ipspace IP address: 150.110.93.104 Path: /vol/vol0/vfguen1 Path: /vol/gue_v1 UUID: f72b65ca-91d3-11d7-9936-00a09800fb90 filer1></pre>

20.1.2 VFiler Setup

The vfiler now has all required resources for operation. VFILER setup and configuration follows the same process to that of a standard filer.

NOTE: NIS configuration information should not be completed prior to CIFS setup. If NIS configuration is invoked prior to CIFS setup, the assumption is made that the filer will only be used for NFS, and the /etc/passwd file required for CIFS translation of Windows domain ID's to local filer ID's is not automatically created. This setup sequence for multi-protocol filers, has been confirmed by NetApp support, and should be adhered to in the following process.

To setup the vfiler for DNS and NIS services, complete the following steps:

Action	Example	Results
Execute vfiler setup command. Be certain to have all the required DNS and NIS domain and server information.	<pre>vfiler run <vfilername> setup</pre>	<pre>filer1> vfiler run vfiler1-guen1 setup ===== vfiler1-guen1 The administration host is given root access to the filer's /etc files for system administration. To allow /etc root access to all NFS clients enter RETURN below. Please enter the name or IP address of the administration host: Do you want to run DNS resolver? [n]: y Please enter DNS domain name []: ny.ssmb.com You may enter up to 3 nameservers Please enter the IP address for first nameserver []: 146.128.168.73 Do you want another nameserver? [n]: 146.128.201.218 Please answer "y" or "n". Do you want another nameserver? [n]: Do you want to run NIS client? [n]: y Please enter NIS domain name []: ue.ssmb.com Please enter list of preferred NIS servers [*]: 162.124.248.195 Mon Jun 2 12:18:12 EDT [vfiler1-guen1@filer1: nis_poll:warning]: Bound to preferred NIS server 162.124.248.195 Changing password for user root on vfiler vfiler1-guen1 New password: Retype new password: User <root> added. filer1> Mon Jun 2 12:18:36 EDT [vfiler1- guen1@filer1: rc:info]: User <root> added</pre>
Verify DNS server connectivity	<pre>vfiler run <vfilername> dns info</pre>	<pre>filer1> vfiler run vfiler1-guen1 dns info ===== vfiler1-guen1 DNS is enabled IP Address State Last Polled Avg RTT Calls Errs ----- 146.128.168.73 UP 0 0 0 Default domain: ny.ssmb.com Search domains: ny.ssmb.com ssmb.com filer1></pre>

Action	Example	Results
Verify NIS serve connectivity	<code>vfiler run <vfilername> nis info</code>	<pre>filer1> vfiler run vfiler1-guen1 nis info ===== vfiler1-guen1 NIS domain is ue.ssmb.com NIS group cache has been disabled IP Address Type State Bound Last Polled Client calls ----- 162.124.248.195 PREF ALIVE YES Mon Jun 2 13:17:11 EDT 2003 0 filer1></pre>

20.1.3 CIFS Setup

As per current SOE NetApp Standards, all management functions of filers, and now including vfilers, will be carried-out from Wintel workstations. For this reason, all vfilers must run the CIFS protocol, in order to map the management directory (*vfilername*lc\$) to Wintel management workstations.

To configure CIFS on VFILER, please complete the following steps:

Action	Example	Results
Execute vfiler cifs setup command. Be sure to have your AD OU administrator create the domain computer account in the "servers" container prior to running CIFS setup. Be certain to have all the required WINS and Active Directory domain and server information available. NOTE: If using Webadmin 3.0 to add computer names, the account used to add the computer to the domain must be the same account used to complete the CIFS setup.	<code>vfiler run <vfilername> cifs setup</code>	<pre>filer1> vfiler run vfiler1-gue cifs setup ===== vfiler1-gue Enable CIFS access to the filer by a Windows(tm) PC Your filer is currently only visible to PCs on the same net. Do you wish to make the system visible via WINS? [no]: y Please enter next WINS IP address (or return to end list) []: 146.128.201.56 Please enter next WINS IP address (or return to end list) []: 150.110.234.253 Please enter next WINS IP address (or return to end list) []: This filer is currently licensed for both CIFS and NFS. The filer will be configured as a Multiprotocol filer. It can be changed to a NTFS-only filer by reissuing the CIFS setup command. The default name of this filer will be 'VFILER1-GUE'. Do you want to modify this name? [no]: n CIFS supports three types of user authentication: 1. Windows Domain authentication. 2. Windows Workgroup authentication using the filer's user accounts. 3. /etc/passwd and/or NIS based authentication. What type of authentication will this filer use? [1]: 1</pre>

