



CATE Storage and Backup Engineering
Citi Architecture & Technology Engineering

NetApp Data ONTAP 8

CSBE-10059-SG

Solutions Guide

Version 4.3

1 Version Control

Version	Date	Author	Reviewed	Comments
0.1	18-APR-2011	Richard Sopp	Stephen Hopkins	Initial draft
1.0	06-MAY-2011	Richard Sopp	CTI NAS Ops	Initial production release
2.0	27-JAN-2012	Richard Sopp	Tim Armstrong Stephen Hopkins	ONTAP 8.0.2 update
4.0	31-MAR-2014	Richard Sopp	Tim Armstrong David Lee Jason Palmer	ONTAP 8.1.4 update
4.0.1	01-MAY-2014	Richard Sopp	Jason Palmer	Minor updates
4.0.2	12-MAY-2014	Richard Sopp	n/a	RBAC and bug updates
4.0.3	15-MAY-2014	Richard Sopp	n/a	References update
4.0.4	20-MAY-2014	Richard Sopp	n/a	Updated for new SMB2 bugs
4.1	25-AUG-2014	Richard Sopp	Stephen Hopkins	ONTAP 8.1.4P4 update
4.1.1	22-SEP-2014	Richard Sopp	Jacek Nogala	Production Release
4.2	01-OCT-2014	Richard Sopp	Jackson Tan	ONTAP 8.1.4P5 update
4.2.1	20-OCT-2014	Richard Sopp	Jacek Nogala	Production Release
4.2.2	12-NOV-2014	Tan Jackson	Jacek Nogala	Uplift from v4.2.1 – added ONTAP 8.1.4P6 details
4.2.3	19-NOV-2014	Tan Jackson	n/a	RFP
4.2.4	09-DEC-2014	Richard Sopp	n/a	EMC SRM integration
4.2.5	19-DEC-2014	Richard Sopp	CTI NAS Ops	Production Release
4.2.6	12-MAR-2015	Richard Sopp	n/a	Revised guidelines regarding the use of SMB 2.0
4.2.7	18-MAR-2015	Richard Sopp	n/a	Revised CPU threshold when SMB 2.0 is acceptable
4.2.8	15-MAY-2015	Tan Jackson	n/a	Updated screen shots for Service Processor to version 1.3.3P2
4.2.9	18-MAY-2015	Richard Sopp	n/a	Windows 2012 DC update
4.3	03-JUN-2015	Richard Sopp	n/a	Windows 2012 DC support RFP

2 Legal Notice

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

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

3 Contact

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

4 Latest Versions

It is recommended that the latest version of the document be used from the CATE Storage Engineering website <https://catecollaboration.citigroup.net/domains/platstor/gdse/Pages/default.aspx>. Newer versions may include certification on additional platforms, or comments on issues discovered after deployment.

5 Regional Reviewers

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

Region	Reviewed by	Version Reviewed
Asia Pacific	No response received	n/a
EMEA	Jason Palmer [NAS L3 team]	4.0
Japan	No response received	n/a
Latin America / Mexico	No response received	n/a
North America	No response received	n/a

Table 1 - Regional stakeholder review

6 Table of Contents

1	Version Control	2
2	Legal Notice	3
3	Contact	3
4	Latest Versions	3
5	Regional Reviewers	3
6	Table of Contents	4
6.1	Index of Tables	6
6.2	Index of Figures	7
7	Introduction	8
8	Assumptions, Conventions, Definitions, Abbreviations and References	9
8.1	Assumptions	9
8.2	Issuing commands	9
8.3	Documentation conventions	9
8.4	How to use this document	9
8.5	Definitions and abbreviations	10
8.6	References	10
9	General Design	11
10	Getting Started	13
10.1	Infrastructure requirements	13
10.2	Hardware platform requirements	13
10.3	Software installation	13
10.4	Scalability	15
10.5	New installations	15
10.6	Upgrading from a previous release	15
10.7	Unsupported configurations	15
11	Best Practices	16
11.1	ONTAP SOE configuration: options and registry settings	16
11.2	Network	16
11.2.1	Single-mode (Active/Passive) Interface Groups	17
11.2.2	Multi-mode (Etherchannel) Interface Groups	18
11.2.3	10GB Ethernet	19
11.2.4	Router Pair Connectivity and Cluster Configuration	22
11.2.5	Service Processor (SP) / Remote LAN Management (RLM)	23
11.2.6	The Management Port (e0M)	27
11.2.7	The Alternate Control Path port (e0P)	30
11.2.8	Host and rc file configuration	31
11.2.9	Filer NetBIOS Naming Standards	32
11.2.10	Domain Name and Host Resolution	33
11.2.11	Time Synchronization and Time Zones	35
11.2.12	Windows Internet Name Service (WINS)	36
11.2.13	Appropriate Byte Counting	38
11.3	Storage	39
11.3.1	Data ONTAP Clustering (HA-Pairs)	39
11.3.2	Disk Shelves	39
11.3.3	Disks	45
11.3.4	Aggregates	49
11.3.5	Volumes	55
11.4	File Service Protocols – CIFS & NFS	58
11.4.1	CIFS	58
11.4.2	NFS	72
11.4.3	Multi-protocol Volumes	88
11.5	Data Protection (Point in time copies, replication and backup)	94
11.5.1	Snapshots (point in time disk based copies)	94
11.5.2	Snapmirror (asynchronous replication)	96
11.5.3	Backup	102

11.6	Advanced features – vFilers, Snapvault, Snaplock and Lockvault	106
11.6.1	Vfiler (multistore)	106
11.6.2	Snapvault.....	113
11.6.3	Snaplock.....	118
11.6.4	Lockvault.....	118
11.7	Storage Optimization Technologies – Dedupe, compression, flash cache, file folding, file blocking.....	146
11.7.1	Deduplication	146
11.7.2	Compression	154
11.7.3	Snapshot ‘file-folding’	155
11.7.4	File Blocking	156
11.7.5	Flash Cache	158
11.7.6	Flash Pools.....	158
11.8	Alerting and Monitoring	159
11.8.1	AutoSupport.....	159
11.8.2	SNMP	159
11.9	Reporting.....	163
11.9.1	EMC Vopr SRM	163
12	Security Best Practices	164
12.1	Overview	164
12.2	Security Standards	164
12.3	Information Classification Certification	164
12.4	Anti-Virus.....	165
12.4.1	Summary of features in VirusScan Enterprise for Storage (VSE-S) 1.0:	166
12.4.2	Multiple to Multiple AV Configurations.....	166
12.4.3	VirusScan Enterprise for Storage 1.0 ePO NetApp Policies configuration	168
12.4.4	ONTAP “vscan” Settings	168
12.4.5	NetApp Data ONTAP Anti-Virus Service Help	169
12.4.6	Resetting the flags in the filer’s cache.....	169
12.4.7	VirusScan for NetApp - Virus Detection	170
12.5	Disk Sanitization.....	171
12.5.1	Disk Sanitization Limitations.....	171
12.5.2	Licensing Disk Sanitization.....	171
12.5.3	Sanitizing Disks	171
12.5.4	Reading Disk Sanitization Log Files.....	173
12.6	CITMP Information Security Policy Convergence	174
12.6.1	Information Ownership	174
12.6.2	Identification and Authentication	175
12.6.3	Authorization and Access Control	184
12.6.4	Restricting management physical host access	188
12.6.5	Traceable Access Entitlements	188
12.6.6	Confidentiality and Integrity	189
12.6.7	Incident Detection and Response	189
12.6.8	Administration.....	191
12.7	TCP/IP Port Requirements.....	200
13	Document version control	201
14	Appendices	206
14.1	Appendix A: ePO & VirusScan Enterprise for Storage Configuration.....	207
14.2	Appendix B: NetApp SNMP MIB	208
14.3	Appendix C: Data ONTAP SOE Configuration Standards.....	219
14.4	Appendix D: ONTAP, Shelf and Disk Firmware installation procedures.....	239
14.4.1	Data ONTAP Install Procedure.....	239
14.4.2	Disk Firmware Upgrade Procedure	239
14.4.3	Shelf Firmware Upgrade Procedure	240
14.5	Appendix E: Tape SAN connectivity	245
14.5.1	Observations from testing.....	246
14.6	Appendix F: SOE Vfiler Creation Example	248

14.6.1	Basic Creation	248
14.6.2	SOE Vfiler DR/COB Creation Example	257
14.7	Appendix G: Role Based Access Control Capabilities	266
14.7.1	Overview	266
14.7.2	High-level Administrative User Concepts	267
14.7.3	Low-level Administrative User Concepts	269
14.7.4	RBAC Deployment Script	282
14.7.5	RBAC Upgrade Script	286
14.7.6	RBAC Downgrade / Revert Script	291
14.7.7	Application Infrastructure Integration RBAC Scripts	296
14.8	Appendix H Snapvault – further reading	297
14.8.1	Detailed Snapvault Snapshot Schedule Configuration	297
14.8.2	Procedure to recover from an “Out of Snapshots” condition	299
14.8.3	Overview for System Administrators	301
14.8.4	Information for Systems Administrators	302
14.9	Appendix I NFSv4 ACLs – further reading	312
14.10	Appendix J OnCommand Systems Manager	315
14.10.1	System Manager - Overview	315
14.10.2	System Manager - Installation	315
14.10.3	System Manager - Configuration	316

6.1 Index of Tables

Table 1 - Regional stakeholder review	3
Table 2 - Supported port roles	16
Table 3 - Hardware requirements to support 10GB ethernet	19
Table 4 - The /etc/rc file key concepts	31
Table 5 - Recommended name resolution order	33
Table 6 - Default Name Resolution Order	33
Table 7 - Disk shelf firmware supported versions	40
Table 8 – FCAL HBA slot and port assignments for multipath-HA	43
Table 9 - Supported disk drive types	45
Table 10 - Maximum number of disk drives supported in a HA or standalone configuration	46
Table 11 - Drive re-build, sanitization and zeroing times	47
Table 12 - System spare disks: How to calculate number of drives	47
Table 13 - Recommended root volume size	50
Table 14 - Recommended RAID group configuration	51
Table 15 - RAID group sizes for large capacity SATA drives	51
Table 16 - Procedure to remove aggregate snapshot reserve	52
Table 17 – Maximum aggregate size per hardware platform with ONTAP 8.1	53
Table 18 - Windows Domain Controller O/S Interoperability Matrix	59
Table 19 - Windows Authentication and Directory Service Protocol Interoperability Matrix	59
Table 20 - Windows Server and Client platform O/S Interoperability Matrix	60
Table 21 - ONTAP licenses required for VMware over NFS	84
Table 22 - Name translation requirements for multi-protocol volumes	88
Table 23 - Changing file permissions for multi-protocol volumes	89
Table 24 - Suggested Snapshot settings	94
Table 25 - Snapmirror standard schedule	97
Table 26 - Maximum number of concurrent replication operations	99
Table 27 - Citi recommended vFiler limits per Storage Controller	107
Table 28 - SNMP configuration settings	159
Table 29 - Filer to AV scanner ratio	167
Table 30 - Required IP ports	200
Table 31 - SNMP trap list table definitions	208
Table 32 - SNMP MIB matrix	209
Table 33 - ONTAP Mandatory (Security) Configuration Matrix	219

Table 34 - ONTAP Recommended (Technical Reference) Configuration Matrix	226
Table 35 - Registry custom settings	236
Table 36 - Remote Management Port software configuration settings	237
Table 37 - Additional configuration settings	238
Table 38 - Tape SAN Infrastructure configuration steps	246
Table 39 - Tape SAN World Wide Node Name and Port Name considerations	246
Table 40 - Tape SAN Lab test performance benchmarks	247

6.2 Index of Figures

Figure 1 - NetApp NAS Filer Architecture Design	11
Figure 2 - Physical network connectivity example for 10GB Ethernet	21
Figure 3 - General network connectivity example	22
Figure 4 - SP (remote management device) and eOM network connections	23
Figure 5 - Multipath HA configuration	41
Figure 6 - CIFS ABE disabled	70
Figure 7 - CIFS ABE enabled	70
Figure 8 - Backup Simple "Local" NDMP Configuration	102
Figure 9 - Backup Three-party NDMP Configuration	103
Figure 10 - NetApp MultiStore (Vfiler)	106
Figure 11 - Lockvault Implementation	120
Figure 12 - File blocked client message	158
Figure 13 - VSE-S User File Access Data Flow	166
Figure 14 - NetApp_Users and NetApps_A administrators Groups C\$ Security Permissions	186
Figure 15 - NetApp Data ONTAP "/etc/log/auditlog"	188
Figure 16 - NetApp Data ONTAP "adtlog evt" file Event View	191
Figure 17 - Tectia Client GUI Agent Forwarding configuration	195
Figure 18 - Tectia Client GUI Authentication Mechanism configuration	196
Figure 19 - Disk firmware update	240
Figure 20 - RBAC entity concepts	266
Figure 21 - RBAC high level overview	267
Figure 22 - System Administrator Role	269
Figure 23 - File System Administrator Role	270
Figure 24 - CIFS File System Administrator Role	271
Figure 25 - NFS File System Administrator Role	272
Figure 26 - Storage Administrator Role	273
Figure 27 - Security Administrator Role	274
Figure 28 - Performance Administrator Role	275
Figure 29 - Network Administrator Role	276
Figure 30 - License Administrator Role	277
Figure 31 - Backup Administrator Role	278
Figure 32 - Account Administrator Role	279
Figure 33 - Auditor/Monitor Role	280
Figure 34 - EMC Vopr SRM Role	281
Figure 34 - Direct write to primary with NAS replication	304
Figure 35 - Dual application write to both systems	304

7 Introduction

The operating system used by NetApp Filers is Data ONTAP. This guide is a general release solution implementation for all Citi NetApp NAS Filers.

This guide describes how to set up and configure NetApp Filers that run Data ONTAP 8.x.x / 7-mode software. The document is intended for Systems and Storage Administrators responsible for the management of Citi SOE NetApp Filers. The goal of this release is to provide supported and certified software for use on all Citi NetApp NAS platforms.

The solutions guide is split into four main sections:

- Chapter 10 - [Getting Started](#)
 - o Infrastructure requirements
 - o Hardware platform requirements
 - o Software installation
 - o Scalability
 - o Unsupported configurations
- Chapter 11 - [Best Practices](#)
 - o Network
 - o Storage
 - o File Sharing Protocols
 - o Data Protection
 - o Advanced ONTAP features
 - o Storage Optimization
 - o Monitoring and Alerting
- Chapter 12 - [Security Best Practices](#)
 - o Anti Virus
 - o Disk erasure (sanitization)
 - o Management authentication and account administration
 - o Restricting management access
 - o Functional IDs
 - o System volume security
 - o Audit logging
 - o TCP-IP required ports
- Chapter 14 - [Appendices](#)
 - o SNMP
 - o Standard Build Configuration Checklist
 - o Technical extended references

For links to best practices regarding specific use cases follow the hyperlinks below:

1. [VMware on NFS](#) (section [11.4.2.8](#))

This document makes extensive use of hyperlinks to improve navigation by referencing related chapters and by providing links to definitions of terms and acronyms unique to the product. Chapter referencing is used to avoid cluttering high level overview sections with too much detail, allowing the reader to take a deeper dive into certain topics if desired.

8 Assumptions, Conventions, Definitions, Abbreviations and References

8.1 Assumptions

This document assumes the reader is familiar with the basics of NetApp NAS storage, including: hardware platforms, Data ONTAP, vFilers, 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.

8.2 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 based user interface (Systems Manager). Where necessary, the type of interface used will be indicated throughout this document.

8.3 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>

As of ONTAP 7.0 and above, all CATE Storage and Backup Engineering documentation will refer to the versions of the ONTAP operating system by the generation name, for example ONTAP 8. Storage administrators and managers should consult the CATE Storage and Backup Engineering NetApp product matrix at <http://catecollaboration.citigroup.net/domains/platstor/gdse/products/platforms/NASFiler/Filer-ConfigurationMatrix.pdf> for the recommended minimum supported version on ONTAP in a given family. The current CATE Storage and Backup Engineering recommended release of Data ONTAP 8 is 8.1.4P1.

8.4 How to use this document

This document is not intended to be read cover to cover, it is purely a reference document designed to augment the product documentation provided by NetApp; provide focused and detailed insights of storage management when using technologies such as thin provisioning and storage over subscription; and consolidate important reference materials required by Citi's Information Technology Management Policies (CITMP). The best practices portion of this Solutions Guide document has been written with CITMP in mind.

8.5 Definitions and abbreviations

Term or Acronym	Explanation or Definition
Cluster	A pair of Filers configured for High availability – deprecated in favor of HA-Pair
EMS	Enterprise Management Services (SNMP monitoring)
Filer	A single network attached storage device – deprecated in favor of Storage Controller. The term is used in this document to distinguish it from the traditional Windows Server
HA-Pair	A pair of storage controllers configured for High availability
ISA	Infrastructure Security Administrator
Node	An alternate term for a storage controller
Storage Controller	The NetApp Filer “head” unit which contains compute and memory resources; and connects the IP network and disk subsystem

8.6 References

Item Name	Location / Value
CATE Storage Engineering Website	http://catecollaboration.citigroup.net/domains/platstor/gdse/products/default.aspx
CATE Storage Engineering NetApp Standards	http://catecollaboration.citigroup.net/domains/platstor/gdse/products/techstdshmpgs/netappfiler.aspx
ONTAP 8.1.4 Project Scope	https://catecollaboration.citigroup.net/domains/platstor/gdse/products/platforms/NASFiler/12032-PS.pdf
ONTAP 8.1.4 Release Notes	https://catecollaboration.citigroup.net/domains/platstor/gdse/products/platforms/NASFiler/12032-RN.pdf
ONTAP 8.1.4 Solutions Guide	https://catecollaboration.citigroup.net/domains/platstor/gdse/products/platforms/NASFiler/10059-SG.pdf
FAS2040 Release Notes	http://catecollaboration.citigroup.net/domains/platstor/gdse/products/platforms/NASFiler/08043-RN.pdf
FAS3200 series Release Notes	https://catecollaboration.citigroup.net/domains/platstor/gdse/products/platforms/NASFiler/10058-RN.pdf
FAS6200 series Release Notes	https://catecollaboration.citigroup.net/domains/platstor/gdse/products/platforms/NASFiler/10060-RN.pdf
CATE VMware VDI 1.5 Storage Design document	https://catecollaboration.citigroup.net/domains/platstor/gdse/products/Designs/CSBE-12017-DD.pdf
CATE VMware VSI Storage Design document	https://catecollaboration.citigroup.net/domains/platstor/gdse/products/Designs/CSBE-11079-DD.pdf
DNS and NTP configuration standards	http://ipns.ny.smb.com/documents/Product-Development/CA-NTP/GNAST/NTP_design_RFP.doc
NetApp Active / Active Controller Configuration	http://www.netapp.com/library/tr/3450.pdf
NetApp Spare Disk policy	http://media.netapp.com/documents/tr-3437.pdf
NetApp SMB 2.0 Technical Reference	http://media.netapp.com/documents/tr-3740.pdf
Citi Standard Build repository	https://catecollaboration.citigroup.net/domains/platstor/gdse/products/Standard%20Builds/NetAppNASFilersb.aspx

9 General Design

This section explains the general design for the implementation of network-attached storage using filers from NetApp. The diagram below provides a basic implementation example using the FAS3200 series Filer for demonstration purposes.




Figure 1 - NetApp NAS Filer Architecture Design

Filers and hardware high-availability: NetApp filer HA-pairs will be used in production with data being replicated to remote COB/Backup filer or filer HA-pair, where required by the business use case. Each HA-pair / filer will have locally attached storage. Data from multiple production HA-pairs 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 [11.4.3](#)).

Filer Fault Tolerance: Hardware high availability is implemented so as to protect from the failure of one filer head (node). In the case of a node failure, the surviving node in the HA-pair 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: 10Gig Ethernet is the current data center standard (where the network infrastructure supports it) with 1 GbE as the legacy standard. For fault tolerance, two 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 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. CATE Storage Engineering and CATE Host Security Engineering has fully integrated and certified McAfee VirusScan Enterprise (VSE) and VirusScan Enterprise for Storage (VSE-S) for On-Access scanning. Please refer to [Anti-Virus](#) section later in this document.

10 Getting Started

10.1 Infrastructure requirements

In order to deploy a Citi standard NetApp NAS solution, the following infrastructure components are required:

- Storage Controller (Filer) – device serving data via CIFS and/or NFS protocols
 - Active Directory domain and infrastructure – required for user authentication to CIFS shares and administrative authentication to the Systems Manager UI and CLI
 - LDAP domain and infrastructure – required for centralized authentication to NFS data
 - AntiVirus Servers – required for virus protection software to protect data
 - Tivoli (or equivalent) infrastructure – receives SNMP traps generated by the Filer
 - OnCommand Unified Manager server – monitors, reports, alerts and centralized storage management application for NetApp Filers
- Note: When upgrading ONTAP 8.1.4 ensure that OnCommand Unified Manager 5.2P2 (or) later is deployed to prevent interoperability issues with performance monitoring functionality**
- OnCommand Systems Manager server – storage system management via GUI
 - EMC Vopr SRM Collector server – retrieves capacity and performance data from the storage system
 - Syslog server (HP ArcSight / SEMS appliance) – remote recipient of ONTAP's audit and system message logs
 - Clients – Windows, ESX, RedHat and Solaris clients wishing to leverage shared or dedicated network attached storage
 - Network – a common IP network to connect all these infrastructure components
 - Contingency (DR) Filer – replication destination system when DR services are required
 - Backup infrastructure (Networker/NetBackup) – tape based backup solution when backups are required for additional data protection services

10.2 Hardware platform requirements

Not all NetApp storage controllers support ONTAP 8. See the [ONTAP configuration matrix](#) for details:

- FAS2040
- FAS3100 series
- FAS3200 series
- FAS6000 series
- FAS6200 series

10.3 Software installation

Software is obtained from NetApp's Support website (<http://support.netapp.com>). A list of approved software images for ONTAP 8.1.4P6 can be found in section 8 of the [CATE ONTAP 8.1.4P6 Release Notes](#).

CATE Storage and Backup Engineering does not test, qualify or certify releases of system board firmware, opting instead to follow the latest version available from NetApp.

Disk shelf and RLM firmware is tested and certified on a regular basis. Details of the latest certified versions can be found in section [11.3.2.1](#) and [11.2.5](#) respectively.

Unless a critical bug is discovered that requires a mandatory upgrade to a certain release of disk drive firmware, CATE Storage and Backup Engineering will only qualify the latest drive firmware included in currently certified release of ONTAP.

Many issues relating to system performance and stability are resolved by implementing the latest firmware levels. It's important to keep system board, shelf and disk firmware levels up to date to minimize issues.

10.4 Scalability

Starting with ONTAP 8.0.1 and above, CATE Storage and Backup Engineering supports NetApp Filer hardware configurations up to **100% of the maximum supported raw storage** as listed in the product specifications on the <http://support.netapp.com> site. However there are scaling limitations with the use of NDMP that may artificially limit the raw storage capacity of your storage system. Storage Administrators are requested to check with their local backup teams to confirm the supported maximum capacity of the storage system permitted to allow a successful NDMP backup schedule. Consult the platform hardware release notes on the [CATE NetApp Filer product page](#) for details.

10.5 New installations

New installations should refer to Appendix C in section [14.3](#) for the operating system configuration options.

10.6 Upgrading from a previous release

CATE Storage and Backup Engineering is releasing ONTAP 8.1.4P6 for installation on new systems as well as designating 8.1.4P6 as the converged upgrade path for system currently running ONTAP 7.3.6, 8.0.2 and prior versions of ONTAP 8.1.

10.7 Unsupported configurations

CATE Storage and Backup Engineering has not certified, nor supports the use of the following configurations or features within ONTAP. This list is neither complete nor exhaustive but is designed to answer the regular support enquiries received for the NetApp NAS Filer platform. If a certain configuration is not listed in this document (either as supported or unsupported) then please check with gdse@citi.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)
- Metro-cluster (geographically stretched cluster)
- SyncMirror (maintains two copies of the data in an aggregate by using separate plexes)
- Mixed mode security style volumes (exception: vfiler pseudo root qtrees only¹)

¹ This only applies to vFilers where a qtree is used to provide the root volume to a vFiler. Current best practices recommend using a flexvol to house a vFiler root volume.

11 Best Practices

This section of the document covers best practice configuration settings that relate to the use of ONTAP 8 / 7-mode. For security related configuration information, refer to section [12](#).

11.1 ONTAP SOE configuration: options and registry settings

All Operating System options and customized registry settings are defined in Appendix C in section [14.3](#). Refer to the [Table 33 - ONTAP Mandatory \(Security\) Configuration Matrix](#), [Table 34 - ONTAP Recommended \(Technical Reference\) Configuration Matrix](#) and [Table 35 - Registry custom settings](#) for full details.

11.2 Network

Depending on the hardware series NetApp Filers are capable of supporting diverse network configurations through a collection of multi-purpose and single tasked network interfaces. The following table presents a summary of the different network options available with each hardware platform supported by CATE with ONTAP 8.1.4P6.

Platform Model	Data	Management ²	Console (serial)
FAS2040	Various	e0M	Serial
FAS3140/3160/3170	Various	e0M	RLM preferred
FAS3210/3240/3270	Various	e0M	SP preferred
FAS6030/6040/6070/6080	Various	e0M	RLM preferred
FAS6210/6240/6280	Various	e0M	SP preferred

Table 2 - Supported port roles

For the remainder of this document the term port is used to describe the physical network connection, interface is used in terms of the logical connection and interface group is used for the aggregation of interfaces.

Data connectivity is comprised of two interface types on the NetApp NAS Filer platform: single mode or multi-mode interface groups. Consult the hardware platform release notes on the [CATE NetApp Filer product page](#) for details on which network ports should be configured depending on the interface group type and count required.

For security information regarding the IP port requirements on NetApp Filers see section [12.7](#).

² Refer to section [11.2.5](#) for complete details on the service processor (SP), baseboard management controller (BMC) and section [11.2.6](#) for the dedicated e0M management port

11.2.1 Single-mode (Active/Passive) Interface Groups

- In a basic implementation two Ethernet network interface ports will be used in each Filer.
- The two Ethernet ports are 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 or VLAN) since in a cluster fail-over scenario, one node has to takeover the name and IP address of the other node.
- When configuring basic interface groups, single-mode type groups should be used (active-passive).
- In a clustered configuration all interfaces must have a name unique within the cluster
- All single-mode gigabit interfaces will be called "gev" followed by a number
- The first gigabit interface in a cluster will be called gev1. As mentioned above this interface name must be unique within the cluster

Below is a sample command line for interface group creation.

```
ifgrp create single gev1 e0a e2a
```

11.2.1.1 Interface favoring

For an optimal network configuration it is important to distribute the data traffic over the redundant pair of switches. In ONTAP there is an interface group favoring technique that can be used to achieve this behavior. Of the two network interfaces configured on each Filer one interface will be considered the primary or preferred interface whose network uplink goes straight through the switch and does not cross an ISL. This configuration assumes an active/active switch pair is being used, as opposed to one of the switches being purely passive for redundancy.

Firstly define the default gateway of each Filer in the NetApp cluster to its preferred switch, so that Filer A will have Switch 1 as its default gateway and Filer B will have Switch 2 as its default gateway. Then set the favored interface in the VIF or Interface Group so that the physical link to the preferred switch is the active network interface (the patching schedule will be required from the Data Center implementation in order to complete this stage). The favored interface should be physically connected to the same switch that is also the default gateway. This is achieved using the ifgrp command:

On an ONTAP 8.x/7-mode system use the ifcfg command.

```
Filer> ifgrp favor <interface_name>
```

Also remember to add this configuration to the startup sequence of the Filer in the /etc/rc file.

11.2.2 Multi-mode (Etherchannel) Interface Groups

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 support Link Aggregation Control Protocol (dynamic 802.3ad) however this requires further testing and is not certified for production use within Citi at the time of writing this document.
- The multi-mode interface group 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 / aggregated IP network ports into an interface group:

- All the ports participating in a bundle **MUST** be connected to the same switch.
- 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 controller fail-over scenario; one node has to takeover the name and IP address of the other node.
- Filers that utilize four or more Ethernet ports will have 3 virtual interfaces each. Two interface groups will be created for each of the 2-port bundles using Multiple-mode (All active). All interfaces in a multimode interface group **MUST** be connected to the same network switch. Thereafter, another interface group will be created in Single-mode (Active-Passive) containing the previous 2 interface groups. The result is that only one network interface group (utilizing half the provisioned ports on the Filer) will be active at one time.
- When increasing the number of active ports in a multi-mode interface group, in order to match Citi network standards, the amount of new ports needs to be in the order of two. An odd number of active ports in a multi-mode interface group are not supported.
- In a controller high-availability configuration all interface and interface groups must have a unique name within the controller HA-pair
- All multi-mode gigabit interface groups will be called "mgev" followed by a number
- The first multi-mode gigabit interface group in a HA-Pair will be called mgev1

Below is a sample command line for a multi-mode interface group creation which is then protected by a single mode interface group for and active / passive configuration.

```
ifgrp create multi mgev1 e0a e2a
ifgrp create multi mgev2 e0c e2c
ifgrp create single gev3 mgev1 mgev2
```

11.2.3 10GB Ethernet

The use of 10Gb NICs is supported by NetApp and certified for use by CATE Storage Engineering where 10Gb network infrastructure is deployed. The latest 10Gb cards benefit from a TSO (TCPIP Segment Offload) capability and the TSO functionality is available even when the ports are joined together in active/passive (single-mode) interface group.

11.2.3.1 Required hardware

The following hardware components are required for a 10GB Ethernet configuration per NetApp storage system HA-Pair:

Storage Platform	Hardware components
FAS2000	Not applicable – 10GB not supported
FAS3100	4 dual port 10GbE SFP+ NICs
FAS3200	(NetApp Part Numbers: X1107A-R6 or equivalent) ³
FAS6000	4 dual port 10GbE SFP+ NICs
	(NetApp Part Numbers: X1107A-R6 or equivalent) ⁴
FAS6200	No additional hardware required for basic fault tolerant active/passive interface configuration – see FAS6200 Release Notes

Table 3 - Hardware requirements to support 10GB ethernet

Note the mid-range (FAS3x00) NetApp system do not contain onboard 10GbE ports so all 10GbE connectivity is delivered through the use of high performance network adapters for installation in the PCI I/O expansion slots.



Prior to ordering the dual port 10GbE SFP+ network adapters, information regarding the physical network topology at the installation location is required to ensure the appropriate SFP+ module is specified (either fiber or twin-ax copper). The physical layer of network connectivity may vary by site so it is important to determine whether fiber or twin-ax copper infrastructure is available so the correct SFP+ modules are ordered. The expectation is the majority of sites will use fiber connectivity for 10GbE cabling.

NetApp dual port 10GbE network adapters feature a single ASIC servicing both ports on the card. In order to ensure ASICs are not overloaded and to provide an appropriate level of hardware component resiliency, two network adapters per storage controller (Filer) are required.

11.2.3.2 Network Adapter IO Expansion Slot Assignments

NetApp details the IO expansion slot assignment best practices in it's Hardware Universe application (<http://hwu.netapp.com/Adapter/Index>). Refer to the 'priority' and 'priority slot assignment' columns in the tool for expansion slot assignment recommendations.

As slot assignment best practices vary between storage controller hardware platforms, network adapter models and ONTAP releases there are too many permutations to fully capture in this design document.

³ At the time of writing this document X1107A-R6 NICs were the latest 10GbE cards available from NetApp and were supported with ONTAP 7.3.2 and above. Later versions of 10GbE cards may require a higher version of ONTAP in order to support the newer hardware. Ensure the appropriate due diligence is performed when ordering hardware to check for changes in the 10GbE NIC part number and ONTAP dependencies.

⁴ At the time of writing this document X1107A-R6 NICs were the latest 10GbE cards available from NetApp and were supported with ONTAP 7.3.2 and above. Later versions of 10GbE cards may require a higher version of ONTAP in order to support the newer hardware. Ensure the appropriate due diligence is performed when ordering hardware to check for changes in the 10GbE NIC part number and ONTAP dependencies.

11.2.3.3 ONTAP Network Settings to support 10GB Ethernet

11.2.3.3.1 Trunking

From each Filer head trunk two interfaces in a 'single-mode' / failover (active/passive) Interface Group. Take one port (interface A) on the first network adapter and another port (interface B) on the second network adapter and combine them together in a single mode interface group or VIF. Interface A will be patched to Switch 1 whereas interface B will be patched to Switch 2 in order to provide a highly available network configuration, protecting the environment against multiple forms of component failure. This is achieved using the ifgrp command:

```
Filer> ifgrp create single <ifgrp_name> <interface1_name> <interface2_name>
```

11.2.3.3.2 Interface failover policy

As there is no second level interface groups in this configuration there is no need to enable the failover.link_degraded ONTAP option. Instead keep this parameter at its default 'off' value. This is configured using the ifgrp.failover.link_graded option.

```
Filer> options ifgrp.failover.link_degraded off
```

11.2.3.3.3 Maximum Transmission Unit (MTU) size

Each 10GbE physical port in a VIF or Interface Group, as well as the logical VIF or Interface Group configuration must specify the default MTU of 1500. This is achieved by the ifconfig command and setting the mtusize attribute for a port to 1500.

```
Filer> ifconfig <interface_name> mtusize 1500
```

Jumbo frames must not be used in implementations using this design as the Cisco General Shared Services (GSS) network is configured with the default MTU of 1500. To avoid transmission fragmentation with jumbo frame enabled hosts communicating over a 1500 MTU network all hosts are recommended to use the default MTU size.

Where Citi IP Fabric is deployed the MTU allowed by the network is 8800. However increasing the MTU size on the storage array yields a marginal benefit at best and as such 1500 will continue as the standard MTU configuration in the IP Fabric. Changes to this default value, whether increase or decrease, will require approval from CATE Storage and Backup Engineering.

11.2.3.3.4 Flow Control

Each 10GbE physical port in a VIF or Interface Group should be configured with flow control set to off. This is achieved by the ifconfig command and setting the flowcontrol attribute for a port to none.

```
Filer> ifconfig <interface_name> flowcontrol none
```

11.2.3.4 Network Switch Settings to support 10GB Ethernet

In accordance with NetApp networking best practices the switch ports should also have flow control disabled. The command to do this will vary by switch and operating system type. Consult your regional networking support team for details

11.2.3.5 Interface naming conventions

- In a storage HA-Pair configuration all interfaces must have a name unique within the cluster
- All 10GbE interfaces will be called “tgev” followed by a number
- The first single-mode gigabit interface in a cluster will be called tgev1

11.2.3.6 Network connectivity example

This section describes the 10GbE network connectivity design for NetApp NAS storage. The following diagram uses a FAS3240AE cluster as an example. The hardware platform in your implementation may be different, as well as the network switch models and the PCIe IO expansion slot assignment preferences may also differ.




Figure 2 - Physical network connectivity example for 10GB Ethernet

11.2.4 Router Pair Connectivity and Cluster Configuration

The following diagram depicts the interconnectivity and relationship between a NetApp NAS Filer HA-Pair and a common switch/router pair network configuration. The diagram demonstrates an example of a 2x1GB multimode interface group with a single-mode interface group on top for active/passive switch failover.




Figure 3 – General network connectivity example

From the diagram you can see that the first Filer is connected to both the first and second router and the second Filer is connected to the second and first routers. The wording of this statement is a clue to how the default gateways should be configured on an HA-Pair Filer configuration.

The default gateway should match the switch with the active interface(s) connected and the interface group connected to the default gateway should be the favored interface and defined in the `/etc/rc` file accordingly. Therefore if there is an issue with one router only the active interfaces on one Filer in the cluster will be affected. Refer to section [11.2.8](#) for further details on the appropriate entries in the `/etc/rc` and `/etc/hosts` files.

11.2.5 Service Processor (SP) / Remote LAN Management (RLM)

The Service Processor is the next generation management interface introduced with new Filer models, starting with the FAS3200 and FAS6200 series, replacing the RLM seen on the FAS3000, FAS3100 and FAS6000 series. The networking requirements are the same whether an RLM or SP is installed on the NetApp array so for the remainder of this section where the term SP is used it can be replaced with RLM unless the two terms cannot be used interchangeably and at which point the differences will be noted.

The SP and e0M management interfaces share the same physical RJ45 port on the center of the storage controller's rear panel which is indicated by the wrench or spanner icon above it. The two interfaces can share the same port due to a small network switch embedded in on the controller's system board.




Figure 4 - SP (remote management device) and e0M network connections

Starting with ONTAP 8.1 CATE restrictions on connecting the Service Processor or the RLM to Citi's management network have been lifted due to control enhancements introduced in Data ONTAP.

11.2.5.1 The Purpose of the Service Processor

The SP enables you to access the storage system remotely to diagnose, shut down, power-cycle, or reboot the system, regardless of the state of the storage controller. It is powered by a standby voltage, which is available as long as the system has input power to at least one of the system's power supplies. The SP monitors the system temperatures, voltages, currents, and fan speeds. When an environmental sensor has reached an abnormal condition, the SP logs the abnormal readings, notifies Data ONTAP of the issue, and sends alerts and "down system" notifications as necessary through an AutoSupport message, regardless of whether the storage system can send AutoSupport messages.

The SP performs Field Replaceable Unit (FRU) reporting, FRU tracking; monitors and reports on internal sensors; reports on and can set LED states; and actively manages fan speed. It has a nonvolatile memory buffer that stores up to 4,000 system events in a system event log (SEL) to help you diagnose system issues.

The SP only uses SSH and does not support Telnet or RSH. It provides access to the serial console in a secure manner, whereas some terminal servers still rely on telnet. Therefore, in Citi, the SP is recommended interface for accessing the system console. After retrofitting SP connectivity to existing storage controllers the serial console terminal server interface must be decommissioned.

11.2.5.2 Network requirements

This section details the network requirements associated with using the Service Processor or RLM. As the e0M interface shares the same physical port it also shares same network requirements:

- 2 physical network connections per HA-pair (one per storage controller)
- 4 IP addresses per HA-pair (two per storage controller, one for the SP and one for e0M)
- Port speed of 100 mbps (minimum) or 1 gbps (preferred)
- Data center management network address space / subnet that complies with RFC 1918
- Both nodes of the storage HA-pair must have their management interfaces on the same management subnet
- The management subnet **must** be different to the default route subnet on the Filer (check the /etc/rc file for the 'route add default' statement to determine the default route subnet)
- Data center management network must be routable
- The DNS name for the management interface (e0M) and the service processor / rlm should be significantly different from the DNS name for the data interfaces to prevent clients from attempting to connect to the management interface
- Data center management network must allow access to the following infrastructure services:
 - o Citi Tivoli trap hosts via SNMP
 - o Citi Mail relays via SMTP
 - o Citi Domain Name Resolution via DNS
 - o Citi Time Servers via NTP
 - o Internet access for call home via HTTPS using Citi's web proxy servers
 - o Peer connectivity to the HA partner management port e0M over port 4444
 - o Inbound SSH connectivity on port 22 from management hosts on the internal Citi network

Note: If retrofitting SP / RLM connectivity on an existing NetApp storage controller there may not be enough physical connectivity or IP addresses on the management network infrastructure for that location as these network requirements may not have been factored into the original system installation request.

11.2.5.3 Account entitlements and limitations for the SP

Due to the limited amount of resources on the service processor it does not contain the relevant software components in order to authenticate with Active Directory. Therefore only local accounts can be used to login to the Service Processor. As the Service Process is deemed a management interface of last resort, and as such is not intended for day to day system administration, it is expected that that use of privileged interactive functional IDs such as root with appropriate password ownership controls will be an acceptable process when required.

Only select accounts have the entitlement that allows them to log into the SP. By default all accounts in the ONTAP built-in 'Administrators' group inherit this privilege. Other accounts are granted this privilege through Citi's ONTAP RBAC model via the ONTAP 'login-sp' capability. The two Citi roles which inherit this ability are the System Administrator and the Storage Administrator.

11.2.5.4 ONTAP parameters defining SP / RLM behavior

There are three ONTAP options which define the behavior of the Service Processor or RLM. The options will have slightly different names depending on which generation of remote management device is installed in your storage controller. Simply replace the 'sp' characters in the option name with 'rlm' characters to get the equivalent parameter for the older FAS models.

- `sp.autologout.enable`
This option controls whether the automatic logout inactivity timer is enabled for the SP. It is enabled by default in ONTAP and is configured as part of the Citi ONTAP build defined in section [14.3](#).
- `sp.autologout.timeout`
This option controls the duration of the automatic logout inactivity timer for the SP. The default duration is 60 minutes and is configured to a different value as part of the Citi ONTAP build defined in section [14.3](#).
- `sp.ssh.access`
This option defines which hosts are permitted to access the SP interface as part of an initial validation of authorized access prior to authentication with entitled credentials. The default value for this setting is '*', meaning all hosts on the network are authorized to access. As part of the Citi ONTAP build this setting needs to be configured with the same host restrictions as the option 'ssh.access' as defined in section [14.3](#).

11.2.5.5 Configuring the SP / RLM for first use

The SP can be configured for first time use within Data ONTAP using the '`sp setup`' command. For older hardware platforms with an RLM installed use the '`rlm setup`' command. Prior to running this command the following information is required by the setup process:

- The static IP address assigned to the SP in whitepages
- The subnet mask of the management network to which the SP is connected
- The default gateway of the management network to which the SP is connected

Before initiating the setup process ensure the SP is running at least the recommended minimum firmware version as specified in the CATE Release Notes for your particular NetApp hardware platform. Use the '`sp status`' command within ONTAP to check the firmware version. For older hardware platforms with an RLM installed use the '`rlm status`' command.

```
ontap> sp status
Service Processor      Status: unconfigured
Firmware Version:     1.3.3P2
Mgmt MAC Address:     00:A0:98:14:15:1E
```

One final check before commencing the SP setup process is to validate that the call home options 'autosupport.to' and 'autosupport.mailhost' are configured correctly according to the Citi build standard described in section [14.3](#).

The setup process is a simple five question procedure:

- Confirm you want to configure the SP
- Confirm whether you want to use DHCP to configure the SP; Answer: No
- Provide the statically defined IP address of the SP interface
- Provide the netmask of the management network to which the SP is connected
- Provide the gateway address of the management network to which the SP is connected

The following example uses 'x', 'y' and 'z' to represent the relevant network details. The network responses you provide during your setup process will differ according to your environment.

```
ontap> sp setup
The Service Processor (SP) provides remote management capabilities
including console redirection, logging and power control.
It also extends AutoSupport by sending
additional system event alerts. Your AutoSupport settings are used
for sending these alerts via email over the SP LAN interface.
Would you like to configure the SP? y
Would you like to enable DHCP on the SP LAN interface? n
Please enter the IP address of the SP [ ]: xxx.xxx.xxx.xxx
Please enter the netmask of the SP [ ]: yyy.yyy.yyy.yyy
Please enter the IP address for the SP gateway [ ]: zzz.zzz.zzz.zzz
```

Once the setup process is configured confirm the settings have been implemented correctly using the 'sp status' command. For older hardware platforms with an RLM installed use the 'rlm status' command.

```
ontap> sp status
Service Processor           Status: Online
Firmware Version:          1.3.3P2
Mgmt MAC Address:          00:A0:98:14:15:1E
Ethernet Link:              up, 100Mb, full duplex, auto-neg complete
Using DHCP:                 no
IPv4 configuration:
IP Address:                 xxx.xxx.xxx.xxx
Netmask:                    yyy.yyy.yyy.yyy
Gateway:                    zzz.zzz.zzz.zzz
IPv6 configuration:         Disabled
```

11.2.5.6 SP / RLM Remote Support Agent (RSA)

The SP and RLM both features an embedded piece of software called the Remote Support Agent which, if enabled, could allow a NetApp Technical Support Engineer to run a short pre-defined set of log and information gathering commands on the storage controller.

At this **time this feature is not approved for use** by the Citi Third Party Access Working Group.

In order to validate this feature has not been enabled the following command must be run on the Service Processor itself. In other words this command does not exist in ONTAP.

```
sp> rsa show
Remote Support Agent is disabled.
```

11.2.6 The Management Port (e0M)

All models in the FAS3x00, FAS6x00 series and the FAS2040 feature an e0M interface that uses the same physical RJ45 port as the SP or RLM located on the storage controller's rear panel which is indicated by the wrench or spanner icon above it. The e0M port allows for the separation of data and management traffic on the storage system and supports all the management protocols which are available on other interfaces on the storage system (SSH, HTTPS and SNMP, as well as DNS, SMTP and NTP).

As the e0M interface shares the same physical interface as the SP it shares the same network requirements as described in section [11.2.5.2](#).

11.2.6.1 The Purpose of the Management Interface (e0M)

The purpose of the e0M port is to provide a dedicated management interface to the storage controller in order to provide a mechanism to segregate management and data traffic. As this interface is certified by CATE Storage and Backup Engineering with ONTAP 8.1 and above it is now the preferred interface by which to manage the storage controller.

However it is important to note that the e0M is a low bandwidth interface that should be configured for management traffic only. Running data traffic over the e0M interface can cause performance degradation and routing problems. The following best practices are defined in order to prevent running data traffic over the e0M.

In addition to providing SSH access to the storage controller the e0M interface also allows for the hardware assisted HA-pair failover feature to be leveraged. Storage controllers are usually deployed in pairs for hardware high availability and are interconnected via a dual path heart beat cable (previously referred to as the cluster interconnect). There is a finite level of granularity to the command set that can pass down this interconnect and so the hardware assisted failover feature builds on this to allow one storage controller to detect hardware issues on it's partner quicker and more proactively than waiting for timers to expire with the traditional interconnect. This means that while the failover time will remain the same, the amount of time it takes to detect a chronic hardware failure will be significantly decreased and thereby lessen user impact.

11.2.6.2 Account entitlements for e0M

As e0M is fully integrated into ONTAP it can leverage Active Directory as it's storage administrator authentication mechanism therefore storage administrators are not required to maintain local accounts within ONTAP in order to use the e0M interface.

The e0M interface works with the same login entitlements as other ONTAP interfaces such that only the login-ssh RBAC capability is required obtain the right to login. Management host level authorization is configured through the ONTAP option 'ssh.access' as defined in the Citi ONTAP build (see section [14.3](#) for details). As previously stated authentication is performed by Active Directory.

11.2.6.3 ONTAP parameters defining e0M behavior

There are four ONTAP options which define the behavior of certain aspects of the e0M interface.

- `cf.hw_assist.enable`
The option defines whether hardware assisted storage controller failover is utilized. It is enabled by default in ONTAP and is configured as part of the Citi ONTAP build defined in section [14.3](#).
- `cf.hw_assist.partner.address`
This option defines the IP address of the partner system's e0M interface. This is usually detected and configured automatically by ONTAP during system boot but can be configured manually when performing a first time setup. If a host name is inadvertently entered it will be converted to an IP address and stored as such in this parameter. The default value is the IP address of the partner's e0M interface and is configured as part of the Citi ONTAP build defined in section [14.3](#).
- `cf.hw_assist.partner.port`
This option defines the IP port which each e0M interface uses to communicate to its counterpart on the partner storage controller. The default value for this setting is '4444' and is configured as part of the Citi ONTAP build defined in section [14.3](#).
- `interface.blocked.mgmt_data_traffic`
This option defines whether high bandwidth data traffic such as SnapMirror, SnapVault, CIFS, NFS iSCSI and NDMP are permitted to use the low bandwidth management port. The default value for this setting is 'off' if the system has been upgraded from a version of ONTAP prior to 8.0.2, and 'on' if upgraded from 8.0.2 or later. It is configured as part of the Citi ONTAP build defined in section [14.3](#).

11.2.6.4 Special considerations for blocking replication and backup traffic over e0M

In a single path SnapMirror transfer between the source and the destination, it is possible to block SnapMirror traffic from e0M on the source storage system by using the `interface.blocked.mgmt_data_traffic` option described above; however, you cannot block SnapMirror traffic from e0M on the destination by using the `interface.blocked.mgmt_data_traffic` setting. In order to block traffic on e0M, all the storage systems must use a consistent configuration. Storage administrators must ensure that routes to the destination storage systems do not use e0M, and that neither SnapMirror nor NDMP is configured to use an IP address that is assigned to e0M on a destination storage system.