Action	Example	Results
		The filer will use Windows Domain authentication. Enter the Windows Domain for the filer []: nam.nsroot.net It is highly recommended that you create the local administrator account (VFILER1-GUE\administrator) for this filer. This account enables access to the filer from Windows when domain controllers are not accessible. Create the VFILER1-GUE\administrator account? [yes]: y Set the password for VFILER1-GUE\administrator: Retype the password: Tue Jun 3 09:30:25 EDT [vnfiler1-gue@filer1: rc:info]: User <administrator> added NAM.NSROOT.NET is a Windows 2000(tm) domain. In order to create this filer's domain account, you must supply the name and password of an administrator account with sufficient privilege to add the filer to the NAM.NSROOT.NET domain. Please enter the Windows 2000 user [Administrator@NAM.NSROOT.NET]: rs01146@nam.nsroot.net Password for rs01146@nam.nsroot.net: CIFS - Logged in as rs01146@nam.nsroot.net. CIFS - Updating filer account 'cn=vfiler1-gue,ou=servers,ou=storageeng,ou=testing,dc=nam,dc=nsroot,dc=net' CIFS - Connecting to domain controller. Tue Jun 3 09:30:48 EDT [vfiler1-gue@filer1: rc:info]: Connection with DC \\NAMDCNY002 established Tue Jun 3 09:30:48 EDT [vfiler1-gue@filer1: nbns_timer:info]: All CIFS name registrations complete for local server Welcome to the NAM (NAM.NSROOT.NET) Windows 2000(tm) domain. CIFS local server is running. filer1>

Action	Example	Results
Verify AD domain connectivity	<pre>vfiler run <vfilername> cifs testdc</pre>	<pre>filer1> vfiler run vfiler1-gue cifs testdc ===== vfiler1-gue Using Established configuration Current Mode of NBT is H Mode Netbios scope "" Registered names... VFILER1-GUE < 0> WINS VFILER1--GUE < 3> WINS VFILER1-GUE <20> WINS NAM < 0> WINS Testing all Primary Domain Controllers found 19 unique addresses found PDC NAMDCNJ003 at 150.110.187.80 found PDC NAMDCNY002 at 162.124.192.156 found PDC NAMDCHT001 at 168.109.40.63 found PDC NAMDCNY999 at 162.124.200.76 found PDC NAMDCNY003 at 162.124.190.181 found PDC NAMDCWE001 at 169.188.125.48 Testing all Domain Controllers found 19 unique addresses found DC NAMDCNJ002 at 150.110.177.155 found DC NAMDCNJ003 at 150.110.187.80 found DC NAMDCNJ006 at 150.110.188.172 found DC NAMDCNJ004 at 150.110.188.170 found DC NAMDCNJ001 at 150.110.117.155 found DC NAMDCNJ005 at 150.110.188.171</pre>
Verify CIFS shares available on vfiler	<pre>vfiler run <vfilername> cifs shares</pre>	<pre>filer1> vfiler run vfiler1-gue cifs shares ===== vfiler1-gue Name Mount Point Description ---- - ETC\$ /vol/vol0/vfguen1/etc Remote Administration BUILTIN\Administrators / Full Control HOME /vol/vol0/vfguen1/home Default Share everyone / Full Control C\$ /vol/vol0/vfguen1 Remote Administration BUILTIN\Administrators / Full Control</pre>
Remove the ETC\$ and HOME CIFS shares as per SOE NetApp Standards	<pre>vfiler run <vfilername> cifs shares - delete <sharename></pre>	<pre>filer1> vfiler run vfiler1-gue cifs shares - delete ETC\$ ===== vfiler1-gue filer1> filer1>vfiler run vfiler1-gue cifs shares - delete HOME ===== vfiler1-gue</pre>

Action	Example	Results
Verify CIFS shares available on vfiler	<code>vfiler run <vfilername> cifs shares</code>	<pre>filer1> vfiler run vfiler1-gue cifs shares ===== vfiler1-gue Name Mount Point Description ----- C\$ /vol/vol0/vfguen1 Remote Administration BUILTIN\Administrators / Full Control</pre>

20.1.4 Create /etc/exports File

To facilitate consistent exports of NFS volumes across reboots it is necessary to create or edit the “/etc/exports” file and add the path of the volumes you wish to export.

When rebooting, Data ONTAP executes the `exportsfs -a` commanding the `etc/rc` file to automatically export all resources in the `/etc/exports` file.

In the `/etc/exports` file, you should also specify the access restrictions with which clients can mount each resource.

NOTE: By default the filer SETUP routine creates the “/etc/exports” file and adds all volumes available to the filer. The system administrator should edit this file to add or remove NFS exports restrictions, and to add or remove additional exports paths.

To edit the `/etc/export` file, the following steps must be completed:

Action	Example	Results
Export the volumes created in the previous step.	<code>vfiler run <vfilername> exportfs -p</code>	<pre>vfiler run vnetapps01-gue exportfs -p anon=0 /vol/vol0/vfguen1 ===== vfiler1-gue vfiler run vnetapps01-gue exportfs -p anon=0 /vol/gue_v1</pre>
Verify exported file systems	<code>vfiler run <vfilername> exportfs</code>	<pre>filer1> vfiler run vfiler1-gue exportfs ===== vfiler1-gue /vol/vol0/vfguen1 -anon=0 /vol/gue_v1 -anon=0</pre>

NOTE: At this point VFILER and NFS are setup on the filer and production ready. NFS workstations are now ready to mount exported file systems.

20.1.5 Client NFS Mount Connectivity

To verify client connectivity and specified restrictions, the following steps must be executed:

Action	Example	Results
From Unix workstation, create a mount point (directory) and execute the mount command to the appropriate vfiler resource.	<code>mount vfilername:exported resource mountpoint</code>	<pre>bash-2.03# mount vfiler1- gue:/vol/gue_v1 /vfguen1_v1</pre>

Action	Example	Results
Verify mount command	<code>mount</code>	<pre> bash-2.03# mount / on /dev/dsk/c0t0d0s0 read/write/setuid/intr/largefiles/onerror=panic /dev=800000 on Tue Jun 3 08:58:25 2003 /proc on /proc read/write/setuid/dev=36c0000 on Tue Jun 3 08:58:24 2003 /dev/fd on fd read/write/setuid/dev=3780000 on Tue Jun 3 08:58:26 2003 /etc/mnttab on mnttab read/write/setuid/dev=3880000 on Tue Jun 3 08:58:27 2003 /var/run on swap read/write/setuid/dev=1 on Tue Jun 3 08:58:27 2003 /tmp on swap read/write/setuid/dev=2 on Tue Jun 3 08:58:29 2003 /home/jm50049 on gues2:/export/guegs2dg1vol01/home/jm50049 remote/read/write/setuid/soft/bg/intr/dev=39000 02 on Tue Jun 3 11:17:00 2003 /opt/PDolvm on gues2:/export/guegs2dg1vol02/shared/opt/PDolvw m remote/read/write/setuid/soft/bg/intr/dev=39000 03 on Tue Jun 3 11:17:21 2003 /opt/ADBEacrbt on gues2:/export/guegs2dg1vol02/shared/opt/ADBEac rbt remote/read/write/setuid/soft/bg/intr/dev=39000 09 on Tue Jun 3 16:29:55 2003 /vfquen1_v1 on vfiler1-gue:/vol/gue_v1 remote/read/write/setuid/dev=390000a on Thu Jun 5 13:34:49 2003 </pre>
Change directory to vfiler mount. Execute a dir listing	<pre> cd /mountpoint ls -ltr </pre>	<pre> bash-2.03# cd /vfquen1_v1 bash-2.03# ls -ltr total 2728872 -rw----- 1 root other 4294967296 May 8 15:47 4gfile -rw----- 1 root other 1073741824 May 13 12:10 1gfile -rw-r--r-- 1 root other 125 Jun 6 09:05 burtfile -rw-r--r-- 1 root other 125 Jun 6 09:11 burtfile2 -rw-r--r-- 1 root other 125 Jun 6 09:11 burtfile3 -rw-r--r-- 1 root other 125 Jun 6 09:11 burtfile4 -rw-r--r-- 1 root other 125 Jun 6 09:11 burtfile5 -rw-r--r-- 1 root other 125 Jun 6 09:11 burtfile6 -rw-r--r-- 1 root other 125 Jun 6 09:11 burtfile7 -rw-r--r-- 1 root other 125 Jun 6 09:11 burtfile8 -rw-r--r-- 1 root other 125 Jun 6 09:11 burtfile9 bash-2.03# </pre>

20.1.6 Snapmirror Setup (Asynchronous Replication)

Snapmirror is a NetApp Data Ontap software option that allows asynchronous replication of a volume or qtree to a remote filer.

Once a snapmirror relationship is established, for the vfiler system qtree, a baseline transfer initializes the mirror to create a replica of the vfiler on the destination filer. The entire vfiler including its network and security settings, data, permissions and options are mirrored as one. Snapmirror functions are carried-out from the host filer.

NOTE: Prerequisites for snapmirror operations should be implemented prior to executing the steps outlined below, these include:


Create volumes on DR/COB filers using the same paths and names as the production resources.

Ensure DNS resolution of local and remote filers.

License snapmirror on local and remote filers.

Enable snapmirror on local and remote filers.