11.2.6.5 Management interface configuration standards

The e0M interface is configured using the same 'ifconfig' command which many other interfaces on the storage controller are configured by; however there are notable additions to the configurations parameters for e0M:

- Disable dynamic DNS (no_ddns parameter)
- Disable WINS resolution (-wins parameter)
- Disable negotiated failover of the e0M interface (-nfo parameter)

This configuration for the e0M needs to be applied in two places; the running configuration and the startup configuration. To change the running configuration use 'ifconfig' command similar to the following example, where xxx.xxx.xxx.xxx represents the statically defined IP address to be assigned to e0M and yyy.yyy.yyy.yyy represents the mask for the subnet:

```
ontap> ifconfig e0M no_ddns xxx.xxx.xxx.xxx netmask yyy.yyy.yyy.yyy mtusize
1500 mediatype auto -wins -nfo
```

To edit the startup configuration follow the following procedure:

1. Mount the C\$ Windows share of the storage controller and navigate to /etc. Make a backup copy of both the hosts and rc files (example: rc.YYYY-MM-DD.bak).
2. Append the IP address and hostname of the e0M interface the to /etc/hosts file and save the updated file:

```
xxx.xxx.xxx.xxx      aaaaaaaaaa-e0M
```

 Where xxx.xxx.xxx.xxx is the IP address assigned for the e0M interface, and aaaaaaaaaa represents the hostname of the physical storage controller.
3. In the section of the /etc/rc file which contains other ifconfig statements append the configuration for e0M to the end of the ifconfig section, where yyy.yyy.yyy.yyy represents the mask of the subnet to which e0M is connected:

```
ifconfig e0M `hostname`-e0M no_ddns netmask yyy.yyy.yyy.yyy flowcontrol
full -wins -nfo
```

11.2.7 The Alternate Control Path port (eOP)

The 3200 and 6200 series, as well as the FAS2040, include an on-board out-of-band disk subsystem management port, designated eOP. Although connectivity of the eOP port to the NetApp SAS shelves is not required for the storage system to serve data, implementation of the private ACP network between all controllers and SAS shelves in a storage system is strongly recommended for system resiliency, hardware recovery, disk enclosure log collection, and the ability to remotely manage the disk enclosure thereby reducing the need for a physical visit to the data center to reboot a 'wedged' enclosure (for example).

The recommend approach when configuring eOP is to use the system default domain, 192.168.0.0/22 (255.255.252.0). If you are configuring ACP for disk shelves attached to an active/active configuration, you must supply the same ACP domain name and network mask for both systems. Do not connect the ACP port to a routed network, and do not configure switches or hubs between the ACP port and the designated Ethernet port. Doing so is not supported and will cause interruptions in service.

To check the domain configured for the ACP use the following command:

```
Filer> storage show acp
```

```
Alternate Control Path: Enabled
Ethernet Interface:    eOP
ACP Status:           Active
ACP IP Address:        192.168.0.75
ACP Subnet:            192.168.0.1
ACP Netmask:           255.255.252.0
ACP Connectivity Status: Full Connectivity
```

Shelf Module	Reset Cnt	IP Address	FW Version	Module Type	Status
3d.11.A	000	192.168.3.35	01.31	IOM3	active
3d.11.B	000	192.168.3.59	01.31	IOM3	active
3a.01.A	000	192.168.2.183	01.31	IOM3	active
3a.01.B	000	192.168.2.197	01.31	IOM3	active

If the ACP private network is not configured to use the system default domain of 192.168.0.0 then use the `acpadmin configure` command to reconfigure the eOP interface as shown in the following example where the incorrect domain 198.15.1.0 has been used (typically as a result of bug 382555):

```
Filer> acpadmin configure
```

```
The Alternate Control Path Management process is already enabled. Do you want
to reconfigure it? [n]: y
Enter the network interface you want to use for the Alternate Control Path
Management [eOP]: eOP
A network domain is required for the Alternate Control Path Management
interface.
Specify a domain not currently in use, and enter it in the form "192.168.1.0"
Enter the domain for Network Interface eOP [192.168.0.0]: 192.168.0.0
Enter the netmask for Network Interface eOP [255.255.252.0]: 255.255.252.0
[acpadmin configure] ACP processing new configuration settings.
delete host 192.15.1.55: not in table
delete host 192.15.1.61
add host 226.225.224.223: gateway 192.168.0.75
delete host 226.225.224.223
arp: Failed. 192.15.1.88 is not on a directly attached subnet.
arp: Failed. 192.15.1.89 is not on a directly attached subnet.
arp: Failed. 192.15.1.90 is not on a directly attached subnet.
arp: Failed. 192.15.1.91 is not on a directly attached subnet.
Tue Apr 26 13:54:34 EDT [11c2-fas3240-01: nbt.nbns.registrationComplete:info]:
NBT: All CIFS name registrations have completed for the local server.
```

Once the NetBIOS registration confirmation is shown the reconfiguration process is complete and the new setting can be confirmed with the **storage show acp** command.

For full details on the ACP network and cabling requirements refer to the CATE [DS4243](#) and [DS2246](#) Release Notes, as well as section 11.2.8 regarding the /etc/rc and /etc/hosts files.

11.2.8 Host and rc file configuration

The /etc/hosts and /etc/rc files contain important details relating to the network configuration of your storage controller. The following sections highlight the high level best practices:

11.2.8.1 /etc/rc

The rc file contains critical network configuration information required to define the network settings on the storage controller during the boot process. The following table defines the key elements in the rc file and provides examples where appropriate. The table does not represent an exhaustive list of the contents and settings of the rc file. The order of contents within the /etc/rc is critical to ensure appropriate storage controller behavior and to avoid any unnecessary system panics. Note not every ONTAP command can be run at startup due to the order in which services are started during boot. It is not recommended to deviate from the following types of commands without review by CATE Storage Engineering and NetApp.

Task	Purpose	Example
Set hostname	Defines the hostname of the system	hostname 11c2-fas3240-01
Create interface groups ⁵	Define and configure network interface groups	ifgrp create single gev1 e0a e2a
Configure vlans	Define and configure vlan settings	
Configure interfaces ⁶	Define and configure the network interfaces on the storage controller (note ONTAP will reference the /etc/hosts file to obtain the IP address associated with `hostname`-gev1)	ifconfig gev1 `hostname`-gev1 mediatype auto netmask 255.255.255.0
Configure vfilers	Prevent vFilers which are currently stopped from restarting if the system reboots	vfiler stop 11c2-fas3240-v01
Default route	Defines the default route – the IP address of the network routers gateway interface to which the Filer is attached (see section 11.2.4 for best practices in a cluster configuration)	route add default 10.40.60.2 1
DNS domain	Define the DNS domain name for the system	options dns.domainname nam.nsroot.net
Enable DNS	Enables the DNS daemon	options dns.enable on
Disable NIS ⁷	Disable the NIS daemon	options nis.enable of
Enable core files	In the event that a storage system panics the savecore commands allows a memory dump to be saved to /etc/crash on next boot	savecore
Other commands	Other commands required to enforce settings which are traditionally not persistent across reboots	priv set -q diag;setflag smb_enable_2_1 0;priv set

Table 4 - The /etc/rc file key concepts

Under no circumstance should the /etc/rc file contain ifconfig statements for the losk interface. Doing so may cause the storage controller to panic.

⁵ Your system may require more than one entry depending on the network configuration requirements and the number of interface groups.

⁶ Your system may require more than one entry depending on the network configuration requirements and the number of interface groups.

⁷ Following the NIS to LDAP migration program, NIS is no longer supported by CATE Platforms and Storage Engineering with ONTAP 8 and above

11.2.8.2 /etc/hosts

The hosts file maps IP address to host names. The file may include entries that were automatically created by ONTAP (localhost-stack, localhost-bsd, localhost-sk) and these should not be edited unless instructed to do so by NetApp support.

In order to ensure the network configuration is persistent across reboots each defined network interface or interface group should have an entry in the /etc/hosts file that corresponds to an ifconfig statement in the /etc/rc file. The entry for interface e0P in the /etc/hosts file should always be commented out with a '#' character at the start of the line.

11.2.9 Filer NetBIOS Naming Standards

The NetBIOS protocol standards limits device names to 15 characters.

11.2.9.1 Primary (Production) Site

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

Example: NASNY01 or NASLDN01.

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

11.2.9.2 Backup / Disaster Recovery Site

To distinguish primary filers from disaster recovery filers, "DR" will 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.

11.2.10 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 to resolve names in the order specified in the */etc/nsswitch.conf* file.

Citi Name Resolution Order	
Map	Services in each Map
hosts	files DNS
Passwd	files LDAP
netgroups	files LDAP
group	files LDAP
shadow	files LDAP

Table 5 - Recommended name resolution order

For example, using the Citi name resolution order shown above:

- 1) Initially, the filer attempts to resolve a netgroup names locally by searching the */etc/netgroup* file (designated by "files" in the table below and in the */etc/nsswitch.conf* file)
- 2) If name is not resolved, the filer tries LDAP (if LDAP is enabled)

Note: LDAP does not provide host name resolution at Citi and hence only the local */etc/hosts* file and DNS should be used for name resolution. For this reason the default hostname search should always be modified to remove NIS from the "hosts" entry.

By default ONTAP creates NIS entries in the */etc/nsswitch.conf* file during system setup (see [Table 6](#)). As NIS is now a retired product within Citi, all NIS entries should be removed from the */etc/nsswitch.conf* file as part of the system build and configuration process.

The default name resolution file supplied by NetApp is shown below:

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 6 - 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, LDAP configuration must be completed later manually. Starting with ONTAP 8.0.1 and above NIS should not be enabled by the operator as part of the initial filer setup process. **Filers must be configured to use the global DNS settings as described in http://ipns.ny.ssmb.com/documents/Product-Development/CA-NTP/GNAST/NTP_design_RFP.doc. The order in which these IP addresses should be configured at each site can be found in Appendix C of the above document. If the site at which the 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.

To ensure interoperability with the AnyCast DNS infrastructure ensure that the option [ip.fastpath.enable](#) is disabled on all NetApp NAS Filers. Failure to disable the ip.fastpath.enable option may result in communication issues between the Filer and Active Directory due to DNS lookups of the domain SRV records failing as the lookup transaction connection gets continually reset by the DNS server. The message log extract below is indicative of a system suffering from this issue:

```
[<filer_name>: Auth00:info]: Lookup of
_ldap._tcp.dc._msdcs.<domain_name>.nsroot.net failed with DNS server
<dns_server_ip_address>: Connection reset by peer.
```

The other benefit of disabling the ip.fastpath.enable option is it decreases the ability for ARP spoofing and session hijacking attacks.

11.2.11 Time Synchronization and Time Zones

11.2.11.1 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.

There are some notable changes in the configuration of NTP in ONTAP between the 7.x and 8.x releases. Consult the [CATE ONTAP 8.0.2 Release Notes](#) for details.

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. As of ONTAP 7.2 or above this value defaults to zero and is now considered obsolete.
- As of ONTAP 8.0 and above the time protocol defaults to ntp and rtc and rdate values have become obsolete.
- As of ONTAP 8.0 and above the time synchronization schedule and scheduling window parameters have become obsolete.

Refer to Appendix C of http://ipns.ny.ssmb.com/documents/Product-Development/CA-NTP/GNAST/NTP_design_RFP.doc for the order in which the time server IP addresses should be configured at each site.

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 recommended NTP configuration settings have been moved to section [14.3](#). Follow the [options timed.*](#) hyperlink to jump straight to the page.

11.2.11.2 TimeZones

CATE Storage and Backup Engineering recommend that the local time zone be used for the geographic location at which the Filer is installed. For example, Filers in California should be configured PST8PDT. If the time zone in which the Filer is located supports Daylight Savings Time (DST), then the Filer should be configured with the time zone information that supports DST. Contact gdse@citi.com if you are unsure which time zone value to use for your Filer.

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

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

11.2.12 Windows Internet Name Service (WINS)

The release date of ONTAP 8.1.4 in Citi aligns with the date at which a collaborative effort to retire WINS is initiated between a number of different CATE and CTI teams. Windows client and server builds have been updated to remove dependencies on WINS for name resolution and to have an appropriate configured DNS suffix search order. Historically WINS server configuration has always been a stage in the CIFS server setup process which has now been updated with the Citi ONTAP 8.1 build to no longer require this step. However tWINS is required to be unconfigured in the existing environment. The purpose of this section is to describe the best practices with regards to disabling WINS for name resolution.

One of the challenges with WINS is that there is little to no way of determining which client systems are using it from an infrastructure perspective. As name resolution occurs prior to clients connecting to the storage system there is no way to tell from a storage system perspective what name resolution mechanism was used. The WINS server itself has very little logging capability outside of the database management and health check features. As such, removing WINS dependencies from a storage system is best performed in a controlled and measured approach using phases to implement the WINS configuration removal process.

- Disable WINS on a per vFiler / CIFS server level
- Disable WINS at the physical Filer / CIFS server level
- Disable WINS on all network interfaces on the storage controller

11.2.12.1 How to disable WINS at the vFiler / physical Filer level

Disabling WINS on a per vFiler or physical Filer (vFiler0) level is achieved by configuring the WINS servers list with a null string:

```
ontap> options cifs.wins_servers ""
```

This change takes effect immediately and is persistent across reboots.

If WINS disablement is planned to be implemented outside of an ONTAP upgrade then the CIFS daemon needs to be stopped and restarted for the change to take effect. Note this change will cause an interruption to service for Windows users accessing network shares on the vFiler where the CIFS daemon is restarted.

The following commands are used to stop and restart the CIFS service:

```
ontap> cifs terminate
ontap> cifs restart
```

11.2.12.2 How to disable WINS at the storage controller network interface level

Disabling WINS at the network interface level will prevent the storage system from performing NetBIOS broadcasts on the subnet. This is achieved by disabling WINS on each configured interface group or per network interface if not participating in an interface group.

1. Use ifgrp status to determine which interface groups are configured

```
ontap> ifgrp status
default: transmit 'IP Load balancing', Ifgrp Type 'multi_mode', fail 'log'
gev1: 1 link, transmit 'none', Ifgrp Type 'single_mode' fail 'default'
```

2. Disable WINS on each interface group of the running configuration and check the status

```
ontap> ifconfig gev1 -wins
ontap> ifconfig gev1
gev1: flags=0x26f4c863<UP, BROADCAST, RUNNING, MULTICAST, TCPCKSUM, NOWINS> mtu
1500
```

3. Disable WINS on each interface group in the system startup configuration. This step requires editing the /etc/rc file. Ensure best practices are followed to make a backup copy of the file before editing. For each 'ifconfig' statement of an interface group in the /etc/rc file add the '-wins' parameter to the 'ifconfig' statement. If the 'ifconfig' statement also specifies the 'partner' parameter ensure the '-wins' is entered prior to the 'partner' parameter.
4. Disable WINS on non-data interfaces e0M and e0P in the running configuration. Note older FAS platforms may not be equipped with an e0P interface.

```
ontap> ifconfig e0M -wins
ontap> ifconfig e0P -wins
ontap> ifconfig e0M
e0M:
flags=0x6b4c867<UP, BROADCAST, RUNNING, MULTICAST, TCPCKSUM, NOWINS, MGMT_PORT>
mtu 1500
ontap> ifconfig e0P
e0P:
flags=0x6b4c867<UP, BROADCAST, RUNNING, MULTICAST, TCPCKSUM, NOWINS, ACP_PORT>
mtu 1500 PRIVATE
```

5. Disable WINS on non-data interfaces e0M and e0P in system startup configuration. Note older FAS platforms may not be equipped with an e0P interface. This step requires editing the /etc/rc file. Ensure best practices are followed to make a backup copy of the file before editing. For the 'ifconfig' statement of e0M and e0P in the /etc/rc file add the '-wins' parameter to end of the 'ifconfig' statement.

11.2.12.3 Recommended approach to prioritizing systems for WINS retirement

The following list is offered as a suggestion as to how to prioritize which systems to retire WINS from in your environment starting with least impactful.

1. NFS only systems
2. Archive systems
3. DR systems
4. Production systems

11.2.13 Appropriate Byte Counting

ONTAP 8.1 introduces a networking feature called Appropriate Byte Counting (ABC) which is an algorithm implemented in the TCP stack of Data ONTAP that enables you to improve TCP performance and security. The effects on the overall workload stack of this new setting are undetermined at this time and require further testing. As such CATE Storage and Backup Engineering in conjunction with CATE Networks are recommending this feature be disabled upon upgrading the storage system to ONTAP 8.1.4. Refer to section [14.3](#) for the recommended setting for configuration option `ip.tcp.abc.enable`.

11.3 Storage

11.3.1 Data ONTAP Clustering (HA-Pairs)

A cluster consists of a pair of NetApp filers that are connected and configured 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 to use each other's NVRAM (Operational Memory) or NV-MEM (where applicable), 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
```

11.3.2 Disk Shelves

NetApp provides disk capacity in units known as a Disk Shelf. There are three models: the legacy fiber channel DS14 3U 14 drive FCAL enclosure, a 4U chassis known as a DS4243 and a 2U chassis known as a DS2246; both of which are SAS shelf types containing 24 SAS drives⁸. Consult the respective [DS4243](#) and [DS2246](#) release notes or, where applicable, the appropriate CATE design document for information on when and where to use each shelf type. This section presents the minimum and recommended certified versions of shelf firmware and presents the best practices on cabling and attaching disk shelves to the storage controller (Filer head). For further information storage administrators are strongly advised to review the [DS4243](#) and [DS2246](#) release notes. Mixing of the FCAL and SAS shelf types on the same system is not recommended.

Note: As of ONTAP 8.1 ESH2 FCAL controllers are no longer supported. Storage systems containing ESH2 controllers cannot be upgraded to ONTAP 8.1 and must be retired as part of a tech refresh / data migration activity.

⁸ Also available in half populated 12 drive configurations. Variances from full or half populated configurations may be available on request but are only available if the NetApp sales team is permitted to order such a configuration from manufacturing.

11.3.2.1 Disk Shelf Firmware

Shelf Type	Module type	Minimum required version	Latest recommended version
DS14	AT-FCX	38.00	38.00
DS14	ESH4	14.00	14.00
DS2246	ACP	02.06	02.08
DS2246	IOM6	0151	0162
DS4243	ACP	01.31	01.41
DS4243	IOM3	0160	0172
DS4246	ACP	02.06	02.08
DS4246	IOM6	0151	0162

Table 7 - Disk shelf firmware supported versions

Firmware updates to ESH as well as IOM3 and IOM6 modules and their associated ACP modules are non-disruptive (i.e they do not interrupt client access to data during the upgrade) where multipath HA is deployed. For single path HA configurations such as the FAS2040, a cluster failover is required to upgrade the module firmware.

On systems running a version of ONTAP prior to 8.0.4, in order to avoid bug 526941, ensure the ACP firmware checksum files are not copied to the `/etc/acpp_fw` directory when updating the ACP firmware. The checksum files will match the following pattern: `/etc/acpp_fw/*.FVF`. If they are present on the system delete them from the `/etc/acpp_fw` directory before upgrading to a version of ONTAP later than 8.0.4 in order to avoid a resource starvation condition that may panic the storage system. Once ONTAP is upgraded these firmware checksum files can be copied to the system again.

Firmware updates to AT-FCX modules have historically been disruptive as client IO traditionally gets disrupted for about 2 minutes. Starting with AT-FCX firmware release 37.00 upgrades to later versions of firmware are no-disruptive only if each of the following conditions are met:

- ONTAP 7.3.1.1P6 or above is installed on the Filer
- The connectivity between the Filer and the disk shelf is a dual or multipath configuration
- The current version of AT-FCX firmware installed on the system is 37.00 or above

11.3.2.2 Multipath HA

Multipath HA Storage provides an additional path from the storage controller (Filer) to the disk shelves for active/active configurations for improved resiliency and performance. The following diagram demonstrates the high level component cable interconnect configuration.




Figure 5 - Multipath HA configuration

11.3.2.2.1 Benefits of Multipath HA

While Multipath HA has been the standard on the NetApp hardware platforms that support it since about 2007 it is important to remember the benefits derived from this feature. Multipath HA reduces the need for takeover by the partner Filer by being able to:

- Sustain a FC Loop or SAS path failure
- Storage HBA failure provided the SAS paths or FC Loops are on different HBAs
- Shelf Controller/Module failure
- Dual Inter-Shelf Cable failures
- Provides performance enhancement:
 - Reducing the need for takeover minimizes inconsistent performance due to an over-burdened partner node

In addition to improved resiliency multipath-HA allows the storage layer to attempt load balancing in order to even out the I/O traffic across the SAS adapters. Each disk is assigned a primary path through a single SAS adapter and all I/O traffic for that disk uses the primary path. Over time the storage layer measures the I/O traffic on all the primary paths and it uses that information to periodically rebalance the load. It rebalances the load by changing the disk primary paths to use different adapters. The goal is to get a similar amount of I/O load on the different SAS adapters that can communicate with the same disks. The periodic rebalancing happens after a certain amount of I/O traffic has been transferred on the system. It rebalances at these points 100GB, 200GB, 400GB, 800GB and so on up to 10TB where it then reaches a steady state and rebalances after every 10TB of data. The storage layer may rebalance around other events as well such as disk insertions or removals. Due to the nature of disk reservations the storage layer cannot dynamically load balance on a per I/O basis and it also cannot switch primary paths very often.

11.3.2.2.2 When to use multipath-HA

Multipath-HA is the CATE standard for all systems to reduce the potential for cluster failover events, except for the FAS2040 with the DS2246 or DS4243 enclosures which only supports single path HA due to limitation in the number of SAS ports on the processor controller modules⁹. Multipath-HA is also recommended on systems with multistore (vFilers) enabled as the cluster failover time is longer on systems hosting more than one vFiler.

11.3.2.2.3 Pre-requisites to installing multipath-HA hardware

Software Disk Ownership (SANOWN) must be enabled prior to installing the additional hardware. The procedure for enabling SANOWN is outlined in the NetApp tech article [“How does disk ownership of NetApp Shared Storage work and what are the requirements”](#).

11.3.2.2.4 How to confirm a multipath configuration

Once the multipath option has been configured a secondary path to the disk is available, as shown in the example below.

```
Filer> storage show disk -p
```

PRIMARY	PORT	SECONDARY	PORT	SHELF	BAY
1c.50.0	A	1b.50.0	B	50	0
1c.50.1	A	1b.50.1	B	50	1
1c.50.2	A	1b.50.2	B	50	2
1c.50.3	A	1b.50.3	B	50	3
1c.50.4	A	1b.50.4	B	50	4
1c.50.5	A	1b.50.5	B	50	5
1c.50.6	A	1b.50.6	B	50	6
1c.50.7	A	1b.50.7	B	50	7

11.3.2.3 SAS Port / Stack speeds – Mixing SAS shelves in the same SAS stack

Unlike the old ESH modules the speeds on the SAS expanders are not set with a mechanical switch and so they only way for a SAS stack to suffer a degradation in the performance of each SAS LANE is to mix a 3GB SAS shelf (DS4243) with a 6GB SAS shelf (DS2246) in the same stack. To this end neither NetApp nor CATE recommend mixing SAS shelf types in the same stack.

⁹ Considering the requirements of the entry level FAS2040 model and the improved reliability offered by the SAS architecture of the DS4243, the reduction in path resiliency was a reasonable tradeoff with the improved disk subsystem stability provide through SAS.

11.3.2.4 FC and SAS adapter PCI slot assignments

Due to the increased number of different hardware configurations with the introduction of the FAS3200 and FAS6200 series, detailed reference materials on SAS adapter PCI slot assignments, SAS stack counts and SAS HBA quantity requirements have been relocated to the release notes for each hardware platform. Use the links below to access the relevant documentation:

- [FAS2040](#)
- [FAS3200 series](#)
- [FAS6200 series](#)

The following table lists the recommended slot assignments for legacy FC HBAs based on hardware platform and the number of FC loops. This table is maintained here for reference purposes only.

Hardware Platform	Shelf Type	Number of loops (per controller)	Port / Slot Assignments			
			Primary		Redundant	
			Local	Partner	Local	Partner
FAS3040	DS14	1	0a	0b	0c	0d
		2	1a	1b	1c	1d
FAS3070	DS14	1	0a	0b	0c	0d
		2	1a	1b	1c	1d
FAS3140	DS14	See section 7.1.2 of the FAS3100 Release Notes				
FAS3160	DS14	See section 7.1.2 of the FAS3100 Release Notes				
FAS3170	DS14	See section 7.1.2 of the FAS3100 Release Notes				
FAS6030	DS14	1	0a	0b	0c	0d
		2	5a	5b	5c	5d
FAS6070	DS14	1	0a	0b	0c	0d
		2	5a	5b	5c	5d

} use all 8 ports for 2 loops

} use all 8 ports for 2 loops

} use all 8 ports for 2 loops

} use all 8 ports for 2 loops

Table 8 – FCAL HBA slot and port assignments for multipath-HA

11.3.2.5 Port Load Balancing

The storage layer load balancing attempts to even out the I/O traffic across the FC or SAS adapters. Each disk is assigned a primary path through a single adapter and all I/O traffic for that disk uses the primary path. Over time the storage layer measures the I/O traffic on all the primary paths and it uses that information to periodically rebalance the load. It rebalances the load by changing the disk primary paths to use different adapters. The goal is to get a similar amount of I/O load on the different adapters that can communicate with the same disks. The periodic rebalancing happens after a certain amount of I/O traffic has been transferred on the system. It rebalances at these points 100GB, 200GB, 400GB, 800GB and so on up to 10TB where it then reaches a steady state and rebalances after every 10TB of data. The storage layer may rebalance around other events as well such as disk insertions or removals. Due to the nature of disk reservations the storage layer cannot dynamically load balance on a per I/O basis and it also cannot switch primary paths very often.

11.3.2.5.1 Monitoring and Alerting

Unified Manager 5.x does not have a mechanism to capture whether Multipath-HA is installed in the NAS environment.

The current SNMP MIB for Tivoli and Netcool includes rules which process the traps generated by disk multipath events (available in ONTAP 7.2.4P6 or above). Refer to [SNMP MIB section](#) for details.

In addition, review the Filer message logs for the following entry that would indicate a degraded state in the multipath-HA configuration:

```
Tue Aug 21 15:31:46 EDT [filer1: config.NotMultiPath:warning]: Disk 0c.29 and other disks on this loop are not multipathed and should be for improved availability.
```

Finally, running the 'storage show disk -p' command will demonstrate that the secondary link is unavailable (by default the remaining link becomes the primary link):

```
filer1> storage show disk -p
PRIMARY PORT SECONDARY PORT SHELF BAY
-----
0c.16 B 1 0
0c.17 B 1 1
0c.18 B 1 2
0c.19 B 1 3
0c.20 B 1 4
...
...
...
0d.16 A 0b.16 B 1 0
0d.17 A 0b.17 B 1 1
0d.18 A 0b.18 B 1 2
0b.19 B 0d.19 A 1 3
0d.20 A 0b.20 B 1 4
...
...
...
```

11.3.2.6 Port / Loop speeds

ONTAP 7.2 introduced support for DS14MK4 shelves and 4GB line speed FC drives. All MK4 shelves are shipped set to 2GB speed. To remove the requirement for production down-time ensure all 4GB enabled shelves are set to 4GB before your storage system is commissioned for production use. Use the `sysconfig -a` command to check the link data rate of the disk loop. If a DS14MK4 shelf set to 2GB is introduced into a DS14MK4 loop where all the other shelves are set to 4GB, the loop as a whole will throttle down to 2GB. ONTAP 7.3.2 and above introduced support for the DS4243 which utilizes 3GB SAS lanes. Unlike the ESH modules the speeds on the SAS expanders are not set with a mechanical switch.

11.3.3 Disks

CATE supports the following drive types within NetApp disk shelves. For information regarding the procedure to erase (sanitize) disks within a DS14, DS2246, DS4243 or DS4246 shelf refer to section [12.5](#).

Shelf Model	Topology	Size	Speed	Minimum ONTAP (CATE)	Use Case	CATE Support
DS2246	SAS	600gb	10K	7.3.5.1	Home / Groups dirs VMware VDI	Current standard / preferred
DS4243	SAS	600gb	15K	7.3.3	VMware VSI Home / Groups dirs	Current standard / preferred Legacy
DS4243	SAS	450gb	15K	7.3.3	VMware VSI Home / Groups dirs	Current standard / preferred for FAS3210 Legacy
DS4243	SAS	300gb	15K	7.3.3	CTI Grid	Current standard / preferred
DS4243	BSAS (SATA)	3tb	7.2K	8.1.4	Archive	Current standard / preferred
DS4246	BSAS (SATA)	3tb	7.2K	8.1.4	Archive	Current standard / preferred
DS14	FC	600gb	15K	7.3.2	Home / Groups dirs	Legacy
DS14	FC	450gb	15K	7.2.5.1	Home / Groups dirs	Legacy
DS2246	SAS	450gb	10K	7.3.5.1	Home / Groups dirs	Legacy
DS14	(S)ATA	1tb	7.2K	7.2.4P6	Archive	Legacy
DS4243	BSAS (SATA)	2tb ^{10 11}	7.2K	8.0.1	Archive	Legacy
DS4243	BSAS (SATA)	1tb	7.2K	7.3.3	Archive	Legacy
DS4246	BSAS (SATA)	2tb ^{11 12 13}	7.2K	8.1.4	Archive	Legacy
DS4246	BSAS (SATA)	1tb ¹²	7.2K	8.1.4 ¹³	Archive	Legacy

Table 9 - Supported disk drive types

The minimum recommended disk drive firmware bundle is dated 2014-04-25.
The associated disk qualification package is datecode 20130711.

¹⁰ 2TB drives are recommended for use as part of a 64-bit aggregate as 32-bit aggregates would lead to artificially small RAID groups due to the 16TB size limitation

¹¹ Refer to the CATE Release Notes for known issues with drive size selection when creating aggregates with 2TB drives

¹² 1TB and 2TB DS4246 configurations supported for DS4243 systems requiring expansion after the DS4243 has reach EOA

¹³ DS4246 is supported from ONTAP 8.1.4 and above although 1TB drives in other enclosures were supported prior to this ONTAP release

11.3.3.1 Maximum drive counts per platform

The following table lists the supported physical drive count limitation per hardware platform in a clustered or standalone configuration. The third column shows the NetApp supported drive count and the fourth column lists CATE Storage and Backup Engineering's recommended maximum number of drives¹⁴. Consult the hardware platform release notes for further details.

Hardware Platform	Shelf Type ¹⁵	Max # of disk drives (multipath-HA support) - NetApp	Max # of disk drives (multipath-HA support) - Citi
FAS2040 ¹⁶	DS424x or DS2246	108	84
FAS3040	DS14	336	252
FAS3070	DS14	504	378
FAS3140	DS14	420	308
FAS3140	DS424x or DS2246	420	312
FAS3160	DS14	672	504
FAS3160	DS424x or DS2246	672	504
FAS3170	DS14	840	630
FAS3170	DS424x or DS2246	840	624
FAS3210	DS424x or DS2246	240	240
FAS3220	DS424x or DS2246	480	480
FAS3240	DS424x or DS2246	600	600
FAS3250	DS424x or DS2246	720	720
FAS3270	DS424x or DS2246	960	960
FAS6030	DS14	252	182
FAS6040	DS14	840	630
FAS6040	DS424x or DS2246	840	624
FAS6070	DS14	302	224
FAS6080	DS14	840	630
FAS6080	DS424x or DS2246	1176	864
FAS6210	DS424x or DS2246	1200	1200
FAS6220	DS424x or DS2246	1200	1200
FAS6240	DS424x or DS2246	1440	1440
FAS6250	DS424x or DS2246	1440	1440
FAS6280	DS424x or DS2246	1400	1400
FAS6290	DS424x or DS2246	1440	1440

Table 10 - Maximum number of disk drives supported in a HA or standalone configuration

¹⁴ Starting with the FAS3200 and FAS6200 series CATE has aligned its drive count support policy with NetApp's

¹⁵ DS424x means both DS4243 and DS4246 disk enclosure models

¹⁶ FAS2040 is an exception to this table as it is a single path HA configuration. Figures quoted include the 12 internal SAS drives in the FAS2040 chassis.

11.3.3.2 Disk performance characters for maintenance activities

In addition to the publicized I/O metrics for each drive type the following performance characteristics are collected when new drive types and sizes are certified: rebuild time, sanitization time and zeroing time. The table below summarizes these attributes for the most recently certified drive types were collected using a test environment under light load. For worse case scenarios consider doubling the figures listed below.

Topology	Size	Speed	Rebuild Time (Full)	Rebuild Time (Rapid)	Zeroing Time	Sanitization Time (1 cycle)	Sanitization Time (7 cycle) ¹⁷
FC and SAS	300gb	15K	2.0 hrs	1.0 hrs	1.5 hr	5.75 hrs	-
SAS	450gb	10k	3.8 hrs	2.5 hrs	2.3 hrs	7.75 hrs	-
FC and SAS	450gb	15K	2.3 hrs	1.5 hrs	2.2 hrs	7 hrs	-
SAS	600gb	10k	4.4 hrs	3.8 hrs	2.6 hrs	9.5 hrs	-
FC and SAS	600gb	15K	2.8 hrs	1.8 hrs	2.5 hrs	8.5 hrs	21.5 hrs
(S)ATA	1tb	7.2K	9.3 hrs	5.5 hrs	4.3 hrs	14 hrs	-
(S)ATA / BSAS	1tb	7.2K	9.3 hrs	5.5 hrs	4.3 hrs	11 hrs	-
(S)ATA / BSAS	2tb	7.2K	18.5 hrs	12.8 hrs	5.6 hrs	17 hrs	121.25 hrs
(S)ATA / BSAS	3tb	7.2K	27.8 hrs	18.3 hrs	7.0 hrs	-	160.1 hrs

Table 11 - Drive re-build, sanitization and zeroing times

11.3.3.3 Spare Disks

Spare disks are global. Spare disks can be used for both the creation of new volumes (sometimes referred to as “unassigned spares”) as well as replacement of failed drives (sometimes referred to as “dedicated global spares”). As of ONTAP 7.3 and above CATE Storage and Backup Engineering is adopting the spare disk best practices as outlined in [NetApps 3437 Technical Report](#) document. The current best practices recommend the same number of spare disks to be used whether your system contains SAS or SATA drives and are listed in the table below. Exceptions to these rules are listed below the table.

Number of DS2246 or DS4243 Shelves	Number of Disks	Recommended number of spares
1	28	2
4	112	3
7	168	3
14	336	4
21	504	6
42	1008	12

Table 12 - System spare disks: How to calculate number of drives

¹⁷ NetApp provides sanitization times based on a single cycle configuration. Citi’s secure media disposal guidelines require a seven cycle configuration when sanitizing disks. This column has been added to provide CTI NAS operations a more realistic time line for the sanitization process. ‘-’ indicates no test result data is currently available for this drive type

The above approach ensures the minimum of two spares disks are available in order to support the functionality of the disk maintenance center (see section [11.3.3.4](#)). The one exception to this rule is single shelf systems such as the FAS2040 model:

- Single shelf systems will only have 1 global spare per system. For clustered single shelf systems when the disk shelf is split into A and B sides to provide a clustered infrastructure, one spare disk must be allocated to each side of the cluster.

Additional behavioral aspects of spare disks and their associated best practices are listed in the bullet points below:

- 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.
- For systems that contain differ sizes or types of disks ensure that the above best practice is applied pro rota across the storage system. For example if a SATA system was 40% 500GB drives and 60% 1TB drives, apportion the number of spare disks for each size of disk accordingly so that 40% of the spares pool would be 500GB drives and 60% would be 1TB SATA. A similar principle would apply is SAS and SATA were mixed in the same system.
- 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, CATE Storage Engineering 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.

11.3.3.4 Maintenance Center

The purpose of Maintenance Center is to improve storage reliability by reducing the number of unnecessary disk returns to NetApp due to transient errors. As discussed in section [11.3.3.3](#), ensure there are a minimum of two spares disks on the system before enabling maintenance center.

Use the following process to enable Maintenance Center and refer to Appendix C: Data ONTAP SOE Configuration Standards for details:

```
options disk.maint_center.enable on
```


11.3.4 Aggregates

11.3.4.1 An introduction to aggregates and flexible volumes

With ONTAP 7 NetApp introduced the flexible volume (flex-vol) and aggregate technology. Compared to the original NetApp volume type, referred to as a traditional volume (trad-vol), which were a group of disks brought together as a volume and assigned a RAID level (a direct relationship between a disk and the volume in which it was a member); aggregates and flex-vols provided a layer of abstraction 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. These technologies continue to play a significant role in the underlying architecture of ONTAP 8, especially with the introduction of 64-bit aggregates (see section [11.3.4.6.1](#)).

11.3.4.1.1 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 8.0, 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.

11.3.4.2 Support for Traditional Volumes

CATE will not support Traditional Volumes on a system running ONTAP 8 or above.

Where CTI Global NAS operations require refresh of an older system using traditional volumes to a system running ONTAP 8, a CATE design review and approval is required.

11.3.4.3 Considerations when 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 FAS2000 series that only has a small number of drives available. Therefore the FAS2000 series 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 (thick provisioned).
- 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. CATE Storage and Backup Engineering recommends setting the root volume size at double 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 for ONTAP 8	CATE Recommended Root Volume Size for ONTAP 8
FAS2040	132GB	264GB
FAS3140	134GB	268GB
FAS3160	232GB	464GB
FAS3170	250GB	500GB
FAS3210	145GB	290GB
FAS3220	250GB	500GB
FAS3240	195GB	390GB
FAS3250	250GB	500GB
FAS3270	250GB	500GB
FAS6030	250GB	500GB
FAS6040	250GB	500GB
FAS6070	250GB	500GB
FAS6080	250GB	500GB
FAS6210	250GB	500GB
FAS6220	250GB	500GB
FAS6240	250GB	500GB
FAS6250	250GB	500GB
FAS6280	250GB	500GB
FAS6290	250GB	500GB

Table 13 - Recommended root volume size

- The recommended volume sizes listed above include the snapshot reserve for the root volume such that a 500GB root volume would have 475GB useable capacity with a snap reserve of 5%.
- As of ONTAP 8.1 the default snap reserve for the root volume is 5%. The default snap reserve size must be used.
- 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.

11.3.4.4 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. 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.

11.3.4.5 Aggregate Sizing Standards

For implementations where the root volume is to be housed on a dedicated aggregate, then the aggregate should comprise of the appropriate number of disks in a RAID-DP configuration. See the notes in section [11.3.4.3](#) with regards to the required 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 (SAS or SATA) you are using.

Drive Type	Filer Model / Series	Recommended RAID Type	Max # of disks per RAID Group
SAS	FAS2040	See FAS2000 Series Release Notes	
	FAS3200	RAID-DP	18 (16+2)
	FAS6200	RAID-DP	18(16+2)
(S)ATA	FAS2040	See FAS2000 Series Release Notes	
	FAS3200	RAID-DP	16 (14+2)
	FAS6200	RAID-DP	16 (14+2)

Table 14 - Recommended RAID group configuration

For VMware on NFS NetApp NAS storage aggregate configuration and sizing information consult the on [CATE VMware VSI Storage Design document](#) for VSI and [CATE VMware VDI 1.5 Storage Design document](#).

To prevent unnecessary fragmentation of the file systems and possible system performance degradation, NetApp does not recommend allocating aggregates at over 90% utilization rates.

11.3.4.5.1 Special considerations for RAID configuration when using large SATA drives

When creating 32-bit aggregates up to their maximum capacity of 16TB, there are some special considerations for a balanced RAID group configuration when using large capacity SATA drives which are shown in the following table.

Drive Capacity	Aggregate Type	RAID group 1	RAID Group 2	RAID Group 3 *
1 TB SATA	32-bit	10+2	9+2	n/a
2 TB SATA	32-bit	9+2	n/a	n/a
1 TB SATA	64-bit	14+2	14+2	14+2
2 TB SATA	64-bit	14+2	14+2	14+2

Table 15 - RAID group sizes for large capacity SATA drives

- Due to the inefficient use of disks creating a 32-bit aggregate with 2TB SATA drives is not recommended however it is acknowledged that there are volume snapmirror limitations between a 32-bit and 64-bit aggregate (see [CATE release notes](#) for ONTAP 8.0 and section [11.5.2](#) for further details) and so 32-bit aggregates comprised of 2TB SATA drives are permitted on an exception basis in this instance. With the arrival of ONTAP 8.1 it is possible to replicate between volumes on a 32-bit and 64-bit aggregate.
- As indicated by the * symbol some hardware platforms may not be able to support a third RAID group as the maximum supported size of a 64-bit aggregate varies by hardware platform due to memory constraints..

11.3.4.5.2 How aggregate capacity is calculated

ONTAP 8 leverages the changes that were introduced in ONTAP 7.3 regarding how aggregate capacity is calculated. The first is that parity disks no longer count towards the cost of the space in the aggregate. The second change is that the 'write size' of the disk will be used when calculating supported aggregate capacity and not the raw size (or marketing size) of the disk.

11.3.4.5.3 Aggregate Snapshots / Snap Reserve

Starting with ONTAP 8.1 an aggregate retains 0% of available space as an aggregate level snapshot reserve (prior to ONTAP 8.1 the default value was 5%). The CATE Storage and Backup Engineering standard has always been a 0% aggregate snapshot reserve.

For legacy systems which may be running with an aggregate snap reserve greater than zero, the aggregate snap reserve can be removed by issuing the following commands.

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 16 - Procedure to remove aggregate snapshot reserve

11.3.4.6 Aggregate Creation

Aggregates are created either using the Systems Manager or from the CLI using the `aggr create` command. Refer to the man pages for complete details of this command. By default with ONTAP 8.0.x all aggregates are created as 32-bit aggregates. Starting in ONTAP 8.1 all aggregates are created as 64-bit by default.

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 the best approach for management of disk errors. As such select a disk count of a certain capacity rather than selecting individual drives when creating an aggregate.

A storage controller can support up to 100 aggregates irrespective of whether they are 32-bit or 64-bit aggregates.

11.3.4.6.1 64-bit Aggregates

ONTAP 8.0 and above introduces support for 64-bit aggregates. In ONTAP 8.0 all aggregates are created as 32-bit by default unless the `-B 64` parameter is used. In ONTAP 8.1 all aggregates are created as 64-bit by default. There are special considerations for replication and data migration when using 64-bit aggregates (see both the [ONTAP 8.0 CATE release notes](#) and [ONTAP 8.1 CATE release notes](#) for full details). These should be reviewed and considered before deploying 64-bit aggregates with ONTAP 8 in your environment. Unless 64-bit aggregates are required for large pools of SATA storage with ONTAP 8.0.x the CATE recommendation is to withhold the wide scale adoption of 64-bit aggregates until the in place aggregate conversion feature is available as part of ONTAP 8.1. With the arrival of ONTAP 8.1 the in place conversion process is available as part of a pilot program. Those interested in a pilot should contact gdse@citi.com.

Unlike 32-bit aggregates which scale up to 16TB, the maximum size of a 64-bit aggregate varies by hardware platform due to the relationship between inode tables and the memory capacity on a storage system. The following table is designed to assist in determining the appropriate hardware platform for a given contiguous aggregate capacity requirement.

Filer Model	Maximum aggregate size (TB)
FAS2040	50
FAS3140	75
FAS3160	90
FAS3170	105
FAS3210	75
FAS3220	90
FAS3240	90
FAS3250	105
FAS3270	105
FAS6040	105
FAS6080	162
FAS6210	162
FAS6220	162
FAS6240	162
FAS6250	162
FAS6280	162
FAS6290	162

Table 17 – Maximum aggregate size per hardware platform with ONTAP 8.1

There is no method to convert a 32-bit aggregate to a 64-bit aggregate with ONTAP 8.0.x. With the arrival of ONTAP 8.1 the in place conversion process is available as part of a pilot program. Those interested in a pilot should contact gdse@citi.com.

11.3.4.7 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.

11.3.4.8 Aggregate Utilization Alerting

Starting with Data ONTAP 8.0.1P5 NetApp have introduced a hidden registry key that enables the ability to generate SNMP alerts on based on aggregate utilization levels. This feature is not enabled by default in Data ONTAP. For details on how to configure this feature as well as setting custom aggregate utilization thresholds see section [11.8.2.1](#).

11.3.5 Volumes

11.3.5.1 Naming convention

11.3.5.1.1 Primary Filer

Since multiple departments will be sharing filers, CATE Storage and Backup Engineering 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 in the name (like apps or data etc.). We have also taken into consideration the fact that one department could have multiple volumes on the same filer. From the wider reporting perspective this will allow all volumes to be unique names in the environment, assisting in simplifying the reports generated by a larger environment.

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

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

11.3.5.1.2 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.

11.3.5.2 Volume Creation

A volume is a logical file system on a Filer striped across all the disks in the aggregate. Volumes are created either from Systems Manager or by the `vol add` command on the CLI. Refer to the man pages for complete details of this command.

- ONTAP will support up to 500 flex-vols on a Filer (1000 in a cluster) on most Filer models with the exception of the FAS2040 and FAS3210 which only support up to 200. As of ONTAP 8 and above the CATE support model aligns with NetApp's figures. During a cluster failover situation a single node in a cluster can support up to 1000 flex-vols. The basic rule of cluster failover performance is the fewer volumes you have in a cluster the faster the cluster failover event will be. Refer to section [11.5.3](#) for the impact of NDMP throughput on scalability calculations.
- 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.
- 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
```

11.3.5.3 Volume Sizing

NetApp storage controllers today generally employ 450GB and 600GB SAS disk drives, as well as 1TB, 2TB and 3TB SATA drives, all certified for use.

For CIFS and NFS file system protocols the maximum volume size is limited to the maximum file system size the OS can write. In both cases the maximum file system size supported by the different CATE certified platforms is much larger than the standard volume size used in Citi today.

When creating a new volume, one must take into account the requirements for backup (see section [11.5.3.2](#)) and replication (see section [11.5.2.3](#)) in determining the optimal size for a volume.

11.3.5.4 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.

11.3.5.5 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. In order to avoid known issues Storage administrators are advised not to create cloned volumes which are larger than the parent volumes (see bug [628774](#) for details).

11.3.5.6 Volume Utilization Alerting

Starting with Data ONTAP 8.0.1P5 NetApp have introduced a hidden registry key that corrects for a known condition where SNMP traps are truncated when multiple volumes concurrently experience a volume utilization threshold breach. This new setting is not enabled by default and needs to be configured. For details on how to configure this feature as well as defining volume utilization thresholds see section [11.8.2.1](#).

11.3.5.7 Volume Maximum Directory Size

The main purpose of the maxdirsize limit is performance. Scans through directories are linear. On average, half of the directory size will need to be scanned for each lookup, which can be painfully slow. WAFL grows directory data structures as used.

This option is useful for environments where system users may grow a directory to a size that starts impacting system performance. When a user tries to create a file in a directory that is at the limit, the system returns a ENOSPC error and fails the create. NetApp recommends not to set this option too large. By default with ONTAP 6.5 and later the default maxdirsize per volume is equivalent to 1% of the total system memory in KB.

To prevent performance issues in the directory it is not recommended setting maxdirsize to a value that is greater than 2% of the total system memory in KB.

Maximum directory size is configured on a per volume basis using the command vol options <vol_name> maxdirsize <value>. See NetApp KnowledgeBase article [3011562](#) for further details.

11.4 File Service Protocols – CIFS & NFS

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 (hundreds 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.

11.4.1 CIFS

ONTAP has the flexibility to meet the requirements of the Citi Windows Security Model and as such implements the configuration to restrict the browsing of CIFS shares by anonymous (unauthenticated) users as well as the signing of SMB traffic. Refer to section [14.3](#) for the [cifs.restrict_anonymous](#) option and [cifs.signing.enable](#) / [cifs.smb2.signing.required](#) for SMB v1 and v2 signing respectively.