To establish vfiler snapmirror replication for COB/DR, the following steps must be executed:

Action	Example	Results
Map a drive to <i>hostfiler</i>\\C\$ share. Using your favorite editor: edit or create the <i>/etc/snapmirror.allow</i> file on the local filer Or Using the snapmirror options: "snapmirror.access" replace "legacy" with the target filer hostname. This replaces the need for the "snapmirror.allow" file	Use favorite text editor	 Or <pre>filer1> options snapmirror.access host=filer2</pre>
From the target/remote filer: Map a drive to <i>targetfiler</i>\\C\$ share. Using your favorite editor: edit or create the <i>/etc/snapmirror.conf</i> file on the local filer. The snapmiror.conf file is used to keep the snamirrored storage units up to date based on schedules. Please reference the NetApp ONTAP Data Protection Guide for schedule options.	Use favorite text editor	

Action	Example	Results
Create and initialize the snapmirrors. NOTE: Make certain to restrict volume access on the target volume before initializing snapmirror	<pre>snapmirror initialize -S sourcefiler:path targetfiler:path</pre>	<pre>filer2> snapmirror initialize -S filer1:/vol/vol0/vfguen1 filer2:/vol/vol0/vfguen1 Transfer started. Monitor progress with 'snapmirror status' or the snapmirror log. filer2> filer2> snapmirror initialize -S filer1:gue_v1 filer2:gue_v1 Transfer started. Monitor progress with 'snapmirror status' or the snapmirror log. filer2></pre>
Check snapmirror status to ensure file transfers are active.	<pre>snapmirror status</pre>	<pre>filer2> snapmirror status Snapmirror is on. Source Destination State Lag Status filer1:/vol/vol0/vfguen1 filer2:/vol/vol0/vfguen1 Uninitialized - Transferring (148 KB done) filer1:gue_v1 filer2:gue_v1 Uninitialized - Transferring (259 MB done)</pre>

20.2 SOE Vfiler DR/COB Creation Example

20.2.1 Vfiler Inventory

When it comes to moving the vfiler structure from one filer to another, we need to have all the technical information about the vfiler's configuration. The following inventory should be collected:

Paths; list of all the data paths (resources) assigned to the vfiler.

IP Address; vfiler's IP address. The IP address of the production vfiler will not be necessary. As part of the COB/DR setup Storage Integration will assign an IP address designated during the 3DNS service request.

Quotas; if necessary.

Use the "vfiler status -a *vfilername*" command to display a full list of vfiler information.

Use the "ifconfig -a" command to display all network interface information.

Outputs of these commands are as follows:

```
filer1> vfiler status -a vfiler1-gue
vfiler1-gue      running
  ipspace: default-ipospace
  IP address: 150.110.93.104
  Path: /vol/vol0/vfguen1
  Path: /vol/gue_v1
  UUID: 65465b8a-91d4-11d7-9936-00a09800fb90
  Protocols allowed: 3
  Allowed: proto=rsh
  Allowed: proto=nfs
  Allowed: proto=cifs
  Protocols disallowed: 0
filer1>
```

```
filer1> ifconfig -a
e0: flags=8042<BROADCAST,RUNNING,MULTICAST> mtu 1500
    ether 00:a0:98:00:fb:90 (auto-unknown-cfg_down)
e4: flags=1108042<BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 1500
    ether 02:a0:98:00:fb:90 (auto-unknown-cfg_down) flowcontrol full
    trunked vif0
e5: flags=1948043<UP,BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 1500
    ether 02:a0:98:00:fb:90 (auto-1000sx-fd-up) flowcontrol full
    trunked vif0
lo: flags=1948049<UP,LOOPBACK,RUNNING,MULTICAST,TCPCKSUM> mtu 1500
    inet 127.0.0.1 netmask 0xff000000 broadcast 127.0.0.1
    ether 00:00:00:00:76:68 (VIA Provider)
vif0: flags=1948043<UP,BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 1500
    inet 150.110.93.101 netmask 0xffffffff broadcast 150.110.93.255
    inet 150.110.93.104 netmask 0xffff0000 broadcast 150.110.255.255
    partner vif1 (not in use)
    ether 02:a0:98:00:fb:90 (Enabled virtual interface)
```

20.2.2 DR/COB Filer Configuration

The most important rule for the COB/DR filer is that it must mimic the same environment as the Production filer. From the COB/DR network, we must ensure the following:

- Access to the same NIS servers
- Access to the same DNS servers
- Access to the same WINS servers
- Access to the same AD domain
- Same volume structure as Production filer.
- Invoke “vfiler setup -r” command to re-create the production vfiler from the snapmirror volumes.
- May need to run setup
- May need to run CIFS setup.