11.4.1.1 Windows Interoperability

ONTAP is designed to interact with the Windows platform in order to provide access to files for Windows clients and also participate in Active Directory to authenticate users to their data. In order to ensure the availability of data to Windows clients storage administrators should refer to the information in this section with regards to minimum ONTAP versions required to support the latest Windows platforms and features.

Due to known issues Storage Administrators are required to disable SMB 2.1 upon upgrading the storage controller to ONTAP 8.1.4P1. See the Known Issues section of the [ONTAP 8.1 CATE release notes](#) for details on bug 798842 and refer to section [14.3](#) for the procedure to disable SMB 2.1 manually. Note that with the arrival of ONTAP 8.1.4P4 and subsequent ONTAP 8.1.4 patch releases, SMB 2.1 is disabled by default and there is no need to implement the workaround documented in section [14.3](#).

For NetApp systems running versions of ONTAP prior to 8.1.4 where SMB 2.0 was disabled due to bug 645251, the upgrade to ONTAP 8.1.4P1 will not completely address the performance issue associated with SMB2 signing due to a subsequent performance related bug 817844 introduced as a side effect of the fix for bug 645251. The current recommendation is where systems running older versions of ONTAP have SMB2 already disabled; SMB2 should remain disabled upon upgrade to 8.1.4P1 or later. Also the upgrade process to 8.1.4P1, P4, P5 or P6 may inadvertently re-enable SMB2 (see bug 649716). Therefore the value of option `cifs.smb2.enable` should be checked on each vFiler post upgrade.

Subsequently with the release of ONTAP 8.1.4P4 and above the performance related bug 817844 is addressed. However it is recommended to keep SMB2 disabled upon upgrade to 8.1.4P4, P5 or P6 in order to prevent unexpected premature exhaustion of storage system resources, known to occur in Citi's virtualized desktop environment. Any systems which were upgraded to 8.1.4P1 prior to the release of 8.1.4P4 (or later) must upgrade to the P4 or most current patch release available from CATE.

Upon upgrading to ONTAP 8.1.4 or later Storage Administrators are reminded to enable the SID compression support for the Active Directory computer account for each vFiler. Consult the [Data ONTAP Active Directory SID Compression Interoperability Tech Note](#) for the procedure on how to enable SID compression.

11.4.1.1.1 Domain Controllers (Infrastructure)

In order to participate in an Active Directory domain where all the domain controllers are running the same version of Windows, the following ONTAP minimum versions are required:

Windows Domain Controller Version	Minimum ONTAP Version	Current Recommended ONTAP 8 Version
2008 SP1	7.2.5.1P6	8.0.2P4
2008 SP2	7.2.5.1P6	8.0.2P4
2012	7.3.6 ¹⁸	8.1.4P6

Table 18 - Windows Domain Controller O/S Interoperability Matrix

There are other functional interoperability aspects to ensuring successful user authentication and access to data when participating in a Windows domain. These relate to security policies that are defined within the domain and protocol versions for authentication traffic. The table below shows the release of ONTAP when the feature was initially supported, the latest recommended minimum version that takes into account any bug fixes or enhancements required to support changes made by Microsoft to the feature in question, and finally the current recommended release of ONTAP.

Feature	Supported from ONTAP Version	Minimum Recommended ONTAP Version	Current Recommended ONTAP 8 Version
NTLMv2	6.5	7.3.6	8.1.4P6
LDAPS	7.1.2.1P1	7.3.6	8.1.4P6

Table 19 - Windows Authentication and Directory Service Protocol Interoperability Matrix

11.4.1.1.1.1 Important changes due to Windows 2012 Domain Controllers

With the arrival of Windows 2012 servers, the Citi Windows server build has been updated to disable NetBIOS over TCP. As a result Data ONTAP may experience issues communicating with a Windows 2012 (or later) domain controller which can culminate in the NetApp File losing its ability to communicate with the Active Directory domain. Data ONTAP does not autonegotiate between NetBIOS and SMB over TCP when authenticating with Domain Controllers and as such the following setting is required to be configured on each vFiler actively running CIFS services.

```
ontap> options cifs.netbios_over_tcp.enable off
```

Once NetBIOS over TCP is disabled in Data ONTAP the existing connections to Domain Controllers need to be reset in order to switch to using SMB. Neither disabling NetBIOS over TCP nor resetting connections to domain controllers requires CIFS to be restarted on the vFiler. These changes can be made with CIFS services remaining active and no interruption to Windows file sharing services. In order to reset the connections to domain controllers run the following command:

```
ontap> cifs resetdc
```

For environments where pre Windows 2012 Domain Controllers are still deployed, Data ONTAP will successfully communicate with these devices using SMB over TCP and as such makes this configuration change backwardly compatible.

¹⁸ Requires workaround to disable [SID compression](#)

11.4.1.1.2 Windows Server and Client platforms (CIFS Clients)

In order for various Windows client and server platforms to connect successfully to NetApp NAS storage devices the following recommended version of ONTAP should be running on the system. It is highly likely that a storage system will have a mixture of all Windows client and server platforms connecting to it at any one time and so the recommended release will always target the version of ONTAP that can support the widest diversity of Windows platforms with the least amount of interoperability issues.

With Citi's Windows Security Model requiring mandatory signing of Windows Server to Windows Server communication (where a NetApp Filer running the CIFS protocol is considered to be a Windows Server for the sake of SMB signing), SMB signing cannot be disabled independently of completely disabling a given version of the SMB protocol. Failure to comply with this best will result in a denial of service for CAP virtual desktops.

Windows Platform Version	Minimum ONTAP 8 Version	Current Recommended ONTAP 8 Version
Windows XP	8.1.4P1	8.1.4P6
Windows Vista	8.1.4P1	8.1.4P6
Windows 7	8.1.4P1	8.1.4P6
Windows Server 2003	8.1.4P1	8.1.4P6
Windows Server 2008	8.1.4P1	8.1.4P6

Table 20 - Windows Server and Client platform O/S Interoperability Matrix

11.4.1.2 Configuring CIFS

The setup of the CIFS daemon on a NetApp storage controller is one of the fundamental steps in preparing the device for use on the Citi network. It allows for the Filer to participate in Active Directory which in turn provides the authentication mechanism for users to access their data as well as administrators to login to the system for storage device management purposes.

All NetApp Filers will participate in Active Directory. Authentication via local accounts for data access (a standalone or workgroup configuration) is not supported by CATE. The goal of this section is to provide the steps necessary for a storage administrator to successfully add a NetApp Filer to an Active Directory domain using the 'cifs setup' process. As CIFS is defined at the vFiler level these steps apply to both physical systems and vFilers

11.4.1.2.1 Pre-requisites

The following list details the pre-requisite activities required to be completed before adding a NetApp Filer to Active Directory:

1. A common routable IP network connecting the Filer to DNS servers and Active Directory Domain Controllers
2. The CIFS protocol license installed on the storage controller
3. The IP address and subnet mask of the primary data interface on the storage system. Contact your Active Directory support group and ensure the subnet range in which the IP address resides is defined in Active Directory Sites and Services, so that when the Filer is added to the domain it receives configuration information that directs it to use localized domain controllers that are designed to optimize the user authentication request time.
4. The name of the Active Directory domain and OU (Organizational Unit), including the OU's canonical path name (example ou=servers,ou=storageEng), to which the Filer will be added

5. An Active Directory user account and password with privileges to create computer objects in the OU specified in the previous step and the right to add devices to the domain
6. The ONTAP setup process has been completed and DNS is configured:
 - a. DNS resolver is configured
 - b. The DNS domain name for the Filer is configured (usually the same as the Active Directory domain name)
 - c. At least two name servers defined
 Use the `dns info` command to check these values
7. The DNS servers must be able to resolve the SRV record of the Active Directory domain in order to provide the Filer with a list of Domain Controller IP addresses when the storage system is being joined to the domain.
8. Starting with ONTAP 8.1 within Citi, there is no longer a requirement to define WINS servers when performing a CIFS setup
9. NTP (network time protocol) time servers are configured on the storage controller and the system time is synchronized with the external time sources

11.4.1.2.2 The CIFS setup process

Data ONTAP provides a convenient interview process in which most of the data gathered from the previous section is entered to configure the CIFS daemon on the storage controller. The high level steps are listed below for convenience followed by an example.

1. Configure the system as a multiprotocol filer (recommended)
2. Configure the NetBIOS host name for the storage system
3. Add the Filer to Active Directory
4. Set the local Windows Administrator account password

The following example walks through the CIFS setup process and includes a common response error message when Active Directory cannot enumerate the user rights to provide a list of candidate OUs to which to add the Filer.

```
11c2-3240-v02> cifs setup
Do you want to setup CIFS? [y]: y
This process will enable CIFS access to the filer from a Windows(R) system.
Use "?" for help at any prompt and Ctrl-C to exit without committing changes.

Your filer does not have WINS configured and is visible only to
clients on the same subnet.
Do you want to make the system visible via WINS? [n]: n
A filer can be configured for multiprotocol access, or as an NTFS-only
filer. Since multiple protocols are currently licensed on this filer,
we recommend that you configure this filer as a multiprotocol filer

(1) Multiprotocol filer
(2) NTFS-only filer

Selection (1-2)? [1]: 1
CIFS requires local /etc/passwd and /etc/group files and default files
will be created. The default passwd file contains entries for 'root',
'pcuser', and 'nobody'.
Enter the password for the root user []: *****
Retype the password: *****
The default name for this CIFS server is '11C2-3240-V02'.
Would you like to change this name? [n]: n
Data ONTAP CIFS services support four styles of user authentication.
Choose the one from the list below that best suits your situation.

(1) Active Directory domain authentication (Active Directory domains only)
(2) Windows NT 4 domain authentication (Windows NT or Active Directory domains)
(3) Windows Workgroup authentication using the filer's local user accounts
(4) /etc/passwd and/or NIS/LDAP authentication
Selection (1-4)? [1]: 1
```

```

What is the name of the Active Directory domain? [namdev.nsrootdev.net]:
namdev.nsrootdev.net
In order to create an Active Directory machine account for the filer,
you must supply the name and password of a Windows account with
sufficient privileges to add computers to the NAMDEV.NSROOTDEV.NET
domain.
Enter the name of the Windows user [Administrator@NAMDEV.NSROOTDEV.NET]:
rs01146@nsrootdev.net
Password for rs01146@nsrootdev.net: *****
CIFS - Logged in as rs01146@nsrootdev.net.
Setup was unable to retrieve a list of joinable containers
(organizational units) from Active Directory, therefore a list of
selectable options cannot be provided. Please enter the distinguished
name of the container that you would like the filer to join below.
There is no need to add the domain name portion,
'DC=namdev,DC=nsrootdev,DC=net', of the distinguished name.
Enter the name of the organizational unit [CN=Computers]:
ou=Servers,ou=StorageEng
CIFS - Starting SMB protocol...
It is highly recommended that you create the local administrator
account (11C2-3240-V02\administrator) for this filer. This account
allows access to CIFS from Windows when domain controllers are not
accessible.
Do you want to create the 11C2-3240-V02\administrator account? [y]: y
Enter the new password for 11C2-3240-V02\administrator: *****
Retype the password: *****
Wed May 4 11:32:55 EDT [11c2-3240-v02@11c2-fas3240-02:
nbt.nbns.registrationComplete:info]: NBT: All CIFS name registrations have
completed for the local server.
Welcome to the NAMDEV.NSROOTDEV.NET (NAMDEV) Active Directory(R) domain.
Wed May 4 11:33:03 EDT [11c2-3240-v02@11c2-fas3240-02:
cifs.startup.local.succeeded:info]: CIFS: CIFS local server is running.

```

11.4.1.2.3 Diagnosing issues in the CIFS setup process

The purpose of this section is to provide recommended courses of action when Active Directory provides certain responses back to the Filer when attempting to join the storage controller to the domain.

1. Active Directory user rights cannot be enumerated to provide a list of candidate OUs in which the Filer's machine account can be created.

Setup was unable to retrieve a list of joinable containers (organizational units) from Active Directory, therefore a list of selectable options cannot be provided. Please enter the distinguished name of the container that you would like the filer to join below. There is no need to add the domain name portion, 'DC=namdev,DC=nsrootdev,DC=net', of the distinguished name.

The OU path will need to be provided manually. Use the canonical OU path name gathered as part of pre-requisite step 4. For example the full path to the OU maybe 'OU=servers,OU=storageEng,DC=namdev,DC=nsrootdev,DC=net' however ONTAP already know the domain (DC) portion of the path so you only need to provide 'OU=servers,OU=storageEng'. Again this is an example so you path and domain details will differ

2. A computer account with the same name as the NetBIOS name of the Filer already exists in the Active Directory domain, either because the name is being re-used (re-running CIFS setup) or the computer account was pre-created in the domain using Active Directory tools.

An account that matches the name '11C2-FAS3240-01' already exists in Active Directory: 'cn=11c2-fas3240-01,ou=servers,ou=storageeng,dc=namdev,dc=nsrootdev,dc=net'. This is normal if you are re-running CIFS Setup. You may continue by using this account or changing the name of this CIFS server.

Either follow the on screen options to re-use the existing computer account or cancel the CIFS setup process or follow the on screen process to change the NetBIOS name of the Filer. If you choose not to use the computer account that is already in the domain and that account is not being used by another system you should ask your Active Directory administrator to remove the account if you do not have rights to do so yourself via A/D tools.

3. The IP address of the primary interface of the Filer is on a subnet that is not registered as part of Active Directory's Sites and Services tool as the pre-requisite step 3 had not been completed.

This filer does not seem to be a member of any Active Directory Site (see Site definitions using the 'Active Directory Sites and Services' tool.) CIFS works best in an Active Directory domain when its subnet is part of a defined Site that contains localized domain controllers. Specify how CIFS should treat Active Directory Site information.

(1) Ignore Active Directory Site information when choosing domain controllers

(2) Choose a site name from the list of available sites

Selection (1-2)? [1]: 2

Choose a Site name from the following list of Sites.

(1) Warren

(2) Warren-Temp

(3) none of the above

Selection (1-3)? : 1

In the above example a temporary workaround was implemented as the subnet had not been assigned to a site in Active Directory. The storage administrator had contacted the Active Directory administrator to confirm the appropriate site to which the Filer should be localized (in this case it was the site 'Warren'). This will allow the rest of the storage controller's setup and configuration processes to be completed. The Active Directory administrator should then update the Sites and Services information for the domain with the subnet information. Once this update has been replicated around the domain the '`cifs reset dc`' command should be run on the Filer to undo the workaround and inherit the site information from the domain.

11.4.1.2.4 Post domain join housekeeping tasks

In order to meet Citi's segregation of duties requirements, Domain Administrators must be removed from the local Administrators group on the Filer. This can be achieved in one of two ways:

1. Via computer manager MMC
 - a. On your Windows desktop point Computer Manager at the Filer using the UNC path of the system (`\\<FilerName>`)
 - b. Expand System Tools / Local users and groups
 - c. Click on the 'Groups' container
 - d. Right click 'Administrators' group and select 'Properties'
 - e. Select the entry '<Domain>\Domain Admins' from the list and then click the 'Remove' button (where <Domain> should be replaced with the name of the A/D domain)
 - f. Click the 'OK' button to close the Administrators Properties window
 - g. Close Computer Manager MMC

2. Via the ONTAP command line
 - a. Login to the Filers CLI
 - b. Remove Domain Administrators using the useradmin command
`useradmin domainuser delete "<Domain>\Domain Admins" -g Administrators`
 (where <Domain> should be replaced with the name of the A/D domain)

Once Domain Administrators are removed from the local administrators group, if the storage system is running a version of Data ONTAP prior to 8.1.3 then the Active Directory computer object which represents the Filer or vFiler is required to have the SID compression attribute disabled.

In order to complete this task, follow the procedure documented in section 7.2 of the [Data ONTAP Active Directory SID Compression Interoperability Tech Note](#).

If this task is not completed for a system running a version of ONTAP prior to 8.1.3 and Windows 2012 Domain Controllers are deployed in the environment then CIFS users will be prevented from accessing their data on the storage system (a Windows Error 5: access denied error is returned).

Upon upgrading to ONTAP 8.1.4 or later Storage Administrators are reminded to enable the SID compression support for the Active Directory computer account for each vFiler. Consult the [Data ONTAP Active Directory SID Compression Interoperability Tech Note](#) for the procedure on how to enable SID compression.

11.4.1.3 Share Naming Recommendation

Since clustering is being used on the production filers 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.

11.4.1.4 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.

11.4.1.5 SMB 2.0

Server Message Block (SMB) 2.0 is the next version of the Common Internet File System (CIFS)/SMB protocol. CIFS can also be considered as SMB 1.0. SMB 2.0 had the first major redesign considering the needs of the next generation of file servers and clients. These needs include a redesign for modern networking environments such as wide area networks (WANs), possible high-loss networks, time-outs, high latency, and so on.

The information prevented in this section is the basics required to implement SMB 2.0 on a NetApp NAS Filer. For a detailed technical presentation on SMB 2.0 refer to [NetApp TR-3740](#).

11.4.1.5.1 When should SMB 2.0 be enabled

From a client perspective SMB 2.0 is enabled in Windows Vista and later, and Windows Server 2008 (W2K8) and above releases. CATE Platform Engineering builds for these releases have SMB 2.0 enabled by default and as such Platform and Storage Engineering support the use of SMB 2.0 between the aforementioned clients and NetApp NAS Filers running ONTAP 7.3.3 or above. **Due to bugs that have been discovered since the initial SMB 2.0 enabled release of ONTAP, SMB 2.0 should only be enabled on systems where the monthly average CPU utilization on a storage controller is at or below 30% in order to prevent unexpected premature exhaustion of storage system resources which is known to occur in Citi's virtualized desktop environment^{19 20}. Please refer to section [11.4.1.1](#) for a summary of known SMB2 related issues.**

¹⁹ Use NetApp's OnCommand Unified Manager 'Controller CPU usage graph' report to make this determination

²⁰ Note this is not considered a CATE Standard and is included here purely in an effort to introduce a level of stability into the environment

11.4.1.5.2 Configuring SMB 2.0

In situations which meet the SMB 2.0 permitted use requirements detailed in section 11.4.1.5.1, SMB 2.0 is enabled at the physical Filer and vFiler levels. Enabling it on just the physical Filer will not mean it's enabled on all the vFilers hosted on that system.

SMB 2.0 is configured through the use of five ONTAP level options. These are discussed in detail in section [14.3](#) together with their default and recommended values as well as a description of each option. All settings with the exception of `cifs.smb2.enable` are hidden starting with ONTAP 8.1 meaning that to check the current configured value the option name must be requested specifically to return the configured value and the option value will not be returned with the simple 'options' list query. The descriptions of these settings can be accessed directly through the use of the following links:

[cifs.smb2.client.enable](#)
[cifs.smb2.durable_handle.enable](#)
[cifs.smb2.durable_handle.timeout](#)
[cifs.smb2.enable](#)
[cifs.smb2.signing.required](#)

11.4.1.5.3 Configuring SMB 2.0 Signing Optimization

Beginning with ONTAP 8.1.4P4, a new SMB 2.0 signing optimization feature has been introduced to parallelize the serial signing operations which caused a performance bottleneck in earlier ONTAP releases (see bug 817844 for details). Please note that enabling this feature is recommended as part of the same maintenance window where the ONTAP upgrade occurs as enabling the optimization is disruptive for CIFS clients. This parameter only needs to be configured if the storage system meets the SMB 2.0 permitted use requirements detailed in section 11.4.1.5.1.

The process to configure this optimization feature varies between ONTAP releases. With ONTAP 8.1.4P6 and later use the following steps to enable this feature:

1. Stop the CIFS service on all vfilers (including vfiler0) on the system
2. Enable the signing optimization by setting option `cifs.smb2.signing.multiprocessing` to 1 (it will show as 'enabled' rather than '1' in the options output)
3. Verify the value for options `cifs.smb2.signing.max_threads` is set to `n-1` (where `n` is the number of cores on the storage controller and the number of cores is the value returned in the "Processors" line of a 'sysconfig -a' output.)
4. Start the CIFS service on all vfilers (including vfiler0) on the system which were running CIFS prior to the start of this procedure

For the limited number of systems which are running 8.1.4P4 follow the steps below to enable this feature. Note the recommended value of the signing optimization feature (options `cifs.smb2.signing.max_threads`) will vary by hardware platform. Refer to section 14.3 for the recommended value for this setting. By default the SMB2 multi-processor signing optimization is disabled when upgrading from releases of ONTAP prior to 8.1.4. Follow the process below (irrespective of the version of ONTAP you are upgrading from) to ensure a consistent and predictable experience:

1. Stop the CIFS service on all vfilers (including vfiler0) on the system
2. Disable the signing optimization by setting options `cifs.smb2.signing.max_threads` to 0
3. Verify the value for options `cifs.smb2.signing.max_threads` is set to 0

4. Set the signing optimization to a value equal to the number of cores on the filer minus 1 (note you can determine the recommended value for the hardware platform by the entry for options `cifs.smb2.signing.max_threads` in section 14.3). The number of cores is the value returned in the "Processors" line of a `'sysconfig -a'` output.
5. Verify the value for options `cifs.smb2.signing.max_threads` is set to number defined in step 4.
6. Start the CIFS service on all vfilers (including `vfiler0`) on the system which were running CIFS prior to the start of this procedure

To explain the behavior of the `cifs.smb2.signing.max_threads` parameter further, when its value is set to zero the signing operation executes in the main code path (the CIFS processing domain) which is the default behavior for ONTAP prior to 8.1.4P1. When the value is set to 1 the signing handling operation is split out into a separate processing domain, the exempt domain, which prevents the signing process from gating other CIFS operations. Increasing the value of `cifs.smb2.signing.max_threads` further allows more cores to concurrently participate in the SMB2 signing handling process.

11.4.1.5.4 Configuring SMB 2.1

Starting with ONTAP 8.1, SMB 2.1 is enabled by default. There are no SMB 2.1 specific options which control the behavior of SMB 2.1 because 2.1 is an extension of the original 2.0 standard. Therefore the ability to disable SMB 2.1 and SMB 2.1 signing is tied to the following respective SMB 2.0 settings.

[cifs.smb2.enable](#)
[cifs.smb2.signing.required](#)

While there are no SMB 2.1 specific ONTAP configuration options there are still ways to control the behavior of SMB 2.1 on the storage system which is important considering the implications of bug 798842 in ONTAP 8.1.4P1.

Due to known issues Storage Administrators are required to disable SMB 2.1 upon upgrading the storage controller to ONTAP 8.1.4P1. See the Known Issues section of the [ONTAP 8.1 CATE release notes](#) for details on bug 798842 and refer to section 14.3 for the procedure to disable SMB 2.1 manually. Note that with the arrival of ONTAP 8.1.4P4 (or later version) SMB 2.1 is disabled by default and there is no need to implement the workaround documented in section 14.3.

11.4.1.5.5 Support for SMB 1.0 clients

SMB 1.0 clients will be able to communicate with an SMB 2.0 enabled Filer due to the protocol autonegotiation at the start of the CIFS session between the client and the Filer.

You can use the updated `cifs sessions` command in ONTAP 7.3 and above to determine which version of SMB a client is using to communicate with the Filer.

`cifs sessions -p smb` will show all SMB 1.0 client communication
`cifs sessions -p smb2` will show all SMB 2.0 client communication and also demonstrate whether the transactions are signed with a SHA256 key

11.4.1.6 SMB 3.0

NetApp Data ONTAP 7-mode does not support SMB 3.0 and the vendor has no plans to implement SMB 3.0 in the 7-mode architecture.

11.4.1.7 Access Based Enumeration (ABE)

Data ONTAP 7.2 and above provide storage system support for Access based Enumeration (ABE), a shared resource security feature introduced in Microsoft Windows Server 2003 Service Pack 1. This feature allows administrators to control the display of folders and other shared resources according to a user's access rights. This functionality is covered in this document to make storage administrators aware of this feature. CATE Storage Engineering is not mandating the use of this option on every CIFS share. It is intended for use for businesses who wish to use this functionality.

Conventional share properties allow you to specify which users (individually or in groups) have permission to view or modify shared resources. However, they do not allow you to control whether shared folders or files are visible to users who do not have permission to access them. This could pose problems, if the names of shared folders or files describe sensitive information, such as the names of customers or new products under development.

Access-based Enumeration extends share properties to include the enumeration of shared resources. When ABE is enabled on a CIFS share, users who do not have permission to access a shared folder or file underneath it (whether through individual or group permission restrictions) do not see that shared resource displayed in their environment. ABE therefore enables you to filter the display of shared resources based on user access rights.

In addition to protecting sensitive information in your workplace, ABE enables you to simplify the display of large directory structures for the benefit of users who do not need access to your full range of content. ABE can increase worker productivity. End users see only the files and folders that they need for their responsibilities, rather than spending time looking through lists of inaccessible folders and files. Administrators can be more productive because they don't have to help users navigate through complex shared folders.

Access-based Enumeration is enabled using the ONTAP CLI (it is not available in the Filerview GUI and cannot be enabled via the Microsoft Management Console). ABE is not enabled on all CIFS shares by default and needs to be enabled on each share by the Storage Administrator as the Storage Administrator is only admin with access to the ONTAP CLI. ABE is enabled at the CIFS share level using the following command (this example modifies an existing share – the same parameter can be used with the 'cifs shares –add' command).

```
nas3050gdse01> cifs shares -change <share_name> -accessbasedenum
```

It is disabled using the following command:

```
nas3050gdse01> cifs shares -change <share_name> -noaccessbasedenum
```

Changes to ABE are implemented immediately and are reflected on a Windows client when a refresh operation is performed in the Windows Explorer. Storage Administrators and SAs considering implementing ABE should ensure the user community is educated on the forthcoming system change to avoid a wave of invalid data deleted helpdesk calls and restore data requests.

If a file system is shared out at two different entry points, ABE will have to be enabled at both shares for the user experience to be consistent. For example, if /vol/data was shared and /vol/data/business/share is also shared out (the second share being in a child directory of the original share), then both shares would required ABE to be enabled to make the user directory browsing experience consistent.

Other items to consider when implementing ABE include:

1. When enabling ABE on production CIFS shares, ABE will also need to be enabled in DR. If you choose to configure your DR environment whereby CIFS shares are not enabled until the failure to COB occurs then be sure to update any scripted processes with the new ABE attribute.
2. In environments where you have Unix and Windows users accessing the same data you only need to enable ABE when you have CIFS users accessing an NTFS security file system. Enabling ABE on a file system with unix security may cause unexpected results in the users Explorer window.

The following example shows a share where a user only has access to four directories. The screen on the left is the normal Windows view and the screen on the right has ABE enabled. System Administrators familiar with Novell will recognize Access-Based Enumeration.




Figure 6 - CIFS ABE disabled




Figure 7 - CIFS ABE enabled

To determine if ABE has been enabled on a share use the 'CIFS shares' command as shown below.

```
Filer1> cifs shares rs01146
```

Name	Mount Point	Description
----	-----	-----
rs01146	/vol/rs01146	test

... access based enum supported

11.4.1.8 CIFS multiplexing

In order to operate more efficiently, some applications such as Citrix and IIS combine many CIFS sessions into the same TCP connection. These connections are known as multiplexed connections. The 'cifs.max_mpx' option in ONTAP determines how many CIFS sessions can be multiplexed into a single TCP connection. When that value is reached, creation of new CIFS sessions within that TCP connection will fail and those users may not be able to access files via CIFS. Other users may experience increased latency until existing CIFS sessions are closed.

The option 'cifs.max_mpx' is configured at both the physical Filer and vFiler levels and requires a termination and restart of the CIFS daemon for changes to take effect. The recommended process is to terminate CIFS, change the setting and then restart CIFS.

Increasing this value of 'cifs.max_mpx' does not impact the Filer's performance but it is known that increasing this value consumes more resources on the client side. With a desktop minimum specification of Windows XP, and server minimum specification of Windows 2003, the associated hardware for these O/S releases is expected to have enough resource capacity to comfortably accommodate this change.

Due to the diversity of Windows client platforms present on Citi's network a value for 'cifs.max_mpx' has been defined in section [14.3](#) that provides the best level of interoperability.

11.4.1.9 CIFS TCP Window Size

In comparing the relative CIFS performance on the different NAS platforms within Citi, the default setting for the CIFS TCP Window Size in Data ONTAP is relatively small. In an effort to offer a more comparable performance between NAS platforms, especially for CIFS transitions which occur over a long geographical distance, CATE Storage and Backup Engineering in conjunction with CATE Networks has tested and determined the viability of changing the default CIFS TCP Window Size in Data ONTAP.

By default this configuration parameter is hidden in Data ONTAP meaning that to check the current configured value the option name must be requested specifically to return the configured value and the option value will not be returned with the simple 'options' list query. The new recommended, non-default, value for this setting is defined in section [14.3](#) under the configuration parameter `cifs.tcp_window_size`.

11.4.2 NFS

11.4.2.1 NFS Overview

NFS is one of the file level access protocols supported by NetApp's Data ONTAP operating system; both NFSv3 and NFSv4 are supported although use of NFSv4 is not recommended by NetApp in production workloads prior to ONTAP 7.3. The NFS v4 functionality certified by CATE in ONTAP 7.3.3 is also available as part of ONTAP 8.0.x.

While most of the NFS configuration in this section is applicable to both NFSv3 and NFSv4 deployment, NFSv4 has additional configuration requirements and these are detailed in the NFSv4 subsection.

11.4.2.2 General NFS Configuration

Like vFiler, NFS is an optional NetApp feature that requires a license purchase. The current certified version of ONTAP 8.0 supports both NFS v3 and NFSv4 (NFSv4 is now certified by CATE Storage and Backup Engineering on ONTAP 7.3.3 or above), and like Oracle Solaris NFS server supports most of the NFS standards export options.

If you are implementing a new system as part of a tech refresh or migration that will host NFS data referencing CIFS-level permissions to grant UNIX-mapped access through NFS, CATE Storage Engineering recommends you implement the following setting:

```
filer1> options nfs.ntacl_display_permissive_perms on
```

This feature needs to be configured on each Filer or Vfiler where data is accessed using the method described above.

In order to optimize the Filers network configuration when using NFS ensure the [NFS receive window size for TCP](#) is configured as shown below for each physical Filer and vFiler in the environment. NFS clients will be required to either unmount / remount or reboot in order to establish a connection with the Filer at the new receive window size. Also ensure the TCP flow control settings for NFS are configured correctly (see section [14.3](#) for a complete list of all NFS settings).

For all NFS use cases `filer1> options nfs.tcp.recvwindowsize 65940`

LDAP netgroups may be used to select which servers have read-only, read-write or root access to the NFS export. Local netgroups are permitted in ONTAP using the `/etc/netgroup` file, it is preferable for netgroups to be held in LDAP as otherwise all filer netgroups on each filer or vfiler must be configured and kept in sync manually.

In order to reduce lookups to the name services NFS authentication is cached on the filer, this is especially important for NFS exports with read-write access control as this authentication is checked on **every** NFS IO. The use of caching on the filer means that a server will not immediately gain or lose access after a change to a netgroup, the data in cache must expire or the information in the cache must be manually purged before new values are available.

The following NFS options control the time data resides in cache and should be left at the default values:

- `nfs.export.harvest.timeout`²¹ (Default value: 1800 – Time in seconds that an unused successful access response will be cached)
- `nfs.export.neg.timeout`²² (Default value: 3600 – Time in seconds that an “access denied” response will be cached)
- `nfs.export.pos.timeout` (Default value: 36000 – Time in seconds that a successful access will be cached)

It is possible to manually flush the NFS authentication cache for all exports using the “`exportfs -f`” command. This command may also be used to flush the authentication cache for a specified export by supplying the export name with the `exportfs` command.

When using netgroups the “strict” option should be configured by using the options command shown below:

```
options nfs.netgroup.strict      on                (ONTAP default is off)
```

When the netgroup strict option is enabled all netgroups must be preceded by “@”, this prevents the filer trying to interpret netgroups as a hostname and hence increasing the lookup overhead when expanding netgroups.

There is an order of precedence for all export options and hence the same server should not be present in more than one netgroup used to limit read-only or read-write access when used on the same export unless this precedence is well understood. An example of this is where host-a is defined in both netgroup1 and netgroup2, if the export below is used access to host-a will only be read-only despite it appearing in the read-write netgroup as “ro=” has a higher precedence than “rw=”.

```
/vol/testvol      -sec=sys,ro=@netgroup1,rw=@netgroup2,nosuid
```

Additional restrictions:

- You can specify up to 255 names with the `rw` option. The names can be host names or subnets; you cannot specify netgroups.

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 `exportfs` 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.
 - o This method is for exporting resources that you want to export after each filer reboot. NOTE: This is the only method for SOE NetApps NFS support.
- Use the `exportfs` 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.
 - o This method is for exporting resources temporarily because the `exportfs` command is not persistent across filer reboots.

²¹ In the ONTAP MAN pages for ‘na_options’ this value is listed incorrectly as 3600. The correct value is the one shown in this document.

²² In the ONTAP MAN pages for ‘na_options’ this value is listed incorrectly as 1800. The correct value is the one shown in this document.

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

```
filer1> rdfile /vol/vol0/etc/exports
/vol/vol0 -rw=canonball.ny.ssmb.com,root=canonball.ny.ssmb.com
#/vol/philtest -sec=sys,ro,rw=roothost:@FIADM,nosuid,root=roothost
/vol/philtest -sec=sys,rw=roothost:@FIADM,nosuid,root=roothost
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
ifgrp 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
```

11.4.2.3 NFSv4 configuration and use

NFSv4 is the latest standard for NFS and is certified for use in ONTAP 7.3.3 and above (including ONTAP 8.0.x). NFSv4 has improvements in the following areas:

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

NFSv4 is currently being certified in a phased approach and the functionality detailed in this section relates to phases 1 through 3 of the NFSv4 certification project. Additional NFSv4 functionality will be added in later phases of the project and this solution guide will be updated as this work is completed.

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

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

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

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

- NFSv4 is now stateful (like CIFS) however the integration of NFS locking into NFSv4 actually improves the robustness of the protocol compared with NFSv3. It would also be incorrect to conclude that the move to a stateful protocol will adversely impact clients during failover as can be the case with in certain circumstances with CIFS.
- There are no longer ancillary protocols such as locking; all NFS related functions are now contained within the same protocol.
- IP port 2049 is used for all NFSv4 traffic.
- NFSv4 now makes use of compound operations with multiple operations being sent in the same request, thus reducing traffic. It should be noted that whilst this may speed up some workloads, for those workloads with a high component of non-IO operations (i.e. meta-data operations) there may be a reduced performance compared with NFSv3.
- NFSv4 is now much closer in functionality to the CIFS protocol with the addition of extra operations such as the “open” call.
- UDP is not supported for NFSv4 hence only NFSv4 over TCP is now supported.
- The NFSv4 standard states that all NFSv4 implementations must now support Kerberos authentication. Previously in the NFSv3 standard support for Kerberos authentication was optional, support for Kerberos is now mandatory in NFSv4. Whilst Kerberos availability in an NFSv4 implementation by a vendor is now mandatory AUTH_SYS may still be used by clients, this is set on the servers export permissions.
- When Kerberos is used as the authentication message the maximum number of groups that a Unix user can be a member of increases from 16 to 32, this can be further increase to 256 if required by setting “nfs.max_num_aux_groups” to the required value.

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

- The “-actual=” flag on an NFSv4 export is no longer supported in ONTAP. If a filer or vfiler has the “-actual=” flag set on an export and NFSv4 is enabled on the vfiler then the export of this area will fail.
- All storage providing NFSv4 must be configured to use LDAP or files and directories will appear to be owned by “nobody”.
- The “root” user and groups must be defined in the local passwd and group files on the filer or root ownership will not display correctly.

As the “-actual” flag is no longer supported when NFSv4 is enabled the mount point used by clients will also change. Thus enabling NFSv4 on a filer or vfiler when “-actual” flags are used for existing exports will necessitate changing not only the export entries but also the client mount definitions as well.

Previously not all filers/vfilers that supported NFSv3 were configured to use LDAP, many of these deployments either used IP addresses instead of netgroups or less commonly these netgroup were defined locally on the filers. As the storage, instead of the servers, now convert UID/GID’s into a user or group names (and also the reverse) the use of LDAP is now effectively mandated for all storage providing NFSv4 if these translations are to succeed.

As the “root” user and group are not defined in LDAP (they are usually defined on the server in its local files) the storage must have the root user and group defined in its own local passwd and groups files. If these root accounts are not defined then a file or directory with a root user ownership or having the root group will display as “nobody” as the lookup in LDAP will fail for the “root” user/group.

The following options should be set as follows on ONTAP to enable NFSv4 support:

```

nfs.v4.acl.enable      on
nfs.v4.enable          on
nfs.v4.id.domain       nfsdomain
nfs.v4.read_delegation on
nfs.v4.write_delegation on
nfs.v4.acl_preserve    on ←Recommended with ONTAP 8.0.2P4 or later (bug 438464)
                        Prior to ONTAP 8.0.2P4 this feature should be disabled

```

Note: "nfs.v4.acl_preserve" is a hidden value and has a default of "off", with the option at its default value changing a file's permissions will erase any ACL set on the file. Setting this option to "on" changes the behavior of ONTAP so that changing a file's Unix permissions via tools such as `chmod` will not clear a file or directory's ACLs. Setting the value to 'on' is not recommended unless running ONTAP 8.0.2P4 or above in the ONTAP 8.0 code line.

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

NFSv4 ACLs should be enabled for all filers using NFSv4 and "nfs.v4.acl_preserve" should be set to "on" only on Filers running ONTAP 8.0.2P4 or above in the ONTAP 8.0 code line.

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

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

Read delegation

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

Write delegation

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

If no delegations are held by any client then the client should work in a similar way to NFSv3 with regard to caching with the polling of modification times.

It is also worth noting that if an application has a lock on a file when it gives up its lock it must also flush any modified data.

As delegations are functions of the file open call within NFSv4 they will only be used when specifically coded in an application. The way in which a file is being used should determine whether delegations are requested as it may be inappropriate for some patterns of file e.g. access from multiple clients. For critical data it is recommended that write delegation is not used to ensure that minimal data is lost should a client crash before flushing data written to a file. Read delegation may not be appropriate where another client is periodically writing to the same file due to the need to revoke delegations.

Due to a bug in ONTAP 8.0.1 (BURT 458144) the value of "nfs.v4.acl.enable" must be set to "on" on both the base filer and vfiler or ACLs will not be visible to clients. Once the storage system is upgraded to ONTAP 8.0.2P4, or later, it is no longer required to enable this feature on the hosting (base) Filer.

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

Both NFSv4 client and storage must have the same value for the NFSv4 Domain ID, if these values do not match every file will appear to have an owner and group of "nobody".

Whilst the client now passes user details as strings the owner and group set on files and directories in WAFL still uses numeric UID and GID values., This ensure compatibility with NFSv3 clients. If a UID or GID of the file or directory cannot be translated to a known name in LDAP (or passwd/group files) then the value that cannot be translated will be displayed as "nobody" on the NFSv4 client rather than simply displaying the UID/GIG value as was the case with NFSv3. If all owner and group fields show as "nobody" check the NFSv4 domain value on the client and storage match, if only some of the files show as either owned by "nobody" or have "nobody" as their group it's likely that the owners of group cannot be found in LDAP or local passwd/group files.

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

11.4.2.3.1 NFSv4 ACLs

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

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

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

(See [14.9](#) Appendix I NFSv4 ACLs – further reading)

From Solaris 10 the same file looks as follows:

```
# ls -lv testfile
-rw-rw----+ 1 root      nasgrp9101      16 Oct 26 09:07 testfile
      owner@:rw-p--aARwC-s:-----:allow
      owner@:--x-----o-:-----:deny
      group@:rw-p--aARwC-s:-----:allow
      group@:--x-----o-:-----:deny
      everyone@:-----a---c---:-----:allow
      everyone@:rxp---ARW-Cos:-----:deny
```

As can be seen whilst the permissions are the same they display completely differently.

In ONTAP 8.01 a maximum of 400 ACE entries may be contained within each ACL.

The “nfs.v4.acl.enable” option set to “on” allows a client to view and manipulate ACLs, **it does not affect whether existing ACLs are obeyed** i.e. setting the “nfs.v4.acl.enable” option to off does not disable NFSv4 ACLs when checking file protections as the option name implies.

If a file has an ACL defined and then “nfs.v4.acl.enable” is set to off the existing ACL will still be obeyed but this ACL will be invisible to the client and cannot be changed (unless ACLs are re-enabled). This can result in a user being granted or denied access to a file when the Unix permissions seem to imply that they should, or should not, have access, this is a result of the hidden ACL on the file. Hence without visibility of ACLs an SA cannot be certain of whether access will be allowed or denied for a particular file or directory as the ACL component of the file protection is functioning but cannot be seen by the client.

The same is also true of inherited ACLs; even with “nfs.v4.acl.enable” set to “off” files or sub-directories created in a directory where an inherited ACL is defined on the parent directory will still inherit the defined ACL, neither of which will be visible to clients.

For these reasons “nfs.v4.acl.enable” should always be set to “on” where NFSv4 is used.

Whilst not strictly required in order to configure support of NFSv4 ACLs on NAS storage below is a brief discussion of some of the issues when using NFSv4 ACLs.

The support for NFSv4 ACLs differs from client operating system to client operating system, for example whilst Solaris 10 supports NFSv4 ACLs natively RHEL 5.4/5.5 requires an additional module (nfs4-acl-tools) to view and manipulate ACLs. Also the operating system tools such as cp, tar, cpio etc on platforms that natively support NFSv4 ACLs (such as Solaris) allow ACLs to be retained when copying or backing up files, this unfortunately is not true of Linux. RHEL uses Posix ACLs for its own file systems hence its native tools are Posix, not NFSv4, ACL aware and hence while ACLs may be queried and set using the `nfs4_getfacl` and `nfs4_setfacl` RHEL native commands such as cp, tar, cpio etc will strip off all NFSv4 ACLs from target files (they will not remove them from source files).

Care must be taken when using inherited ACLs. If an inherited ACE is set on a directory then files created in this directory will not have the users default Unix permission mask set on files, only the inherited ACL will be set. Therefore not only must the required inherited ACE be set but the required Owner, Group and Everybody Unix mask permissions must also be set as inherited ACEs.

When using inherited ACLs it is easy for a user to create a file or sub-directory using the parent directories permissions mask only to have all access removed from the newly created file as soon as the inherited ACL is added unless the inherited ACLs are set appropriately.

If deny ACL rules are used care must be taken as the order of the ACE entries in the ACL will determine the resulting permissions. All ACE's are processed starting with the rules displayed at the top and working down the list of ACE, the first match with the user is used to determine the access to the file or directory.

For example a rule that denies User_A all access to a file that is beneath another rule that grants read access to EVERYONE will not prevent User_A getting read access to the file, If however the order of the rules are reversed then User_A will be denied access as the first rule found is now the deny rule rather than the allow rule.

11.4.2.3.2 NFSv4 and NFSv3 coexistence

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

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

- NFSv3 access to files or directories with NFSv4 ACLs set will obey these ACLs but NFSv3 clients will not be able to view or change these ACLs.
- Creating files in sub-directories from NFSv3 in a parent directory with inherited NFSv4 ACLs defined will caused the created files to have the ACLs attached to the files/sub-directories in the same way as for an NFSv4 client file create.
- NFSv3 and NFSv4 locks are visible to each protocol and locks on one protocol are honored on the other.
- NFSv3 cannot perform byte range locking hence it will lock whole files, NFSv4 can perform byte range locking. Partial byte range locking in NFSv4 will prevent an NFSv3 gaining a lock, the reverse is also true.
- If an NFSv3 client crashes without releasing its locks then the lock will remain and must be cleared manually as for any other NFSv3 lock, this may impact NFSv4 clients. An NFSv4 clients locks will automatically be removed if an NFSv4 client crashes without releasing its locks. (NFS locks may be listed and manually released using the ONTAP "lock" command (see `na_lock` in NetApp documentation for more details) care must be taken as this command views and manipulates both CIFS and NFS locks.
- As files and directories have owners and groups defined on files using numeric UID/GID's in WAFL both NFSv3 and NFSv4 can access the same files successfully.

11.4.2.3.3 Useful NFSv4 information

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

- Solaris
 - o NFSv4 domain ID is set in the file “/etc/defaults/nfs”
 - o Set NFSv4 ACLs using `chmod`
 - o View NFSv4 ACLs using “`ls -lv`” or “`ls -lV`” (“`ls -ldV`” or “`ls-ldv`” for directories)
- Linux
 - o NFSv4 domain ID is set in the file “/etc/idmapd.conf”
 - o Requires module `nfs4-acl-tools` to manipulate NFSv4 ACLs
 - o Set NFSv4 ACLs using “`nfs4_setfacl`”
 - “`nfs4_setfacl`” cannot set empty ACE’s.
 - o `nfs4_getfacl` and `nfs4_setfacl` do not except wildcards for the filename.
 - o View NFSv4 ACLs using “`nfs4_getfacl`”
 - o Using Posix “`setfacl`” or “`getfacl`” commands will return errors as NFSv4 does not support the Posix ACLs used by these commands.
 - o OS native client copy and backup tools do not retain NFSv4 ACLs.