20.2.3 Special Setup – Volume FSID update

In order to allow transparent redirection from the NFS clients to the COB/DR vfiler, we need to be sure the FSID of the volume on the production vfiler is the same as the COB/DR vfiler. To change the volume fsid on the COB/DR filer complete the following:

From the Production filer:

```
Enter priv set advanced
Enter vol read_fsid <volname>
Result: Volume data has a FSID of 0x76005881
Record FSID number
Enter priv set admin
```

From the COB/DR filer:

```
Enter priv set advanced
Enter vol restrict <volname> (should be same as production volume).
```

NOTE: You cannot restrict root volume while the filer is up, all other volumes can be restricted when the filer is up but users cannot access the volume.

```
Enter vol read_fsid <volname>
Result: Volume data has a FSID of 0x50008ddd
Enter vol rewrite_fsid <volname> 0x76005881
Result: Volume data has had its FSID changed from 0x50008ddd to 0x50008ddd
Enter vol online <volname>
Enter priv set admin
```

20.2.4 Recreate Production Vfiler on DR/COB

To re-create the production vfiler in COB/DR, the following steps must be executed:

Action	Example	Results
From COB/DR filer; Check snapmirror status Verify snapmirror(s) are active.	<code>snapmirror status</code>	filer2> <code>snapmirror status</code> Snapmirror is on. Source Destination State Lag Status filer1:/vol/vol0/vfguen1 filer2:/vol/vol0/vfguen1 Snapmirrored 00:00:14 Idle filer1:gue_v1 filer2:gue_v1 Snapmirrored 00:00:12 Idle
From COB/DR filer; Quiesce the vfiler snapmirror volumes	<code>Snapmirror quiesce /vol/vol0/vfiler qtree</code> <code>Snapmirror quiesce /vol/datavol</code>	filer2> <code>snapmirror quiesce /vol/vol0/vfguen1</code> snapmirror quiesce: in progress This can be a long-running operation. Use Control - C (^C) to interrupt. snapmirror quiesce: /vol/vol0/vfguen1 : Successfully quiesced filer2> <code>snapmirror quiesce gue_v1</code> snapmirror quiesce: in progress This can be a long-running operation. Use Control - C (^C) to interrupt. snapmirror quiesce: gue_v1 : Successfully quiesced
From COB/DR filer; break the vfiler snapmirror volumes in order to make the volumes writeable.	<code>Snapmirror break /vol/vol0/vfiler qtree</code> <code>Snapmirror break /vol/datavol</code>	filer2> <code>snapmirror break gue_v1</code> Thu Jun 5 16:12:35 EDT [rc:error]: quota: /vol/gue_v1 or a quota tree on it referred to a non-existent vfiler. quota: quotas have been turned off. snapmirror break: Destination gue_v1 is now writable. filer2> <code>snapmirror break /vol/vol0/vfguen1</code> snapmirror break: Destination /vol/vol0/vfguen1 is now writable.
From Production filer; Halt Production Filer to simulate disaster. Or If halting the filer is not an option... STOP the VFILER in question.	<code>halt</code> or <code>vfiler stop vfiler1-gue</code>	filer1> <code>halt</code> CIFS local server on vfiler vfiler1-gue is shutting down... Thu Jun 5 16:08:34 EDT [vfiler1-gue@filer1: rc:info]: Connection with \\NAMDCNY002 terminated Thu Jun 5 16:08:34 EDT [vfiler1-gue@filer1: rc:info]: CIFS local server has shut down... CIFS local server on vfiler vfiler0 is shutting down... Thu Jun 5 16:08:34 EDT [filer1: rc:info]: Connection with \\NAMDCSOL004 terminated Thu Jun 5 16:08:39 EDT [filer1: rc:error]: No WINS servers are responding; still trying Thu Jun 5 16:08:39 EDT [filer1: rc:info]: CIFS local server has shut down... Thu Jun 5 16:08:39 EDT [filer1: kern.shutdown:notice]: System shut down because : "halt". Program terminated ok

Action	Example	Results
From COB/DR filer; Re-create Production vfiler	<pre>vfiler create vfilername -r /vfilerqtree</pre>	<pre>nasrut1001dr> vfiler create vnetapps01-gue -r /vol/vol0/vfguen1 Thu Jun 5 16:17:43 EDT [vnetapps01-gue@nbns_timer:info]: All CIFS name registrations complete for local server Thu Jun 5 16:17:51 EDT [vnetapps01-gue@rc:info]: Lookup of _ldap._tcp.nyc._sites.NAM.NSROOT.NET failed with DNS server 146.128.201.218: Can't assign requested address. Thu Jun 5 16:17:58 EDT [vnetapps01-gue@rc:warning]: Cannot connect to server. Socket connect error: Can't assign requested address Thu Jun 5 16:17:58 EDT [vnetapps01-gue@rc:info]: CIFS - can't make TCP connection with \\NAMDCNY002 Thu Jun 5 16:18:13 EDT [vnetapps01-gue@rc:error]: Can't connect socket to WINS server: Can't assign requested address Thu Jun 5 16:18:13 EDT [vnetapps01-gue@rc:error]: Can't connect socket to WINS server: Can't assign requested address Thu Jun 5 16:18:16 EDT [vnetapps01-gue@rc:warning]: CIFS - Could not find a domain controller. CIFS local server is running. Thu Jun 5 16:18:16 EDT [vnetapps01-gue@nis_poll:error]: Cannot send NIS binding request to 162.124.248.195 : Can't assign requested address Thu Jun 5 16:18:16 EDT [vnetapps01-gue@nis_poll:info]: NIS server(s) not available Thu Jun 5 16:18:16 EDT [vnetapps01-gue@nbns_timer:info]: All CIFS name registrations complete for local server NFS server is running. export: Can't find address for host prometheus.ny.ssmb.com vnetapps01-gue running ipspace: default-ipospace IP address: 150.110.93.104 [unconfigured] Path: /vol/vol0/vfguen1 Path: /vol/gue_v1 UUID: 65465b8a-91d4-11d7-9936-00a09800fb90 Protocols allowed: 3 Allowed: proto=rsh Allowed: proto=nfs Allowed: proto=cifs Protocols disallowed: 0 The IP addresses of vfiler vnetapps01-gue are still unconfigured. To use these addresses, you must configure them using ifconfig, and make those settings persistent by adding the ifconfig command lines to /etc/rc nasrut1001dr> Thu Jun 5 16:18:31 EDT [vnetapps01- gue@rc:notice]: Vfiler vnetapps01-gue restored. nasrut1001dr></pre>
From COB/DR filer; Add IP address resource from COB/DR network to vfiler	<pre>Vfiler add vfilername -i ipaddress</pre>	<pre>nasrut1001dr> vfiler add vnetapps01-gue -i 192.168.30.144 Thu Jun 5 16:22:49 EDT [vnetapps01- gue@SMBBrowser:error]: Can't connect socket to WINS server: Can't assign requested address Thu Jun 5 16:22:59 EDT [rc:notice]: IP address 192.168.30.144 added to vfiler "vnetapps01-gue" nasrut1001dr></pre>

Action	Example	Results
From COB/DR filer; Add the IP address as an alias to the COB.DR filer network interface	<pre>ifconfig interface alias ipaddress netmask xxx.xxx.xxx.yyy</pre>	<pre>nasrut1001dr>ifconfig gev1 alias 192.168.30.144 netmask 255.255.255.0 ifconfig (warning): route already exists (probably due to another interface on the same subnet.) nasrut1001dr></pre>
From COB/DR filer; Check the vfiler status using the <code>-r</code> option. Ensure: vfiler is running IP address assigned is up		<pre>nasrut1001dr> vfiler status -r vfiler0 running ipspace: default-ip-space IP address: 192.168.30.142 Path: / UUID: 00000000-0000-0000-0000-000000000000 vnetapps01-gue running ipspace: default-ip-space IP address: 150.110.93.104 [unconfigured] IP address: 192.168.30.144 Path: /vol/vol0/vfguen1 Path: /vol/gue_v1 UUID: 65465b8a-91d4-11d7-9936-00a09800fb90 nasrut1001dr></pre>
From COB/DR filer; Check the vfiler's NIS info	<pre>vfiler run vfilername nis info</pre>	<pre>vfiler run vnetapps01-gue nis info ===== vnetapps01-gue NIS domain is ue.ssmb.com NIS group cache has been disabled IP Address Type State Bound Last Polled Client calls ----- 162.124.248.195 PREF ALIVE YES Fri Jun 6 10:53:03 EDT 2003 0 nasrut1001dr></pre>
From COB/DR filer; Check the vfiler's DNS info	<pre>vfiler run vfilername dns info</pre>	<pre>nasrut1001dr> vfiler run vnetapps01-gue dns info ===== vnetapps01-gue DNS is enabled IP Address State Last Polled Avg RTT Calls Errs ----- 146.128.201.218 UP Fri Jun 6 09:29:22 EDT 2003 4292851098 109 24 146.128.168.73 UP 0 24 24 Default domain: nam.nsroot.net Search domains: nam.nsroot.net nsroot.net nasrut1001dr></pre>
From User Workstation; Unmount the production vfiler exported <i>path</i>	<pre>umount mounthname</pre>	<pre>bash-2.03# umount /vfguen1_v1</pre>

Action	Example	Results
From User Workstation; Re-mount the COB/DR vfiler exported <i>path</i>	mount vfilename:/path	bash-2.03# mount vnetapps01-gue:/vol/gue_v1 /vfguen1_v1
From User Workstation; Verify mount; list mount directory	ls /mountpoint	bash-2.03# ls -ltr /vfguen1_v1 total 2728872 -rw----- 1 root other 4294967296 May 8 15:47 4gfile -rw----- 1 root other 1073741824 May 13 12:10 1gfile -rw-r--r-- 1 root other 125 Jun 6 09:05 burtfile -rw-r--r-- 1 root other 125 Jun 6 09:11 burtfile2 -rw-r--r-- 1 root other 125 Jun 6 09:11 burtfile3 -rw-r--r-- 1 root other 125 Jun 6 09:11 burtfile4 -rw-r--r-- 1 root other 125 Jun 6 09:11 burtfile5 -rw-r--r-- 1 root other 125 Jun 6 09:11 burtfile6 -rw-r--r-- 1 root other 125 Jun 6 09:11 burtfile7 -rw-r--r-- 1 root other 125 Jun 6 09:11 burtfile8 -rw-r--r-- 1 root other 125 Jun 6 09:11 burtfile9 bash-2.03#