11.4.2.4 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
```

11.4.2.5 Using LDAP for user information

The default Citi Unix authentication is LDAP and Kerberos which replaces NIS+. Only the LDAP component of this solution is required for NetApp filers and LDAP provides the same functionality to NIS+ from a filer perspective.

LDAP cannot be configured from the filer “`setup`” command and hence must be configured manually after the initial configuration. LDAP configuration requires knowledge of both the local LDAP infrastructure servers and also the location of the user, password and netgroup information within the LDAP namespace. The location of this information within the LDAP namespace may differ from region to region and hence the location must be obtained from the local LDAP administrators before the configuration of the filer to use LDAP begins.

Starting with ONTAP 8.1.3 the `ldap.fast_timeout.enable` setting is introduced and by default this feature is enabled. In addition the `ldap.retry_delay` feature is introduced as requires changing from its default setting to a value which has been determined appropriate for Citi’s LDAP environment. Refer to section [14.3](#) for the best practice value for this setting.

The following options should be set on the filer or vfiler to be configured to use LDAP:

ldap.enable	← Enable LDAP
ldap.base	← Default point in LDAP namespace to start looking for data
ldap.base.group	← Point in LDAP namespace to find group data
ldap.base.netgroup	← Point in LDAP namespace to find netgroup data
ldap.base.passwd	← Point in LDAP namespace to find user data
ldap.servers	← LDAP infrastructure IP addresses
ldap.servers.preferred	← LDAP server used in preference to others
ldap.timeout	← Timeout for query before ONTAP tries the next LDAP server
ldap.fast_timeout.enable	← Controls the timeout behavior of LDAP client operations (ONTAP 8.1.3 and above)
ldap.retry_delay	← Sets the time, in seconds, to wait before trying to connect to LDAP servers after a failed attempt (ONTAP 8.1.3 and above)

The example below configures all the required non-default values for a filer to use LDAP

```
options ldap.enable          on
options ldap.base            dc=namux,dc=nsrootplt,dc=com:ONELEVEL
options ldap.base.group      ou=group,ou=nas,dc=namux,dc=nsrootplt,dc=com:ONELEVEL
options ldap.base.netgroup    ou=netgroup,ou=nas,dc=namux,dc=nsrootplt,dc=com:ONELEVEL
options ldap.base.passwd     ou=people,ou=PeopleRoles,ou=nas,dc=namux,dc=nsrootplt,dc=com:ONELEVEL
options ldap.servers          "162.124.148.188 150.110.142.188"
ldap.servers.preferred       162.124.148.188
ldap.timeout                 2
ldap.fast_timeout.enable     on
ldap.retry_delay             12
```

The LDAP servers will be queried in the order supplied in the ldap.servers list therefore any LDAP servers local to the filer should be first in the list. The ldap.servers.preferred should include all ldap servers to be used that are physically local to the filer.

When defining multiple ldap servers for either the "ldap.servers" or "ldap.servers.preferred" entries these should be entered within quotes and space separated. If entries are entered using a different separator, such as a colon, these entries will work for LDAP lookup but the filer will not failover to the remaining LDAP servers in the event that the first ldap server in the list is not available.

In the example above only one level below the indicated point will be searched as indicated by the "ONELEVEL" option on the end of the LDAP location. If no search limits such as "ONELEVEL" are specified the default, "SUBTREE", search criteria will be used. Any "SUBTREE" lookup will search all entries below the point indicated. Selecting the "onelevel" criteria limits the amount of searching performed in the LDAP namespace for unknown (or known) entries. Depending on the layout of the local LDAP namespace it may be possible that "SUBTREE" searches are required if data is located within multiple layers in the namespace and this should be determined during consultation with the local LDAP namespace administrators.

If a search point relevant to a particular namespace query is not defined then the "ldap.base" location will be used, specific data search points should be provided where possible. If these data specific search points are provided then LDAP searches will start at the relevant point in the LDAP namespace indicated by the specific LDAP entry e.g. for netgroups the relevant entry is "ldap.base.netgroup". It is likely that if only "ldap.base" is defined then a "subtree" search will be required but this will be inefficient and to be avoided if at all possible as searches will be unable to go directly to the location holding the information.

The following general information also applies to the use of LDAP within Citi:

- The Citi LDAP configuration is RFC 2307 compliant and hence the default ONTAP LDAP schema may be used without modification.
- No usermap information is available within the Unix LDAP and hence the filer /etc/usermap.cfg file can be used if required.
- The default LDAP tcp port (389) is used by the Citi LDAP servers for LDASP traffic.
- Connection to the LDAP server is not encrypted and hence SASL and SSL are not required.
- Anonymous access to the Citi LDAP servers is allowed hence no access accounts need be configured on the filer.

11.4.2.6 LDAP testing

It is possible to check the correct operation of LDAP using the advanced commands shown below:

```
*> getXXbyYY getpwbyname_r <username>    ← Returns users full details from LDAP
*> getXXbyYY getpwbyuid_r <UID>           ← Returns users full details from LDAP
*> getXXbyYY getgrlist <username>         ← Returns users group membership details from LDAP
*> getXXbyYY getgrbygid <GID>             ← Returns groups name from LDAP
*> getXXbyYY getgrbyname <group name>     ← Returns groups GID from LDAP
*> getXXbyYY netgrp <netgroup name> <hostname> ← Tests if hostname is a member of the netgroup in LDAP
```

Note: The user must be in the advanced privilege mode to access these commands. This command will also query NIS if this is configured hence the commands output will only be meaningful for LDAP testing if only LDAP is configured on the filer/vfiler.

11.4.2.7 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.

This should be modified to reflect the use of LDAP. Only the local files and dns should be queried for hostname lookup as LDAP is not configured to provide host lookups within Citi. An example for LDAP is shown below:

ldap example:

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

If NIS is not longer required it should be removed from the filer configuration including the nsswitch.conf file.

11.4.2.8 VMware on NFS

The purpose of this section of the Solutions Guide is to describe the operational best practices associated in implementing NFS storage for ESX server (VMware). In order to optimize the Filers network configuration when using NFS ensure the [NFS receive window size for TCP](#) is configured as shown below for each physical Filer serving VMware images. Also ensure the TCP flow control settings for NFS are configured correctly (see section [14.3](#) for a complete list of all NFS settings).

```
filer1> options nfs.tcp.recvwindowsize 65940
```

VMware does not support the use of the NFSv4 protocol.

11.4.2.8.1 System Licenses

The following ONTAP licenses are used as part of the VMware using NAS NFS configuration. Actual licenses in production deployment may vary depending on the implementation model and components used (for example if no storage DR is required to provide contingency to DR ESX hosts then the snapmirror license would not be required).

License Name	Purpose
a_sis	De-duplication
cifs	Filer management and authentication to system volume
cluster	High availability
multistore	Allows for the use of vFilers
nfs	File serving protocol license to ESX servers

Table 21 - ONTAP licenses required for VMware over NFS

11.4.2.8.2 Default volume security style

With the exceptions of Vol0, all volumes on the NAS filer will have **unix security**. The management of the Filer will be made simpler by defining the default volume security style as unix to avoid creating CIFS security style file systems for VMware:

```
options wafl.default_security_style unix
```

If this setting is not implemented then for each volume created for use with VMware the qtree security style will need to be defined.

```
qtree security /vol/<vol_name> unix
```

11.4.2.8.3 Aggregate and RAID group configuration

Please refer to the two VMware design documents for a demonstration of the aggregate and RAID group configuration. Refer to section [Error! Reference source not found.](#) for links to the latest documents.

11.4.2.8.4 Volume configuration, sizing and naming convention

The terms data-store [used in ESX] and volume [used in NAS] refer to the same object where VMware over NFS is concerned. The two terms can (and most likely will be) used interchangeably.

From a storage administration perspective, when provisioning volumes for VMware there are two types of volumes: one used as swap space and one used to store the virtual disks that make up the VM guest. The volume configuration will be the same irrespective of use, but the sizing (including space management policies) and naming conventions will differ.

11.4.2.8.4.1 Configuration

All volumes will be created using NetApp's thin provisioning process as storage utilization rates will vary dynamically over time.

When creating a volume use the '-s none' option to specify there is no space guarantee on the volume.

Ensure each volume has the 'no_atime_update' option set to on.

```
filer1> vol options <vol_name> no_atime_update on
```

Ensure quotas have not been enabled on the volumes. Use the 'quota status' command to confirm this.

The final volume configuration option relates to snapshots. See section [11.4.2.8.9](#) for snapshot (point-in-time copies) configuration for each volume.

11.4.2.8.4.2 Sizing

Volume sizes are defined in each of the specific VMware (VDI and VSI) design documents. Refer to section [Error! Reference source not found.](#) for links to the latest design documents. Automatic volume capacity increase (known as autosizing) must not be used as part of the VMware configuration.

11.4.2.8.5 NFS export configuration for VMware

A NFS export policy is used to prevent unauthorized hosts from accessing a mount point on the NAS Filer.

For a VMware over NFS implementation the VM Kernel IP address for each ESX host must be added to the export policy of every data-store / volume that the ESX cluster needs to have access to. Irrespective of the number of hosts in the initial ESX cluster, the IP address allocation for all the potential hosts in the ESX cluster should be known at system installation time. Where possible storage administrators should be provided with the list of all VM Kernel IP addresses for the fully populated ESX cluster at the time when the file system is originally exported. This will prevent changes to the NFS export configuration as subsequent hosts are added to the ESX cluster.

Based on current ESX deployment strategy it is assumed that there will be an average of ten (10) ESX hosts in a cluster. The export policy example show below is based on that assumption and is applied to the volume esxswd01.

In essence the export policy is allowing the VM Kernel port read/write access as Root to the NFS file system as ESX performs all I/O operations to the NFS file system as Root. The export policy also disables the use of setuid and setgid executables as well as preventing mknod commands.

This is achieved by the 'nosuid' parameter in the export policy and is required for a secure build. The export policy relies on the Unix SYS AUTH authorization mechanism.

```
filer1> exportfs
/vol/esxswd01 -sec=sys,rw=10.40.42.10:10.40.42.11:10.40.42.12:
10.40.42.13:10.40.42.14:10.40.42.15:10.40.42.16:
10.40.42.17:10.40.42.18:10.40.42.19,root=10.40.42.10:10.40.42.11:
10.40.42.12:10.40.42.13:10.40.42.14:10.40.42.15:10.40.42.16:
10.40.42.17:10.40.42.18:10.40.42.19,nosuid
```

Note: only IP addresses can be used when exporting an NFS mount point for ESX. Testing has shown the use of netgroups instead of IP addresses in the export policy does not work with ESX. In any case the use of IP addresses is more efficient as it removes the need for latency sensitive DNS lookups.

The following command was used to apply the above export policy to the volume:

```
filer1> exportfs -p sec=sys,rw=10.40.42.10:10.40.42.11:10.40.42.12:
10.40.42.13:10.40.42.14:10.40.42.15:10.40.42.16:
10.40.42.17:10.40.42.18:10.40.42.19,root=10.40.42.10:10.40.42.11:
10.40.42.12:10.40.42.13:10.40.42.14:10.40.42.15:10.40.42.16:
10.40.42.17:10.40.42.18:10.40.42.19,nosuid /vol/esxswd01
```

11.4.2.8.6 CIFS shares configuration for VMware

CIFS shares are not required when configuring volumes for use as VMware NFS datastores. CIFS should only be licensed on the storage system to allow administrative access to the root volume on the Filer.

11.4.2.8.7 Deduplication for VMware

Dedupe (ASIS) is configured at the volume (datastore) level and executes as a scheduled process.

Before writing any data to a volume or creating a snapshot policy, storage administrators must activate dedupe on each volume that will contain a non-swap space datastore (remember swap space is not being deduped in this implementation).

Refer to section [11.7.1](#) of this document for all procedures and processes related to implementing deduplication.

Remember that dedupe is a scheduled operation and so storage utilization rates will rise over the course of the business week and contract at the weekends when the deduplication operations runs.

11.4.2.8.8 Replication (snapmirror) for VMware

At the time of writing this document there were no infrastructure replication requirements for VMware storage. Per the VMware team all replication for DR is to be performed at the O/S level in the VMware guest or at the application level. A snapmirror license on the storage system is not required.

11.4.2.8.9 Point in time copy (snapshot) for VMware

Snapshots are not to be configured on VMware NFS datastores. While the ability to offer point in time disk copies may seem appealing initially, there is considerable risk to system availability when the change rate of the VMs in the data store result in the storage volume running out of capacity. After a joint review between VMware and Storage teams the decision was made to disable the use of storage level snapshots across the entire VMware environment (both VDI and VSI).

```
filer1> snap sched <vol_name> 0 0 0
```

In addition to configuring the snapshot schedule, the snapshot reserve needs to be configured for each NFS volume. ESX needs to be aware of the total amount of space available on a volume, rather than what is assigned to the active file system and what is reserved for snapshots because ESX does not have the embedded intelligence to interpret this vendor specific volume presentation methodology (nor should it be required to have it). Therefore, for each NFS volume, the snap reserve needs to be set to zero:

```
filer1> snap reserve <vol_name> 0
```

This completes the snapshot configuration requirements to support VMware on NFS.

11.4.2.8.10 Snapvault (backup) for VMware

Snapvault is not a certified backup approach for VMware environments running on NAS. Guest level backups are the current supported data protection strategy.

11.4.2.8.11 Expanding storage and existing aggregates

When expanding the storage system in order to provide additional physical capacity to an existing aggregate, new writes will be biased towards the new disks added to the system. To prevent performance issues it is recommended that an aggregate reallocate process be run against each aggregate contained that has recently been expanded. Use the following command to achieve this:

```
filer1> reallocate start -A <aggr_name>
```

11.4.2.8.12 Provisioning example

Provisioning is described in the specific VMware (VDI and VSI) design documents. Refer to section [Error! Reference source not found.](#) for links to the latest design documents.

11.4.2.8.13 Monitoring VMware NFS Filers

Unified Manager and Performance Advisor are required for monitoring Filers that serve data to ESX clusters over NFS.

11.4.3 Multi-protocol Volumes

There is a common misunderstanding that in order to access data concurrently by both NFS and CIFS the security style must be set to Mixed, this is not true. The protocol used to access the data (NFS/CIFS) and the security style are totally independent.

A data area using the Mixed security style will change the security style on each individual file **every** time a files permissions are changed. The files security style will be set to match the native security style of the changing OS (i.e. Unix or NTFS). The effective security permissions of each file may therefore change dramatically as a files access permissions change between either being native to a particular OS and being translated, this is the reason for it not being recommended by NetApp.

The security style used on a volume or qtree should match the predominant access that is writing data i.e. “Unix” style for NFS (v3 or v4) and “NTFS” for Windows data.

The use of the “Mixed” security style is not certified for use.

In order to access data from either Unix (NFS) or Windows (CIFS) the user trying to access data must provide their credentials to the filer for it to match against to a particular file or directories security settings in order to determine the level of access granted. Depending on the access different mechanisms will be used.

The security settings on a particular file or directory will depend on the security style set on the filer. Each file or directory can only have **either** a Unix **or** a Windows security protection entry; it cannot have both.

Client Accessing Data	Security Style of Data	
	Unix	Windows
Unix	No translation required	Unix username converted to equivalent Windows user and compared
Windows	Windows username converted to equivalent Unix user and compared	No translation required

Table 22 - Name translation requirements for multi-protocol volumes

In order for Unix users to access data protected by Windows security or vice versa a translation must occur from the native roles and credentials to one native to the security style. It is also worth understanding that Unix and Windows security mechanisms work in completely different ways so that even when Unix or Windows security translations occur the resulting access to file will not always be identical under both schemes.

Whilst the NetApp filer has the ability to map specific Unix users to specific Windows users (and vice versa) via the NetApp usermap.cfg configuration file the assumption is that the Unix and Windows usernames will be identical, if this is the case then no entries are required in the usermap.cfg file for user mappings as the default mapping will be used. This means that a Windows user ab12345 will be matched to a Unix user ab12345 and will be assumed to be the same person. **If a match is not found in the other OS then no access will be granted unless a specific entry has been manually added to the /etc/usermap.cfg file on the filer to map that user to a specific account.**

11.4.3.1 Unix access to data secured using Windows security style

The following steps take place when a Unix users attempts to access data protected using the windows security style.

- Unix user presents their User ID (UID) and Group ID's (GID's) to filer via NFS
- Filer converts user UID to Unix username by looking up UID in NIS passwd table
- Filer looks up this username in Windows active directory and obtains users windows roles for this username
- Access to the file is granted on the roles retrieved via Windows active directory in conjunction with the CIFS ACLs on the file system or file.

Note: GID and UID's are integer numbers used to represent the various users the groups defined on the Unix system.

11.4.3.2 Windows access to data secured using Unix security style

The following steps take place when a Unix users attempts to access data protected using the windows security style.

- Windows user presents their Windows active directory account username
- Filer looks up the windows username in the NIS passwd table and returns the matching Unix users UIS and primary GID
- Filer looks up the username in the NIS group table and returns all GID's granted to this username
- Access to the file is granted based on the UID and GID's granted to the username

11.4.3.3 Changing file permissions

When a user access data held in a non-native security style then there are also changes to the modifications that a user can make to the file access permissions as shown below:

Client Accessing Data	Security Style of Data	
	Unix	Windows
Unix	File permissions (or ACLs for NFSv4) changed via normal Unix tools	Cannot change file permissions or ACLs
Windows	Can only change file permissions if a cifs option is changed from ONTAP default.	File permissions changed via security tab on file properties

Table 23 - Changing file permissions for multi-protocol volumes

Enabling the `cifs.preserve_unix_security` option allows the modification of Unix file permissions by a Windows client (ONTAP 7.2 or above).

NetApp note that many Windows applications incorrectly interpret ACL's when reading files that have Unix security and that when the application saves the file the original Unix permissions may be lost.

NetApp recommend enabling `cifs.preserve_unix_security` for files in a qtree with Unix security where files have Unix permissions and are edited using Windows applications. This option is enabled as shown below.

```
options cifs.preserve_unix_security on
```

Where Unix files are accessed from Windows `cifs.preserve_unix_security` should be enabled on the filer or vfiler.

Setting this option also changes the way a Unix security style share is seen by Windows. With `cifs.preserve_unix_security` disabled (the only option prior to ONTAP 7.2) a Unix share would display as a FAT volume, enabling this option changes the volume type seen from Windows to NTFS. This change to NTFS also means that files now have a security tab in their properties screen and hence a file with Unix permissions may be manipulated from Windows. Setting Unix security values via the Windows security tab is not user friendly and is best avoided if possible to avoid incorrect security settings on the file.

11.4.3.4 Default file permissions

When files are created they are created with a default set of permissions, these permissions depend on the security style of the directory into which the file is to be created.

11.4.3.4.1 Unix security style

Files are created according to the umask setting on the users Unix server session. As Windows does not have the concept of the umask a default umask setting must be created on the CIFS share when it is created (or existing shares modified), without this setting all files are created with Owner, Group and Others having read, write and delete access. Having files writeable by all users is in conflict with the Citi security standards and hence must be avoided.

In order to restrict the Group and Other access to files in line with Citi Unix security standards a umask of 022 should be set, this still gives the Owner read, write and delete access but restricts Group and Other access to read (The default value for the umask is 000 on a NetApp filer which allows world read and write).

The digits of the umask refer to the default permissions as shown below:

- 0 - read, write and execute
- 1 - read and write
- 2 - read and execute
- 3 - read only
- 4 - write and execute
- 5 - write only
- 6 - execute only
- 7 - no permissions

The first digit of the umask sets the default Owners permissions, the second the Group and the third digit the Others access. In the recommended settings of 022 this gives Owners a value of 0 (read, write and execute) and both Group and Others a value of 2 (read and execute only).

The command below shows setting the umask on a CIFS share, this may either be done when a CIFS share is originally created on the filer or later as shown in the example.

```
filer> cifs shares testunix
```

Name	Mount Point	Description
testunix	/vol/eng/testunix	test unix
	everyone / Full Control	

Note: there is no umask value hence the default umask value of 000 is in effect.

```
filer> cifs shares -change testunix -umask 022
1 share(s) have been successfully modified
```

```
filer> cifs shares testunix
```

Name	Mount Point	Description
testunix	/vol/eng/testunix	test unix
	... umask=22	
	everyone / Full Control	

Umask value of 022 set

Using the “umask” switch the same umask value is applied to both files and directories. If too restrictive a umask is used the resulting file permissions may stop Windows users accessing the directories.

An alternative mechanism on ONTAP 7.2 and above is where both the file and directory umask values are set independently as shown below. If a value is not set then the default umask value of “000” will be used.

```
filer> cifs shares -change testunix -file_umask 022
```

```
filer > cifs shares testunix
```

Name	Mount Point	Description
testunix	/vol/eng/testunix	
	... file_umask=22	
	everyone / Full Control	

File umask set to 022

```
filer > cifs shares -change testunix -dir_umask 022
```

```
filer > cifs shares testunix
```

Name	Mount Point	Description
testunix	/vol/eng/testunix	
	... dir_umask=22	
	... file_umask=22	
	everyone / Full Control	

Directory umask set to 022

Shows value from previous command (above)

As with the “umask” these switches may either be used to modify an existing share or when creating a new share. The “-dir_umask” and “-file_umask” settings override any “-umask” settings on the share.

As the “umask” options are not available from the normal Windows mechanisms to create CIFS shares all CIFS shares referencing data using Unix Security Style must be either created on the filer itself via the “cifs shares” with a umask switch command or have the umask set on any existing share via “cifs shares -change”.

11.4.3.4.2 NTFS Security Style

When using the NTFS security style the file always inherits its permissions from the directory in which it resides. Whilst the windows permissions may be changed via CIFS connections it is not possible from Unix NFS connections. Attempting to use the Unix “chmod” command to change file permissions will result in an error message as shown below:

```
$ chmod 0-r aaa.txt
chmod: WARNING: can't change aaa.txt
```

11.4.3.5 General Configuration

In order to provide a secure translation of privileges the following options should be set on the filer.

```
options wafl.default_unix_user "pcuser" (ONTAP default is "pcuser")
options wafl.default_nt_user "" (This is the ONTAP default)
```

The following changes should be added to the /etc/usermap.cfg file to increase security; this file is only used when translating Windows users to Unix users and vice versa.

- ***\root => ""** – Stops a Windows user called “root” gaining root privileges for Unix areas.
- **"" <= administrator** – Stops a Unix user called “administrator” being mapped to the windows administrator account for Windows areas.

Note: The first match found from the top of the file in the /etc/usermap.cfg file is used for any mapping.

As the standard NetApp management platform is Windows and CTI standard security style for the filers configuration data area is the Windows security style all access to these configuration areas will use Windows security mechanisms.

The `wafl.nt_admin_priv_map_to_root` option defines whether a Windows users in the administrator group will be mapped to the root account where data is protected with the Unix security style and is recommended to be enabled allowing a Windows administrator to access Unix security style data as the root user. [The setting of this option is currently being reviewed by CATE Storage Engineering].

It is important to note that for environments that have not implemented ONTAP RBAC that storage administrators whose accounts are members of the local administrators group inherit root privileges to NFS data when this option is configured to on. This is one of the ways in which RBAC helps achieve segregation of duty.

11.4.3.6 NFSv4 and CIFS multi-protocol access

When configuring NFSv4 and CIFS access to the same data area the requirements for security styles and mechanisms used to translate user credentials are the same as for NFSv3 and CIFS as discussed previously in this document. There are however differences when NFSv4 ACLs are used.

Despite CIFS supporting its own ACLs that are similar to NFSv4 ACLs these are neither visible nor modifiable from Windows clients; however as with NFSv3 any access to a file will still be determined by these hidden ACLs. When viewed from Windows only the standard user, group and other permission bits will be displayed not NFSv4 ACLs, this means that a Windows user may be allowed or denied access to a file even when the displayed Windows permissions would appear to deny or allow to the user.

An NFSv4 client viewing a Windows area will see the same permissions as would an NFSv3 client; the NFSv4 client is not able to see or manipulate the NTFS ACLs despite the large similarity between NFSv4 and Windows ACLs.

An additional complication is that the ACL permissions required to allow (or deny) access to a file to a Windows user is different to those required for a Unix user. Hence in order to permit both Windows and Unix access to a file protected by NFSv4 ACLs the permissions must be expanded as shown below:

NFSv4 ACL's required to grant CIFS read access are (rnc)

- Read_data (r)
- Read_xattr (n)
- Read_acl (c)

NFSv4 ACL's required to grant CIFS write access are (rwaTnNc)

- Read_data (r)
- Write_data (w)
- Append_data (a)
- Write_attributes (T)
- Read_xattr (n)
- Write_xattr (N)
- Read_acl (c)

The NFSv4 ACL requirements to grant access to CIFS clients is not documented in any NetApp or IETF standards documents; it is not yet apparent whether these requirements are unique to NetApp interoperability between NFSv4 and CIFS or whether it is also applicable to multi-protocol implementations from other vendors.

11.5 Data Protection (Point in time copies, replication and backup)

11.5.1 Snapshots (point in time disk based copies)

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.

On volumes where snapshots have been enabled, every directory in the filer's active file system contains a directory named ~snapshot (for CIFS) and .snapshot (for NFS), 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.

11.5.1.1 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 and above *one volume can have up to a maximum of 255 snapshots*, however **CATE Storage Engineering 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 hourly.0 and hourly.1. 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 nightly.0 to nightly.21. 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 n in nightly. n is adjusted such that the higher the value of n, 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 24 - 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 ([11.5.3.5](#)) for details.

11.5.1.2 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 default snapshot reserve size for any new volume created on the storage system is inherited from the snapshot reserve size of the root volume. If you plan to implement a fixed snapshot reserve size across a Filer then defining it on the root volume will simplify snapshot management.

CATE Storage and Backup Engineering recommends that the snapshot reserve value should be decreased in fixed amounts of 5% (for example 20 to 15, 15 to 10), but realizes this approach may not suit all business requirements. Operations teams are encouraged to implement snapshot management policies that optimize their storage environment, while meeting any business requirements for point-in-time copies.

It should also be noted that there may be occasions where the rate of change in the data is high enough to warrant an increase in the snapshot reserve. The temptation to increase the snap reserve should be metered with an approach that reduces the amount of online snaps, and thereby relieves the pressure to increase the snap reserve. CATE Storage and Backup Engineering does not recommend using more than **25%** of the active volume for snap reserve as you are not making appropriate use of the active file system. If your requirements for a snap reserve are greater than 25% you should consider moving the snap shots to a secondary/achieve tier using snapvault. Note that a snapvault transfer will require a baseline image to be created on the target.

11.5.1.3 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](#)
- Using the Previous Versions tab in the Properties dialog box from Windows Explorer

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

When viewing NFS file systems from Unix clients, the snapshot directory is listed as '.snapshot', as opposed to '~snapshot' for CIFS.

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).

11.5.1.4 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.

11.5.1.5 Snapshot Creation Timestamp

Due to the default naming convention of snapshots and the behavior of various client operating systems it is sometimes difficult to determine the time each snapshot was created by just browsing the snapshot directory.

For Linux and Unix use `ls -l`.

From a Windows client when viewing a directory contents in Explorer with a 'Detailed' view, ensure to add the column "Date Accessed". The "Date Accessed" column will depict the time and date at which the snapshot was created.

11.5.2 Snapmirror (asynchronous 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`.

11.5.2.1 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. Consult section [11.5.2.3](#) 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. On larger Filers with many volumes it may not be possible to run at a 10 minute schedule. Storage Administrators may have to allow for more time between asynchronous replications to avoid CPU and memory resources from being overwhelmed with snapmirror processes. Approval should be obtained from the business prior to making any scheduling change as by implementing a longer replication duration the storage administration is changing the recovery point objective of the destination (DR) system.

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. 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, CATE Storage Engineering 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 25 - Snapmirror standard schedule

11.5.2.2 Snapmirror Replication Throttling

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.

On a storage system with many replicated volumes, defining a rate at which to throttle each snapmirror transfer can become tedious. With ONTAP 7.2 and above it is possible to throttle all snapmirror transfers using the following options. **It's important to note that these options also affect SnapVault transfers.**

`options replication.throttle.enable on` enables global network throttling of snapmirror and snapvault transfers. The default value for the options is 'off'.

`options replication.throttle.incoming.max_kbs <value>` specifies the maximum total bandwidth used by all the incoming (applied at destination) snapmirror and snapvault transfers, specified in kilobytes/sec. The default value for this option is 'unlimited', which means there is no limit on total bandwidth used. Valid transfer rate values are 1 to 125000. This option is valid only when the option `replication.throttle.enable` is on.

`options replication.throttle.outgoing.max_kbs <value>` specifies the maximum total bandwidth used by all the outgoing (applied at source) snapmirror and snapvault transfers specified in kilobytes/sec. The default value for this option is 'unlimited', which means there is no limit on total bandwidth used. Valid transfer rate values are 1 to 125000. This option is valid only when the option `replication.throttle.enable` is on.

The 'per transfer' throttle (as defined earlier) from the snapmirror.conf will still be applied. When both are per transfer and system wide throttling are configured, then throttling at system wide will be applied only if combined bandwidth used by all the relationships goes above the system wide throttling value.

11.5.2.3 Scaling concurrent replication processes

The purpose of the following table is to provide capacity planning data in terms of the number of concurrent replication operations supported per hardware platform. For the purposes of capacity planning a replication event is considered to be a snapmirror (both volume and qtree) and snapvault operation. A volume SnapMirror replication, qtree SnapMirror replication, or SnapVault replication consists of two operations, one operation on the source side of the transfer and the other on the destination side. Therefore, if a storage system is the source of one replication and the destination of another replication, it uses two replication operations. Likewise, if a storage system is the source and destination of the same replication, it uses two replication operations.

Each value listed in the following table indicates the maximum number of concurrent replication operations allowed for a specific storage system, when using a single type of replication. The values are exclusive, and not cumulative. For example, if you are using the maximum number of volume SnapMirror replication operations for a storage system as a source, you cannot use any more replication operations of any type for the storage system.

The amount of concurrent replication processes varies depending on whether the NearStore license is installed on the storage system.

Filer type	Without NearStore License						With NearStore License					
	Volume SnapMirror		Qtree SnapMirror		SnapVault		Volume SnapMirror		Qtree SnapMirror		SnapVault	
	Src	Dst	Src	Dst	Src	Dst	Src	Dst	Src	Dst	Src	Dst
2040	50	50	64	64	64	64	50	100	160	160	160	160
3040	50	50	64	64	64	64	50	100	160	160	160	160
3070	50	50	64	64	64	64	50	100	320	320	320	320
3140	50	50	64	64	64	64	50	100	80	80	320	320
3160	50	50	64	64	64	64	50	100	160	160	320	320
3170	50	50	64	64	64	64	50	100	192	192	384	384
3210	50	50	64	64	64	64	50	100	60	60	120	120
3220	50	50	64	64	64	64	50	100	160	160	320	320
3240	50	50	64	64	64	64	50	100	160	160	320	320
3250	50	50	64	64	64	64	50	100	160	160	320	320
3270	50	50	64	64	64	64	50	100	160	160	320	320
6040	100	100	96	96	96	96	100	200	192	192	384	384
6080	150	150	128	128	128	128	150	300	256	256	512	512
6210	150	150	128	128	128	128	150	300	256	256	512	512
6220	150	150	128	128	128	128	150	300	256	256	512	512
6240	150	150	128	128	128	128	150	300	256	256	512	512
6250	150	150	128	128	128	128	150	300	256	256	512	512
6280	150	150	128	128	128	128	150	300	256	256	512	512
6290	150	150	128	128	128	128	150	300	256	256	512	512

Table 26 - Maximum number of concurrent replication operations

Src = Source Dst = Destination n/s = NearStore License not supported on platform

11.5.2.4 Snapmirror Compression

ONTAP 7.3.3 and above includes a snapmirror compression process. SnapMirror compression works by running data through a software based compression engine and then transmitting the data to a target system(s) where it is uncompressed, reordered and written to disk. WAN optimization devices, on the other hand, in a similar fashion to SnapMirror compression, compressing as well as eliminating redundant data within the packet stream. The major difference being that WAN optimization devices sit outside the storage so they can handle traffic from multiple applications concurrently. One other difference is that WAN optimization devices are designed to handle issues with high latency, poor line quality between sites.

They are a few key rules of thumb in determining whether to use snapmirror compression:

1. If WAN optimization devices are already deployed at each end of the network link do not enable snapmirror compression as it will just create an unnecessary overhead on the storage controller and WAN optimization device.
2. If you do not have WAN optimization devices and are only looking to optimize SnapMirror bandwidth utilization – use SnapMirror compression.
3. If your network links have high latency / packet loss – use a WAN optimization device. Then you get compression and network 'stabilization.'

During testing conducted in CATE labs compression rates varied from 14:1 to 24:1, with a median average of 20:1. Results will vary over time, workload and the data set update being replicated. Tests showed CPU load increased 2% to 3% for a short period of time (less than a minute) while compression was conducted. Load increase may differ based on other file serving workloads.

Snapmirror compression can be enabled on a per replication relationship basis and requires editing of the `/etc/snapmirror.conf` file. As ONTAP defines replication at either the volume or qtree level this means you have fairly granular control over which replicated file systems will have snapmirror compression enabled.

The high-level steps to configuring snapmirror compression are listed below followed by an example:

1. Backup the `/etc/snapmirror.conf` file on the destination system by creating a copy of the file
2. Edit the `/etc/snapmirror.conf` file on the destination system:
 - a. Create a connection name expression for the source and destination systems
 - b. Define the source and destination file system relationship using the connection name expression
 - c. Enable compression on the relationship by adding the `compression=enable` statement
3. Save the updated `/etc/snapmirror.conf` file

To disable compression on that snapmirror relationship simply remove the `'compression=enable'` statement from that line.

Now let's work through an example of converting an existing snapmirror relationship to a compression enabled relationship.

1. Backup the `/etc/snapmirror.conf` file on the destination system by creating a copy of the file
2. Consider an existing snapmirror relationship defined in `/etc/snapmirror.conf` of the destination storage system as shown below to which you want to enable snapmirror compression:
`naslab01:vol1 naslab01dr:vol1dr - 1-59/10 * * *`
3. Edit the `/etc/snapmirror.conf` file using a text editor such as notepad
4. Firstly comment out the existing relationship by adding the `'#'` character to the start of the line
`# naslab01:vol1 naslab01dr:vol1dr - 1-59/10 * * *`
5. Create a multipath snapmirror relationship between the source and destination storage systems by adding the following line to the file:
`lab01=multi(naslab01,naslab01dr)`

Where:

lab01 is the connection name for the multipath snapmirror relationship

'multi' defines the snapmirror relationship as a multipath relationship

naslab01 is the source system and naslab01dr is the destination system

The connection name could actually be any imagined word, but lab01 is used as a shortened version of the source system name because it makes the output of a snapmirror status command easier to interpret and understand the name of the source system (as you will see later in this example). Removing the nas prefix from the source system name when defining the connection name ensures the connect name is not confused with an absolute reference to the source system

6. Define the replication relationship using the connection name rather than the name of the source system and enable compression on the relationship:

```
lab01:vol1 naslab01dr:vol1dr compression=enable 1-59/10 * * *
```

Where:

lab01 is the connection name for the multipath snapmirror relationship

vol1 is the name of the replication source volume on the source storage system

naslab01dr is the name of the destination storage system

vol1dr is the name of the replication destination volume on the destination storage system

compression=enable activates the compression function for the snapmirror transfer

1-59/10 * * * is the cron job syntax for the replication scheduler

7. Save the changes to /etc/snapmirror.conf and close the text editor

After waiting the allotted time for at least one scheduled transfer to complete (depending on your defined interval), confirm the status of the replication job using the snapmirror status command.

```
naslab01dr>snapmirror status
```

Snapmirror is on.

Source	Destination	State	Lag	Status
lab01:vol1	naslab01dr:vol1dr	Snapmirrored	00:09:11	Idle

As you will notice using the naming convention suggested in the above example you can still determine the source storage system for this replication transaction, which in this case is lab01.

While Snapmirror compression may seem appealing to environments that do not have WAN optimization devices installed there are some challenges with managing and monitoring it that storage operational owners should be aware of.

1. The only way to determine if a snapmirror relationship is enabled is by reviewing the /etc/snapmirror.conf file and search for the 'compression=enable' expression. Compression state is not shown in a snapmirror status command and will only show in a snapmirror status -l command if the transfer is actually in process when the command is run. Unified Manager is not aware of a compression enabled snapmirror transfer.
2. Savings derived from the snapmirror compression are shown in the /etc/log/snapmirror file. However there are no tools or scripts available at the time of writing this documentation by which you can calculate a network bandwidth saving over a period of time. An example of the log entry is shown below

```
dst Mon Jul 26 21:11:00 EDT lab01:vol1 naslab01dr:vol1dr Request (Scheduled)
dst Mon Jul 26 21:11:04 EDT lab01:vol1 naslab01dr:vol1dr Start
dst Mon Jul 26 21:11:08 EDT lab01:vol1 naslab01dr:vol1dr End (136 KB,
Compression 21.5 : 1)
```

To disable compression on a particular snapmirror relationship simply remove the 'compression=enable' statement from that line. Using the above example this is achieved as shown below. The use of the multipath syntax for the snapmirror relationship will not impede the operation of a normal (uncompressed) snapmirror schedule.

```
lab01:vol1 naslab01dr:vol1dr - 1-59/10 * * *
```

11.5.3 Backup

Due to the different backup products used globally within Citi, this section will only include configuration information that is specific to the NAS device itself, or backup infrastructure details that are product neutral. Support groups are advised to read through the section to familiarize themselves with the topology and NDMP requirements. NAS operations groups should document their specific backup implementations as part of their CITMP level III documentation.

11.5.3.1 NDMP Overview

Network Data Management Protocol (NDMP) controls backups and restores on the NetApp 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 up to 40TB of data because it keeps data from flowing across the network, preventing network congestion.




Figure 8 - 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*.

The third and more common approach within Citi is to use Remote NDMP to Tape over the IP network.

11.5.3.2 NDMP Considerations in storage controller scalability

At some point within the lifecycle of the data on the storage system it will require backing up. The most common form of NAS backup in Citi is NDMP. There are scaling limitations with the use of NDMP that may artificially limit the raw storage capacity of your NetApp storage system. Storage Administrators are requested to check with their local backup teams to confirm the supported maximum capacity of the storage system permitted to allow a successful NDMP backup schedule. Remember that production systems often replicate to a DR system where the data is then backed up via NDMP so these scaling requirements apply to both the production system as well as the DR system.




Figure 9 - Backup Three-party NDMP Configuration

A three way NDMP backup is where the tape drives on one Filer are used to backup other Filers, over the network. This may be considered, for redirecting NDMP backups, temporarily, if a local-NDMP configuration has been put in place and the tape drives on a Filer are down for troubleshooting/maintenance.

With Remote NDMP to Tape, NAS data is backed up over the network to a NetBackup Enterprise Media Server. Given the advantages and disadvantages of the above configuration options, the preferred configuration is to use Remote NDMP to tape. In this case, it is recommended to use the primary Gig-E interface on the Filer for backup. Local NDMP should also be considered, for Filers that have a large backup requirement that cannot be satisfied by backing up over a single GIG-E network interface.

11.5.3.3 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.

11.5.3.4 Configuring the tape devices

Log on to the NetApps file server.

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

```
d7-fas3140-01> sysconfig -t

Tape drive (L-11-2-MDS9222-1:1-9.133) IBM      03592E06
rst01l - rewind device,          format is: 900GB J1A comp
nrst01l - no rewind device,      format is: 900GB J1A comp
urst01l - unload/reload device,  format is: 900GB J1A comp
rst01m - rewind device,          format is: 1500GB E05 comp
nrst01m - no rewind device,      format is: 1500GB E05 comp
urst01m - unload/reload device,  format is: 1500GB E05 comp
rst01h - rewind device,          format is: 640GB E06
nrst01h - no rewind device,      format is: 640GB E06
urst01h - unload/reload device,  format is: 640GB E06
rst01a - rewind device,          format is: 1920GB E06 comp
nrst01a - no rewind device,      format is: 1920GB E06 comp
urst01a - unload/reload device,  format is: 1920GB E06 comp

Tape drive (L-11-2-MDS9222-1:1-7.134) IBM      03592E05
rst03l - rewind device,          format is: 300GB J1A
nrst03l - no rewind device,      format is: 300GB J1A
urst03l - unload/reload device,  format is: 300GB J1A
rst03m - rewind device,          format is: 900GB J1A comp
nrst03m - no rewind device,      format is: 900GB J1A comp
urst03m - unload/reload device,  format is: 900GB J1A comp
rst03h - rewind device,          format is: 500GB E05
nrst03h - no rewind device,      format is: 500GB E05
urst03h - unload/reload device,  format is: 500GB E05
rst03a - rewind device,          format is: 1500GB E05 comp
nrst03a - no rewind device,      format is: 1500GB E05 comp
urst03a - unload/reload device,  format is: 1500GB E05 comp

...      ...      ...      ...      ...      ...      ...
...      ...      ...      ...      ...      ...      ...
...      ...      ...      ...      ...      ...      ...
...      ...      ...      ...      ...      ...      ...

Tape drive (L-11-2-MDS9222-1:1-18.137L1) IBM    03592J1A
rst303l - rewind device,          format is: 300GB
nrst303l - no rewind device,      format is: 300GB
urst303l - unload/reload device,  format is: 300GB
rst303m - rewind device,          format is: 300GB
nrst303m - no rewind device,      format is: 300GB
urst303m - unload/reload device,  format is: 300GB
rst303h - rewind device,          format is: 300GB
nrst303h - no rewind device,      format is: 300GB
urst303h - unload/reload device,  format is: 300GB
rst303a - rewind device,          format is: 600GB comp
nrst303a - no rewind device,      format is: 600GB comp
urst303a - unload/reload device,  format is: 600GB comp
```

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

```
d7-fas3140-01*> mt -f rst01a status
Tape drive: IBM      03592E06
Status: offline
Format: 900GB J1A comp
fileno = -1 blockno = -1 resid = 0
```


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.
```

11.5.3.5 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.

11.5.3.6 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.

11.6 Advanced features – vFilers, Snapvault, Snaplock and Lockvault

11.6.1 VFiler (multistore)

11.6.1.1 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 10 - NetApp MultiStore (VFiler)

11.6.1.2 Hosting Filer / 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 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.

As of ONTAP v7.2 and above NetApp supports up to 65 usable vFilers on a single filer node [includes vFiler0 in the count] and up to 130 vFilers in a HA-pair 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 CATE Storage Engineering.

Filer Memory Size	NetApp	Citi
over 2 GB	65	65
Exception: FAS2040	16	16

Table 27 - Citi recommended vFiler limits per Storage Controller

As you provision more vFilers within a HA-pair the time taken to perform a cluster failover will increase as more resources will need to be processed during the failover event. Storage administrators should consider this behavior when planning to drastically increase the number of vFilers per storage controller in the environment.

11.6.1.3 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.

11.6.1.3.1 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.

11.6.1.3.2 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. 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.smb.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.

11.6.1.4 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.

11.6.1.5 VFiler Limitations

The following list describes the limitations of vFiler:

- VFiler management is limited to individual filer nodes. There is no centralized management utility.
- VFilers can only be managed from the command line and from Systems Manager
- 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 Citi on the NetApp NAS Filer platform.

11.6.1.6 Data ONTAP tasks allowed on VFilers

Data ONTAP allows a limited command set 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 LDAP/NFS configuration and “cifs setup” for CIFS configuration
- Configuration of vfiler LDAP use
- Managing quotas. “quota”
- Preparing vfilers for virus protection: “vscan”

11.6.1.7 SOE Vfiler Naming Standards

11.6.1.7.1 Naming Guidelines

NetApp recommends the following guidelines:

- The hostname can contain up to 31 alphanumeric ASCII characters.
 - o NOTE: Citi 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: Citi 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.

11.6.1.7.2 SOE Production VFilers

In cooperation with Storage Integration, CATE Storage Engineering determined the Citi 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

11.6.1.8 SOE Vfiler Production Configuration

11.6.1.8.1 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.

11.6.1.8.2 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 a /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*.

11.6.1.8.3 SOE Vfiler 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 interface group NIC for the vfiler.
 - **NOTE: Create VFILER with IP alias before assigning IP alias to the host filer interface group. This allows us to create vfilers without having to temporarily disable the interface group.**
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

11.6.1.8.4 SOE VFiler Create Example

1. create qtree for vfiler configuration files
`qtree create /vol/vol0/nasrut110v01tfi`
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 nasrut110v01tfi -i 150.110.93.104 /vol/vol0/ nasrut110v01tfi`
3. assign designated data volume(s) to vfiler
`vfiler add nasrut110v01tfi /nasrut110v01tfi _v1`
4. create IP address alias on host filer interface group 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`

11.6.1.8.5 SOE HA-pair Considerations

- Number of Citi supported vFilers is defined in [Table 27](#)
- 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 an HA-Pair must have a partner interface.**

11.6.1.8.6 Vfiler Graceful Shutdown Sequence

In order to avoid authentication resources from being over consumed and an extreme degradation in NetApp storage system performance, it is imperative that a vFiler be shutdown in a particular manner as described below. This procedure must be followed when either a full DR test is invoked (implying access to production is revoked for the duration of the test) or if a vFiler is being decommissioned as part of a data migration / tech refresh.

At a high level, before the vFiler is stopped, ensure the following occurs:

- File protocols are terminated (CIFS and NFS)
- Name resolution services (NIS, LDAP and CIFS) are turned off

Therefore the process to gracefully shutdown a vFiler should resemble the following:

1. cifs terminate -t 0
2. nfs off
3. options ldap.enable off
4. options nis.enable off
5. options dns.enable off
6. vfiler stop <vfiler_name>

Failure to follow the above procedure will incur a service outage if many vfilers are shutdown at the same time.

11.6.1.9 SOE Vfiler DR/COB Configuration

11.6.1.9.1 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.

11.6.1.9.2 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.

11.6.2 Snapvault

SnapVault is a replication-based disk-to-disk backup and is delivered as a licensed feature of the Data ONTAP operating system. It creates full point-in-time backup copies on disk. It uses a block incremental methodology whereby after the first full backup (referred to as a baseline copy), it transfers and stores only new or changed blocks. This minimizes data transfer over the network. Each “block incremental” backup is a full backup copy but only the new or changed blocks are added to the target storage system (often referred to as the Snapvault Secondary). With ONTAP 8.0.2 and above the snapvault destination can be assigned to a vFiler. This feature was not supported in ONTAP 8.0.1 but was supported in earlier versions of ONTAP 7.

11.6.2.1.1 License Requirements

Snapvault requires two licenses keys to be installed, one of the source storage controller (referred to as the snapvault primary) and one on the destination storage controller (referred to as the snapvault secondary). Licenses can be purchased at the time the system is procured or at a later date within the life of the system.

11.6.2.1.2 Enabling Snapvault

To enable SnapVault you must execute this command on all the primary and secondary filers:

```
primary_1> options snapvault.enable on
.
primary_N> options snapvault.enable on
secondary> options snapvault.enable on
```

11.6.2.1.3 Access Control

Access control for SnapVault data transfer requests between SnapVault primary and secondary filers is controlled by the “snapvault.access” option.

To grant unrestricted access on all the primary and secondary filers:

```
primary_1> options snapvault.access all
.
primary_N> options snapvault.access all
secondary> options snapvault.access all
```

Optionally, to grant access on a primary filer only to specific secondary filers:

```
primary_1> options snapvault.access host=<secondary_1>,<secondary_2>,...
```

Optionally, to grant access on the secondary filer only to specific primary filers (for data restore)

```
secondary> options snapvault.access host=<primary_1>,<sv_primary_2>,...
```

For this access control mechanism to work correctly the filer must be able to resolve the specified host names to IP addresses via either a local entry in the /etc/hosts file on the filer or via DNS. Alternatively, literal IP addresses (e.g. “192.168.0.1”) may be specified instead of host names to eliminate the need for hostname to IP address resolution.

11.6.2.1.4 Configuring Snapvault

11.6.2.1.4.1 About setting up SnapVault Backup

Ensure that the appropriate licenses have been installed on the source and destination system before proceeding (see section [11.6.2.1.1](#)).

On filer primary storage systems, use filer console commands to activate the SnapVault primary license, and specify the SnapVault secondary storage host.

On the SnapVault secondary storage system, use the filer console commands to license and enable SnapVault, specify the primary storage systems to back up, and start the initial snapshot backup.

On the primary filer storage systems, schedule times for local SnapVault snapshots to occur. And on the SnapVault secondary storage system, schedule times for those primary snapshots to be backed up to secondary storage.

11.6.2.1.4.2 Configuring a filer primary storage system for SnapVault

On each filer primary storage system to be backed up to the SnapVault secondary storage system, log in to that filer's console and carry out the following steps:

1. Set up the SnapVault primary license on each primary storage system to be backed up. In the filer console, enter the following command:
`license add sv_primary_license`
2. Enable SnapVault on each primary storage system to be backed up. In the filer console enter the following command:
`options snapvault.enable on`
3. Use the options snapvault.access command to specify the name of the SnapVault secondary storage system to back up to. Enter the following command:
`options snapvault.access host= snapvault_secondary`

The filer must be able to resolve the host name (snapvault_secondary) to an IP address in the /etc/hosts file, or else the filer needs to be running DNS. You can also use the literal IP address instead of the host name.

11.6.2.1.4.3 Configuring the SnapVault secondary storage system

Configure the SnapVault secondary storage system. Carry out the following steps:

1. Set up the SnapVault secondary license. In the filer console of the SnapVault secondary system, enter the following command:
`license add sv_secondary_license`
2. Enable SnapVault. In the filer console, enter the following command:
`options snapvault.enable on`
3. Use the options snapvault.access command to specify the names of the primary storage systems to back up and restore. Enter the following command:
`options snapvault.access host= snapvault_primary1, snapvault_primary2 ...`

4. For each qtree on the filer primary storage systems to be backed up, use the snapvault start command line to execute an initial baseline copy of the qtree from the primary to the secondary storage system. On each command line, specify the primary storage filer, volume, and qtree; and the secondary storage host, SnapLock volume, and qtree. Use the -S prefix to indicate the source qtree path. For example, to start a baseline copy of qtrees tree_a, tree_b, and tree_c to a SnapLock Compliance volume called sv_vol, use the following commands:

```

snapvault start -S filer_a:/vol/vol1/tree_a sv_filerb:/vol/sv_vol/tree_a
snapvault start -S filer_a:/vol/vol1/tree_b sv_filerb:/vol/sv_vol/tree_b
snapvault start -S filer_a:/vol/vol1/tree_c sv_filerb:/vol/sv_vol/tree_c

```

11.6.2.1.4.4 Creating Snapvault Jobs

SnapVault protects at the granularity of qtrees on the primary filers. You need to configure SnapVault with the list of primary filer qtrees to protect and to associate these primary filer qtrees with secondary filer qtrees into which the primary qtree data will be stored. This is the purpose of the “snapvault start” command on the secondary filer.