Notes:

Because of the inherent characteristics of the NFS protocol, it is necessary to un-mount and then re-mount the exported resource on the client workstation when failover occurs.

On the other hand, CIFS failover is seamless and requires no user interaction. Mapped drives are offline only for the duration of the failover process, and return online immediately thereafter.

NFS automounter feature is not a feasible alternative for VFILER and NFS because automounter expects both file servers online in parallel. In a production, high availability, implementation automounter is most suitable. However, for active/passive failover implementations the use of 3DNS name resolution fits best.

20.2.5 Vfiler Production Failback

We've now proven the feasibility of vfiler failure and recovery. In order to failback to the production vfiler, the following steps must be executed:

Action	Example	Results
From COB/DR filer; sys Stop vfiler	Vfiler stop vfilename	nasrut1001dr> vfiler stop vnetapps01-gue vnetapps01-gue stopped nasrut1001dr>
Reboot production filer	Boot	ok boot Intel Open Firmware by FirmWorks Copyright 1995-2001 FirmWorks, Network Appliance. All Rights Reserved. Firmware release 4.0_i1 Press Del to abort boot, Esc to skip POST Memory size is 1024 MB Testing SIO Testing LCD Probing devices Testing 1024MB

Action	Example	Results
From COB/DR filer; Resync snapmirror of vfiler resources.	Snapmirror resync path	<pre> nasrut1001dr> snapmirror resync gue_v1 The resync base snapshot will be: nasrut1001dr(0050385317)_gue_v1.112 These older snapshots have already been deleted from the source and will be deleted from the destination: nasrut1001dr(0050385317)_gue_v1.111 Are you sure you want to resync the volume? y Fri Jun 6 11:42:33 EDT [rc:notice]: snapmirror: resynching volume gue_v1 to netapps01:gue_v1 using nasrut1001dr(0050385317)_gue_v1.112 as the base snapshot. Volume gue_v1 will be briefly unavailable before coming back online. Fri Jun 6 11:42:33 EDT [config_async_thread:notice]: CIFS - disabled for volume gue_v1 Fri Jun 6 11:42:33 EDT [config_async_thread:warning]: vfiler vnetapps01-gue: storage unit /vol/gue_v1 now offline Volume gue_v1 is now restricted. Fri Jun 6 11:42:34 EDT [config_async_thread:notice]: Reverting volume gue_v1 to a previous snapshot. Fri Jun 6 11:42:34 EDT [config_async_thread:notice]: vfiler vnetapps01-gue: storage unit /vol/gue_v1 now online nasrut1001dr> snapmirror resync /vol/vol0/vfguen1 The resync base snapshot will be: nasrut1001dr(0050385317)_vol0_vfguen1-dst.36 Resync may alter the data in this qtree. Are you sure you want to resync the qtree? y Fri Jun 6 11:46:09 EDT [rc:info]: snapmirror: resync of /vol/vol0/vfguen1 to netapps01:/vol/vol0/vfguen1 : success Transfer started. Monitor progress with 'snapmirror status' or the snapmirror log. nasrut1001dr> </pre>
From COB/DR filer; Verify snapmirror status.	Snapmirror status	<pre> nasrut1001dr> snapmirror status Snapmirror is on. Source Destination State Lag Status netapps01:/vol/vol0/vfguen1 nasrut1001dr:/vol/vol0/vfguen1 Snapmirrored 00:00:33 Idle netapps01:gue_v1 nasrut1001dr:gue_v1 Snapmirrored 00:01:34 Idle nasrut1001dr> </pre>
From production filer; Verify vfiler status	Vfiler status -r	<pre> netapps01> vfiler status -r vfiler0 running ipspace: default-ipspace IP address: 150.110.93.101 IP address: 150.110.93.102 IP address: 150.110.93.103 Path: / UUID: 00000000-0000-0000-0000-000000000000 vnetapps01-gue running ipspace: default-ipspace IP address: 150.110.93.104 Path: /vol/vol0/vfguen1 Path: /vol/gue_v1 UUID: 65465b8a-91d4-11d7-9936-00a09800fb90 netapps01> </pre>