For example, let's say we have two primary filers, *toaster* and *fridge*. The primary filer *toaster* has qtrees */vol/proj/bread* and */vol/proj/butter* and the primary filer *fridge* has qtrees */vol/proj/juice* and */vol/proj/beer*. The plan is to protect these four qtrees with volume *sv_proj* on the SnapVault secondary filer secondary. To get this going, you would enter the following four commands on secondary:

```

secondary> snapvault start -S toaster:/vol/proj/bread /vol/sv_proj/bread
secondary> snapvault start -S toaster:/vol/proj/butter /vol/sv_proj/butter
secondary> snapvault start -S fridge:/vol/proj/juice /vol/sv_proj/juice
secondary> snapvault start -S fridge:/vol/proj/beer /vol/sv_proj/beer

```

These commands will create the qtrees on the secondary filer (the volume *sv_proj* must already exist), start an initial baseline (level-0) transfer (aka “initialize”) from the primary filers to the secondary filer and store the configuration persistently in the system registry. To confirm that data is being transferred from the primary filers to the secondary filer the “snapvault status” command can be used:

```

secondary> snapvault status
Snapvault secondary is ON.
Source Destination State Lag Status
toaster:/vol/proj/bread secondary:/vol/sv_proj/bread Uninitialized -
Transferring (3 MB done)
toaster:/vol/proj/butter secondary:/vol/sv_proj/butter Uninitialized -
...
...
Transferring (2 MB done)
fridge:/vol/proj/juice secondary:/vol/sv_proj/juice Uninitialized -
Transferring (1 MB done)
fridge:/vol/proj/beer secondary:/vol/sv_proj/beer Uninitialized -
Transferring (1 MB done)

```

To view the primary filer configuration, use the “snapvault status -c” command:

```

secondary> snapvault status -c
Snapvault secondary is ON.
/vol/sv_proj/bread source=toaster:/vol/proj/bread kbs=unlimited tries=2
/vol/sv_proj/butter source=toaster:/vol/proj/butter kbs=unlimited tries=2
/vol/sv_proj/juice source=fridge:/vol/proj/juice kbs=unlimited tries=2
/vol/sv_proj/beer source=fridge:/vol/proj/beer kbs=unlimited tries=2

```

11.6.2.1.4.5 Initial SnapShot Schedule Configuration

On both the filer primary storage systems and the SnapVault secondary storage system, configure a snapshot schedule. To set up a snapshot schedule, complete the following steps:

1. **On the filer primary storage systems:** On each filer primary storage system that contains qtrees to be backed up to a SnapVault secondary storage system, use the `snapvault snap sched` command to schedule sets of SnapVault snapshots on each volume containing the backed-up qtrees.

For each set of snapshots, specify volume name, snapshot base name (for example: "sv_hourly," or "sv_nightly," and so on), number of SnapVault snapshots to store locally, and the days and hours to execute the snapshots. For example:

```
snapvault snap sched vol1 sv_weekly 1@sat@19
snapvault snap sched vol1 sv_nightly 2@mon-fri@19
snapvault snap sched vol1 sv_hourly 11@mon-fri@7-18
```

Note: When specifying the SnapVault snapshot base name, avoid using "hourly," "nightly," or "weekly." Such naming will conflict with the non-SnapVault `snap sched` snapshots.

2. **On the SnapVault secondary storage system:** For each SnapVault volume snapshot set that you scheduled on your filer primary storage systems (see Step 1), enter the following `snapvault snap sched -x` command to schedule a set of transfers to and subsequent snapshots of the SnapVault secondary storage system:

```
snapvault snap sched -x sec_vol snap_name count[@ day_list] [@ hour_list]
```

count is the number of snapshots you want to retain for this set.

@day_list is a comma-separated list that specifies the days on which a new snapshot is created for this set.

@hour_list specifies the hours at which a new snapshot is created for this set.

Example: The following commands schedule transfers and snapshots every Saturday at 8 p.m., every Monday through Friday at 8 p.m., and every hour, Monday through Friday from 7 a.m. to 7 p.m., respectively:

```
snapvault snap sched -x sv_vol sv_weekly 8@sat@20
snapvault snap sched -x sv_vol sv_nightly 7@mon-fri@20
snapvault snap sched -x sv_vol sv_hourly 11@mon-fri@7-19
```

3. **On the SnapVault primary and secondary storage system:** To turn off the regular snapshot schedule, enter the command:

```
snap sched volume 0 0 0
```

It is usually unnecessary and wasteful of storage resources to have both regular scheduled snapshots (those scheduled using the `snap sched` command) and SnapVault scheduled snapshots (those scheduled using the `snapvault snap sched` command). If you have both schedules, you can turn off and clear the regular snapshot schedule to save resources.

See section [14.8.1](#) for a detailed `snapvault` snapshot schedule configuration example.

11.6.2.1.5 User-Triggered Snapvault Snapshot Creation

Forced creation of SnapVault snapshot is possible if it is not the schedule time for the snapshot to be created. The snapshot creation will proceed just as if triggered by the snapshot scheduler. This means that primary filers will simply create snapshots, whereas secondary filers will kick-off the transfers from the primary filers, if their schedules have been set up to do so, before creating a snapshot.

On the primary filers, enter the command:

```
toaster> snapvault snap create proj sv_hourly
fridge> snapvault snap create proj sv_hourly
```

Then, on the secondary filer, enter the command:

```
secondary> snapvault snap create sv_proj sv_hourly
```

11.6.2.1.6 System Monitoring

It is important to monitor the operation of the systems to ensure that backups are proceeding correctly. Error conditions such as failed/slow network connections and out of disk space conditions can prevent the scheduled completions of backup/mirroring operations from succeeding. Also, there is a potential for operator error or malicious activity causing the backups/mirroring operations to fail. Most of the monitoring activity discussed in this section can be accomplished using the latest CATE Storage Engineering certified version of NetApp Unified Manager.

11.6.2.1.6.1 Status

`snapmirror status` and `snapvault status` display the status of each relationship. The meanings of the status keywords are described here:

- *Idle*: The relationship is currently not involved in a transfer
- *Transferring*: The relationship is currently involved in a transfer
- *Pending*: The relationship has pending transfer
- *Aborting*: The relationship is aborting the currently active transfer
- *Quiescing*: The relationship is waiting for other transfers in the secondary volume to complete, OR the relationship could be involved in a rollback operation
- *Resyncing*: The relationship is currently involved in a resync operation

11.6.2.1.6.2 Recommended Monitoring

Outlined here is what to monitor to ensure successful backup/mirroring operations.

On a **daily** basis:

- On the SnapVault Secondary, do a `snapvault status` and ensure that the lag times are appropriate
- If no transfers are active at the time, the status should be *Idle*. Meanings of the other statuses are described above.
- Do a `snapmirror status` on the “Tertiary” system and ensure that lag times are appropriate.

On a **weekly** basis:

- Check the free disk space on the log volume and the secondary volumes and make sure that it is adequate for the coming weeks.

Look at the current `/vol/vol0/etc/log/snapmirror` log file on the “Tertiary” Nearstore for errors (Retry or Abort) indicating that the transfers may have encountered problems.

11.6.3 Snaplock

Snaplock is a licensed feature of the Data ONTAP operating system that creates nonrewritable and non eraseable volumes to prevent files from being altered or deleted until a set retention date. Snaplock is available in two variants: Compliance, for stringent government requirements; and Enterprise, for adhering to organizational best practices in a trusted administrator model.

NetApp has re-introduced support for SnapLock with ONTAP 8.1 and above (7-mode) as it was not available in ONTAP 8.0. Many of the features relating to SnapLock are used in the LockVault (SnapVault + SnapLock) solution, and so in an effort to avoid duplication on content please refer to the following section.

11.6.4 Lockvault

LockVault is a disk-based regulatory compliance solution for unstructured data. It delivers a regulatory solution by making backups compliant and by saving only block-level incremental changes. Block incremental daily Snapshot copies are backed up to secondary storage (using SnapVault) and protected against modification or deletion until a specified retention date (using SnapLock). In addition, a ComplianceJournal (file log) tracks changes between Snapshot copies and stores them on a WORM volume, so the log cannot be modified either.

In addition to the information in this chapter, refer to section [14.8](#) for further Lockvault related content.

11.6.4.1 Overview

WORM data resides on SnapLock storage volumes that are administered much like regular storage volumes. SnapLock volumes operate in WORM mode and support standard file system semantics. Data on a SnapLock volume can be created and committed to WORM state by transitioning the data from a writable state to a read-only state. Marking a currently writable file as read-only on a SnapLock volume commits the data as WORM. This commit process prevents the file from being altered or deleted by applications, users, or administrators. Data that is committed to WORM state on a SnapLock volume is immutable and cannot be deleted. The only exceptions are empty directories and files that are not committed to a WORM state. Additionally, once directories are created, they cannot be renamed.

In order for a magnetic WORM storage system to enforce retention in a regulated environment it must make any time related retention decisions using a secure time base. A secure time base is an essential component to ensure that retained data cannot be deleted prematurely simply by subverting the regular clock of the storage system. It is also important, however, that the storage system provide a time source for general operation that can be synchronized to a common reference time used inside the data center.

In order to meet these requirements SnapLock provides a secure Compliance Clock (CC) that is implemented in software and runs independently of the regular system clock. The CC will track the regular system clock and remain accurate but has a tightly defined drift rate which prevents the CC from being subverted by abrupt changes to the system time. The CC will display significant inertia to any changes in the system time and will take time to "catch up". This catching up also applies to **system downtime** because when the system is down the CC is not being updated.

11.6.4.1.1 Required Components

The following hardware and software components comprise the solution that is described in this document.

- **FAS system with NearStore personality license** – ATA based storage system that provides WORM compliant storage via integrated controls. (Note: there may be some occasions, to be reviewed on a case by case basis, where SnapLock and/or LockVault functionality is required but with the performance offered by Fiber Channel drives due to application I/O requirements).
- 'LockVault' – The software solution which includes modifications to SnapLock and SnapVault to provide a regulatory compliant backup/archive solution.
- SnapVault – Product that performs that actual snapshot type backup of the source file system for compliant backup.
- WORM Log – A log file that is WORM compliant (cannot be deleted or modified) that records system activities. One can append to this log, but existing records cannot be altered or removed.
- SnapLock – NetApp software that provides the integrated controls to make the volumes WORM compliant.
- SnapMirror – NetApp software that provides a second copy of the data and the WORM log on a second NearStore system also setup for WORM compliance.
- Compliance Clock – A secure secondary clock within Data ONTAP that verifies that the system time is accurate against internal security protocols to verify accurate time for file retention and prevents deletions of files that are unexpired. This feature is enabled when the SnapLock license is installed in ONTAP.

11.6.4.1.2 Applicable Environments

The NetApp solution is a product that is best suited for applications that require compliant file system support that can be centrally managed. Applications in scope also require high access rates combined with requirements to meet regulatory compliance in data storage.

- Simplifies and increases the reliability of data backup and recovery
- Significantly reduces the time to recover data, vs. recovery from traditional tape sources
- Performs incremental, forever backups to use CPU, network and storage resources more efficiently
- Stores each backup as a space-optimized full copy by sharing common file blocks across all backups
- Improves the availability of data
- Provides automated operation to lower the cost of backup administration
- Meets the requirements for regulatory compliant backups and archive for file systems

11.6.4.2 Installation

11.6.4.2.1 Basic Configuration

The basic configuration is one or more filers running Data ONTAP 7.3 or later as primaries with a NearStore licensed system running a SnapVault and SnapLock licensed version of Data ONTAP. Data from the filers is backed up on the destination using SnapVault. A regulatory compliant configuration also includes a second destination to which the first destination system is mirrored using Volume SnapMirror. The second destination system is also running a SnapVault and SnapLock licensed version of Data ONTAP.

The following picture shows the typical configuration using SnapVault and SnapMirror. Note, that both the SnapVault target volume(s) and the WORM log volume are SnapMirrored, since both the logs and the data are necessary for compliance.




Figure 11 - Lockvault Implementation

11.6.4.2.1.1 Primary Filer Configuration

Primary filer configuration is no different than for a standard SnapVault configuration. Nothing special needs to be done just because the secondary systems (Nearstore1 and Nearstore2 in the figure) have WORM archive capabilities.

11.6.4.2.1.2 Secondary NearStore Configuration

There are four major differences between configuring a regular SnapVault secondary and configuring a LockVault secondary.

- Initializing the compliance clock
- Creating the SnapVault target volumes as worm volumes
- Creating a worm log volume and configuring the worm log volume.
- Setting appropriate retention periods for the worm target and log volumes

11.6.4.2.1.3 Tertiary NearStore Configuration

There are some differences between the way the “Tertiary” NearStore and the Secondary NearStore are configured. Fundamentally, with Volume SnapMirror between two filers, the secondary volumes are initially created as regular volumes and take on the properties of “worm” volumes after the SnapMirror initialize is complete. Also the minimum, maximum, and default snapshot retention periods are inherited from the SnapMirror primary volumes. Therefore to configure the “tertiary” systems you just create regular volumes to mirror the SnapVault target volumes and the log volume. You must also license SnapLock compliance on the tertiary system.

Note: The system time on all LockVault filers and NearStore systems should be set and verified before and upgrades to Data ONTAP are performed.

11.6.4.2.2 Licensing Snaplock

SnapLock can be licensed in one of two different flavors depending on the regulatory environment and customer requirements. At the present time these licenses are mutually exclusive and cannot co-exist in the system simultaneously. The behavior of a SnapLock volume, however, does not depend on the type of the installed license.

11.6.4.2.2.1 Snaplock Compliance

SnapLock Compliance (SLC) is the strictest form of SnapLock and it designed for use in strictly regulated environments such as those regulated by the SEC. SLC not only enforces the WORM protection of files to the clients but also places tight restrictions on the ability of the storage administrator with regard to any operations that may result in modification or erasure of retained WORM records. Most notably this includes being unable to destroy a volume. This can be thought of as an "untrusted administrator" mode of operation.

In order to create and utilize SLC volumes you must install an optional license for the "snaplock" feature of Data ONTAP. This license is installed using the following procedure from the filer administrative console.

```
filer> license add <code>
A snaplock site license has been installed.
SnapLock(tm) Compliance enabled.
```

Because of the strict limitations placed on the storage administrator with SLC volumes it is recommended that customers perform basic evaluation testing using the more flexible SnapLock Enterprise. This can eliminate the risk of inadvertent long term or permanent allocation of storage resources (e.g. indestructible volumes) during the evaluation period.

When using the integrated SnapLock / SnapVault deployment described in this chapter of the Solutions Guide, you **MUST** use SnapLock Compliance.

11.6.4.2.2.2 Snaplock Enterprise

SnapLock Enterprise (SLE) is a more flexible form of SnapLock designed for use in less strictly regulated environments. SLE enforces the WORM protection of files and snapshots to all clients but leaves the storage administrator free to manage the system with very few restrictions. The storage administrator is permitted to perform operations that may result in the loss of data such as volume destruction.

```
filer> license add <code>
A snaplock site license has been installed.
SnapLock(tm) Enterprise enabled.
```

SnapLock Enterprise (SLE) is **NOT** appropriate for use in strictly regulated environments such as those governed by SEC Rule. 17a-4 as it uses a "trusted administrator" model of operation and provides storage administrators with the ability to destroy retained content.

11.6.4.2.3 Initializing the Compliance Clock

The first step is to initialize the compliance clock. It is the recommendation of CATE Storage Engineering to not create any volumes before configuring the compliance clock. First, ensure that the regular clock is set correctly. You can verify the setting of the regular clock by entering date on the filer console. Then initialize the compliance clock as follows.

Until the CC is initialized and running, SnapLock will not allow deletion of any retained content and will also prevent other retention operations such as the creation and retention of WORM snapshots.

The Compliance Clock (CC) can be initialized only once for the system so **extreme care** should be exercised at install time to ensure that the clock time is set correctly. Make sure that the system time and time zone are correctly set and everything should go smoothly as the CC can only be set from the system time to prevent any data input errors.

The procedure for setting the CC is as follows.

```
filer> snaplock clock initialize

*** WARNING: YOU ARE INITIALIZING THE SECURE COMPLIANCE CLOCK ***

You are about to initialize the secure Compliance Clock of this
system to the current value of the system clock. This procedure
can be performed ONLY ONCE on this system so you should ensure
that the system time is set correctly before proceeding.

The current local system time is: Wed Feb 4 23:38:58 GMT 2004

Is the current local system time correct? y

Are you REALLY sure you want initialize the Compliance Clock? y

Compliance Clock: Wed Feb 4 23:39:27 GMT 2004
```

The current CC time can be queried as follows.

```
filer> snaplock clock status
Compliance Clock: Wed Feb 4 23:42:39 GMT 2004
```

11.6.4.2.4 Enabling Snapvault

To enable SnapVault you must execute this command on all the primary and secondary filers:

```
primary_1> options snapvault.enable on
.
.
primary_N> options snapvault.enable on
secondary> options snapvault.enable on
```

11.6.4.2.5 Access Control

Access control for SnapVault data transfer requests between SnapVault primary and secondary filers is controlled by the “snapvault.access” option.

To grant unrestricted access on all the primary and secondary filers:

```
primary_1> options snapvault.access all
.
.
primary_N> options snapvault.access all
secondary> options snapvault.access all
```

Optionally, to grant access on a primary filer only to specific secondary filers:

```
primary_1> options snapvault.access host=<secondary_1>,<secondary_2>,...
```

Optionally, to grant access on the secondary filer only to specific primary filers (for data restore)

```
secondary> options snapvault.access host=<primary_1>,<sv_primary_2>,...
```

For this access control mechanism to work correctly the filer must be able to resolve the specified host names to IP addresses via either a local entry in the /etc/hosts file on the filer or via DNS. Alternatively, literal IP addresses (e.g. “192.168.0.1”) may be specified instead of host names to eliminate the need for hostname to IP address resolution.

11.6.4.3 Configuration

CATE Storage Engineering recommends the following configuration options:

- The NearStore storage systems should be dedicated for regulatory compliance and not used for other purposes.
- At least two spare disks should be assigned.
- Aggregates should be configured as discussed in section [11.3.3](#) and use RAID-DP.
- Redundant network connections will be used on the NearStore systems

Retention periods for volumes used to meet SEC guidelines should be set to the following:

- Minimum Retention Period should be set to 3 years.
- Maximum Retention Period should be set to 7 years.
- Default Retention Period should be set to ‘minimum’.

11.6.4.3.1 Creating Snaplock Aggregates

Unlike snaplock using traditional volumes, the snaplock configuration is defined at the aggregate level for flexible volumes (however the retention periods are still defined at the volume level). Snaplock aggregates are created in a similar way to regular aggregates except they require the `-L` parameter with the `aggr create` command. Use the compliance attribute to define the snaplock licensed applied to the aggregate.

You can only create Snaplock Aggregates using the CLI as the `-L` option is not available in Filerview.

The following example will create a Snaplock Compliance aggregate, with 18 disks in the aggregate configured as RAID-DP.

```
aggr create aggr2 -L compliance -t raid_dp 18
```

Use the `aggr status` command to display the properties of each aggregate, thereby demonstrating which aggregates are configured as Snaplock Aggregates.

```
filer1> aggr status
Aggr State      Status      Options
aggr0 online     raid_dp, aggr  root
aggr1 online     raid_dp, aggr  snaplock_compliance
```

11.6.4.3.2 Creating Snaplock Flexible Volumes

As previously mentioned, snaplock volumes inherit their compliance type from the aggregate in which they are created. As such a snaplock flexible volume is created in the same way as a regular flexible volume.

Unlike Lockvault in ONTAP 6.5, the `-L` parameter attribute `compliance_log` has been deprecated. Refer to section [11.6.4.3.3](#) on how to assign the worm log volume.

11.6.4.3.3 Assigning the WORM log Volume

One of the volumes needs to be defined as the WORM log volume that SnapVault will know which volume is to be used for logging.

This option name has changed in ONTAP 7.1 or above, compared to the Lockvault PVR implementation in ONTAP 6.5.

```
snapvault.lockvault_log_volume <vol_name>
```

11.6.4.3.4 Setting Volume Retention Parameters

Now set appropriate retention periods for the worm target and log volumes. You will need to set appropriate minimum, maximum, and default retention periods for the volumes.

BE CAREFUL WHEN SETTING RETENTION PERIODS. UNEXPIRED FILES CANNOT BE DELETED AND VOLUMES (AND THEREBY AGGREGATES) CONTAINING UNEXPIRED FILES CANNOT BE DESTROYED AND REUSED FOR OTHER PURPOSES.

The three retention periods can be set as follows:

```
vol options <volname> snaplock_minimum_period <period>
vol options <volname> snaplock_default_period <period>
vol options <volname> snaplock_maximum_period <period>
```

Note: you must set the retention periods on EACH of the SnapLock compliance volumes. SnapLock volumes provide options for configuring a very basic retention policy for both WORM files and snapshots. These options include a default, minimum and maximum retention periods for the volume. A simple way to check the retention setting on each volume is to use the `vol options` command:

```
vol options <volname>
```

It is good practice to configure these retention parameters before bringing a SnapLock volume into service to ensure that the desired retention policy is correctly reflected. The default values, especially for SLC volumes, as mandated by some regulatory bodies, are usually NOT what the customer would choose to deploy.

The default retention period on the secondary volumes is used by SnapVault to establish the retention period for snapshots if no “-o retention_period” is specified in the `snapvault snap sched` or if a manual `snapvault update` is performed. The default retention period on the log volume is used to establish the retention period for the log files on the volume. The minimum and maximum retention periods are used to bind the specified retention periods. The minimum and maximum retention periods are also used to bound the retention period for the dump image and log which may be generated during a `snapvault -r (resync)`. The retention period for the `snapvault -r dump` image is set by finding the largest retention period for any file in the dataset, subject to the minimum and maximum settings.

11.6.4.3.4.1 Period Specification Syntax

The syntax for specifying a retention period is very simple. Periods can be specified in increments of days, months or years and will be applied in a calendar correct manner. That is, for example, a WORM file created on 1st February with one-month retention period will expire on 1st March. The syntax for a retention period is `<count>{d|m|y}` with some examples being “5d” for five days, “18m” for eighteen months and “6y” for six years.

11.6.4.3.4.2 Minimum Retention Period

The minimum retention period specifies the minimum allowable retention period for any file or snapshot committed to WORM state on the SnapLock volume. This option is useful in regulatory environments to ensure that applications or users do not intentionally or unintentionally assign non-compliant retention periods to retained records. Any file or snapshot committed to WORM state with a retention period less than this minimum will automatically have this minimum retention period assigned. The minimum retention period takes precedence over the default retention period.

The procedure for setting the minimum retention period is as follows (example 6 month period):

```
filer> vol options wormvol snaplock_minimum_period 6m
```

If not explicitly configured, the minimum retention period is 0.

11.6.4.3.4.3 Maximum Retention Period

The maximum retention period specifies the maximum allowable retention period for any file or snapshot committed to WORM state on the SnapLock volume. This option is useful in regulatory environments to ensure that applications or users do not intentionally or unintentionally assign excessive retention periods to retained records. Any file or snapshot committed to WORM state with a retention period greater than this maximum will automatically have this maximum retention period assigned. The maximum retention period takes precedence over the default retention period.

The procedure for setting the maximum retention period is as follows (example 3 year period):

```
filer> vol options wormvol snaplock_maximum_period 3y
```

For SnapLock Compliance volumes, the default retention period is equal to the maximum retention period of 30 years. If either the maximum retention period or the default retention period are not changed, nor are explicitly configured, WORM files that received the default retention period cannot be deleted for 30 years.

11.6.4.3.4.4 Default Retention Period

The default retention period specifies the retention period that will be assigned to any file or snapshot committed to WORM state on the SnapLock volume without an explicitly assigned retention period. This option is useful to ensure that if users or applications fail to assign a retention period that a default period will be assigned.

The procedure for setting the default retention period is as follows (example 1 year period):

```
filer> vol options wormvol snaplock_default_period 1y
```

The default retention period also accepts the values "min" and "max" which instruct the storage system to use the current minimum or maximum retention periods respectively as the default retention date. These can be specified as follows:

```
filer> vol options wormvol snaplock_default_period min
filer> vol options wormvol snaplock_default_period max
```

It is important to take note of the default value of this parameter when not explicitly configured, especially with SnapLock Compliance volumes.

For SnapLock Compliance volumes, the default retention period is equal to the maximum retention period of 30 years. If either the maximum retention period or the default retention period are not changed, nor are explicitly configured, WORM files that received the default retention period cannot be deleted for 30 years in respect to SEC mandate.

For SnapLock Enterprise volumes this parameter defaults to "min" which is, by default, zero.

11.6.4.3.4.5 Extending the Retention Period

As of ONTAP 7.1 and above you can extend the retention period of files on a LockVault volume. This is achieved with the `snapvault snap retain` command.

This command has changed from pre-7.1 releases of ONTAP. In earlier versions of ONTAP the command to extend the retention period on a volume was `snap retain`.

The following example would set the retention date for Snapshot “wormsnap” of volume “wormvol” to 5 years from the current time:

```
filer> snapvault snap retain wormvol wormsnap 5y
```

Snapvault snap retain has the following characteristics:

- Can be executed on a per volume basis
- Will **fail** if the new expiration date is **earlier than** the existing expiration date

11.6.4.4 Operational Procedures

11.6.4.4.1 Capacity Planning

The following are aspects of capacity planning:

- SnapVault secondary storage system volume size
- SnapVault log volume size
- Number of qtrees backed up to each volume
- SnapVault schedule

11.6.4.4.1.1 Guidelines for the secondary storage system volume size

The following are conservative guidelines you can use to estimate the size of SnapVault secondary storage system volumes for growth over one year, depending on how often you back up data.

- Three percent (3%) growth every day
- Five percent (5%) growth every week
- Ten percent (10%) growth every month
- Seven percent (7%) monthly growth rate, compounded to 100 percent growth every year

Daily backup example: Assuming a starting baseline transfer of 1 GB and 100 percent growth of new data for the year (1 GB). The changed size due to daily growth if changes are made every weeknight is as follows:

Size due to daily growth = (number of days x .03) - (new data)

= (250 x .03) - 1 GB = 7.5 GB - 1 GB = 6.5 GB of change

Therefore, for 1 GB of source data, you need 1 GB baseline + 1 GB new data + 6.5 GB of changed data = 8.5 GB secondary system volume. In other words, you must be able to scale the secondary system volume to about 8 times the current data set size within one year.

Weekly backup example: Using the same assumptions as the daily backup example, the changed size due to weekly growth if changes are made every week is as follows:

Size due to weekly growth = (number of weeks x .05) - (new data)

= (52 x .05) - 1 GB = 2.6 GB - 1 GB = 1.6 GB of change

Using weekly backups, for 1 GB of source data, you need 1 GB baseline + 1 GB new data + 1.6 GB of changed data = 3.6 GB, or a scaling factor of a little over 3.5 times the current data set size with one year.

11.6.4.4.1.2 Guidelines for the log volume size

NetApp recommends the following guidelines to estimate the log volume size:

- Baseline transfer = number of inodes per volume x 256 bytes where 256 bytes is an estimated size of the log entry.
- Incremental transfers = number of inodes x .03 x 250 snapshots x 256 bytes, where .03 represents the rate of change.

Use these guidelines to estimate the size of the log data associated with each SnapLock volume and create a log volume equal to the total of the combined estimated sizes. You can always grow the size of the flexible volume to compensate for growth in the WORM log data.

11.6.4.4.1.3 How many qtrees should be SnapVaulted to each volume?

Use the following guidelines when backing up qtrees to a volume using SnapVault:

- Because of performance constraints, a maximum of 16 primary qtrees should be backed up to a single secondary volume.
- To gain optimum performance, the number of qtrees backed up to a volume should be approximately six.
- If you are backing up volumes to qtrees using SnapVault, a maximum of 16 volumes should be backed up to a volume.

Note: You can mix backed up volumes and qtrees on the same secondary volume, if desired.

Note: All SnapVault snapshots to the same target volume must have the same name to setup a single SnapVault job from multiple sources.

11.6.4.4.1.4 Snapvault schedule

CATE Storage Engineering recommends that the SnapVault schedule be set nightly. Other schedules can be set on individual SnapVault schedules based on business need.

Based on the initial 1AM recommendation for Project EARP – CATE Storage Engineering recommends the Primary filer schedule is set to 1AM and the Secondary R200 Target be set to 2AM. This may change based on business and technical requirements. Database Dumps may also require a larger window.

The only guidelines for scheduling SnapVault transfers are to not overload the primary storage system or the network and to make sure that the scripts writing data to the primary filer are completed. You must consider the overall load and time to complete all transfers.

Filers and NearStore systems have limits on the number of concurrently transfers they allow. You can schedule more than the concurrent transfer limit for your storage system. If you do schedule more than the limit of concurrent transfers for your filer or NearStore system, the transfers that are over the limit are deferred and then automatically started as previously scheduled transfers finish.

11.6.4.4.2 Configuring Snapvault

11.6.4.4.2.1 About setting up SnapVault Backup

Setting up SnapVault backup means preparing the primary storage systems and SnapVault secondary storage systems to perform backup tasks.

You must have separate SnapVault licenses for the primary storage and the secondary storage to use SnapVault backup.

On filer primary storage systems, use filer console commands to activate the SnapVault primary license, and specify the SnapVault secondary storage host.

On the SnapVault secondary storage system, use the filer console commands to license and enable SnapVault, specify the primary storage systems to back up, and start the initial snapshot backup.

On the primary filer storage systems, schedule times for local SnapVault snapshots to occur. And on the SnapVault secondary storage system, schedule times for those primary snapshots to be backed up to secondary storage.

11.6.4.4.2.2 Configuring a filer primary storage system for SnapVault

On each filer primary storage system to be backed up to the SnapVault secondary storage system, log in to that filer's console and carry out the following steps:

1. Set up the SnapVault primary license on each primary storage system to be backed up. In the filer console, enter the following command:
`license add sv_primary_license`
2. Enable SnapVault on each primary storage system to be backed up.
In the filer console enter the following command:
`options snapvault.enable on`
3. Use the options snapvault.access command to specify the name of the SnapVault secondary storage system to back up to. Enter the following command:
`options snapvault.access host= snapvault_secondary`

The filer must be able to resolve the host name (snapvault_secondary) to an IP address in the /etc/hosts file, or else the filer needs to be running DNS. You can also use the literal IP address instead of the host name.

11.6.4.4.2.3 Configuring the SnapVault secondary storage system

Configure the SnapVault secondary storage system. Carry out the following steps:

1. Set up the SnapVault secondary license. In the filer console of the SnapVault secondary system, enter the following command:
`license add sv_secondary_license`
2. Enable SnapVault. In the filer console, enter the following command:
`options snapvault.enable on`
3. Enable the SnapVault log volume. Enter the following command:
`snapvault.lockvault_log_volume <vol_name>`
 vol_name is the SnapLock log volume.
4. Use the options snapvault.access command to specify the names of the primary storage systems to back up and restore. Enter the following command:
`options snapvault.access host= snapvault_primary1, snapvault_primary2 ...`
5. For each qtree on the filer primary storage systems to be backed up, use the snapvault start command line to execute an initial baseline copy of the qtree from the primary to the secondary storage system. On each command line, specify the primary storage filer, volume, and qtree; and the secondary storage host, SnapLock volume, and qtree. Use the -S prefix to indicate the source qtree path. For example, to start a baseline copy of qtrees tree_a, tree_b, and tree_c to a SnapLock Compliance volume called sv_vol, use the following commands:
`snapvault start -S filer_a:/vol/vol1/tree_a sv_filerb:/vol/sv_vol/tree_a`
`snapvault start -S filer_a:/vol/vol1/tree_b sv_filerb:/vol/sv_vol/tree_b`
`snapvault start -S filer_a:/vol/vol1/tree_c sv_filerb:/vol/sv_vol/tree_c`

11.6.4.4.2.4 Creating Snapvault Jobs

SnapVault protects at the granularity of qtrees on the primary filers. You need to configure SnapVault with the list of primary filer qtrees to protect and to associate these primary filer qtrees with secondary filer qtrees into which the primary qtree data will be stored. This is the purpose of the “snapvault start” command on the secondary filer.

For example, let's say we have two primary filers, *toaster* and *fridge*. The primary filer *toaster* has qtrees */vol/proj/bread* and */vol/proj/butter* and the primary filer *fridge* has qtrees */vol/proj/juice* and */vol/proj/beer*. The plan is to protect these four qtrees with volume *sv_proj* on the SnapVault secondary filer secondary. To get this going, you would enter the following four commands on secondary:

```
secondary> snapvault start -S toaster:/vol/proj/bread /vol/sv_proj/bread
secondary> snapvault start -S toaster:/vol/proj/butter /vol/sv_proj/butter
secondary> snapvault start -S fridge:/vol/proj/juice /vol/sv_proj/juice
secondary> snapvault start -S fridge:/vol/proj/beer /vol/sv_proj/beer
```

These commands will create the qtrees on the secondary filer (the volume `sv_proj` must already exist), start an initial baseline (level-0) transfer (aka “initialize”) from the primary filers to the secondary filer and store the configuration persistently in the system registry. To confirm that data is being transferred from the primary filers to the secondary filer the “snapvault status” command can be used:

```
secondary> snapvault status
Snapvault secondary is ON.
Source Destination State Lag Status
toaster:/vol/proj/bread secondary:/vol/sv_proj/bread Uninitialized -
Transferring (3 MB done)
toaster:/vol/proj/butter secondary:/vol/sv_proj/butter Uninitialized -
...
...
Transferring (2 MB done)
fridge:/vol/proj/juice secondary:/vol/sv_proj/juice Uninitialized -
Transferring (1 MB done)
fridge:/vol/proj/beer secondary:/vol/sv_proj/beer Uninitialized -
Transferring (1 MB done)
```

To view the primary filer configuration, use the “snapvault status -c” command:

```
secondary> snapvault status -c
Snapvault secondary is ON.
/vol/sv_proj/bread source=toaster:/vol/proj/bread kbs=unlimited tries=2
/vol/sv_proj/butter source=toaster:/vol/proj/butter kbs=unlimited tries=2
/vol/sv_proj/juice source=fridge:/vol/proj/juice kbs=unlimited tries=2
/vol/sv_proj/beer source=fridge:/vol/proj/beer kbs=unlimited tries=2
```

11.6.4.4.2.5 Initial SnapShot Schedule Configuration

On both the filer primary storage systems and the SnapVault secondary storage system, configure a snapshot schedule. To set up a snapshot schedule, complete the following steps:

1. **On the filer primary storage systems:** On each filer primary storage system that contains qtrees to be backed up to a SnapVault secondary storage system, use the snapvault snap sched command to schedule sets of SnapVault snapshots on each volume containing the backed-up qtrees.

For each set of snapshots, specify volume name, snapshot base name (for example: “sv_hourly,” or “sv_nightly,” and so on), number of SnapVault snapshots to store locally, and the days and hours to execute the snapshots. For example:

```
snapvault snap sched vol1 sv_weekly 1@sat@19
snapvault snap sched vol1 sv_nightly 2@mon-fri@19
snapvault snap sched vol1 sv_hourly 11@mon-fri@7-18
```

Note: When specifying the SnapVault snapshot base name, avoid using “hourly,” “nightly,” or “weekly.” Such naming will conflict with the non-SnapVault snap sched snapshots.

2. **On the SnapVault secondary storage system:** For each SnapVault volume snapshot set that you scheduled on your filer primary storage systems (see Step 1), enter the following `snapvault snap sched -x` command to schedule a set of transfers to and subsequent snapshots of the SnapVault secondary storage system:

```
snapvault snap sched -x sec_vol snap_name count[@ day_list] [@ hour_list]
```

count is the number of snapshots you want to retain for this set.

@day_list is a comma-separated list that specifies the days on which a new snapshot is created for this set.

@hour_list specifies the hours at which a new snapshot is created for this set.

Caution: Scheduling Snapvault snapshots of a SnapLock Compliance volume automatically creates WORM snapshots. Before scheduling SnapVault snapshots, ensure that the SnapLock default retention period is set correctly or that you explicitly set the retention period when you schedule SnapVault snapshots. Snapshot base names on the primary and secondary systems must match, but snapshot times and number of stored snapshots can differ. The `-x` parameter causes SnapVault to copy new or modified data from the primary qtrees to their associated qtrees on the secondary storage system. After all the secondary qtrees on the specified volume have been updated, SnapVault then creates a snapshot of this volume for archiving.

Example: The following commands schedule transfers and snapshots every Saturday at 8 p.m., every Monday through Friday at 8 p.m., and every hour, Monday through Friday from 7 a.m. to 7 p.m., respectively:

```
snapvault snap sched -x sv_vol sv_weekly 8@sat@20
snapvault snap sched -x sv_vol sv_nightly 7@mon-fri@20
snapvault snap sched -x sv_vol sv_hourly 11@mon-fri@7-19
```

3. **On the SnapVault primary and secondary storage system:** To turn off the regular snapshot schedule, enter the command:

```
snap sched volume 0 0 0
```

It is usually unnecessary and wasteful of storage resources to have both regular scheduled snapshots (those scheduled using the `snap sched` command) and SnapVault scheduled snapshots (those scheduled using the `snapvault snap sched` command). If you have both schedules, you can turn off and clear the regular snapshot schedule to save resources.

See section 14.8.1 for a detailed snapvault snapshot schedule configuration example.

11.6.4.4.3 User-Triggered Snapvault Snapshot Creation

Forced creation of SnapVault snapshot is possible if it is not the schedule time for the snapshot to be created. The snapshot creation will proceed just as if triggered by the snapshot scheduler. This means that primary filers will simply create snapshots, whereas secondary filers will kick-off the transfers from the primary filers, if their schedules have been set up to do so, before creating a snapshot.

On the primary filers, enter the command:

```
toaster> snapvault snap create proj sv_hourly
fridge> snapvault snap create proj sv_hourly
```

Then, on the secondary filer, enter the command:

```
secondary> snapvault snap create sv_proj sv_hourly
```

11.6.4.4.4 User-Triggered Snapvault Snapshot Creation

Verify or Display the SnapVault snapshot schedule

The `snapvault snap sched` command without any arguments can be used to verify the current settings on the secondary:

```
secondary> snapvault snap sched
xfer sv_proj sv_hourly 1@12,20 retention_period=365d
```

Verify or Display the SnapVault status

The `snapvault status` command with the `-s` argument can be used to verify the current status on the secondary:

```
snapvault status -s
```

Example:

```
secondary> snapvault status -s
Snapvault secondary is ON.
Volume      Snapshot      Status      Schedule
-----
sv_proj      sv_hourly      Active      1@12,20
```

11.6.4.4.4.1 Snapshot rotation for SnapLock volumes

Normally, SnapVault uses the count field of the SnapVault schedule on the secondary storage system to determine the number of snapshots that are retained. Once the count has been reached, the oldest retained snapshot is deleted when new snapshots are added. For SnapLock Compliance volumes, older WORM snapshots cannot be deleted until their retention period has expired. It is the storage administrator's responsibility to delete WORM snapshots once they expire; they are not deleted automatically.

11.6.4.4.2 Snapshot names for SnapLock volumes

Snapshots created by SnapVault on SnapLock Compliance volumes use a different naming scheme from that used by snapshots for regular volumes. Snapshot names for SnapLock volumes use the time and date of the snapshot creation as a suffix. As of ONTAP 7.3.2 and above the time and date is no longer taken from the compliance clock but rather the normal system clock instead. This reduces user confusion when listing the snapshots should the compliance clock be lagging the system clock due to prolonged system downtime or a compliance volume being in an offline state for a prolonged period of time. It does not affect the retention period of the volume, which is still enforced by the compliance clock, but it does prevent any user confusion regarding when the snapshot was actually taken. The following is the format of the snapshot name:

snapname.yyyymmdd_hhmmss_zzz

snapname is the snapshot name specified in the schedule

yyymmdd is the year, month, and day

hhmmss is the hour, minute, and second

zzz is the current time zone setting

Example:

%/used	%/total	date	name
-----	-----	-----	-----
2% (2%)	0% (0%)	Feb 04 02:20	svhourly.20040104_120502_GMT
9% (3%)	1% (0%)	Feb 04 02:15	svhourly.20040104_180601_GMT

11.6.4.4.3 Retention period for SnapVault created WORM snapshots

You can configure the schedules on the SnapVault secondary storage system to create WORM snapshots with a retention period that you specify. For WORM snapshot creation to occur, the volume must exist on the secondary storage system and it must be a SnapLock volume.

If you change the retention period in the schedule, WORM snapshots created under the previous schedule retain their retention periods and WORM snapshots created under the changed schedule use the new retention period.

11.6.4.4.4 SnapVault default setting for the retention period

By default, when you configure a SnapVault snapshot schedule for a SnapLock Compliance volume on the secondary storage system, the snapshots created for that volume are WORM snapshots. SnapVault uses the retention period set by the `snaplock_default_period` vol option as the retention period for the WORM snapshots. You should ensure that the retention periods you configured when creating a SnapLock Compliance volume are correct.

The default retention period is 30 years. Be sure to reset the default retention period unless you want all of the created SnapVault WORM snapshots to have a 30-year retention period.

11.6.4.4.4.5 Specifying WORM snapshot retention periods

You can specify retention periods for different WORM volume snapshots. To specify the retention period for WORM volume snapshots, complete the following step.

Example: WORM snapshots created on the secondary storage system in the sv_proj SnapLock volume are retained for 360 days from the time of their creation. Snapshots are created at noon and 8 p.m. every day.

```
snapvault snap sched -x -o retention_period=360d sv_proj sv_hourly 1@12,20
```

Enter the following command:

```
snapvault snap sched -x -o retention_period= period volname snapname count@
day_list@ hour_list
```

period is the retention period in days (d), months (m), or years (y).

volname is the name of the volume where the snapshot resides.

snapname is the name of the snapshot.

count is the number of snapshots you want to retain for this set.

@day_list is a comma-separated list that specifies the days on which a new snapshot is created for this set.

@hour_list specifies the hours at which a new snapshot is created for this set.

11.6.4.4.5 SnapMirror

This section describes the process of replicating SnapLock volumes using the SnapMirror feature of Data ONTAP.

A SnapMirror destination volume created as a SnapLock Compliance volume has the following restrictions:

- You cannot use the `snapmirror resync` command to return a former destination volume and resynchronize its content with the source.

Note: This restriction applies to volume SnapMirror relationships only. Resynchronization is supported for qtree SnapMirror relationships.

- You cannot reinitialize a SnapLock Compliance destination volume because WORM data on the destination volume must be protected from modification or erasure. If you break the SnapMirror relationship between the source and destination volumes using the `snapmirror break` command, Data ONTAP prevents you from ever reestablishing the SnapMirror relationship with the same SnapLock Compliance destination volume.
- You can initialize a new empty SnapLock Compliance destination volume.

A SnapMirror destination volume created as a SnapLock Enterprise volume is the same as a destination volume created for a non-SnapLock volume. You can perform all the same administrative tasks on a SnapLock Enterprise destination that you can perform on a non-SnapLock destination.

11.6.4.4.5.1 Creating a SnapLock Enterprise SnapMirror relationship

Creating a SnapMirror relationship from a SnapLock Enterprise source volume to a SnapLock Enterprise destination volume is the same as creating a SnapMirror relationship for regular volumes.

To create a SnapMirror relationship from a SnapLock Compliance volume to a SnapLock Compliance destination volume, complete the following steps:

1. On the administration host, use the `license` command to confirm that the SnapVault secondary storage device has the SnapMirror feature licensed.
2. On the administration host, use the `license` command to confirm that the destination device, either a filer or NearStore storage, has the features licensed:
 - SnapMirror
 - SnapLock Compliance
3. On the SnapVault secondary storage system console, use the options `snapmirror.access` command to specify the host names of filers that are allowed to copy data directly from the source filer. For example:

```
options snapmirror.access host=d_filerA
```

4. On the administration host, create or edit the `snapmirror.conf` file on the destination to specify the SnapLock Compliance volume source and destination.

Example: The following entry specifies asynchronous mirroring from `vol1` of `s_filer` and `vol2` of `d_filer` using the default replication arguments.

```
s_filer:vol1 d_filer:vol2 - * * * *
```

5. In both the source filer (SnapVault secondary storage system) and destination console, use the `snapmirror on` command to enable SnapMirror on the source and destination systems.
6. On the destination filer console (assuming a snaplock aggregate called `aggr1` already exists), create the destination volume using the `vol create` command.

Example: `vol create vol2 -s volume aggr1`

Unlike the ONTAP 6.5 PVR release of Lockvault where the `-L` option was used when initializing the SnapMirror transfer, in 7.1 the `-L` option is used to denote the Snaplock type at the aggregate level and the volume inherits this attribute. With ONTAP 7.1 and above, `snapmirror initialize` is constructed in such a way as to allow an initialize on a WORM volume as long as the volume is empty.

7. On the destination filer console, mark the volume as restricted using the `vol restrict` command.
8. On the destination filer console, create an initial complete (baseline) copy of the source on the destination and start the mirroring process using the `snapmirror initialize` command (see important note above regarding changes to syntax since ONTA 6.5 LockVault PVR).

Result: After successfully completing the baseline transfer, Data ONTAP converts the destination volume to a SnapLock Compliance destination volume before bringing it online.

Example: `snapmirror initialize -S s_filer:vol1 d_filer:vol2`

Note: After the destination volume is converted to a SnapLock Compliance destination volume, it will always be a SnapLock Compliance volume and cannot be changed.

11.6.4.4.6 System Monitoring

It is important to monitor the operation of the systems to ensure that backups are proceeding correctly. Error conditions such as failed/slow network connections and out of disk space conditions can prevent the scheduled completions of backup/mirroring operations from succeeding. Also, there is a potential for operator error or malicious activity causing the backups/mirroring operations to fail. Most of the monitoring activity discussed in this section can be accomplished using the latest CATE Storage Engineering certified version of Unified Manager.

11.6.4.4.6.1 Status

`snapmirror status` and `snapvault status` display the status of each relationship. The meanings of the status keywords are described here:

- *Idle*: The relationship is currently not involved in a transfer
- *Transferring*: The relationship is currently involved in a transfer
- *Pending*: The relationship has pending transfer
- *Aborting*: The relationship is aborting the currently active transfer
- *Quiescing*: The relationship is waiting for other transfers in the secondary volume to complete, OR the relationship could be involved in a rollback operation
- *Resyncing*: The relationship is currently involved in a resync operation

11.6.4.4.6.2 Recommended Monitoring

Outlined here is what to monitor to ensure successful backup/mirroring operations.

On a **daily** basis:

- On the SnapVault Secondary, do a `snapvault status` and ensure that the lag times are appropriate
- If no transfers are active at the time, the status should be Idle. Meanings of the other statuses are described above.
- Do a `snapmirror status` on the “Tertiary” system and ensure that lag times are appropriate.

On a **weekly** basis:

- Check the free disk space on the log volume and the secondary volumes and make sure that it is adequate for the coming weeks.
- Look at the current `snapvault.yyyymmdd_zzz` log file in the `/etc/log` directory on the worm log volume. Grep for words Abort and Retry which sometimes indicate trouble. Generally when Abort appears in the log file, the reason for the abort is shown. Many aborts are caused by operator errors such as not having the proper facilities enabled or volumes/qtrees are incorrectly configured. Some are caused by downed source systems. The presence of transient aborts in the log does not necessarily mean that any action needs to be taken, however the cause of repeated aborts should be investigated. An occasional Retry is not a cause for concern, however large numbers of retries may indicate an overloaded system or network problems.
- Look at the current `/vol/vol0/etc/log/snapmirror` log file on the “Tertiary” Nearstore for errors (Retry or Abort) indicating that the transfers may have encountered problems.
- Verify that snapshot retention periods on the snapvault secondary Nearstore are correct. Use `snap list -l` to view the retention periods.
- Verify the Compliance Clock is in sync between the secondary and tertiary NearStore systems.

11.6.4.4.6.3 Remediation for Problems

One of the more common problems is that a scheduled transfer will fail due to network problems, a primary filer down for maintenance, or a lack of disk space on the target Nearstore. Once the problem is corrected you can rerun the transfer. To rerun the transfer at a time later than the originally scheduled time, you can perform a snapvault update. If the snapshot on the primary from which you want to transfer the data is still the latest snapshot for snapvault, you can just do a snapvault update. If a later snapshot has already been taken, you should specify the snapshot on the source from which you want to transfer the data with the '-s' option. The snapshot of interest on the source would only have been created if the primary were up at the scheduled time for the snapshot to be taken. The snapshot would only still be present if either the number of snapshots to be retained in the snapvault snap sched command on the primary had been large enough for it to be retained, or if the transfer had started and failed. If the transfer had started and failed, the snapshot would still be present due to a 'softlock'. To find the snapshot of interest, do a snap list on the primary filer for the source volume from which the transfer should have occurred, and find the snapshot that was taken for SnapVault just prior to the scheduled transfer. If it is not present, you will not be able to re-do the transfer and you must just rely on the data from subsequent transfers.

11.6.4.4.7 WORM log configuration

SnapVault uses a SnapLock log volume in which the log files are kept. The volume is similar to a WORM volume, but has the added feature of supporting append-only operations. This allows accurate record keeping but does not allow previous events to be overwritten.

For regulatory compliance, it is not acceptable for logs to be unreliable. It is required that SnapVault operation logs help in accurately determining the exact series of events that occurred during that timeframe. In addition, it is also required that once generated, over-writing or deleting contents or the log files themselves should not be possible. The contents of log files cannot be overwritten or deleted, nor can the log file itself be deleted. The log files must accurately record the exact series of events that occur during a timeframe. In order to protect the log files from being destroyed or tampered with, we require that they be created in a special volume, which is very similar to a WORM volume, but in addition it supports "append only" operations. This special SnapLock log volume must be created before you begin any initial (level 0) SnapVault transfers.

11.6.4.4.7.1 Enabling the SnapVault log volume

Unlike the ONTAP 6.5 PVR release of Lockvault where a SnapLock log volume was created with the -L compliance_log option, for versions of ONTAP 7.1 and above the setting no longer exists. Instead create a Snaplock Compliance volume and ensure the step in the following section is completed before you begin an initial (level 0) SnapVault transfer.

11.6.4.4.7.2 Option to set the log volume

Once the SnapLock log volume is created and ready for use, you need to configure the user option setting on your SnapVault secondary system.

```
options snapvault.lockvault_log_volume <vol_name>
```

Setting this option will only succeed if you provide the name of the volume that was created earlier. Only after you have set this option, will your SnapVault transfers proceed successfully.

All SnapVault logs will be created in this log volume. As you must have noted, this is a system-wide option. So any SnapVault operation on a WORM volume on the system will be logged to this volume.

11.6.4.4.8 Snapvault Log Information

All SnapVault logs will be available under the `etc/log/` directory on the user configured SnapLock log volume. These log files being part of a worm volume will inherit the current setting of the SnapLock default retention period option. Users cannot delete these files until the retention period expires.

Ensure that you set the default retention period on the log volume appropriately.

SnapVault uses two types of log files to record events:

- SnapVault Operations Log File
- SnapVault Transfer Log File

11.6.4.4.8.1 Operations Logs Location

These logs will be available under the `etc/log/` directory of the SnapLock log volume. The format for naming these log files will be `"snapvault.YYYYMMDD_ZZZ"` where:

YYYY is the year
 MM is the month (01-12)
 DD is the date (01-31)
 ZZZ is the current time zone setting

The timestamp denotes the time of creation of the file, and will be generated using the compliance clock, not the system clock.

The operations logs have a weekly rotation policy. A new log file will be created every Saturday morning at midnight.

Note that SnapVault will not try and delete any of these log files itself and will, in fact, be incapable of deleting them until the specified retention period of the logs has been satisfied. SnapVault operations that were active when the log files are being rotated, will log to the most recent log file.

11.6.4.4.8.2 Transfer Logs Location

The transfer logs are specific to every SnapVault relationship and a new log file will be created at the beginning of each SnapVault transfer, in the appropriate directory for that SnapVault relationship.

There will be one directory under the `etc/logs/` directory on the SnapLock log volume per SnapVault secondary volume. The format for naming the directory will be `<secondary_volume_name>_<volume_uuid>`. The volume UUID is guaranteed to be globally unique for NetApp volumes. This should help in identifying volume renames etc.

Under this directory, there will be a separate sub-directory for every snapvaulted qtree with its own name, `<secondary_qtree_name>_<treeid>_<inodenum>_<gennum>`. Again, the tree id, inode number and generation number help in uniquely identifying qtrees within a volume.

Under the qtree directory, there will be a monthly sub-directory with the format “YYYYMM” for its name. The timestamp denotes the time of creation of this directory, and will be generated using the compliance clock, not the system clock.

Underneath this directory, every transfer will generate its own unique log file. The file-history log files will be named using the format “snapvault_filelog.YYYYMMDD_HHMMSS_ZZZ” where:

YYYY is the year
MM is the month (01-12)
DD is the date (01-31)
HH is the hour (00-23)
MM is the minute (00-59)
SS is the second (00-59)
ZZZ is the current time zone setting

The timestamp denotes the time of creation of this file, and will be generated using the compliance clock, not the system clock.

A sample path for SnapVault file-history logs could be:

```
/etc/log/vault_ed0ad520-5f40-11d8-91ca00a09800dcba/
users_19_1435628_1286197/200401/snapvault_filelog.20040123_000500_PST
```

Similar to SnapVault logs, these transfer log files will never be deleted by SnapVault itself. The user will have to delete these, once their retention period has expired.

11.6.4.4.8.3 Transfer Log Contents

Each log file contains a header describing the transfer followed by zero or more entries describing the contents of the transfer. Note that it is possible for a log file to contain only a header if, for example, a snapvault update were performed when no data on the primary had changed since the last snapvault update. A log file will contain only a header for a rollback transfer.

The headers at the beginning of each log file contain:

- A fixed field: START TRANSFER, or RESTART TRANSFER
- Date/Time of when the transfer started
- Sourcefiler:sourcepath
- Destinationfiler:destinationpath
- Date/Time of source filer snapshot.

Note that this date/time will be interpreted according to the time zone settings on the secondary system, not the primary system.

Example of a header log record:

```
** Start Transfer
** Started @ Mon Feb 16 13:39:43 EST 2004
** From sourcefiler:/vol/compat_data/myqtree1
** To destinationfiler:/vol/logs/mult1
** Source Snapshot created @ Mon Feb 16 13:42:44 EST 2004
```

The header is followed by 0 or more log entries. The following types of log entries are included:

- Create File (regular, special or stream)
- Delete File (regular, special or stream)
- Create Directory
- Delete Directory
- Rename From (regular file or directory)
- Rename To (regular file or directory)
- Modify File (regular, special, or stream file, but not directory)
- Modify Attributes (regular, special, or stream file, or directory)

Format of log entry:

<date/time> <action taken> <length of path in single byte chars> <relative file path>

The date/time will be the mtime or ctime from the source system, not the date/time of the backup. The action taken will be one of the “types of log entries” shown above. The path name displayed will be relative to the base of the tree that is mirrored, not the absolute path name. This relative path name is the same on the primary and the secondary.

Examples of log entries:

```
Mon Jan 26 9:37:14 PST 2004 Create File 17 ./docs/newdoc.doc
Fri Jan 23 11:40:45 PST 2004 Rename From 19 ./docs/firstdoc.doc
Mon Jan 26 11:40:45 PST 2004 Rename To 17 ./docs/olddoc.doc
```

The following notes are included to clarify the exact circumstances under which the above log entries will be created and the contents of some of the fields in the log entries:

- The initial transfers (at SnapVault start) of all files and directories are logged as ‘Create File’ and “Create Directory” respectively.
- A Modify Attributes log entry is not included if a Modify File is logged for a file (regular or stream).
- A Modify Attributes log entry for a directory is included if the attributes such as the permissions have been changed or if any entries in the directory have been changed.
- Creations, deletions, or modifications of ACLs are not explicitly logged. If an ACL for a file or directory has changed, a Modify Attributes for the files and directories associated with the ACL is logged.
- A Rename is logged as two log entries: a ‘Rename From’ and a ‘Rename To’. The ‘Rename From’ and associated ‘Rename To’ appear together in the log.
- When a file is renamed, either a Modify Attributes or a Modify File log entry is included. Which one is included depends on whether actual data has been modified. If a Modify File log entry is included, a Modify Attributes is not included.
- When a directory is renamed, a Modify Attributes log entry is also included.
- Since streams have a different syntax for path names, (basefilename:streamfilename), the user can tell a stream file entry from a regular file entry by looking at the path name.
- The date/time in the log entry for a Create, Delete, or Modify Data is the mtime. We use the mtime because mtime refers to the modification time of data and these three log types are generated when data is modified.
- The date/time in the log entry for a Rename or Modify Attributes is the ctime. We use the ctime because ctime refers to the modification time of the container for the data, not the data itself. Renames and Modify Attributes fit this description.
- Modify Attributes entries are included for regular, special, or stream files and directories if any attributes of the file or directory, including the access time (atime) on the primary have changed. As a result, the user may see many Modify Attributes entries for files that have not been changed except that they have been accessed. Accesses include reading files and traversing directories. If the user does not want Modify Attributes log entries for accesses, the no_atime_update option to the vol options filer console command may be set on the primary filer to suppress the changing of the atime field on the files on the source filer. Note that the setting of the no_atime_update option on the secondary filer has no effect on the entries that will appear in the file log.

Stream files do not have their own date/time. Therefore the date/time shown in the log entry for a stream file is the date/time of the base file for the stream file. The mtime of the base file is displayed if on this transfer the base file was created or its data was modified. The ctime of the base file is displayed if on this transfer only the attributes of the base file (such as the addition of a stream file) had changed.

11.6.4.4.9 Disaster Recovery

By now you already know how to setup compliant backups using SnapVault. Once your SnapVault secondary system is configured, you will probably want to keep a plan ready, just in case disaster strikes. We will talk about the different disaster scenarios, and outline some steps to recover from them in the quickest possible way.

What disaster scenarios are we dealing with here?

We will deal with disaster scenarios that are very specific to SnapVault primary and secondary systems. More specifically, we will deal with scenarios that lead to an unplanned outage on one or more SnapVault primary systems themselves or the secondary system itself.

The goal here is to provide storage administrators with a plan to recover from such scenarios, using a well-defined procedure, such that you can restore your SnapVault configuration back to normal. And of course we will always try to use the procedure that takes the shortest amount of time.

11.6.4.4.9.1 Primary system is temporarily unavailable for client access

If one of your primary systems becomes unavailable, and you need to access the data on that primary, then you can directly point your clients to access data in a read-only fashion from the snapvault secondary system. This procedure will be independent of snapvault.

However, if you need read-write access to the primary data, then you will need to first make the required qtree(s) read-write on the secondary system. To do so, you will need to follow a certain procedure:

- Turn snapvault off on the secondary system
- Turn snapmirror off on the secondary system
- Convert the secondary qtree(s) to a snapmirrored qtree(s)
- Turn snapmirror and snapvault back on
- Use snapmirror quiesce and break to make the qtree read-write

You can now direct your clients to the secondary qtree(s) and they should be able to read-write to the qtree(s).

Snapvault scheduled operations will ignore this qtree, since you converted it to a snapmirrored qtree. The reason you have to do the convert is to be able to use the snapmirror commands to make the qtree read-write.

Now after a while, your primary system is up and available for use.

If you want to retain the data that was newly created on the secondary, you will first need to transfer these new files over to the primary system. Snapvault provides you with a restore command to be able to transfer the entire secondary qtree to a new qtree on a primary. The restore command is essentially a snapvault initial transfer from the secondary to a primary. You cannot run incremental restore transfers using snapvault.

Note that the secondary qtree will still be read-write, for the duration of the restore. It is the administrator's responsibility to restrict user changes to the secondary qtree starting from just before the restore command is kicked-off. Any changes made to the secondary qtree after the snapshot for restore is created (at the beginning of the transfer) will not be transferred to the primary.

Once the restore transfer is complete, you need to re-sync the newly created qtree and the secondary qtree, using the snapvault start -r command on the snapvault secondary. This will re-configure the snapvault relationship using the new qtree on the primary and the secondary qtree.

If you want to ignore the data that was newly created on the secondary since the break, you can skip the restore command. And directly execute the snapvault start -r command using the existing primary qtree and secondary qtree.

The start -r command will first archive all the changes made to the secondary after the qtree was made read-write. And then it will re-configure the snapvault relationship.

From this point on, your snapvault updates will continue as per your configured schedules.

11.6.4.4.9.2 Primary system is irrecoverable

If the primary system is irrecoverable, then you will need to configure a new primary system. After setting up appropriate volumes and snapvault licenses etc. you will have to snapvault restore all the desired qtrees from the secondary system to the new primary. Once all restores are complete, you will have to run snapvault start -r command on the snapvault secondary for each of the primary-secondary relationships – specifying possibly a different hostname than before for the new primary system.

Once this has been done, your snapvault updates should continue as per your configured schedules.

11.6.4.4.9.3 Secondary system is irrecoverable

If your secondary system crashes and is irrecoverable, then your best chance to recover from this disaster is if you had earlier setup volume snapmirror relationships from the secondary system to a tertiary system. This configuration not only protects your snapvault relationships, but also keeps remote copies of your retained backup snapshots.

In the absence of the tertiary, you will have to start your snapvault relationships all over from scratch – i.e. you will need to do baseline transfers for all primary qtrees to the new secondary system. Even worse is the fact that you could lose all your backup snapshots on the old secondary in this disaster.

The secondary data and state of snapvault qtrees is preserved when you snapmirrored the secondary volume to the tertiary system. Hence, the quickest approach out of this disaster scenario is to re-configure the tertiary system to run like the snapvault secondary. Here are the steps you will need to follow on the tertiary to do that:

- First you will need to make the snapmirror replica volumes on the tertiary system read-write. The snapmirror break command must be used for this purpose.
- Now license the tertiary system as a snapvault secondary.
- Check the status of your snapvault qtrees using snapvault status. The snapvault qtrees are preserved on the volume snapmirror replica, along with their state.
- Configure the snapvault relationships on the tertiary system, using snapvault start. This step adds the snapvault configuration entries on the tertiary system. Note that no transfers will be kicked off when you execute this command.
- Next configure the snapvault options on the tertiary, as you had before on the secondary.
- Finally set the snapvault schedules on the tertiary, as you had before on the secondary.

Finally your tertiary should be ready to run as a snapvault secondary. The schedules will trigger the next updates for you. These updates, depending upon your snapmirror schedules will be incremental transfers based upon one of the recent successful updates on the snapvault secondary.

The primary systems will automatically report snapvault relationships with the qtrees on the tertiary system. You will also see the old snapvault relationships to the old secondary system. After observing a couple of successful scheduled updates, you could release the resources tied down on the primary systems for the old secondary system by using the snapvault release command.

This configuration helps avoid the time-consuming initial transfer from the primary systems, to a newly setup secondary system.

Note that once you fail-over your snapvault relationships from the secondary to the tertiary system, you cannot easily fail-back to the old secondary. For sake of this discussion, lets say the old secondary system is called S, and the tertiary system that is now running as the secondary for snapvault relationships is called T. Before the disaster, your configuration looked like:

Primary systems → S → T

After your fail-over, the configuration looks like:

Primary systems → T

The old secondary system can only be used again by setting it up as a snapmirror replica system. You will need to carry out initial transfers from T to S. After which your configuration will look like:

Primary systems → T → S

This way you are back to the original setup, and your new secondary (T) is now protected by the system (S).

Because of the slightly high complexity of this fail-over procedure, it is advisable to weigh the benefits of doing this in scenarios where the secondary system is only temporarily unavailable.

11.6.4.4.10 Destroying Snaplock Volumes

SnapLock Compliance volumes constantly track the retention information of all retained WORM files and WORM snapshots. Data ONTAP does not allow you to destroy any SnapLock volume that contains unexpired WORM content. Data ONTAP does allow you to destroy SnapLock Compliance volumes when all the WORM files and WORM snapshots have passed their retention dates, that is, expired.

Note: You can destroy SnapLock enterprise volumes at any time.

To destroy a SnapLock volume, complete the following step.

```
vol destroy <vol_name>
```

If there are any unexpired WORM files in the SnapLock Compliance volume, Data ONTAP returns the following message:

```
vol destroy: Volume volname cannot be destroyed because it is a SnapLock
Compliance volume.
```


Note: In order to delete a SnapLock volume that has just been created, you must cycle the volume through an online and offline cycle in order to delete the volume. The volume must cycle in order for Data ONTAP to determine that there are no files with unexpired retention periods.

11.7 Storage Optimization Technologies – Dedupe, compression, flash cache, file folding, file blocking

11.7.1 Deduplication

NetApp's deduplication technology (formerly called ASIS, now referred to as simply 'De-dupe') is a volume level deduplication process that usually is configured to run on a scheduled basis but can also be run manually via the command line. Data is not de-duped 'on the fly' when it is written to disk but is rather scanned and duplicate blocks replaced by pointers via a scheduled task run on the storage system. Deduplication is a low priority system process that allows the Filer to continue normal data serving and replication operations. It will increase overall CPU utilization while deduplication is running (see section [11.7.1.6](#)). Storing this finger print information while new data is written allows for faster execution of the dedupe process as only the fingers prints are analyzed as opposed to pulling every new block of data from disk for analysis. As of ONTAP 7.3 or above deduplication utilizes 1% to 3% of the available aggregate space to store 'finger print' (metadata) about the data being written to a de-dupe enabled volume.

Refer to the CATE NetApp ONTAP 8.1 Release Notes for known issues when upgrading deduplication enabled systems which have been running an earlier version of ONTAP 8.1.

11.7.1.1 Licensing

Deduplication is a licensed feature of the Data ONTAP operating system. As of ONTAP 8 or above deduplication only requires a single license key (note nearstore_option is no longer required as a pre-requisite)

- a_sis

Best practices mandate that for a clustered Filer **both licenses** be installed on **both nodes in a cluster**.

According to the NetApp account team, most NAS systems procured by Citi in the past 24 months have included the above licenses in the purchase. Storage administrators refer to the license keys provided with the storage system to confirm if valid asis licenses are already in the inventory for each system. The keys may have been provided but never installed on the system.

ONTAP 8.x makes no further changes to the behavior of de-duplication outside of the changes already introduced with ONTAP 7.3 whereby the metadata is stored in the aggregate.

11.7.1.2 Best practices for determining volumes eligible for de-dupe

Outside of VMware on NetApp NFS storage, the other primary candidate for de-dupe is user and group data volumes.

In choosing candidate systems for de-dupe CATE Storage and Backup Engineering has devised the following criteria to assist storage operations teams in their deployment efforts:

1. Supported hardware platforms. Only the following platforms support de-dupe:
 - a. FAS2040 series
 - b. FAS3100 series
 - c. FAS3200 series
 - d. FAS6000 series
 - e. FAS6200 series
2. Storage system utilization:
 - a. Lowest CPU utilization over a sample period of 1 week < 65%
 - b. Snapshot utilization < 50%.
For the purposes of de-dupe viability snapshot utilization is expressed as the amount of space used by snapshots in a volume compared to the total amount of used space in a volume. This applies to both regular and snaplock volumes.
 - c. Aggregate utilization =< 86%.
As of ONTAP 7.3 or above de-dupe needs to build a fingerprint database in the hosting aggregate that can take up 1% to 4% of the total aggregate size. The 86% utilization rate factors in the recommended 10% aggregate free space policy.
 - d. Volume utilization < 98%
2% free space is required in the volume while the de-duplication process is running
3. Storage system use cases. Volume types that are the best candidates for de-dupe:
 - a. User home and group/departmental data CIFS (Windows) volumes on production Filers (as opposed to contingency systems)
 - b. User home and group/departmental data NFS (Unix) volumes on production Filers (as opposed to contingency systems)

This methodology was derived from the de-dupe pilot conducted in the North America during October / November 2008.

11.7.1.3 Best practices for implementing de-dupe

After using the system eligibility and volume criteria outlined in the previous section, use the best practices listed below to implement de-dupe:

1. Install asis licenses on Production and DR Filers
Installation of licenses on the DR Filers will allow you to schedule de-dupe on DR volumes if a business is in a DR state for an extended period of time
2. Ensure your candidate volumes are protected by some form of backup solution to allow for a roll back process in the event of a catastrophic failure
3. Deploy de-dupe across your environment in a phased approach and not all volumes at once
4. Enable de-dupe on each candidate volume with the `sis on /vol/vol_name` command
5. For volumes with existing data use the `sis start -s /vol/vol_name` approach to de-dupe the data already present on the volume. It is recommended that you conduct this stage on a Friday evening or over a weekend to allow the existing data to be processed during non-core business hours. Be careful not to start the de-dupe process too late on a Saturday so that the scan operation will be contending with the RAID scrub processes for system resources.

6. Schedule all subsequent automatic de-dupe operations based on the following standard schedule `sis config -s sun@6-14 /vol/vol_name` for each candidate volume on the system.
7. After the initial de-dupe operation savings will not be immediately shown as the changed blocks are still locked in snapshots. Allow all snapshots and snapvaulted snapshots to rotate one full cycle according to your snapshot schedule policy before determining the space savings with the `df -s -g` command.

11.7.1.4 Commands to manage de-dupe

To enable deduplication on a volume use the following command:

```
filer> sis on /vol/vol_name
```

To begin the deduplication process **on a volume which already contains unde-duplicated data**, determine a time during the week when the Filer is at a respectively low period of workload. Where possible it is recommended to try to perform this on a Saturday to avoid conflicts with the weekly disk scrubbing process on Sunday mornings. Use the following command to de-dupe data on an existing volume with existing data ('-s' forces ONTAP to scan the existing data).

```
filer> sis start -s /vol/vol_name
```

On a new volume where no data exists you can start the de-dupe process manually with the following command (notice the absence of the '-s' parameter).

```
filer> sis start /vol/vol_name
```

Use the 'sis config -s' command to schedule the de-dupe operation to occur automatically. If you just type 'sis config' ONTAP will return the current schedules for all de-dupe enabled volumes. Adding the '-s' parameter allows you to set the schedule.

```
filer> sis config -s [sched]
```

The [sched] portion of the config command can accept parameters in four different formats. However from a standardization perspective we will only use one of these parameters. The schedule is defined in the format [day_list]@[hour_list], similar to how snapshot and snapvault schedules are defined. This is best demonstrated with an example that will also show the recommended best practice for scheduling a de-duplication process:

```
filer> sis config -s sat@8-18 /vol/vol1
```

The above example configures the de-dupe process to run between 8am to 6pm on every Saturday. This configuration is the recommended starting point when considering pilot implementation of dedupe in a production environment. As deployments of de-dupe become more prevalent additional best practices will be provided on scheduling and managing deduplication. It is important to note that in order to preserve system resource only 8 concurrent volume de-duplication operations can take place. Other deduplication operations will be queued until one of the 8 active deduplication jobs completes.

The 'sis status' command is used to report in the status of the de-dupe process. It will be reported in the long (detailed) reporting format when you use the '-l' parameter. CATE Storage Engineering recommends using the '-l' parameter as it provides more information. If you do not specify the name of a particular volume, ONTAP will report the status of every de-dupe enabled volume on the Filer.

```
filer> sis status -l /vol/vol_name
```

An example of the output from the `sis status` command is shown below:

```
filer> sis status -l /vol/esxpse01
Path:                /vol/esxpse01
State:               Enabled
Status:              Idle
Progress:            Idle for 00:14:06
Type:                Regular
Schedule:            mon-sat@3
Last Operation Begin: Tue Jun 10 03:00:00 EDT 2008
Last Operation End:   Tue Jun 10 22:20:37 EDT 2008
Last Operation Size:  30 GB
Last Operation Error: -
```

The status command will report three possible status values in the status field:

- Idle – No de-dupe processes are currently running on the volume
- Pending – A de-dupe process is pending the removal on a resource constraint (typically occurs when more than 8 de-dupe operations are scheduled to run concurrently).
- Active – A de-dupe process is currently running

The status command will report six possible values for the progress field:

- Idle – The de-dupe process has been in an idle state for hh:mm:ss
- Scanned – The de-dupe process was initiated with the '`sis start -s`' parameter and ONTAP is scanning the volume for existing data. The amount of data scanned will be reported.
- Searched – De-dupe is searching for duplicate data. The amount of data searched will be reported.
- Done – The de-dupe operation has saved 'x' GBs of de-duped data which is 'y' % of the total data on the volume designated as a candidate for de-dupe.
- Verified – The de-dupe process is verifying the metadata of processed data blocks. This process will remove unused metadata.
- Merged – Verified metadata is being merged. This process will merge together all verified metadata of processed data blocks to an internal format that supports fast sis operation.

A de-dupe operation can be suspended while it is running using the following command:

```
filer> sis stop
```

It can be restarted again using either the '`sis start`' command or when the next scheduled automated de-dupe process begins. When the deduplication process restarts it will not continue from the point at which it was stopped, but rather it will start anew from the scanning or searching state (as defined above). Note if you stop a de-dupe process in the middle of a scheduled process and the process is defined to execute between a given start time and a given end time then it is possible that the de-dupe process will restart automatically within 5 minutes of issuing the '`sis stop`' command.

De-dupe processes can be disabled on a volume using the following command:

```
filer> sis off /vol/vol_name
```

Scheduled de-dupe processes will not run on a volume where de-dupe has been disabled. It is important to note that disabling dedupe is not the same as returning the data to its original format (complete with duplicate data blocks). It merely means that any new data written to the volume will not have its metadata collected for use in any future de-dupe processing.

If de-dupe is disabled on a volume it can be re-enabled with the '`sis on`' command. Administrators should then run '`sis start -s`' to ensure that any data written to the volume since de-dupe was disabled is scanned.

The final de-dupe command to be discussed in this document is a special command available in the console diag mode. This command should not be used by Storage Administrators without opening a support case with NetApp first. It has been included here to demonstrate there is a process available to return a volume to its original state and backout any changes made by the de-duplication process (consider it an un-de-duplication process for use of a better phrase).

```
filer*> sis undo
```

Storage administrators should use the 'df -s -g' command (discussed below), 'df -A' (to determine free space in the aggregate) and 'vol size' commands (to resize a flexible volume) if considering reverting a volume to its original state.

11.7.1.5 Determining space saved and reporting

There is a special version of the 'df' command that tells the storage administrator how much space has been saved by the deduplication operation. Use the following command to determine how much space has been saved with dedupe (you may use the -m, -g or -h parameters after the -s to list the results in MB, GB or highest measure, respectively):

```
filer> df -s vol_name
```

Here's an example:

```
nas3050gdse01> df -s -g win32
Filesystem      used      saved      %saved
/vol/win32/     93GB     181GB      66%
```

In this example the data on the volume win32 has been deduplicated to 93GB. The 181GB represents the amount of duplicate blocks of data that existed on the volume. The total utilization of the volume without deduplication is calculated as 93GB plus 181GB, yielding 264GB. The %saved is calculated as 181GB / 264GB.

From a Windows or Unix client perspective, the share or mount point space utilization will report the space on the file system consumed after de-duplication. However when querying the size of individual files, the file metadata is maintained so the "on-disk" value is returned as opposed to the deduped size of the file.

The implication of this difference in reported file size (metadata) and space consumed (physical) may mean that changes are required to any chargeback data collection mechanism which walks filesystems to calculate utilized space in order to obtain the correct data.

Operations Manager 3.8.1²³ or above can be used to provide reporting insight into de-duplicated file systems at a regional level.

11.7.1.6 Constraints and limitations when considering de-dupe

Storage administrators should be aware of the following constraints and product limitations when considering deployment of de-dupe in their NetApp NAS environment:

- Deduplication will occur at the volume level only. It is not possible to deduplicate data across the entire aggregate or storage system. That is not to say you cannot deduplicate every volume on the system, it just means you can only single instance duplicate blocks of data within a given volume.
- The finger print data (metadata) used by the deduplication process to record matched blocks will consume between 1% and 3% of the useable space on an aggregate.

²³ Now branded Unified Manager

- Deduplication should not be scheduled to run on volumes during periods of high write activity. The de-dupe process creates a temporary snapshot of a volume to provide a roll back point if an error occurs. If the volume is nearly full at the time this snapshot is taken, then all other writes to the volume will continue to consume space until the dedupe process is near completion and the snapshot is automatically removed by ONTAP.
- Snapmirror destination volumes cannot be de-duplicated as the volume is in a read-only state.
- The data blocks owned by a snapshot cannot be de-duplicated as they are locked in a read-only state.
- Temporary metadata files created on a volume during a de-duplication operation can become locked if a snapshot is scheduled to occur during the dedupe run. The metadata files remain locked until the snapshot copies locking the files are deleted.
- Deduplication on the primary storage is only replicated to secondary (DR) storage when Volume SnapMirror is used. Qtree Snapmirror does not operate at the same logical level so the complete dataset (as opposed to the de-duplicated dataset) is copied to the secondary storage device.
- With ONTAP 7.3 or above both the baseline transfer and all subsequent delta transfers will be deduplicated once sis is enabled on the target volume. Existing SnapVault snapshots will not be deduplicated, only new ones written to the volume once sis is enabled. It will take two successive transfers before any savings are seen on the target volume.
- De-duplication cannot be concurrently run on more than 8 volumes on a Filer. This implies that like other NetApp processes such as snapshots and snapmirror transfers, the efficiency and run duration of the de-dupe process is dependent on the level of change in the dataset between each de-duplication run and the amount of duplicate data blocks in a volume. For volumes which have never been de-duplicated before, the run time of the initial deduplication process will depend on the size of the volume and its utilization rate, as well as the amount of other processes that are running on the storage system and competing for system resources. As more and more volumes on a Filer are de-dupe enabled, administrators are anticipated to spend an increasing amount of time monitoring the efficiency of the de-dupe process and the length of time it takes for the de-dupe process to run on a given volume. The limitation on the number of concurrent de-dupe processes will likely mean administrators will have to manage the de-dupe schedules carefully.
- De-dupe is only supported on certain NetApp hardware platforms and considering the narrowly defined scope of ONTAP 8.0.x deployments within Citi the list has been shortened further to only platforms supporting 8.0.x:
 - It is supported on the FAS2040, FAS3100 series, FAS3200 series, FAS6000 series and FAS6200 series
- Starting with ONTAP 8.1 de-dupe no longer artificially restricts the maximum size of the volume on a given platform. The values are well above the maximum volume size supported by CATE Storage and Backup Engineering in Citi and have been included here for awareness. The maximum volume size supported is determined by the particular FAS model. Refer to the CATE NetApp Hardware Platform release notes for further details on the maximum volume size.
- Only flexible volumes support de-dupe. De-dupe cannot be run on traditional volumes.

- It should be remembered that de-dupe will run as a background process. Where systems have the capacity to leverage under-utilized CPU resources, testing has shown de-dupe processing can increase CPU utilization by 25 to 35%. On systems that are already running with high CPU utilization it is expected that dedupe processes will take longer to complete while ONTAP ensures that data serving and other processes taken higher priority than the de-dupe process.

11.7.1.7 Deduplication during a cluster failover event

During a cluster failover event, deduplication will behave in the following manner:

- Writes to a de-dupe enabled volume will have their finger-prints added to the metadata change log
- No sis commands, including manually and schedule invoked de-duplication processes will run.
- Upon failback, normal A-SIS operations can continue and the updated change log processed.

11.7.1.8 Deduplication during a contingency event

From ONTAP 7.3 and above all volume deduplication metadata is stored in the hosting aggregate. This overcomes two disadvantages experienced with earlier versions of ONTAP: 1) the metadata was replicated each time with the volume snapmirror update, and 2) every snapshot copy maintained the metadata and thus minimizing the space savings.

To overcome the above stated disadvantages, starting with Data ONTAP 7.3 and above, the deduplication metadata has been moved out of the deduplicated volume and into the aggregate. Even though the metadata is not part of the volume, the destination volume is still in a deduplicated state and the data in the destination volume can be accessed just as in a regular volume. However, due to the absence of the metadata in the volume, there are some side effects to keep in mind in failover and failback scenarios described below.

11.7.1.8.1 Failover

In this scenario the destination (DR) volume is made writable for DR purposes. The following list represents the deduplication characteristics for the DR volume:

1. The DR volume is already in a deduplicated state because it is a volume SnapMirror destination volume for a deduplicated primary volume.
2. Any new data written to the DR volume since the failover enjoys space savings within the new data itself but does not enjoy additional space savings with the old data. In other words, deduplication savings are not shared between the new and the old data.
3. If deduplication savings are desired across the entire volume (new and old data), the deduplication metadata must be rebuilt by running the deduplication operation on the entire volume (by using the `sis start -s` command). The volume is available for read and write operations during the deduplication operation phase. The factor driving the decision point on whether or not to rescan all the existing data on DR is how long the DR system is expected to serve data. From experience most DR invocations are failed back at the first available greenzone after production services are restored and thus may not warrant the extra work required in rescanning the old data on the DR system.

11.7.1.8.2 Failback

In this scenario the primary volume is made writable at the production site after a SnapMirror resync operation. It is also assumed that all the new data written at the DR site since failover is replicated from the DR volume to the primary volume. The following are the deduplication characteristics for the primary volume:

1. The primary volume is already in deduplicated state because it is a volume SnapMirror destination volume for a deduplicated DR volume.
2. Any new data written to the primary volume since the failback enjoys space savings with the old data that exists at the time of failover but not with the data written from the DR volume.
3. If deduplication savings are desired across the entire volume, the deduplication metadata must be rebuilt by running the deduplication operation on the entire volume (by using the `sis start -s` command). The volume is still available for read and write operations during the deduplication operation phase.

11.7.1.9 Deduplication during a volume SnapMirror data migration event

In this scenario volume SnapMirror is used to migrate a volume between aggregates within a system or between two different systems (assuming that the source and destination aggregate are of the same type – either both 32-bit or both 64-bit). Upon migration, it is recommended to run the deduplication operation on the entire volume to rebuild the deduplication metadata (by using the `sis start -s` command) for maximum space savings.

11.7.1.10 Deduplication log files

All deduplication event information is written to the following log file:

- `/etc/log/sis`

11.7.2 Compression

Starting with ONTAP 8.1 both inline and post process compression are supported. This section is a placeholder for future content regarding compression. The content will be added to a later draft version of this document.

11.7.3 Snapshot 'file-folding'

File folding is the process of checking the data in the most recent snapshot, and, if it is identical to the blocks in the file currently being saved to the storage device by a CIFS file close event, just referencing the previous snapshot instead of taking up disk space writing the same data in the active file system.

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. This feature was introduced in ONTAP 6.2 to address inefficiencies in how MS Office documents saved changes to disk. Due to the improvements in the MS Office suite and how it handles updating files, plus the arrival of storage level data de-duplication, there is little value to be gained by deploying file-folding.

"snapshot_file_folding" is configured with the options command at a physical Filer or vFiler level. The recommendation is to keep this feature at its default setting 'off'.

```
options cifs.snapshot_file_folding.enable off
```

11.7.4 File Blocking

ONTAP 7.2 and above provide storage system support for simple native file blocking functionality. This feature allows storage administrators to control the types of files being stored on the Filers using file extension. This functionality is covered in this document to make storage administrators aware of this feature. CATE Storage and Backup Engineering is not mandating the use of this option on every Filer or volume. It is intended for use by businesses who wish to use this functionality to control what file types are saved to their shared storage.

11.7.4.1 Overview

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

FPolicy runs per vFiler instance on a physical Filer, therefore each vFiler on a system would need to be configured with File Blocking if the intention was to implement file blocking on the entire physical system. File blocking policies enabled on the physical Filer are not inherited by the vFilers running on the same system.

FPolicy runs independently from the storage system's antivirus vscan facility. All file policies are applied to a client request before virus scanning occurs. Vscan operations are independent of file policies. That is, vscan can open and scan files that have been blocked by file policies.

11.7.4.2 Limitations of native file blocking

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

Native file blocking scans the extension of the file only. It does not scan the file as a whole to determine if it should be blocked if the file extension does not reflect the true content of the file.

The following features are **not available** with native file blocking:

- File and user level blocking event reporting
- Intuitive blocking response to the end user
- Centralized management and reporting of file blocking configurations
- Scan of the file header information to determine whether the file type is the same as the file extension

11.7.4.3 Implementing native file blocking

Use the following steps to implement native file blocking on the NetApp Filer platform. In this example a policy will be created to stop mp3 files from being saved to the volume 'data1'. For all other command regarding File Blocking refer to the "Using Fpolicy" chapter of the "File Access and Protocols Management Guide" from NetApp.

1. Enable fpolicy to enable the file screening process
`Filer> options fpolicy.enable on`
2. Create a policy name and type. Currently the only supported policy type is 'screen'.
 Tip: Make the policy name descriptive such as "mp3blocker".
`Filer> fpolicy create mp3blocker screen`
3. Define the file extensions you intend to block with the policy
`Filer> fpolicy ext inc set mp3blocker mp3`
4. Define the volumes to which you intend to apply the blocking policy. In this example the mp3blocker policy is being applied to volume data1.
`Filer> fpolicy vol inc add mp3blocker data1`
 Policies can also be applied using wildcards to select multiple volumes²⁴.
5. Set the file blocking policy to 'required'. This step is needed when using native file blocking
`Filer> fpolicy options mp3blocker required on`
6. Set the operations and protocols to be monitored by the policy. The following example will enforce the mp3blocker policy on both CIFS and NFS protocols for file create and file rename operations.
`Filer> fpolicy monitor set mp3blocker -p cifs,nfs create,rename`
7. Enable the file blocking policy
`Filer> fpolicy enable mp3blocker -f`
8. Confirm the file blocking policy has been implemented
`Filer> fpolicy show mp3blocker`
 File policy mp3blocker (file screening) is enabled.
 No file policy servers are registered with the filer.
 Operations monitored:
 File create,File rename
 Above operations are monitored for NFS and CIFS
 List of extensions to screen:
 MP3
 List of extensions not to screen:
 Extensions-not-to-screen list is empty.
 Number of requests screened : 0
 Number of screen failures : 6
9. Confirm the volumes to which this policy is applied
`Filer> fpolicy vol inc show mp3blocker`
 List of volume specifications to screen:
 Data1

This completes this example of configuring and enabling native file blocking.

²⁴ Refer to http://now.netapp.com/NOW/knowledge/docs/ontap/rel7251/html/ontap/filesag/accessing/task/t_oc_accs_screening_by_volume.html for details.

11.7.4.4 End user experience

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

An example of the error message experienced by a Windows client when a user tries to copy an mp3 file (test25.mp3) to a volume with the mp3blocker policy enabled (using the example from the previous section) is shown below:




Figure 12 - File blocked client message

11.7.5 Flash Cache

Flash Cache (formerly known as PAM II) is an intelligent cache that combines software and hardware within NetApp storage controllers to increase system performance without increasing the disk drive count. Flash Cache is implemented with software features in Data ONTAP and PCIe-based modules with either 256 or 512GB of Flash memory per module. The capacity of the module is dependent on the hardware platform upon which it is being installed. The modules are controlled by custom-coded Field Programmable Gate Array processors. Multiple modules may be combined in a single system and presented as a single unit.

PCIe slot assignment best practices for Flash Cache can be found in the Release Notes for the [FAS3200](#) and [FAS6200](#) series platforms.

11.7.6 Flash Pools

Formally known as hybrid aggregates, Flash pools are supported with ONTAP 8.1.1 and above. Flash pools allow the use of different tiers of storage, usually SSD and SAS or SSD and SATA in the same storage pool or aggregate.

At this time Flash Pools are not certified for use within Citi.

11.8 Alerting and Monitoring

11.8.1 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 webproxy server so that HTTPS posts can be sent to NetApp.

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 Citi 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. If you wish to receive email notification internally then ensure the appropriate value for an SMTP server for your region is defined in the autosupport.mailhost 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.

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. Refer to the [autosupport section](#) of chapter [14.3](#) for details on autosupport configuration standards.

11.8.2 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 [14.2](#) of this document.

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.

Due to bug 589098, if updating traphosts by adding new traphosts and deleting old traphosts then it is imperative to reinitialize the SNMP daemon using `snmp init 0` followed by `snmp init 1`.

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 init 1</code>	Initialize SNMP daemon

Table 28 - SNMP configuration settings

11.8.2.1 Alerting on disk space utilization

Starting with ONTAP 8.0.1P5 NetApp have introduced two hidden registry options that correct gaps and deficiencies in the ONTAP operating system's ability to alert on disk space utilization events via SNMP. These new parameters are not enabled by default and need to be configured as part of the Citi build. Use the following process to enable these parameters once you have your storage system running ONTAP 8.0.1P5 or later:

1. Set the console to advanced mode

```
Filer1> priv set advanced
```

2. Confirm presence of the new registry parameters

```
Filer1> registry walk options.thresholds
options.thresholds.fsFull=98
options.thresholds.fsNearlyFull=95
options.thresholds.fsReportEachFullVolumeChange=off
options.thresholds.fsReportFullAggregates=off
```

3. Enable the two new threshold variables²⁵

```
Filer1> registry set options.thresholds.fsReportEachFullVolumeChange on
Filer1> registry set options.thresholds.fsReportFullAggregates on
```

4. Confirm new variable values have been applied

```
Filer1> registry walk options.thresholds
options.thresholds.fsFull=98
options.thresholds.fsNearlyFull=95
options.thresholds.fsReportEachFullVolumeChange=on
options.thresholds.fsReportFullAggregates=on
```

5. These registry entries take effect immediately. No reboot of the storage system is required. Updates and changes to the registry are persistent across reboots.

By default ONTAP will use the same utilization thresholds of 95% and 98% when evaluating both aggregate and volume level utilization rates. Optionally the default threshold values can be updated and / or customer threshold values can be created on a per volume or per aggregate basis. Before embarking on such an undertaking take a moment to review the advantages and disadvantages of such a highly customized approach:

- Changing the system level default values for options.thresholds.fsFull and options.thresholds.fsNearlyFull can lead to a large amount of alerts being generated depending on the utilization rates of the volumes on the systems and the values defined for each of these thresholds
- Creating custom volume or aggregate level entries (example exception were you do not wish to follow the system level default settings) involves significant operational overhead
 - o Two registry entries (full and nearly full) will need to be created for each custom volume or aggregate threshold configuration
 - o If you create new volumes and aggregates on a system and you wish for them to inherit a custom setting applied to another volume or aggregate you still have to create two registry entries and apply them to that new volume or aggregate

²⁵ Note the aggregate level alerting requires both registry settings to be enabled (you cannot alert on aggregate utilization without enabling fsReportEachFullVolumeChange first)

- o Registry entries, volume names and aggregate names are case sensitive in ONTAP so as custom registry entries need to include the name of the volume or aggregate the correct case must be used
- o If a volume or aggregate with a custom threshold is renamed the registry is not updated and the renamed file system will use the system level default values instead until such a time as the registry entries are updated to reflect the new file system names.
- o When a volume or aggregate with a custom threshold is deleted from the system, the registry entries relating to that file system still remain and are required to be removed manually. Deleting registry entries has a certain risk associated with it. Storage administrators are recommended to engage NetApp support to the latest procedures when deleting registry entries. It is also recommended that registry entries be removed during a maintenance greenzone.

A best practice implementation would leverage the `options.thresholds.fsFull` and `options.thresholds.fsNearlyFull` values for all volumes on the system but have custom entries for aggregates which represent a larger pool of storage and greater inherent risk should the aggregate run out of space. Additionally aggregates are more persistent than volumes in the overall life of the storage systems, are created / deleted much less than volumes and are almost never renamed.

Use the following process to set a custom threshold for an aggregate²⁶. In this example we will set a threshold for an aggregate called `aggr1` and implement the best practice when 90% utilization is considered a full for an aggregate:

```
Filer1> priv set advanced
Filer1> registry set options.thresholds.aggr1.fsFull 90
Filer1> registry set options.thresholds.aggr1.fsNearlyFull 85
```

Again use the following command to confirm the settings have taken effect:

```
Filer1> registry walk options.thresholds
```

This setting takes effect immediately so if the defined threshold is less than the actual utilization rate then SNMP traps will be generated within 60 seconds of this setting being configured.

The commands shown above can also be used to modify the values of existing settings.

²⁶ The process is similar for a volume. Simply replace the aggregate name with the name of the volume in question.

There are some important behavioral aspects of disk space utilization alerting on NetApp NAS that require further explanation here:

- Aggregate utilization rates report differently depending on whether thick (space guarantee=volume), thin (space guarantee=none) or a combination is used on the same aggregate.
 - o Where an aggregate is comprised of only thick volumes, aggregate utilization is calculated as the sum of allocated capacity of all volumes on the aggregate
 - o Where an aggregate is comprised of only thin volumes, aggregate utilization is calculated as the sum of the utilization rate of all volumes on the aggregate
 - o Where an aggregator consists of both thin and thick volumes a hybrid calculation is performed
- During a cluster failover scenario the node that is in a degraded state (namely the Filer that is failed over) will not send SNMP alerts relating volume or aggregate utilization while the system remains in this state. Alerts will be sent once again when the cluster is returned to normal operational status (a giveback has been performed and the cluster is up and enabled).

11.8.2.2 SNMP alerting with Tivoli NetCool

When there is an alert event, NetApp Filers will send out an SNMP alert to the traphost configured in the system. In order for the traphost to create a problem ticket (ServiceNow) or send out SMS/ Paging alert, it is important to check that a valid entry exists for the host (filers) could be found in [Aperture](#).

For more information on NetCool Impact Architecture, please refer to the following [link](#).

11.9 Reporting

11.9.1 EMC ViPr SRM

EMC ViPR SRM (Storage Resource Manager) is Citi's strategic storage reporting tool and version 3.5.1 or later is compatible with Data ONTAP 7.3.x and above (ONTAP 7-mode only).

To onboard a NetApp Data ONTAP 7-mode storage system to SRM the following information is required to appropriately configure the storage system:

- The ONTAP Role Based Access Controls (RBAC) policy for EMC SRM
- The name of the functional ID (region specific) responsible for collecting capacity and performance data from the storage system
- The IP address of the EMC SRM collector server responsible for collecting capacity and performance data from the storage system

Follow the steps below to achieve the tasks described above:

1. Run the RBAC script shown in section [14.7.7.1](#) to create the roles and group necessary for SRM functional ID
2. Add the EMC SRM functional ID detailed in section [12.6.2.5](#) to the ONTAP user group `grp_app_emc_srm` using the following command:

```
ontap> useradmin domainuser add <domain\userid> -g grp_app_emc_srm
```

3. Obtain the IP address of the EMC SRM collector server from the Citi SRM Support team and add the address to the ssh access list within ONTAP. Note the SSH access list is overwritten each time its modified so remember to capture the existing list contents before updating the list, and then confirm the new contents:

```
ontap> options ssh.access (record the output from this command)
ontap> options ssh.access <old_list>,<emc_srm_server_IP>
ontap> options ssh.access
```

Refer to section [14.3](#) for further details on this setting

12 Security Best Practices

This section only applies to all Citi certified NetApp Filers running ONTAP 8.x..

12.1 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 Citi Information Security Standards ("CISS").

This section, together with the Information Security ("IS") IRCAM/ISRP processes, provide the foundation for the assessment of operational risk associated with production NAS environments.

12.2 Security Standards

Citi Information Security Standards ("CISS") have been used as a guideline to produce this document. See appendix A for a summary of the CISS standards. Please refer to the Information Security Services ("ISS") website at <http://www.citigroup.net/informationsecurityservices/> for more details.