Action	Example	Results
From User Workstation; Unmount the production vfiler exported <i>path</i>	umount mountname	bash-2.03# umount /vfguen1_v1
From User Workstation; Re-mount the COB/DR vfiler exported <i>path</i>	mount vfilename:/path	bash-2.03# mount vnetapps01-gue:/vol/gue_v1 /vfguen1_v1
From User Workstation; Verify mount; list mount directory	ls /mountpoint	bash-2.03# ls -ltr /vfguen1_v1 total 2728872 -rw----- 1 root other 4294967296 May 8 15:47 4gfile -rw----- 1 root other 1073741824 May 13 12:10 1gfile -rw-r--r-- 1 root other 125 Jun 6 09:05 burtfile -rw-r--r-- 1 root other 125 Jun 6 09:11 burtfile2 -rw-r--r-- 1 root other 125 Jun 6 09:11 burtfile3 -rw-r--r-- 1 root other 125 Jun 6 09:11 burtfile4 -rw-r--r-- 1 root other 125 Jun 6 09:11 burtfile5 -rw-r--r-- 1 root other 125 Jun 6 09:11 burtfile6 -rw-r--r-- 1 root other 125 Jun 6 09:11 burtfile7 -rw-r--r-- 1 root other 125 Jun 6 09:11 burtfile8 -rw-r--r-- 1 root other 125 Jun 6 09:11 burtfile9 bash-2.03#
From management workstation; Verify CIFS access to /etc directory	Windows explorer, map to network drive vfilename\\c\$	

21 Appendix J: Role Based Access Control capabilities

21.1 Overview

NetApp have implemented the following Role Based Access Control model with ONTAP 7.1.




Figure 21.1 - RBAC entity concepts

Using the above framework, the following logical rules can be extrapolated:

- Each user can be a member of multiple groups
- Each role can be a member of multiple groups
- Conversely, each group can be assigned multiple roles
- Each capability can be a member of multiple roles
- Conversely, each role can be assigned multiple capabilities

Capabilities are group into four classes:

Login – right to authenticate to the device using a particular method: Telnet, RSH, SSH, Console, HTTP

CLI – right to execute each command available in the CLI. Each and every CLI command is mapped to a capability type. The CLI command can be wild-carded to support all concatenated sub-commands, or grant only the initial command. For example cli-cifs* (wildcard) or just cli-cifs. CLI capabilities do not support the granting of switches or arguments/subcommands on a per argument basis, therefore you cannot distinguish between the right to perform sysconfig –t and sysconfig –d, nor does it distinguish the right to perform cifs shares –add and cifs shares –delete.

API – right to call an ONTAP API (ONTAPI). Unlike CLI rights, APIs can be defined at each level of the API call and therefore support concatenation of all API sub-types

Security – right to perform a privileged system function

This framework offers a reasonably flexible and highly adaptive structure that can be used to meet the majority of role based access controls for administrative duties on the NetApp NAS Filer Platform.

From a CATE perspective the final implementation design follows the KISS principle as the flexibility of the RBAC framework. Therefore roles have been designed to match business functions such that, for example, all CIFS administrator commands are assigned to a CIFS administrator role. Groups can then be used to collate supersets of roles, or where “super users” (complex roles) are required; a user can be assigned to more than one group to achieve the required roll up of distinctive roles.

21.2 High-level Administrative User Concepts

The diagram below provides a high level overview of the RBAC model developed by CATE for the NetApp Filer platform.




Figure 21.2 - RBAC high level overview

21.2.1 Common Capabilities

All administrators have a right to execute a core set of basic, non-intrusive, informational commands:

- Echo
- Help / ?
- Logout
- Man
- Version

21.2.2 System Administrator

A super set of groups (represented by letters) and roles (represented by roman numerals)

a) File Storage Administrator

i) CIFS Administrator

Create/manage qtrees, shares; run volume utilization reports

ii) NFS Administrator

Create/manage qtrees, mount points; run volume utilization reports

b) Network Administrator

Responsible for configuring and diagnosing all network related functions on the storage device, with the exception of IPsec

c) License Administrator

Responsible for the management of software licenses on storage device

d) Storage Administrator

Responsible for provisioning storage pools, managing virtual storage devices, replication processes, performing system software upgrades, running scripted processes and sanitizing disks

e) Backup Administrator

A read only view of tape and ndmp related system information

f) Security Administrator

Responsible for configuring security (secureadmin) and any key management, including ipsec

g) Performance Administrator

A read only performance data-gathering role

21.2.3 Account Administrator

Account administrator is the equivalent of the Infrastructure Security Administrator (ISA). The term 'account' is used in preference to the 'security' as it is more descriptive and applicable to the role. An account administrator has the following capabilities:

- Create accounts
- Modify account description
- Change account password
- Change account permissions / entitlements
- Provide reports on account status and entitlements
- Delete accounts
- Manage groups and roles (RBAC framework)
- Must not be able to perform storage management functions