12.3 Information Classification Certification

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

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

For Confidential and Confidential PII support, all new applications are required to encrypt confidential PII data in transit (refer to Appendix A.1 of [CISS](#)). For Restricted or Authentication data, encryption of data is required both in-flight and at-rest.

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

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

12.4 Anti-Virus

Three options exist for virus scanning on the NT/CIFS space for NetApp 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 that was replaced with the introduction of #3.

CTI employs McAfee AV Suite of products as its SOE standard for virus scanning. With the availability of VirusScan for NetApp (which was subsequently replaced with VirusScan Enterprise for Storage) on-access scanning, the 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 Enterprise for Storage (VSE-S) 1.0 protects data stored on 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 VSE-S server scans the file and takes appropriate action, if necessary.

VirusScan Enterprise for Storage (VSE-S) 1.0 is deployed to a scanner-server via an ePO deployment task. There is no VSE-S SPAC or TPM package to deploy.

VSE-S requires the following prerequisites to be installed on the scanner server:

- McAfee Agent 4.0 (or later)
- McAfee VirusScan Enterprise 8.7i (or later)

Additional documentation provided by Host Security Engineering regarding these other AV components can be found at the following locations:

[Release Notes for McAfee Agent 4.0](#)

[Release Notes for McAfee VirusScan Enterprise 8.7i](#)

[Release Notes for McAfee VirusScan Enterprise for Storage 1.0](#)

[ePO Configuration Standards for McAfee VirusScan Enterprise for Storage 1.0](#)

12.4.1 Summary of features in VirusScan Enterprise for Storage (VSE-S) 1.0:

Refer to [Figure 1 - NetApp NAS Filer Architecture Design](#) 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 VSE-S tasks from the ePolicy Orchestrator console.
 - o View ePolicy Orchestrator tasks in the local VirusScan server VirusScan Console.




Figure 13 - VSE-S User File Access Data Flow

12.4.2 Multiple to Multiple AV Configurations

In previous revisions, AV scanner to Filer best practices provided excellent performance at the cost of underutilized server hardware. Redeploying AV scanners as a farm maintains the current performance standards while reducing the number of deployed servers. VMware virtualization capabilities will provide additional cost savings because the underlying hardware can be multi-purposed depending on business requirements.

Establishing an AV scanner farm provides higher resource utilization through consolidation and, in case of failure, less performance degradation. Larger AV scanner farms can withstand a member going offline in comparison to a 1:1 AV to Filer ratio configuration. Netapp Filers round-robin scanning requests between each AV scanner that is configured for performance and redundancy.

Based upon information gathered from production servers, CATE Storage Engineering recommends the following configurations based upon a filer (not cluster) to scanner ratio:

Netapp Filers ²⁷	AV Scanners
1-4	2+ ²⁸
5-8	4+
9-16	8+
17-32	16-24 ²⁹

Table 29 - Filer to AV scanner ratio

This table is only a guideline based upon performance data obtained from production environments and NetApp best practices. This guideline provides the flexibility to add or remove AV scanners to maximize scanner utilization as needed. AV scanner performance should be monitored using appropriate performance monitoring tools to ensure Filer responsiveness to end-user requests. AV farm members must be similar in performance characteristics and hardware (cpu, memory, network) since round-robin load balancing does not have insight into individual server utilization. VMware is a candidate to host AV scanners and is fully supported by McAfee when using VSE 8.7i and VSE-S 1.0 or above.

Filer traffic statistics is a direct representation to how much work the filer is performing, compared to raw storage or current users. As a rule-of-thumb, 1.5GHz of CPU power and 384MB RAM provides sufficient resources to scan 6MB/second of Netapp filer traffic. For example, if 8 Netapp filers have an aggregate 120MB/second of network traffic, at least 4 AV scanners would be required. Each AV scanner would need a minimum of 1920MB RAM and 7.5 GHz CPU cycles. A typical server might have 2GB RAM and 2 GHz dual-core CPUs.

The following Netapp filer performance metrics should be monitored:

- System utilization from statit
- Network throughput on all interfaces that are in use from statit
- vscan statistics (stats show vscan_stats, stats show vscan_srv)
- CIFS op-mix (stats show cifs:cifs:cifs_op_pct)

The following AV scanner performance metrics should be monitored via the performance monitoring tools available in the environment:

- \\Server\Processor[*]\%Processor Time
- \\Server\Network Interface [*]\Bytes Received/sec
- \\Server\Network Interface [*]\Bytes Sent/sec

The following steps should be followed to migrate a Filer from independent AV scanners to an AV farm:

1. Identify network information for existing AV scanners, farm AV scanners, and Netapp Filers.
2. Configure AV farm scanner ePO organizational containers with additional Netapp Filers.
3. Unconfigure independent AV scanner ePO organization containers by removing the migrated Netapp Filers.

Note: At least two AV scanners should be listed at all times for fault redundancy. This is to ensure redundancy in case of an AV scanner failure. For 200 and 2000 series Filers, a single AV

²⁷ Note: This is per Filer, not Cluster which consists of two Filers.

²⁸ Minimum of 2 AV scanners for fault tolerant requirements.

²⁹ Contact CSBE if AV scanner farm needs to expand past 24 nodes.

scanner may be deployed if all the businesses using the storage device agree to the reduced redundancy of a single AV server.

For reference, please consider the follow resources:

[Netapp Antivirus Sizing Guide](#)

[Netapp Antivirus Best Practices](#)

12.4.3 VirusScan Enterprise for Storage 1.0 ePO NetApp Policies configuration

Citi utilizes McAfee ePolicy Orchestrator to manage anti-virus configurations centrally. The standards regarding this configuration policy are detailed in the Host Security Engineering documentation [ePO Configuration Standards for VSE-S 1.0](#).

12.4.4 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
Virus scanning is enabled
```

```
filer1> Fri Sep 19 12:10:45 EDT [filer1: vscan.enable:info]: CIFS: Virus
scanning has been enabled.
```

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. Disabling mandatory scanning is the preferred options as the data will scanned on the client side as well.

- b. Set vscan options client_msgbox off

Configures the 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 and “vscan options”

Displays vscan settings and provides summary information about scan requests.

```
filer1> vscan
Virus scanning is enabled.
```

Virus scanners(IP and Name)	P/S	Connect time (dd:hh:mm)	Reqs	Fails
10.40.43.150	\\NASMGMTGDSE02 Pri	00:13:50	0	0
10.40.43.155	\\NASMGMTGDSE03 Pri	00:13:38	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, OV?, PIF, POT, PP?, RAR, RTF, SCR, SHS, SMM, SWF, SYS, VBS, VS?, VXD, WBK, WPD, XL?, XML


```
List of extensions not to scan:
Extensions-not-to-scan list is empty.
```

```
Number of files scanned: 7
Number of scan failures: 0
```

```
filer1> vscan options
vscan options timeout:      10 sec
vscan options abort_timeout: 50 sec
vscan options mandatory_scan off
vscan options client_msgbox off
```

12.4.4.1 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 removing.
- Replacing the file types in the extensions list:
vscan extensions set ext[,ext...]” , where “ext” is the extension you want to set.

12.4.5 NetApp Data ONTAP Anti-Virus Service Help

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

```
vscan help
```

12.4.6 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>
```

12.4.7 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: CIFS_VScan: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.

12.5 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.

12.5.1 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 CATE Storage Engineering 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.

12.5.2 Licensing Disk Sanitization

The disk sanitization feature for Data ONTAP 7.0 and above 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.

12.5.3 Sanitizing Disks

You can sanitize any disk that has “spare” status. This includes disks that exist on the Filer 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 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. For versions of ONTAP up to and including 7.1 the default value is 6 cycles. For version 7.2 and above the default value is 1. The maximum value is 7 cycles. **CTI GIS Secure Media Disposal Guidelines require this value to be set at 7.**
 - 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 Citi required seven cycles (for a total of 21 overwrites) to the specified disks, 7.6, 7.7, and 7.8.

```
disk sanitize start -c 7 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 21 pattern overwrite passes of disk 7.3.

```
disk sanitize start -p 0x55 -p 0xAA -p 0x37 -c 7 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 22 passes times 30 minutes, or 760 minutes or 12.75 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
```

12.5.4 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.
```

12.6 CITMP Information Security Policy Convergence

12.6.1 Information Ownership

The head of Distributed Engineering has information ownership of NAS products and all related engineering documents. NAS products are available in the Citi 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://catecollaboration.citigroup.net/domains/platstor/gdse/products/techstdshmpgs/netappfiler.aspx>

12.6.2 Identification and Authentication

12.6.2.1 Verify Session Authentication

NetApp Data ONTAP requires user authentication for every console or Systems Manager system administration session. Users are authenticated by their respective userid.

With ONTAP 7.3 and above admin accounts can fall into two categories: domain based accounts and local accounts. To simplify account administration within ONTAP CATE Storage and Backup engineering recommends using domain based accounts for authentication to storage management services. At the time of writing this document it is recommended that functional IDs should remain as local admin accounts.

12.6.2.2 ONTAP Domain Based Admin accounts

As of ONTAP 7.3 and above Windows domain based accounts are supported for authentication. This requires the storage controller to be added to an Active Directory domain (see section [11.4.1.2](#) for details). While this feature was available in earlier versions of ONTAP it did not support the Citi required mandatory NTLM v2 setting.

To add a domain based user or group use the `useradmin domainuser add` command in ONTAP.

The recommended approach when implementing domain based authentication is to use a domain based group and add it to the local groups within ONTAP to further simplify storage administrators account management.

For the System Administrator role discussed in section [14.7](#) it is recommended that you use the same Windows Global Group discussed in section [12.6.3.6](#) that allows administrators full control over vol0 via CIFS. For all other granular ONTAP roles that are used in your environment a new global group will be required to be created in the Windows domain.

As an example let's review the command necessary to add the existing North America Windows global group with CIFS full control over vol0 to the necessary local ONTAP groups to create domain based authentication to the System Administrator role:

```
useradmin domainuser add nam\NetApp_Admins -g "grp_sys_admins_1,grp_sys_admins_2, grp_sys_admins_3, grp_sys_admins_4,power users"
```

12.6.2.2.1 Behavioral aspects of Domain based admin accounts and ONTAP

1. Domain based accounts can be used for SSH (CLI) and Systems Manager (GUI) authentication. Console and Service Processor authentication cannot delegate to an external authentication service. Where Console and Service Processor authentication is required use of a local Functional ID is recommended. By delegating SSH and GUI authentication to the domain, password management within ONTAP is simplified and the gains from this outweigh the occasional opportunities where access to the system console is required.
2. Domain based account lockouts or account disabling actions are instantaneously detected via ONTAP. Account validation is not only conducted at the start of an ONTAP management session but also frequently during session activities.

12.6.2.2.2 ONTAP 8 diag user

A diagnostic account, named “diag,” is provided with ONTAP. You enable the diagnostic account to perform troubleshooting tasks in the (FreeBSD) systemshell. The diagnostic account and the systemshell are intended only for low-level diagnostic purposes and should be used only with guidance from technical support.

Upon upgrading from ONTAP 7.x Storage Administrators are required to unlock, configure the password for the diag privileged functional ID (using the standard functional ID password management tools and processed) and relock the account.

The diagnostic account is the only account that can be used to access the systemshell, through the advanced command systemshell. The diagnostic account is disabled by default. The account must be enabled and its password set before using it. Neither the diagnostic account nor the systemshell is intended for general administrative purposes.

With guidance from NetApp Technical Support the diagnostic account can be enabled and used, but first requires the shell privilege level to be raised to advanced. The following examples show the procedure to enable the diagnostic account, set the password, login to the systemshell, logout of the systemshell and disable the diagnostic account.

Use the following process to determine the current state of the diagnostic account:

```
11c2-fas3240-01> priv set advanced
11c2-fas3240-01*> useradmin diaguser show
Name: diag
Info: Account for access to systemshell
Locked: yes
```

Use the following process to enable use of the diagnostic account:

```
11c2-fas3240-01*> useradmin diaguser unlock
11c2-fas3240-01*> useradmin diaguser show
Name: diag
Info: Account for access to systemshell
Locked: no
```


Use the following process to set the password for the diagnostic account:

```
11c2-fas3240-01*> useradmin diaguser password

Please enter a new password: *****
Please enter it again: *****
Wed May  4 17:02:02 EDT [11c2-fas3240-01: passwd.changed:info]: passwd
for user 'diag' changed.
```

Use the following process to login to the systemshell (note the system shell is designated by the '%' sign in the prompt):

```
11c2-fas3240-01*> systemshell

Data ONTAP/amd64 (11c2-fas3240-01) (ttyp0)

login: diag
Password: *****

WARNING: The system shell provides access to low-level
diagnostic tools that can cause irreparable damage to
the system if not used properly. Use this environment
only when directed to do so by support personnel.

11c2-fas3240-01% whoami
diag
```

Use the following process to log out of the systemshell

```
11c2-fas3240-01% exit
Logout
```

Use the following process to disable use of the diagnostic account:

```
11c2-fas3240-01*> useradmin diaguser lock

11c2-fas3240-01*> useradmin diaguser show
Name: diag
Info: Account for access to systemshell
Locked: yes
```

12.6.2.3 ONTAP Local Admin accounts

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.

12.6.2.3.1 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 command line interface, service processor or the Systems Manager user interface. All three utilities can be utilized for making configuration changes. As of NetApp Data ONTAP 7.1 and above, 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>
```

12.6.2.3.2 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.
- Passwords must contain a minimum of one uppercase and one lowercase character
- The login name is case-sensitive.
- The login name cannot contain any of the following characters:
* \ : , = / > [< | + ?] ; space

12.6.2.3.3 Automatic disabling of accounts

NetApp Data ONTAP currently has the capability to disable local accounts based on invalid login attempts. The system log also captures invalid or failed login attempts, when a user has 6 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' expanded view command which will report on all accounts on the storage system. For example:

```
useradmin user list -x
```

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 must not be used for day-to-day or automated/scripted device management

12.6.2.3.4 Password Aging

Automated password expiration should be enforced after 90 days and account inactivity expiration should be enforced after 120 days. Password expiration policy features are available from ONTAP 7.1 and above. 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.

Functional IDs may follow a different password ageing cycle. Refer to section [12.6.2.5](#) for further information in Functional IDs.

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 90
```

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 or the 'useradmin user list -x' query. Centralized reporting is available with Operations Manager 3.7 or above via the 'Host Local Users, All' report. 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.

ONTAP does not have the ability to query the age of an administrator password, nor does it warn that passwords are about to expire. Depending on how you log in to the Filer the event that identifies your password as expired will differ:

- For the cli –

```
filer1> version
Permission denied, user nn12345 does not have access to version
filer1> Tue Jul 17 15:29:59 EDT [filer1: passwd.expired:warning]: user 'nn12345' has an expired password and must change it before attempting further actions.
```

12.6.2.4 Role Based Access Controls

With ONTAP 7.1 or above, 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 Citi NetApp support region must move to a segregated administration model after deploying ONTAP 7.1 or later.

In order to implement RBAC, CATE Storage Engineering recommends using the following process:

- Configure the Filer for RBAC using the RBAC scripts discussed in section [14.7.4](#).
- There is no need to deploy RBAC policies on a per vFiler basis as all Filer management is executed from the physical Filer (vFiler0). The only accounts required on a vFiler are the Root and Administrator built-in functional IDs.
- Create ISA accounts using 'useradmin user add <user name> -c "<comment>" -n "<user full name>" -g grp_acc_admins'
- Once the ISA team has become experienced in the use of account administration within ONTAP, modify all the SOE ID accounts currently in the 'Administrators' group so they become 'Storage Administrators'; thereby removing their ability to run the useradmin command 'useradmin user add <user name> -c "<comment>" -n "<user full name>" -g grp_sys_admins_1,grp_sys_admins_2,grp_sys_admins_3,grp_sys_admins_4,Power Users'
- Creating the ISA and Storage Administrator roles meets the Citi requirements for segregation of duty. If you wish to add other users to the granular subset of roles then do so as necessary

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 Citi 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 G: Role Based Access Control Capabilities are defined on the Filer.
- 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.

12.6.2.5 Functional IDs

The purpose of this section is to provide a list of the Functional IDs required to implement Unified Manager within a NetApp NAS Storage environment as well as other functional IDs associated with the day to day processes on a NetApp storage system. The table below provides information on the Functional IDs according to new Functional Account technical requirements recently published by the CTI Information Security Administration (ISA) organization at http://www.citigroup.net/informationsecurityservices/doc/istg/TechnicalRequirementsforFunctionalAccountsv1_2.pdf. Each account is listed, along with its purpose; whether the account is mandatory or optional for a given implementation of OM (depending on use case and product features leveraged); the scope of the account (domain [Windows] or local); whether the account is used interactively or not; whether the account is privileged or not; and finally the recommended password change frequency based upon interactivity status and privilege level.

The actual name of some Windows domain functional accounts may vary when used in a production environment based on Operational best practices and naming conventions employed by the Active Directory ISA group responsible for managing functional IDs for each region.

Data ONTAP today uses five local and one Active Directory domain SOE functional ID's for specific processes:

Account Name	Root
Purpose	System super-user account created when the ONTAP operating system is installed
Required	Yes
Scope	Local – this account is created locally on each storage controller
Interactivity	Interactive – this is an operating system super user account
Privileged	Privileged – this system account has privileged access to each Filer or vFiler
Password Change Frequency	A regular password change is required every 6 months

Account Name	Administrator
Purpose	System super-user account created when the CIFS protocol is enabled in the ONTAP operating system
Required	Yes
Scope	Local – this account is created locally on each storage controller
Interactivity	Interactive – this is an operating system super user account for the CIFS protocol
Privileged	Privileged – this system account has privileged access to each Filer or vFiler
Password Change Frequency	A regular password change is required every 6 months

Account Name	Diag
Purpose	The diagnostic account is the only account that can be used to access the (FreeBSD) systemshell, through the advanced command systemshell. The diagnostic account is disabled by default. Neither the diagnostic account nor the systemshell is intended for general administrative purposes.
Required	Yes
Scope	Local – this account is built-in locally on each storage controller
Interactivity	Interactive – this is an operating system diagnostic user account
Privileged	Privileged – this system account has privileged access to underlying O/S kernel
Password Change Frequency	A regular password change is required every 6 months

Account Name	Dfmproxy
Purpose	Used by NetApp's Unified Manager management software to allow UM to authenticate to the storage system to retrieve performance and system status information
Required	Yes
Scope	Local – this account is created locally on each storage controller
Interactivity	Non-interactive – this is a system account used for backend processes
Privileged	Privileged – this system account has privileged access to each Filer
Password Change Frequency	A regular password change is not required

Account Name	Backup
Purpose	System account used to perform NDMP backups using NDMP hashed passwords
Required	Yes
Scope	Local – this account is created locally on each storage controller
Interactivity	Non-interactive – this is a system account used for backend processes
Privileged	Privileged – this system account has privileged access to each Filer
Password Change Frequency	A regular password change is not required

Account Name	[domain]\Srv_filerVirusScan
Purpose	Allows the McAfee Anti-Virus proxy server to authenticate with the storage device for virus scanning of CIFS (Windows) user data
Required	Yes
Scope	Domain – this domain account is added to the local administrators group on each storage device
Interactivity	Non-interactive – this is a system account used for backend processes
Privileged	Privileged – this system account has privileged access to each Filer
Password Change Frequency	A regular password change is not required

Account Name	[domain]\gnocsisvcsrcm01
Purpose	Allows EMC ViPR SRM to collect capacity and performance data from the NetApp storage system
Required	Yes
Scope	Domain – this domain account is added to the local “grp_app_emc_srm” group on each storage device
Interactivity	Non-interactive – this is a system account used for backend processes
Privileged	Non-Privileged – this system account has elevated but non-privileged access to each Filer
Password Change Frequency	A regular password change is not required
Account owner	Level 3 NAS support
Password Management	Dual controlled password entry with L3 NAS Support and SRM teams

Creation and use of functional id's should follow the best practices described in section [12.6.2.2](#). All Functional Ids audit events are captured in the “/etc/log/auditlog” file in the system volume (/vol/vol0 or \\filername\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 audit logs must be forwarded to HP ArcSight / SEMS via syslog (see section [12.6.7.3](#)).

12.6.3 Authorization and Access Control

12.6.3.1 User Access Controls

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

- System Administration access controls
- File system access controls

12.6.3.2 System Administrator Access Controls

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

12.6.3.3 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.

12.6.3.4 Windows Active Directory Support

From a Windows domain perspective, each filer looks and behaves like a Windows Server including NetBIOS name resolution. 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 2008 Server and Windows 2012 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.

12.6.3.5 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 Windows Computer Manager to manage local user accounts.
- Local groups on filers can be created and modified using Windows Computer Manager or the useradmin group command.

12.6.3.6 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. NetApp Filers are pre-configured with c\$, etc\$ and home shares on the root volume. The etc\$ and home shares are never used and the directories are already accessible 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 Citi 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_Admns", and "Domain\NetApp_Users" respectively.
 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 Windows Computer Manager to 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 12.6.2.2, ***the "administrator" is a system account, its use should be managed with a Change Management system, and should be the only specific 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 Computer Manager to remove "[DOMAIN]/administrators" from the local "Administrators" group.
 - b. Add global group "[Domain]\NetApp_Admns" 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 an Active Directory 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 can be done through the CLI, for example:

```
Filer> cifs access c$ NAM\NetApp_Users "Full Control"
```

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 “[Filer]\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 14 - 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.

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

12.6.3.7 System Administration Entitlements

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

12.6.3.8 File System Access Entitlements

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

12.6.4 Restricting management physical host access

12.6.4.1 System Administration Access Controls

System administration access controls must be implemented whereby Filer management is restricted to a limited number of specific management stations. At a minimum this should consist of the following system count, although some regional teams may require more SOE terminal servers per group depending on the size of the group and the performance of the terminal server:

- 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 package installed
 - o These terminal servers should be specifically used for NAS management and should be configured in Application Terminal Services mode (where supported).
- 3 Unified Manager servers per support group
 - o One installed at the main production site
 - o One installed at the main contingency site
 - o One installed at the same site to conduct performance monitoring

12.6.5 Traceable Access Entitlements

12.6.5.1 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 09:19:21 EDT [netapp02: rc:auditlog]: eg55388:OUT:console shell:netapp02*
Tue Apr 13 09:19:21 EDT [netapp02: rc:auditlog]: eg55388:IN:console shell:
Tue Apr 13 09:19:28 EDT [netapp02: rc:auditlog]: eg55388:OUT:console shell:netapp02*
Tue Apr 13 09:19:28 EDT [netapp02: rc:auditlog]: eg55388:IN:console shell:vscan scanner server
Tue Apr 13 09:19:28 EDT [netapp02: rc:auditlog]: eg55388:OUT:vscan:The following commands are available: for more information
Tue Apr 13 09:19:28 EDT [netapp02: rc:auditlog]: eg55388:OUT:vscan:type "vscan help <command>"
Tue Apr 13 09:19:28 EDT [netapp02: rc:auditlog]: eg55388:OUT:vscan:extensions      off      options scanner
Tue Apr 13 09:19:28 EDT [netapp02: rc:auditlog]: eg55388:OUT:vscan:help      on      reset
Tue Apr 13 09:19:33 EDT [netapp02: rc:auditlog]: eg55388:OUT:console shell:netapp02*
Tue Apr 13 09:19:33 EDT [netapp02: rc:auditlog]: eg55388:IN:console shell:vscan scanners
Tue Apr 13 09:19:33 EDT [netapp02: rc:auditlog]: eg55388:OUT:vscan:No virus scanners are registered with the filer.
Wed Apr 14 09:14:51 EDT [netapp02: rc:auditlog]: eg55388:OUT:console shell:netapp02*
Wed Apr 14 09:14:51 EDT [netapp02: rc:auditlog]: eg55388:IN:console shell:
Wed Apr 14 09:14:51 EDT [netapp02: rc:auditlog]: eg55388:OUT:console shell:netapp02*
Wed Apr 14 09:14:51 EDT [netapp02: rc:auditlog]: eg55388:IN:console shell:priv set
Wed Apr 14 09:14:51 EDT [netapp02: rc:auditlog]: eg55388:OUT:console shell:netapp02*
Wed Apr 14 09:14:51 EDT [netapp02: rc:auditlog]: eg55388:IN:console shell:useradd eg55388
Wed Apr 14 09:14:51 EDT [netapp02: rc:auditlog]: eg55388:OUT:useradmin:New password:
Wed Apr 14 09:14:51 EDT [netapp02: rc:auditlog]: eg55388:OUT:useradmin:Password must contain at least 2 alphabetic characters and
Wed Apr 14 09:14:51 EDT [netapp02: rc:auditlog]: eg55388:OUT:useradmin:at least 1 numeric or special character
Wed Apr 14 09:14:51 EDT [netapp02: rc:auditlog]: eg55388:OUT:useradmin:New password:
Wed Apr 14 09:14:51 EDT [netapp02: rc:auditlog]: eg55388:OUT:useradmin:Password is too short - must be at least 6 characters
Wed Apr 14 09:14:51 EDT [netapp02: rc:auditlog]: eg55388:OUT:useradmin:New password:
Wed Apr 14 09:14:51 EDT [netapp02: rc:auditlog]: eg55388:OUT:useradmin:Password is too short - must be at least 6 characters
Wed Apr 14 09:14:51 EDT [netapp02: rc:auditlog]: eg55388:OUT:useradmin:Too many failures - try later
Wed Apr 14 09:14:51 EDT [netapp02: rc:auditlog]: eg55388:OUT:console shell:netapp02*
Wed Apr 14 09:14:51 EDT [netapp02: rc:auditlog]: eg55388:IN:console shell:
Wed Apr 14 09:14:51 EDT [netapp02: rc:auditlog]: eg55388:OUT:console shell:netapp02*
Wed Apr 14 09:14:51 EDT [netapp02: rc:auditlog]: eg55388:IN:console shell:useradd eg55382
Wed Apr 14 09:14:51 EDT [netapp02: rc:auditlog]: eg55388:OUT:useradmin:New password:
Wed Apr 14 09:14:51 EDT [netapp02: rc:auditlog]: eg55388:OUT:useradmin:Retype new password:
Wed Apr 14 09:14:51 EDT [netapp02: rc:auditlog]: eg55388:OUT:useradmin:User <eg55382> added.
Wed Apr 14 09:14:51 EDT [netapp02: rc:auditlog]: eg55388:OUT:console shell:netapp02*
Wed Apr 14 09:14:51 EDT [netapp02: rc:auditlog]: eg55388:IN:console shell:
  
```

Figure 15 - NetApp Data ONTAP “/etc/log/auditlog”

12.6.5.2 File System Access

CIFS file access entitlements are captured through domain authentication.

12.6.6 Confidentiality and Integrity

12.6.6.1 Approved Anti-Virus

NetApp Filers configured for CIFS 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. Citi SOE certifies McAfee VirusScan as its standard anti-virus software. Refer to section [12.4](#) for more information on Anti-Virus configuration.

12.6.6.2 Support by Approved and Certified Vendor

NetApp is a certified vendor in the current Citi Technology Catalog. CATE Storage and Backup Engineering owns and maintains all current 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://catecollaboration.citigroup.net/domains/platstor/gdse/products/techstdshmpgs/netappfiler.aspx>. The NetApp Data ONTAP operating system is also in the Citi Vulnerability Threat Matrix (VTM) inventory, and is scrutinized for UNIX and Windows vulnerabilities (see section [12.6.8.4](#)).

12.6.6.3 Disposal of failed disks

NetApp filers use 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 Citi "Medium Risk" Confidential information. Currently, all failed media should be processed in accordance Global Data Center Secure Media Disposal policies.

12.6.7 Incident Detection and Response

12.6.7.1 Login Violations

As of version 7.1 and above 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 [12.6.2.3.3](#).

12.6.7.2 Audit Trails

NetApp Data ONTAP captures audit events for both:

- System administration, and
- File system access.

12.6.7.3 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 [12.6.3](#) 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 configure the syslog daemon on the Filer to send audit logs to two HP ArcSight/SEMS syslog appliances. The Data ONTAP SEMS rules are used to configure the HP ArcSight appliance. The modifications to the syslog.conf file on the NetApp Filer is shown below:

Filer Syslog Configuration

/etc/syslog.conf

```
.....
#Start of NetApp Auditlog syslog export
cmdsaudit.auditlog      @10.40.42.49    # Prod SEMS appliance
cmdsaudit.auditlog      @10.40.45.60    # DR SEMS appliance
#End of NetApp Auditlog syslog export
#This will send all entries as they are written to the auditlog to the remote
syslog servers 10.40.42.49 and 10.40.45.60
.....
# Log messages of priority info or higher to the console
# Log messages of priority debug or higher to the messages log - required for
NFS mountd trace logging
*.info                  /dev/console
*.debug                 /etc/messages
.....
```

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

12.6.7.4 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. Auditable events must be set on files and directories on the storage system by setting the level of auditing via Windows Explorer
3. 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 “adtlog.evt” files, and does not relate to any operational concerns of the system. CATE Storage and Backup Engineering 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 16 - NetApp Data ONTAP "adtlog.evt" file Event View

12.6.7.5 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.

12.6.8 Administration

12.6.8.1 Information Security Administration (ISA)

The ISA role is introduced in ONTAP 7.1 and above. Please refer to Appendix G: Role Based Access Control Capabilities for details as to ISA capabilities. CATE Storage and Back Engineering recommends a phased approach to implementing RBAC whereby ISAs are allowed to become acclimated to the product and processes before storage administrators have the rights removed that relate to create/modify/delete users and change other user passwords processes.

12.6.8.2 Security Standards

CATE Storage Engineering in cooperation with CTI 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.

12.6.8.3 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.

12.6.8.4 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 Citi's VTM database as a UNIX system. Please refer to the VTM website for the latest vulnerability resolutions:

<https://stcprod001.nj.smb.com/eii/main/program/index.aspx?programid=231>

12.6.8.5 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 functionality that enables support for SSH and SSL access to the filer system console as part of one of the core features of ONTAP.

12.6.8.6 Secure Filer Access (SecureAdmin)

SecureAdmin is a part of Data ONTAP and 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

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

12.6.8.7 Configuring SSH access to the Filer

This section details the various tasks involved in securing the SSH configuration of the NetApp storage controller.

12.6.8.7.1 Configuring the ONTAP integrated SSH server

Use the following procedure to configure the integrated SSH server in ONTAP:

1. Login to the ONTAP CLI via the system console
2. Temporarily disable SSH by running the command

```
Filer> secureadmin disable ssh
```

3. Initiate the SSH setup operation with the following command

```
Filer> secureadmin setup ssh
```

```
SSH Setup
```

```
-----
```

```
Determining if SSH Setup has already been done before...no
Do you want to proceed? [no] yes
```

```
SSH server supports both ssh1.x and ssh2.0 protocols.
```

```
SSH server needs two RSA keys to support ssh1.x protocol. The host key is
generated and saved to file /etc/sshd/ssh_host_key during setup. The server
key is re-generated every hour when SSH server is running.
```

```
SSH server needs a RSA host key and a DSA host key to support ssh2.0
protocol.
```

```
The host keys are generated and saved to /etc/sshd/ssh_host_rsa_key and
/etc/sshd/ssh_host_dsa_key files respectively during setup.
```

```
SSH Setup will now ask you for the sizes of the host and server keys.
```

```
For ssh1.0 protocol, key sizes must be between 384 and 2048 bits.
```

```
For ssh2.0 protocol, key sizes must be between 768 and 2048 bits.
```

```
The size of the host and server keys must differ by at least 128 bits.
```

4. Specify the following key sizes:
 - a. SSH v1 host key = 2048 bits
 - b. SSH v1 server key = 1536 bits
 - c. SSH v2 host keys = 2048 bits

```
Please enter the size of host key for ssh1.x protocol [768] :2048
```

```
Please enter the size of server key for ssh1.x protocol [512] :1536
```

```
Please enter the size of host keys for ssh2.0 protocol [768] :2048
```

5. Confirm the key sizes:

```
You have specified these parameters:
```

```
host key size = 2048 bits
```

```
server key size = 1536 bits
```

```
host key size for ssh2.0 protocol = 2048 bits
```

```
Is this correct? [yes] y
```

```
Setup will now generate the host keys. It will take a minute.
```

```
After Setup is finished the SSH server will start automatically.
```

6. You will receive a stream of warning messages regards unprotected private key files, with permissions 0777 being too open for the following files:
 - a. /etc/ssh/sshd/ssh_host_key
 - b. /etc/ssh/sshd/ssh_host_rsa_key
 - c. /etc/ssh/sshd/ssh_host_dsa_key

These messages can be ignore for a moment. The keys will be secured in an upcoming step.

7. Once the keys are created you will receive a confirmation message similar to the following:

```
Tue May  3 22:03:16 EDT [11c2-fas3240-02:
secureadmin.ssh.setup.passed:info]: SSH setup is done and ssh2 is enabled.
Host keys are stored in /etc/ssh/sshd/ssh_host_key, /etc/ssh/sshd/ssh_host_rsa_key,
and /etc/ssh/sshd/ssh_host_dsa_key.
```

For other [SSH configuration options](#) refer to the SSH section of chapter 14.3.

12.6.8.7.2 Securing the ONTAP private host keys

The section details to steps necessary to secure the private host keys from the warnings experienced in step 6 of the previous section.

1. From Windows Explorer map a drive to [\\FilerName\c\\$](#) using an account that is a member of the Filer's Local Administrators group (where FilerName is replaced with the host name of the NetApp storage controller)
2. Navigate to \etc directory and then select properties on the \etc\sshd directory
3. Break security inheritance from the parent directories and ensure the local Administrators group has full control of the \etc\sshd directory and its respective sub-directories and files. No other user or group should be included in this ACL. Using a Windows tools such as Cacl the permissions on the sshd directory on system 11c2-fas3240-02 can be shown accordingly:

```
C:\>cac ls \\11c2-fas3240-02\c$\etc\sshd
\\11c2-fas3240-02\c$\etc\sshd BUILTIN\Administrators:(OI)(CI)F
```

12.6.8.7.3 Configuring the Citi Tectia SSH Client for NetApp connectivity

The Citi standard SSH client is Tectia SSH client. Information about the latest version of Tectia SSH client is maintained under the CATE Platforms and Storage SharePoint site at the following link <https://catecollaboration.citigroup.net/domains/platstor/osunix/wpages/ProductDetail.aspx?Name=Tectia&Publish=1>

The latest client version to be qualified with NetApp NAS storage is Tectia version 6.0.12.50.

To ensure compatibility between the Tectia SSH client and the NetApp Filer disable X11 connection forwarding and agent forwarding. From the CLI is this achieved by using the following syntax:

```
C:\>ssh2 -x -a <hostname>
```

Where:

-x disables the X11 connection forwarding

-a disables the agent forwarding

<hostname> is the hostname of the Filer you are trying to connect to.

When using the GUI X11 forwarding is disabled by default, however NetApp Filer administrators will need to disable the agent forwarding option in order to successfully connect to the Filer. This is achieved using the following process and use the screenshot below for reference:

1. Launch Tectia SSH Client GUI
2. Select Edit menu and then the SSH Tectia Configuration item
3. From the Connection Profiles menu in the left pane select the profile you want to configure
4. In the main window select the 'Tunneling' tab
5. At the bottom of the main (right) pane uncheck the "Allow Agent Forwarding" in the Forwarding Options section
6. Click the 'Apply' button

Ensure these options are configured before you begin saving connection profiles to devices in your environment.




Figure 17 - Tectia Client GUI Agent Forwarding configuration

If you are seeing the following message entry appearing in the ONTAP message log then you have an SSH client that is incorrectly configured and is allowing agent forwarding to the Filer:

```
Thu Mar 24 18:15:02 EDT [11c2-fas3240-01: openssh.invalid.channel.req:warning]:
SSH client (SSH-2.0-6.0.12.50 SSH Secure Shell) from 172.24.66.105 sent
unsupported channel request (10, x11-req).
```

Ensure both GUI and CLI SSH sessions are following the best practices outlined above to avoid such an error message. For further information refer to [NetApp Knowledge Base article 2013692](#).

In addition you need to ensure the appropriate authentication mechanism is configured in your SSH client. This is achieved using the following process and use the screenshot below for reference:

1. Launch Tectia SSH Client GUI
2. Select Edit menu and then the SSH Tectia Configuration item
3. From the Connection Profiles menu in the left pane select the profile you want to configure
4. In the main window select the 'Authentication' tab
5. In the main (right) pane uncheck the "Use Defaults" checkbox
6. Select "Password" from the list and using the 'up arrow' button move this method to the top of the list
7. Click the 'Apply' button

Ensure these options are configured before you begin saving connection profiles to devices in your environment.




Figure 18 - Tectia Client GUI Authentication Mechanism configuration

12.6.8.7.4 Configuring the Citi Tectia SSH Client and ONTAP for publickey authentication

When using public key authentication, Tectia SSH client saves keys in the SecSH format. Exporting public keys in this format to the Filer will result in failed public key authentication as the embedded OpenSSH server on the Filer cannot translate the SecSH format.

Use the following procedure to create the necessary key pairs for publickey authentication.

1. Login to the Windows workstation or server from which you want to access the storage system via publickey authentication. Ensure that the account you login to the Windows system with is the account for which you want to create the public key pair.
2. Ensure the Tectia client is installed on the Windows system
3. Generate a 2048-bit RSA and 2048-bit DSA key pair using the following commands on the Windows system:
 - a. `ssh-keygen2 -t rsa -b 2048`
 - b. `ssh-keygen2 -t dsa -b 2048`

```
C:\>ssh-keygen2 -t rsa -b 2048
```

```
You don't have an existing random seed file. Creating random seed.
```

```
Generating 2048-bit rsa key pair
```

```
3 0.o0o.o0o.o
```

```
Key generated.
```

```
2048-bit rsa, rs01146@CTIREGLL3GN738, Tue May 03 2011 22:16:32 -0500
```

```
Passphrase :
```

```
Again      :
```

```
Key is stored with NULL passphrase.
```

```
(You can ignore the following warning if you are generating hostkeys.)
```

```
This is not recommended.
```

```
Don't do this unless you know what you're doing.
```

```
If file system protections fail (someone can access the keyfile),  
or if the super-user is malicious, your key can be used without  
the deciphering effort.
```

```
Private key saved to C:\Users\rs01146\Application
```

```
Data\SSH\UserKeys\id_rsa_2048_a
```

```
Public key saved to C:\Users\rs01146\Application
```

```
Data\SSH\UserKeys\id_rsa_2048_a.pub
```

```
C:\>ssh-keygen2 -t dsa -b 2048
```

```
Generating 2048-bit dsa key pair
```

```
32 .o0o.o0o.o0
```

```
Key generated.
```

```
2048-bit dsa, rs01146@CTIREGLL3GN738, Tue May 03 2011 22:17:15 -0500
```

```
Passphrase :
```

```
Again      :
```

```
Key is stored with NULL passphrase.
```

```
(You can ignore the following warning if you are generating hostkeys.)
```

```
This is not recommended.
```

```
Don't do this unless you know what you're doing.
```

```
If file system protections fail (someone can access the keyfile),  
or if the super-user is malicious, your key can be used without  
the deciphering effort.
```

```
Private key saved to C:\Users\rs01146\Application
```

```
Data\SSH\UserKeys\id_dsa_2048_a
```

```
Public key saved to C:\Users\rs01146\Application
```

```
Data\SSH\UserKeys\id_dsa_2048_a.pub
```

4. Navigate to the file system location provided in the previous step for the two public key files that were generated, in this example the following:
 - a. C:\Users\rs01146\Application Data\SSH\UserKeys\id_rsa_2048_a.pub
 - b. C:\Users\rs01146\Application Data\SSH\UserKeys\id_dsa_2048_a.pub
5. From a Windows client map a drive to the C\$ share of Filer and then open a Windows command line session and change directory to the newly mapped network drive. Then issue the following commands:
 - a. CD to \etc\sshd
 - b. mkdir <user> where <user> is actually the user ID of the public key created in the previous step
 - c. cd \etc\sshd\<user>
 - d. mkdir .ssh
6. Refer to the “Editing public keys generated by SecureCRT and ssh.com clients” chapter of the [ONTAP 8.0 System Administration Guide](#) for the procedure to manually edit each of the public key files mentioned in step 4 into the format required by ONTAP. Place both converted keys into a single file and save the file as authorized_keys2 into directory path created in step 5 (\etc\sshd\<user>\.ssh).

An example of the contents of a authorized_keys2 file is shown here for demonstration purposes. The contents of your file will be different. There are only two lines here: one starting ssh-rsa and one ssh-dsa. The other lines have been wrapped for presentation purposes.

```
ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAQCAQC5fq/d/bBXuRDVIRkLML9vPlfbDb86Q/5fmb5Qam0UX
E6u0CdJ73GEZ+HPgl/tMslDIg16BmzN1MwgQvTo+FFp5j9Nt3T09+V/dMwSdr5WMhjz0FMq5L
JyvzVmvmr6rRCN8i7iQoyqFXyHwEjCwCrC0I0J0W4bApQMs16KQ6c2l8+z0JoNemo1jMHC4t
EAcq2P7y399Rf9E1FJd5fZoptz3THJfSRd5a1wskSxJUVXF939ajm+vhMKWXYQcWlpNbBYDb
S2vJYew/zn5v3/DiC6wqZpQfWUrKX+9bFduuXW1UTqHLfhfJEn+LjAYIectKWlpankvW0AzBS
+DIm7T/
ssh-dsa
AAAAB3NzaC1kc3MAAABAKzn3hKjC7KqBR13jo7ofFBQ+kLFk56HwFJE1+esJSSEXaC2bRwC7
ZfdET9xB5mHL1RJ0qvV2imRZ9DZimaHdn2U/fIWc9hq4cF7Cng1zbc1tDdQF/qKsm//71kCj1
67YisUAsiQCA6GsyL1sKrNGyYmMHI2Xpcxh+fZAQZSi4Jw4a0+F2NmLBC6Th2k0JScEMZqK
4mSmK0izCU3prnmxz0ic7ez/bMYTERybh20soBwFBl2PsB0mJHR02V8M/r027/GmcgSCZ3b6S
Cp1bt5gjyXXzfdqsJ4PACq5ImwAJgpivzcA1jeAUmiCjIuLk4Wzfn0224US+GXRqG57rfcAA
AAVAJGpoJ9858GD03XlRRJlKhDwXNjJAAABAAJKjJmFRuJ9Aww8RmEMlHR2fqWPz3PntRcHSZ
hJZvoqoEkt7PFAl26EeM2bB6y9FCgHl3SSdZ8+1jALGB/ZM8hbrYntfw7f/FCb2y9r5h9aFbv
VMW1bIYo30d7ER44RlsLIJfcWYtNNU0ADETC05rhhK/6yt+eBz+c2NCTXhn4SCRT8/BDTdT27p
g7KEHIPVeG6u3L0boYmGo35wOS3e2qKRnf87D800tEFUgRMZhZ0wZbfyWP4cLGfXzQZUuLHLZ
C0kY5b9unGSF11EJ7c3ikWse35LyQgd7Ts2baA7A9N5a9WqH7jmUsC/lTxDLAGHID/Jr6ET6i
i5Z4EkvX1VTo0AAAEAVvbAy12YVWhDakkf1L7i4wB8IM/LyGkbj+h3AVn90LEUz4kaxMph3W6
Iru8UTpaZIAQ805iDKFWpS3DiPnoGZ/FPKYpu9Y20AnbAQ5UuaaGX7r2IBleeRnPxhF6ZhtZn
S+Dwfm30/g9Cq1svrGwyE3alxnjW0D34NLtlwxYvA5KBIdWvHeelAzpzAwwPnB0G1X46ngkNy
5XlJiJudsNYZK9m3SoIZiKNEGLs2eYIH/05fXUQbfGT/oA1EZ8/IjsVkyXUemRleQ25CCh9j/
j8/0ETNynBdS+YdaoscLKfZSmRDT35IBZB+4WBmFbreSULj16a877OnbyH9EhdB6+Lsg
==
```

12.6.8.8 Configuring SSL on the Filer's embedded management web server

Use the following procedure to configure SSL for the embedded web server used by Systems Manager and other NetApp and third party management tools:

1. Login to the ONTAP CLI via the system console
2. Temporarily disable SSL by running the command
`Filer> secureadmin disable ssl`
3. Initiate the SSH setup operation with the following command
`Filer> secureadmin setup ssl`
4. Answer the interview questions providing details relevant to your environment, ensuring the following key points
 - a. Days until expires is set to 730 (2 years)
 - b. Key length (bits) is set to 2048

```
Do you want to proceed? [no] yes
Country Name (2 letter code) [US]: US
State or Province Name (full name) [California]: New Jersey
Locality Name (city, town, etc.) [Santa Clara]: Warren
Organization Name (company) [Your Company]: Citigroup
Organization Unit Name (division): CATE - Storage and Backup Engineering
Common Name (fully qualified domain name) [11c2-fas3240-02.namdev.nsrootdev.net]: 11c2-fas3240-02.namdev.nsrootdev.net
Administrator email: gdse@citi.com
Days until expires [5475] :730
Key length (bits) [512] :2048
Tue May 3 23:45:45 EDT [11c2-fas3240-02:
secureadmin.ssl.setup.success:info]: Starting SSL with new certificate.
```

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

For other [SSL configuration options](#) refer to the SSL section of chapter [14.3](#).

As the two year duration of the certificate is traditionally half the length of the operational life of the NetApp storage system, the certificate will need to be regenerated towards the end of the second year (and subsequently two years after is necessary). In order to generate new SSL certificates run the following command:

```
Filer> secureadmin setup ssl -f
```

12.7 TCP/IP Port Requirements

The purpose of this section is to provide an overview of the network protocols and ports required for a successful implementation of a NetApp Filer storage environment on the Corporate Network. The following information is especially important if there is a firewall device between the client device and the storage controller.

Port	Type	Protocol	Source	Target	Direction	Purpose
22	TCP	SSH	Client, Management Server or Unified Manager Server	Filer	Inbound	Device Management Secure console
25	TCP	SMTP	Filer	SMTP relay	Outbound	Device Management Email notifications of autosupport messages
111	UDP TCP	NFSv3	NFS Client	Filer	Inbound	Data Access Portmap a.k.a. rpcbind
123	UDP TCP	NTP	Filer	Time server	Outbound	Device Management Time synchronization
137	UDP	CIFS	CIFS Client	Filer	Inbound	Data Access NetBIOS nameserver for CIFS data
138	UDP	CIFS	CIFS Client	Filer	Inbound	Data Access NetBIOS datagram service for CIFS data
161	UDP	SNMP	Filer	Trap Target	Outbound	Device Management SNMP Monitoring of Filers
443	TCP	HTTPS	Client, Management Server or Unified Manager Server	Filer	Inbound	Device Management Secure User Interface
443	TCP	HTTPS	Filer	Proxy Server	Outbound	Device Management Vendor call home information posted via HTTPS
445	TCP	CIFS	CIFS Client	Filer	Inbound	Data Access CIFS over TCP with NetBIOS framing
514	UDP	SYSLOG	Filer	Syslog server	Outbound	Device Management Transferring syslog messages to centralized syslog server
2049	UDP TCP	NFSv3	NFS Client	Filer	Inbound	Data Access Primary NFS service
2049	TCP	NFSv4	NFS Client	Filer	Inbound	Data Access NFS service (includes locking)
4444	UDP	SNMP	Filer	Filer partner	Inbound or Outbound	Device Management High availability
10000	TCP	NDMP	Backup Server or Filer	Filer	Inbound or Outbound	Device Management NDMP backup to tape or NDMP copy to another Filer
10566	TCP	Snapmirror	Filer	Filer	Inbound or Outbound	Device Management Data replication protocol

Table 30 - Required IP ports

13 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
2.9	- Added links to CATE VMware VDI and VSI design documents in section 8.6 - Removed links for outdate VMware NFS design documents in section 8.6 - Removed out dated content statement at the start of 11.4.2.8 - Removed snpamirror and snapvault licenses from Table 21 - ONTAP licenses required for VMware over NFS - Added multi-store license to Table 21 - ONTAP licenses required for VMware over NFS - Added links to VMware design documents in 11.4.2.8.3 - Removed context from section 11.4.2.8.4.2 and replaced with links to the VMware design documents - Removed naming convention section regarding volume names as this has been superseded by a level III PCM operational standard - Updated section 11.4.2.8.6 so that CIFS shares are not required on top of NFS datastores - Updated section 11.4.2.8.8 to reflect storage level replication is not required - Updated section 11.4.2.8.9 to reflect storage level snapshots are not required - Updated section 11.4.2.8.10 to reflect snapvault is not a supported backup strategy by VMware - Updated section 11.4.2.8.12 to refer to VMware design document regarding hot to provision capacity on the storage system - Updated link to CSBE homepage in Section 4 - Added disable Active Directory SID compression attribute task to the post domain join activities in section 11.4.1.2.4 - Added clarification to when nfs.v4.enable should be set to on or off in section 14.3
4.0	- Document update for ONTAP 8.1.4P1 - Added draft document designation - Updated directives in section Upgrading from a previous release - Changed port recommendations in Table 2 - Supported port roles now that e0M, RLM and SP ports are certified; and added FAS3100 series and FAS6000 series systems to the table now that those systems can be upgraded to ONTAP 8.x. - Added FAS3100 and FAS6000 to Table 3 - Hardware requirements to support 10GB ethernet now that those systems can be upgrade to ONTAP 8.x - Removed details change control entries from all document versions prior to 2.9 - Removed link to NetApp System Configuration Guides in section entitled Network Adapter IO Expansion Slot Assignments and replaced with link to NetApp Hardware Universe - Added statement regarding Arista IP Fabric MTU standards to section entitled Maximum Transmission Unit (MTU) size - Revoked non-approved status for service processor / rlm. Added best practices and network requirements to section entitled Service Processor (SP) / Remote LAN Management (RLM) - Revoked non-approved status for management interface e0M. Added best practices and network requirements to section entitled The Management Port (e0M) - Updated Definitions and abbreviations to include new terms - Updated /etc/rc section to include cautions on appropriate order of entries and not to define the losk interface in the /etc/rc file - Added section regarding retirement of Windows Internet Name Service (WINS) - Added DS14 and FCAL legacy architectures to Storage section to provide reference materials for ONTAP 7.x systems upgrading to ONTAP 8.1 - Added 3TB BSAS / SATA drives to Storage section - Added FAS3100 and FAS6000 series platforms to Storage section - Added new FAS3220, FAS3250, FAS6220, FAS6250, FAS6290 to Storage section

Version	Changes
4.0 (Cont...)	- Corrected links CATE VMware VDI and VSI design documents in Aggregates - Removed justification of 0% aggregate snapshot reserve as that is the new default value in ONTAP 8.1 in section Aggregates - Removed comparison of aggregate and volume level snapshots as this is redundant with a 0% aggregate snapshot reserve - Updated aggregate creation with revised NetApp best practices on disk selection in section Aggregates - Updated maximum aggregate size with ONTAP 8.1 in section Aggregates - Replaced references to Filerview with Systems Manager - Added SMB 2.1 oplock leasing bug 798842 to section CIFS - Added SID compression enablement to section CIFS - Added support for Windows 2012 Domain Controllers to section CIFS - Updated minimum ONTAP version for all Windows clients to ONTAP 8.1.4P1 in section CIFS - Removed WINS server and NetBIOS requirement from CIFS SETUP process in section CIFS - Updated SMB 2.0 ONTAP version recommendations to version 8.1.4P1 or later in section CIFS - Added section regarding SMB 2.1 and known issues to section CIFS - Added section regarding option cifs.tcp_window_size to section CIFS - Added section regarding Appropriate Byte Counting feature of ONTAP 8.1 - Added new parameters ldap.fast_timeout.enable and ldap.retry_delay to sections 14.3 and NFS - Added legacy and new FAS3200 and FAS6200 series platforms to section Scaling concurrent replication processes - Removed Vfiler Planning Worksheet and 3DNS request form template from section Vfiler (multistore) as they redundant at this point - Snaplock and LockVault functionality is added in ONTAP 8.1 so merged the existing ONTAP 7.x best practices for these features in the ONTAP 8 document and updated compliance clock command changes for ONTAP 8.1 - Removed references to R200 in section Lockvault - Removed volume size limitations in section Deduplication as they no longer apply in ONTAP 8.1 - Added chapter for compression (details to be added in later draft version of this document) - Added chapter providing current CATE certification support statement regarding Flash Pools - Updated Information Classification Certification with new statement applying CISS to all storage products - Removed Filerview readonly access configuration as Filerview is retired and replaced with Systems Manager - Added reminder to configure diag privileged user account upon upgrade to ONTAP in section ONTAP 8 diag user - Added upper and lower case minimum character requirements to local password policies - Removed srv_sqlserver_soedb functional ID requirement with the retirement of WebAdmin for NAS inventory from section Functional IDs - Removed requirement to use snapshots to retain audit logs on vol0 now that SEMS / HP ArcSight is deployed in section Functional IDs - Removed global inventory WebAdmin server requirement from section System Administration Access Controls - Removed inventory agent for NetApp Filers chapter in section Administration due to retirement of WebAdmin - Added port 4444 requirement for e0M interface to section TCP/IP Port Requirements - Removed appendix detailing the data schema for the NetApp Inventory Agent due to the retirement of WebAdmin and adjusted Appendix letters after D accordingly - Incorporated Lockvault/SnapLock information into appendix from the ONTAP 7.x solutions guide and updated compliance clock commands

Version	Changes
4.0 (Cont...)	- Added section Appendix J OnCommand Systems Manager for further reference information - Added the following configuration parameters to the configuration matrices in section Appendix C: Data ONTAP SOE Configuration Standards - Options rlm.autologout.enable - Options rlm.autologout.timeout - Options rlm.ssh.access - Options sp.autologout.enable - Options sp.autologout.timeout - Options sp.ssh.access - Options cf.hw_assist.enable - Options cf.hw_assist.partner.address - Options cf.hw_assist.partner.port - Options ip.tcp.abc.enable - Options ip.tcp.rfc3390.enable - Options cf.giveback.auto.enable - Options cf.giveback.auto.after.panic.takeover - Options flexscale.rewarm - Options cifs.smb2_1.branch_cache.enable - Options ldap.fast_timeout.enable - Options ldap.retry_delay - Options cifs.W2K_password_change - Options cifs.wins_servers - Options cifs.tcp_window_size - Options interface.blocked.mgmt_data_traffic - Options nfs.always.deny.truncate - Options nfs.nfs_rootonly - Options nfs.rpcsec.trace - Options nfs.v2.enable - Options nfs.v3.snapshot.active.fsid.enable - Options nfs.v4.id.allow_numerics - Options nfs.v4.snapshot.active.fsid.enable - Options nfs.vstorage.enable - Options security.passwd.rules.minimum.lowercase - Options security.passwd.rules.minimum.uppercase - Scope of the configuration parameter has changed for the following settings: - Options cifs.oplock.enable - Options cifs.smb2.enable - Options cifs.smb2.signing.required - Options cifs.wins_servers - Recommended value of the configuration parameter has changed for the following settings: - Options cifs.rpcfd_timeout - Options cifs.W2K_password_change - Options cifs.weekly_W2K_password_change - The following configuration parameters have been removed from the build: - Options cifs.weekly_W2K_password_change replaced with options cifs.W2K_password_change - Added configuration parameters required on the remote management port (the RLM or SP depending on FAS hardware type) - Added new configuration parameters section to handle any additional settings which fall outside of the usual ONTAP options, ONTAP registry and Remote Management Port parameters in Table 37 - Additional configuration settings

Version	Changes
4.0.1	- Corrected minor typos and updated various references of Operations Manager to Unified Manager - Added DS4246 to Table 7 - Disk shelf firmware supported versions - Added DS4246 to Table 9 - Supported disk drive types - Added DS4246 to Table 10 - Maximum number of disk drives supported in a HA or standalone configuration - Added Unified Manager 5.2P2 dependency warning to Infrastructure requirements - Updated reviewers in Table 1 - Regional stakeholder review - Updated disk shelf firmware latest versions to those included in ONTAP 8.1.4 for SAS based shelves in Table 7 - Disk shelf firmware supported versions - Updated links in References - Updated Figure 1 - NetApp NAS Filer Architecture Design - Moved Interface Favoring standard under single-mode interface groups as it applies to all single-mode interface groups irrespective of port speed - Re-ordered Table 9 - Supported disk drive types with preferred drive types first and updated preferred drives for Grid, VMware VDI and VMware VSI use cases - Reword three-party NDMP configuration and added Remote NDMP to tape to section NDMP Considerations in storage controller scalability - Removed references to Windows 2003 Domain Controllers - Removed SOE Web Browser support section as this is now covered in the System Manager release notes. - Added root volume size includes snap reserve statement and mandated root volume snap reserve - Added minimum recommended drive firmware bundle and disk qualification package requirements to section Disks
4.0.2	- Updated RBAC policies to align with ONTAP 8.1.4 features in section 14.7 - Added bug 526941 advisory to section 11.3.2.1 - Removed draft document designation
4.0.3	Updated references section entries
4.0.4	- Added SMB2 related bugs 645251, 817844 and 649716 to section Windows Interoperability - Added support statement for SMB3 to section CIFS - Updated SMB2 enable recommendations in section 14.3 due to bugs 645251 and 817844
4.1	- Added SMB 2.1 changed default behavior (disabled) with ONTAP 8.1.4P4 in section Windows Interoperability - Added SMB 2 state recommendation with ONTAP 8.1.4P4 in section Windows Interoperability - Updated ONTAP recommended versions in Table 18, Table 19 and Table 20 - Updated recommendation on when SMB 2.0 should be enabled in section 11.4.1.5.1 - Added section 11.4.1.5.3 describing the steps to configure the SMB multi-processor signing optimization - Added SMB 2.1 changed default behavior (disabled) with ONTAP 8.1.4P4 in section 11.4.1.5.4 - Updated recommendation for option cifs.smb2.enable in section 14.3 - Added option cifs.smb2.signing.max_threads to section 14.3 - Updated Table 37 regarding default behavior for disabling SMB 2.1
4.1.1	Removed draft document designation

Version	Changes
4.2	- Updated ONTAP recommended versions in Table 18, Table 19 and Table 20 - Updated SMB 2 multi-processor signing optimization enablement feature in section 11.4.1.5.3 - Added option cifs.smb2.signing.multiprocessing to section 14.3
4.2.1	- Removed draft document designation - Added ONTAP RBAC revert scripts for downgrading / backing out an upgrade change to section 14.7
4.2.2	- Uplift from version 4.2.1 – Added ONTAP 8.1.4P6 details - RFT
4.2.3	RFP
4.2.4	- Added section 14.7.7 detailing RBAC requirements for infrastructure applications integration - Added section 14.7.7.1 detailing RBAC requirements for EMC ViPR SRM integration - Changed document designation to Draft - Added placeholder for EMC ViPR SRM functional ID to section 12.6.2.5
4.2.5	- Removed draft document designation - Added EMC ViPR SRM to section 10.1 Infrastructure requirements - Added details of EMC ViPR SRM functional ID to section 12.6.2.5 - Added EMC ViPR SRM server to list of authorized SSH hosts (options ssh.access) in section - Corrected syntax errors with script in section 14.7.7.1 - Added EMC ViPR SRM RBAC diagram to section 14.7.3.13 - Added section 11.9 detailing reporting tools
4.2.6	- Updated best practices around the support and deployment of SMB2 in order to prevent unexpected premature exhaustion of storage system resources, known to occur in Citi's virtualized desktop environment. Changes made to the wording of section entitled Windows Interoperability - Clarified implications of disabling SMB signing independent of disabling actual SMB version in section 11.4.1.1.2 - Introduced new 40% CPU utilization threshold guideline for SMB2 in section 11.4.1.5.1 - Updated recommended value for option cifs.smb2.enable in section 14.3
4.2.7	Updated CPU utilization threshold guideline to 30% for SMB2 in section 11.4.1.5.1
4.2.8	Updated screen output to reflect SP version 1.3.3P2 in section 11.2.5.5
4.2.9	- Added CIFS terminate and restart requirement to disabling WINS outside of an ONTAP upgrade process in section 11.2.12.1 - Added section 11.4.1.1.1 detailing steps to disable NetBIOS over TCP, thereby enforcing SMB over TCP when communicating with Windows 2012 Domain Controllers
4.3	Removed draft document designation

14 Appendices

14.1 Appendix A: ePO & VirusScan Enterprise for Storage Configuration

Please refer to Host Security Engineering's [ePO Configuration Standards for McAfee VirusScan Enterprise for Storage 1.0](#) document for detailed procedures of configuring VirusScan Enterprise and VirusScan Enterprise for Storage for use with NetApp NAS Storage.

14.2 Appendix B: NetApp SNMP MIB

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

For purposes of certification a trap is considered new to ONTAP if it has been added to the MIB since the previous Citi certified release of ONTAP (8.0 or 7.3.6). [Table 31](#) contains the column definitions used in the SNMP MIB matrix ([Table 32](#)) for NetApp MIB version 2.1.2.

Table 31 - SNMP trap list table definitions

ONTAP Version	The earliest Citi 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 than 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	Citi 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 and which severity of events will invoke operations support call out from your regional Citi command center. 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.
Netcool Severity	Citi 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 and which severity of events will invoke operations support call out from your regional Citi command center
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 32 - SNMP MIB matrix

Introduced ONTAP Version	MIB #	OID / Trap #	Trap Name	Description	# Different Message Strings (sub-types)	Default Severity	Citi Tivoli Severity	Citi Netcool Severity	Auto Clear	Clear OID
6.4.x	1.13	2	userDefined	A polling-style trap built using the 'snmp traps' command on the filer.	n/a	Alert	Harmless	Minor	Yes	2
6.4.x	1.13	6	dhmNotice DegradedIO	Disk Health Monitor - Reported a Disk Degraded-I/O Event	5	Notification	Warning	Critical	No	N/A
6.4.x	1.13	7	dhmNotice PFAEvent	Disk Health Monitor - Reported a Disk Predictive-Failure Event	n/a	Informational	Warning	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	Warning	Warning	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	Warning	Warning	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	Warning	Warning	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	Warning	Warning	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	Warning	Warning	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	Warning	Warning	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	Warning	Warning	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	Minor	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	Critical	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	Critical	Yes	26
6.4.x	1.13	26	diskRepaired	The failed disks have been repaired.	n/a	Notification	Harmless	Clear	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	Critical	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	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	Minor	Yes	36
6.4.x	1.13	36	fanRepaired	All fans are repaired.	1	Notification	Harmless	Clear	Yes	36

Introduced		OID / Trap #	Trap Name	Description	# Different Message Strings (sub-types)	Default Severity	Citi Tivoli Severity	Citi 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	Critical	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	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	Major	Yes	46
6.4.x	1.13	46	powerSupply Repaired	Previously-failed power supplies or power rails have been repaired.	1	Notification	Harmless	Clear	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	Minor	Yes	56
6.4.x	1.13	56	cpuOk	CPU utilization has dropped back below 90%.	1	Notification	Harmless	Clear	No	N/A
6.4.x	1.13	62	nvrAmBattery Discharged	The NVRAM battery is fully discharged.	2	Alert	Critical	Critical	No	N/A
6.4.x	1.13	63	nvrAmBattery Low	The charge in the NVRAM battery is low.	8	Critical	Fatal	Critical	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	Critical	Yes	76
6.4.x	1.13	75	clusterNode TakenOver	The partner has taken over for a failed cluster node.	1	Warning	Fatal	Minor	Yes	76
6.4.x	1.13	76	clusterNode Repaired	A cluster node has resumed operation.	1	Notification	Harmless	Clear	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	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	Critical	Yes	86
6.4.x	1.13	86	volume Repaired	All volumes are now under 95% full.	1	Notification	Harmless	Clear	Yes	86
7.1.2.1P1	1.18.2	87	volumes StillFull	At least one volume that was full or nearly full is now OK, but more volumes still exceed the volumeNearlyFull or volumeFull thresholds	1	Information	Warning	Minor	No	N/A
6.4.x	1.13	91	overTemp Shutdown	System temperature is too high to continue operating. The system is shutting down.	3	Emergency	Fatal	Critical	Yes	96
6.4.x	1.13	95	overTemp	System temperature is too high and in the warning level.	1	Warning	Critical	Critical	Yes	96

Introduced		OID / Trap #	Trap Name	Description	# Different Message Strings (sub-types)	Default Severity	Citi Tivoli Severity	Citi Netcool Severity	Auto Clear	Clear OID
ONTAP Version	MIB #									
6.4.x	1.13	96	overTemp Repaired	System temperature has returned to an acceptable value.	1	Notification	Harmless	Clear	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	Critical	Yes	106
6.4.x	1.13	106	shelfRepaired	A previously-reported shelf fault is now corrected.	2	Notification	Harmless	Clear	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	Critical	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	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	Minor	Yes	116
6.4.x	1.13	116	globalStatus Ok	The appliance's overall status returned to normal.	1	Notification	Harmless	Clear	Yes	116
6.4.x	1.13	126	softQuota Exceeded	A user has exceeded their soft quota limit.	1	Notification	Warning	Warning	Yes	127
6.4.x	1.13	127	softQuota Normal	A user is safely back under their soft quota limit.	1	Informational	Harmless	Clear	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	Major	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	Warning	Warning	Yes	136
6.4.x	1.13	136	autosupport Sent	Autosupport was sent successfully.	6	Notification	Harmless	Clear	Yes	136
7.0.3 / 7.0.4	1.17.3	142	upsLinePowerOff	UPS: Input line power has failed and UPS is now on battery	1	Alert	Harmless	Clear	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	Clear	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	Clear	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	Clear	Yes	146
7.0.3 / 7.0.4	1.17.3	146	upsLinePowerRestored	UPS: Input line power has been restored and UPS is now off battery	1	Notification	Harmless	Clear	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	Clear	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	Clear	No	N/A
7.0.3 / 7.0.4	1.17.3	153	appCritical	The application encountered a critical condition	1	Critical	Harmless	Clear	No	N/A

Introduced ONTAP Version	MIB #	OID / Trap #	Trap Name	Description	# Different Message Strings (sub-types)	Default Severity	Citi Tivoli Severity	Citi Netcool Severity	Auto Clear	Clear OID
7.0.3 / 7.0.4	1.17.3	154	appError	The application encountered an error condition	1	Error	Harmless	Clear	No	N/A
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	Clear	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	Clear	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	Clear	No	N/A
7.0.3 / 7.0.4	1.17.3	158	appTrap	The application requires to be debugged	1	Debug	Harmless	Clear	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	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	Major	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	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	Warning	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	Clear	Yes	177
6.4.x	1.13	187	wafDirFull	The directory has been filled to its limit (64MB)	1	Informational	Critical	Major	No	N/A
6.4.x	1.13	192	eccSummary	Memory ECC: number of new correctable ECC errors	1	Alert	Warning	Warning	No	N/A
6.4.x	1.13	195	eccMasked	Memory ECC: High frequency of ECC errors	1	Warning	Critical	Warning	No	N/A
7.0.3 / 7.0.4	1.17.3	204	ftpdError	Ftp daemon: service stopped	2	Error	Warning	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	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	Warning	No	N/A
7.1.2.1P1	1.18.2	222	scsitgtFCP LinkBreak	SCSI Target: Link Break on FCP adapter	2	Alert	Critical	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	Major	No	N/A
7.0.3 / 7.0.4	1.17.3	226	ScsitgtThrottle Notice	SCSI Target: Throttle limit event notification	2	Notification	Warning	Warning	No	N/A

Introduced ONTAP Version	MIB #	OID / Trap #	Trap Name	Description	# Different Message Strings (sub-types)	Default Severity	Citi Tivoli Severity	Citi Netcool Severity	Auto Clear	Clear OID
6.5.x	1.16	237	vifPrimaryLink Failed	Vif: The primary Interface on a Single mode vif has failed	1	Informational	Critical	Critical	No	N/A
6.5.x	1.16	238	vifAllLinks Failed	Vif: All the links of the vif have failed	1	Debug	Critical	Critical	No	N/A
6.5.x	1.16	245	vfStopped	A vFiler has stopped.	1	Warning	Warning	Warning	No	246
6.5.x	1.16	246	vfStarted	A vFiler has started.	1	Notification	Harmless	Clear	Yes	246
6.5.x	1.16	254	vscanVirus DetectedError	VScan: Vscan scanner has detected a virus on the Filer	2	Error	Warning	Major	No	N/A
6.5.x	1.16	255	vscanDis Connection	VScan: There has been a connection to the vscan servers dropped	2	Warning	Warning	Warning	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	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	Minor	No	N/A
6.5.x	1.16	266	vscanServerU pgrade	VScan: The Vscan server has been upgraded	1	Notification	Harmless	Minor	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 waffron has failed. This trap is issued to alert operator that a volume is not accessible and waffron must be started to allow access to it.	2	Alert	Critical	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	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	Major	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	Clear	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	Clear	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	Clear	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	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	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	Warning	No	N/A

Introduced		OID / Trap #	Trap Name	Description	# Different Message Strings (sub-types)	Default Severity	Citi Tivoli Severity	Citi Netcool Severity	Auto Clear	Clear OID
ONTAP Version	MIB #									
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	Warning	No	N/A
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	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	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	Major	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	Major	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	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	Critical	No	N/A
6.5.x	1.16	364	snapmirror SyncFailed	Synchronous SnapMirror failed and went into asynchronous mode.	1	Error	Critical	Critical	Yes	366
6.5.x	1.16	366	snapmirror SyncOk	Synchronous SnapMirror went into synchronous mode.	1	Notification	Harmless	Clear	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	Critical	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	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	Critical	Yes	376
7.0.3 / 7.0.4	1.17.3	376	chassis Temperature OK	The chassis temperature is OK	1	Notification	Harmless	Clear	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	Critical	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	Critical	Yes	386

Introduced ONTAP Version	MIB #	OID / Trap #	Trap Name	Description	# Different Message Strings (sub-types)	Default Severity	Citi Tivoli Severity	Citi Netcool Severity	Auto Clear	Clear OID
7.0.3 / 7.0.4	1.17.3	386	chassisCPU FanOK	All CPU fan(s) are properly functioning	1	Notification	Harmless	Clear	Yes	386
7.0.3 / 7.0.4	1.17.3	391	chassisPower Supplies Failed	Multiple chassis power supplies failed	1	Emergency	Fatal	Critical	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	Critical	Yes	396
7.0.3 / 7.0.4	1.17.3	393	chassisPower Supply Failed	One chassis power supply failed	1	Critical	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	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	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	Clear	Yes	396
7.0.3 / 7.0.4	1.17.3	397	chassisPower SupplyOK	This chassis power supply is functioning properly	1	Informational	Harmless	Clear	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	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	Clear	Yes	406
7.0.3 / 7.0.4	1.17.3	412	chassisFan Degraded	A chassis fan has been degraded	1	Alert	Critical	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	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	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	Critical	Yes	416
7.0.3 / 7.0.4	1.17.3	416	chassisFan OK	All chassis fans are functioning properly	1	Notification	Harmless	Clear	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	Minor	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	Minor	No	N/A

Introduced		OID / Trap #	Trap Name	Description	# Different Message Strings (sub-types)	Default Severity	Citi Tivoli Severity	Citi Netcool Severity	Auto Clear	Clear OID
ONTAP Version	MIB #									
7.2.4	1.19.3	436	dcPasswd ChangeFailed	An attempt to change the filer's password with the domain controller has failed	1	Notification	Critical	Critical	No	N/A
7.3.1.1	2.0.1	437	domain Controller Connected	A CIFS domain controller connection to the filer has been established	1	Information	Harmless	Clear	No	N/A
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	Critical	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	Critical	No	N/A
7.2.4	1.19.3	464	shelfSES Electronics Failed	One or more of the enclosure services devices in a disk shelf has failed. Some shelf designs combine the enclosure monitoring hardware function into the module that provide the storage interface to the shelf. A failure in the enclosure monitoring section of these combined modules does not necessarily indicate a failure in disk or loop or bus operation, which may be able to continue	1	Error	Critical	Critical	No	N/A
7.2.4	1.19.3	467	shelfSES Electronics Info	A previously reported failure of an enclosure services device in a disk shelf has been corrected, or the device has reported information that does not necessarily require customer action	1	Information	Harmless	Clear	No	N/A
7.2.4	1.19.3	473	shelfIFModule Failed	One or more of the storage interface modules in a disk shelf has failed. Some shelf designs combine the enclosure monitoring hardware function into the module that operates the fibre channel loop or SCSI in the shelf. This failure is of the storage interface itself, not a failure of the enclosure monitoring, which may be able to continue. This failure may make one or more disks in the shelf or in the loop or bus unavailable.	1	Critical	Critical	Critical	No	N/A
7.3.1.1	2.0.1	475	shelfIFModule Warning	One or more of the storage interface modules in a disk shelf has a warning condition. This warning condition is of the storage interface itself, not a failure of the enclosure monitoring, which may be able to continue	1	Warning	Critical	Critical	Yes	477
7.2.4	1.19.3	477	shelfIFModule Info	A previously reported failure of a disk shelf interface module has been corrected, or the module has reported information that does not necessarily require customer action	1	Information	Harmless	Clear	No	N/A
7.0.3 / 7.0.4	1.17.3	482	maxDirSize Alert	A directory has reached its maxdirsize limit. Either increase the maxdirsize or clean up the directory	1	Alert	Critical	Critical	No	N/A
7.0.3 / 7.0.4	1.17.3	485	maxDirSize Warning	A directory is getting close to its maxdirsize limit. Either increase the maxdirsize or clean up the directory	1	Warning	Critical	Critical	No	N/A
7.2.4	1.19.3	497	cifsStats ExhaustMem CtrlBlk	A request for a CIFS control block can not be granted. All control blocks are in use	1	Information	Critical	Critical	No	N/A
7.2.4	1.19.3	501	chassisPS RemovedxMin Shutdown	Power Supply Unit is removed from the system. System will be shutdown in x minutes if it is not put back	1	Emergency	Fatal	Critical	No	N/A

Introduced		OID / Trap #	Trap Name	Description	# Different Message Strings (sub-types)	Default Severity	Citi Tivoli Severity	Citi Netcool Severity	Auto Clear	Clear OID
ONTAP Version	MIB #									
7.2.4	1.19.3	505	chassisPSUs Mismatch	Power Supply Units in the system are of different type. Please replace them with correct Power Supply Units	1	Warning	Critical	Critical	No	N/A
7.2.4	1.19.3	511	chassisFan FailxMinShut down	Multiple chassis fan failure. System will shut down in a few minutes if not corrected	1	Emergency	Fatal	Critical	No	N/A
7.2.4	1.19.3	515	chassisPSU wrongInput	Power Supply Unit (PSU) in the system is connected to an incompatible external power source. Please change the power source being used or replace the PSU with one that matches the current power source	1	Warning	Critical	Critical	No	N/A
7.2.4	1.19.3	521	powerSupply FanFailxMin Shutdown	Multiple power supply fans failure. System will shut down in a few minutes if not corrected	1	Emergency	Fatal	Critical	No	N/A
7.3.1.1	2.0.1	532	remoteSystem MgtAlert	System remote management (RLM or BMC) detected a system down event. The events that would generate this trap would be: post_error: Fatal POST error on boot power_loss: Loss of power to system abnormal_reboot: l2_watchdog_reset	1	Alert	Critical	Critical	No	N/A
7.3.1.1	2.0.1	535	remoteSystem MgmtWarning	System remote management (RLM or BMC) detected a system down event. The events that would generate this trap would be: loss_of_heartbeat: Loss of heartbeat to the RLM	1	Warning	Critical	Critical	No	N/A
7.3.1.1	2.0.1	536	remoteSystem Mgmt Notification	System remote management (RLM or BMC) detected a system down event. The events that would generate this trap would be: power_off_via_rlm: This indicates system power off power_cycle_via_rlm: This indicates system power cycle reset_via_rlm: This indicates system reset	1	Notification	Harmless	Clear	No	N/A
7.3.1.1	2.0.1	547	remoteSystem MgmtPeriodic	System remote management (RLM or BMC) detected a system down event. The events that would generate this trap would be: Remote Management Event: type=keep_alive, severity=info, event=periodic_message	1	Information	Harmless	Clear	No	N/A
7.3.1.1	2.0.1	556	remotesystem MgmtTest	System remote management (RLM or BMC) detected a system down event. The events that would generate this trap would be: Remote Management Event: type=test, severity=notice, event=test	1	Notification	Harmless	Clear	No	N/A
7.1.2.1P1	1.18.2	562	diskMultipath OneSwitch	A multipathed disk is only connected to one switch.	1	Alert	Critical	Critical	No	N/A
7.1.2.1P1	1.18.2	563	diskMultipath NoTakeover	Multipath disks or luns could not be detected for the partner. The node will be unable to takeover correctly.	1	Critical	Critical	Critical	No	N/A
7.1.2.1P1	1.18.2	565	diskMultipath Warning	Synchronous mirroring is enabled, even though disks are not multipathed.	1	Warning	Critical	Critical	No	N/A
7.2.4	1.19.3	574	driveDisable Err	A drive has been disabled by the shelf module due to hardware errors. If it is the only drive in the shelf that is disabled, replace the drive	1	Error	Critical	Critical	No	N/A

Introduced		OID / Trap #	Trap Name	Description	# Different Message Strings (sub-types)	Default Severity	Citi Tivoli Severity	Citi Netcool Severity	Auto Clear	Clear OID
ONTAP Version	MIB #									
7.2.4	1.19.3	587	hbaOffline Information	A host bus adapter (HBA) port has been disabled. This state may be caused either by user command or by errors that make the interface unable to communicate with attached devices	1	Information	Critical	Critical	No	N/A
7.3.1.1	2.0.1	606	lunSnap RestoreStatus	Status of the Snap Restore	1	Notification	Harmless	Clear	No	N/A
7.3.1.1	2.0.1	616	lunClone Create	Created a LUN Clone	1	Notification	Harmless	Clear	No	N/A
7.3.1.1	2.0.1	626	lunCloneSplit Start	LUN Clone Split started	1	Notification	Harmless	Clear	Yes	627
7.3.1.1	2.0.1	627	lunCloneSplit Complete	LUN Clone Split completed	1	Information	Harmless	Clear	No	N/A
7.3.1.1	2.0.1	636	flexCloneSplit Start	Flex Clone Split started	1	Notification	Harmless	Clear	Yes	637
7.3.1.1	2.0.1	637	flexCloneSplit Complete	Flex Clone Split completed	1	Information	Harmless	Clear	No	N/A
7.3.1.1	2.0.1	646	volumeClone Create	Created a Volume Clone	1	Notification	Harmless	Clear	No	N/A
7.3.1.1	2.0.1	656	snapAuto Delete	Snapshot Autodeleted	1	Notification	Harmless	Clear	No	N/A
7.3.1.1	2.0.1	666	volume Autogrow	Volume is Autogrown	1	Notification	Harmless	Clear	No	N/A

14.3 Appendix C: Data ONTAP SOE Configuration Standards

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.

The settings in these five tables are considered standard across all NetApp platforms and use cases, and as such should be implemented on every system to meet the objective of simplification and consistency. Where particular option values are based on certain criteria, these criterium are specified accordingly.

Table 33 - ONTAP Mandatory (Security) Configuration Matrix

From ONTAP Release	ONTAP Security Option Name	Description	Value	Applicability	
				Physical	vFiler
6.4	auditlog.enable	Enables/disables the audit logging of commands executed at the console/telnet shell or by using ssh or GUI. The default value 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.	on	Yes	No
7.3	auditlog.readonly_api.enable	Enables/disables the logging of read only API calls. If an API is used only for retrieving information and not for modifying the state of the system, the read-only API is not audited by default. The default value is off . Enabling logging of read-only APIs will result in many more records logged to the auditlog. It should only be enabled at the direction of NetApp Technical Support for diagnostic purposes. In summary this option reduces the level of audit logging.	off	Yes	No
6.4	autologout.console.enable	Enables/disables the autologout of console connections. The default value 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.	on	Yes	No
6.4	autologout.console.timeout	The number of minutes the console is idle after which console connection is disconnected.	60	Yes	No
6.4	autologout.telnet.enable	Enables/disables the autologout of telnet connections. The default value 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.	on	Yes	No
6.4	autologout.telnet.timeout	The number of minutes after which telnet connections are disconnected if autologout.telnet.enable is on. The default value is 60 minutes. Any change to this option requires a logout before it takes effect.	60	Yes	No
6.4	autosupport.content	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 value for this complete . Any change to this option will take effect the next time an autosupport message is generated.	complete	Yes	No

From ONTAP Release	ONTAP Security Option Name	Description	Value	Applicability	
				Physical	vFiler
7.0	autosupport.support.proxy	Dictates the address of the Citi 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.	webproxy.wlb.nsroot.net:8080	Yes	No
7.0	autosupport.support.transport	Dictates the type of transport protocol used to send the autosupport message. The default value is https .	https	Yes	No
6.4	cifs.audit.enable	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 cifs.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 value is off .	off	Yes	Yes
6.4	cifs.audit.file_access_events.enable	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 value is on .	on	Yes	Yes
6.4	cifs.audit.logon_events.enable	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 value is on .	on	Yes	Yes
6.4	cifs.bypass_traverse_checking	The default value is on . When on directories in the path to a file are not required to have the 'X' (traverse) permission. This option does not apply in UNIX qtrees.	off	Yes	Yes
7.2.5.1	cifs.restrict_anonymous	This option is used to control the behavior of ONTAP when a null (anonymous) user attempts to connect to the Filer in order to enumerate the shares on the system. The option cifs.restrict_anonymous replaces cifs.restrict_anonymous.enable as of ONTAP 7.2.5.1 and above. Refer to NetApp knowledge base article 47838 for further details. The default value is 0 .	1	Yes	Yes
7.3.3	cifs.signing.enable	Signing is a security feature provided by the CIFS protocol that is designed to detect and prevent 'man-in-the-middle' intrusion into CIFS communications. This is achieved by calculating a security signature value for every incoming and outgoing CIFS packet. The default value is off .	on	Yes	Yes
7.3.1	cifs.smb2.signing.required	This option is to enforce SMB signing on all SMB 2.0 sessions. The default value is off . This option must be enabled to comply with CISS policy requirements. The filer incurs a CPU performance overhead when signing is enabled. Starting with ONTAP 8.1 this setting also configures support for SMB 2.1 as well as SMB 2.0 (a setflag workaround is required to disable SMB 2.1 independently of SMB 2.0 as part of bug 798842 - See the following /etc/lrc file section)	on	Yes	Yes

From ONTAP Release	ONTAP Security Option Name	Description	Value	Applicability	
				Physical	vFiler
8.1.1	cifs.W2K_password_change	Enables or disables automatic storage system domain account password change. This setting was called cifs.weekly_W2K_password_change in earlier versions of ONTAP prior to 8.1.1. The default value is off .	off	Yes	Yes
6.5	cifs.weekly_w2k_password.change	This option causes a storage system belonging to a Windows Active Directory domain to change its domain password once a week. The password change occurs at approximately 1:00 a.m. on Sunday. The default value is off . This option is replaced by cifs.W2K_password_change in ONTAP version 8.1.1 or later.	off	Yes	Yes
6.4	ftpd.enable	If on the appliance accepts FTP connections on port 21. If off connections on port 21 are refused. The default value is off .	off	Yes	Yes
8.0.1	httpd.admin.enable	Enables HTTP access to FilerView, the graphical user interface for the Filer. The default value is off .	on	Yes	No
7.2	httpd.admin.top-page.authentication	If enabled, the top-level page of FilerView will have authenticated access. The default value is on .	off	Yes	Yes
7.3.3	httpd.admin.ssl.enable	Enables HTTPS access to FilerView. Configured via the secureadmin ONTAP CLI command. The default value is on .	on	Yes	No
6.4	httpd.enable	Enables HTTP access to the filer. Valid values for this option are on or off. The default value for this option is off .	off	Yes	Yes
7.0	ip.ping_throttle.drop_level	Specifies the maximum number of ICMP echo or echo reply packets (ping packets) that the filer will accept per second. Any further packets within one second are dropped to prevent ping flood denial of service attacks. The default value is 150 .	150	Yes	No
7.0	nfs.export.auto-update	This option controls whether automatic updates are performed on the /etc/exports file. If it is not set, then the commands 'vol create', 'vol delete', and 'vol rename' will not automatically rewrite the file. Instead they will syslog the need to edit the file. When volumes are moved between vfilers, automatic updates on the /etc/exports file of the source and destination vfilers are dependent on this option. Turning this option off prevents storage administrators from inadvertently creating NFS exports when a CIFS volume is created. The default value is on .	off	Yes	Yes
6.4	nfs.webnfs.enable	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 .	off	Yes	Yes
6.4	pcnfsd.enable	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 .	off	Yes	Yes
8.1.1	rlm.autologout.enable	Enables or disables the automatic logout of idle RLM SSH sessions (only applies to hardware platforms with the RLM port). The default value for this option is on .	on	Yes	No
8.1.1	rlm.autologout.timeout	Specifies the number of minutes after which RLM SSH idle connections are disconnected if rlm.autologout.enable is on (only applies to hardware platforms with the RLM port). The default value for this option is 60	30	Yes	No

From ONTAP Release	ONTAP Security Option Name	Description	Value	Applicability	
				Physical	vFiler
8.1.1	rlm.ssh.access	Restricts the SSH access to the RLM interface (only applies to hardware platforms with the RLM port). 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 Unified Manager 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>,<dfm_server3_IP>'	See notes	Yes	No
6.4	rsh.enable	Enables the RSH server on the filer. Valid values for this option are on or off. The default value for this option is off .	off	Yes	Yes
7.1	security.passwd.firstlogin.enable	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 off .	on	Yes	Yes
7.1	security.passwd.lockout.numtries	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 4294967295 . Configuring this option to '6' meets the CITMP policy requirements.	6	Yes	Yes
7.3.3	security.passwd.rootaccess.enable	This option controls whether root can have access to the system. A value of off means that root cannot login or execute any commands. This option is reset to on if a user changes root's password, or during a boot without etc/rc. The default value for this setting is on .	on	Yes	No
6.5	security.passwd.rules.enable	This option controls whether a check for minimum length 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).	on	Yes	Yes
7.0	security.passwd.rules.everyone	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 on . security.passwd.rules.enable (above) must have the value on or this option is ignored.	on	Yes	Yes
7.1	security.passwd.rules.history	Defines the number of previous passwords that ONTAP will remember for each administrator account. The default value for this option is 6 .	6	Yes	Yes
7.0	security.passwd.rules.maximum	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. The default value for this option is 256 .	14	Yes	Yes

From ONTAP Release	ONTAP Security Option Name	Description	Value	Applicability	
				Physical	vFiler
7.0	security.passwd.rules.minimum	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.	6	Yes	Yes
7.0	security.passwd.rules.minimum.alphabetic	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.	2	Yes	Yes
7.0	security.passwd.rules.minimum.digit	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.	1	Yes	Yes
8.1.2	security.passwd.rules.minimum.lowercase	This option controls the minimum number of lowercase alphabetic characters ("a" to "z") that a password must contain. This option is ignored if the security.passwd.rules.enable option is not set to on. The default value for this option is 0 .	1	Yes	Yes
7.0	security.passwd.rules.minimum.symbol	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.	0	Yes	Yes
8.1.2	security.passwd.rules.minimum.uppercase	This option controls the minimum number of uppercase alphabetic characters ("A" to "Z") that a password must contain. This option is ignored if the security.passwd.rules.enable option is not set to on. The default value for this option is 0 .	1	Yes	Yes
8.1.1	sp.autologout.enable	Enables or disables the automatic logout of idle SP SSH sessions (only applies to hardware platfroms with the SP port). The default value for this option is on .	on	Yes	No
8.1.1	sp.autologout.timeout	Specifies the number of minutes after which SP SSH idle connections are disconnected if rlm.autologout.enable is on (only applies to hardware platfroms with the SP port). The default value for this option is 60	30	Yes	No
8.1.1	sp.ssh.access	Restricts the SSH access to the SP interface (only applies to hardware platfroms with the SP port). 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 Unified Manager 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>,<dfm_server3_IP>'	See notes	Yes	No

From ONTAP Release	ONTAP Security Option Name	Description	Value	Applicability	
				Physical	vFiler
7.0	ssh.access	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>,<dfm_server3_IP>,<emc_srm_collector_server1_IP>'	See notes	Yes	Yes
6.4	ssh.enable	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 on .	on	Yes	Yes
7.0	ssh.idle.timeout	Timeout value for ssh sessions in seconds. The default value for this option is 600 seconds .	600	Yes	Yes
7.0	ssh1.enable	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 off .	off	Yes	Yes
7.0	ssh2.enable	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 .	on	Yes	Yes
6.4	ssl.enable	Enables or disables the SSL protocol on the filer. Valid values for this option are on or off. The default value for this option is on .	on	Yes	No
7.3	ssl.v2.enable	Enables or disables the SSL v2 protocol on the filer. Valid values for this option are on or off. The default value for this option is on .	off	Yes	No
7.3	ssl.v3.enable	Enables or disables the SSL v3 protocol on the filer. Valid values for this option are on or off. The default value for this option is on .	on	Yes	No
8.0.1	telnet.distinct.enable	Enables making the telnet/SSH and console separate user environments. If it is off, then telnet/SSH and console share a session. The two sessions view each other's inputs/outputs and both acquire the privileges of the last user to login. Valid values for this option are on or off. The starting default value on a factory install for this option is on .	on	Yes	No
6.4	telnet.enable	Enables the Telnet server on the filer. Valid values for this option are on or off. The default value for this option is on .	off	Yes	No
6.4	tftpd.enable	Enables the tftp (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 tftp server allows <i>get</i> requests, but does not allow <i>put</i> requests.	off	Yes	No
7.0	trusted.hosts	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. '<soe_server1_IP>,<soe_server2_IP>,<dfm_server1_IP>,<dfm_server2_IP>'	see notes	Yes	No

From ONTAP Release	ONTAP Security Option Name	Description	Value	Applicability	
				Physical	vFiler
6.4	wafl.default_nt_user	Provides default translation rules for Unix user accessing a Windows security style volume. The default value for this option is "".	""	Yes	Yes
6.4	wafl.default_unix_user	Provides default translation rules for Windows user accessing a unix security style volume. The default value for this option is 'pcuser'.	pcuser	Yes	Yes
6.4	wafl.nt_admin_priv_map_to_root	When turned on, an NT administrator is mapped to the Unix root account. The default value for this option is 'on'.	on	Yes	Yes
8.0.1	webdav.enable	Enables WebDAV access to the filer. Valid values for this option are on or off. The default value for this option is 'on'.	off	Yes	Yes

Table 34 - ONTAP Recommended (Technical Reference) Configuration Matrix

From ONTAP Release	ONTAP Technical Option Name	Description	Value	Applicability	
				Physical	vFiler
6.4	autosupport.enable	Dictates whether recipients defined in autosupport.to will receive autosupport emails.	on	Yes	No
6.4	autosupport.from	Filer Administrator recipient email address, in the standard Citi email address form. Alternatively this email address can be a distribution list. Whichever value you use, it must be a valid entry in the Citi Global Address List (GAL).	<first.last@citi.com>	Yes	No
6.4	autosupport.mailhost	The address of the SMTP server used to send the autosupport email to the recipients defined in the autosupport.to option.	mailhub-nyc3.ny.smb.com	Yes	No
6.4	autosupport.to	Filer Administrator recipient email address, in the standard Citi 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 Citi email system. Due to the clear text nature of autosupport messages, internet email addresses are forbidden in this field	<first.last@citi.com>	Yes	No
6.4	autosupport.support.enable	Dictates whether NetApp receives the autosupport message. The default value is on.	on	Yes	No
7.3.3	cdpd.enable	When this option is set to ON, Cisco Discovery Protocol v1 (CDPv1) Daemon is enabled on all physical network ports so that it starts sending and processing CDPv1 advertisements. The default value is off.	off	Yes	No
6.4	cf.giveback.auto.enable	This option turns on/off automatic giveback. An automatic giveback is invoked when one node of a High Availability (HA) configuration is in takeover mode and the "down" node is repaired and reboots. If enabled the repaired node will boot into Data ONTAP and the node in takeover mode will detect this and initiate a give-back. This feature is only available on flash booted systems. The default value is on.	off	Yes	No
8.1.2	cf.giveback.auto.after.panic.takeover	Enables or disables automatic giveback after a takeover-on-panic event. The default value is on.	off	Yes	No
8.1.4	cf.hw_assist.enable	This option turns the hardware-assisted storage controller takeover functionality on or off (see supporting documentation for details). It requires e0M and SP or RLM to be configured. The default value is on.	on	Yes	No
8.1.4	cf.hw_assist.partner.address	The hardware failure notification is sent to this partner IP address. If hostname is given, it is converted into an IP address. The default value is unconfigured	IP address of e0M interface of partner node	Yes	No
8.1.4	cf.hw_assist.partner.port	The hardware failure UDP SNMP notification is sent to this partner port. The default value is 4444	4444	Yes	No
6.5	cifs.audit.autosave.onsize.enable	When CIFS audit log reaches the value defined cifs.audit.logsize the audit log will be automatically saved and a new log started. The default value is off.	on	Yes	Yes

From ONTAP Release	ONTAP Security Option Name	Description	Value	Applicability	
				Physical	vFiler
7.3.3	cifs.max_mpx	This option controls how many simultaneous operations the filer reports that it can process. An "operation" is each I/O the client believes is pending on the filer including outstanding change notify operations. Clients such as Windows Terminal Server or IIS may require that this number be increased to avoid errors and performance delays. The default value is dependant on the version of ONTAP deployed: 8.0 = 50; 8.0.1 = 255 ; 8.0.2 = 253.	253	Yes	Yes
7.3.3	cifs.netbios_over_tcp. enable	This option enables the use of NetBIOS over TCP, which is the standard protocol used for CIFS prior to Windows 2000. In certain Windows 2000 networks it is desirable to disable that protocol. This option corresponds to the "Enable NetBIOS over TCP" setting in the Windows 2000 Advanced TCP/IP settings tab. If it is set to off, all clients must be Windows 2000 (or above), and only Windows 2000 (or above) domain controllers and virus scanners can be used. The default value is on.	off	Yes	Yes
6.4	cifs.oplocks.enable	CIFS oplocks (opportunistic locks) enable the redirector on a CIFS client in certain file-sharing scenarios to perform client-side caching of read-ahead, write-behind, and lock information. A client can then work with a file (read or write it) without regularly reminding the server that it needs access to the file in question. This improves performance by reducing network traffic. The default value is on. Starting with ONTAP 8.1 this feature also determines whether lease locks are enabled for SMB 2.1 if SMB 2.1 is enabled on the storage system.	on	Yes	Yes
8.0.2	cifs.rpcfd_timeout	Added in ONTAP 8.0.2 to address an issue (bug 430721) where certain misbehaving client applications can cause a temporary service outage by consuming a large portion of the available memory, so that memory can be released. By default the timeout is set to 20 minutes (expressed as seconds in the setting). However regression testing has uncovered bug 490774, whereby if an RPC (ZAPI) request takes longer than the cifs.rpcfd_timeout option it may cause a filer to panic when the long-running request finally completes. The workaround to bug 490774 is the CIFS rpcfd timeout setting needs to be disabled by setting it to a value of '-1' (minus one). Bug 490774 is fixed in ONTAP 8.1.4 and later. Although this feature has CIFS in its name, CIFS does not need to be running or licensed on the system for the NetApp Filer to be exposed to the risk of this bug. The default value is 1200.	-1 (for ONTAP 8.0.2) 1200 (for ONTAP 8.1.4 and above)	Yes	Yes
7.3.1	cifs.smb2.client. enable	This option is to enable/disable SMB 2.0 client capability on the storage system. When this option is enabled, NetApp system initiated connections to Windows domain controllers will attempt to use the SMB 2.0 protocol. In case the Windows domain controller does not support the SMB 2.0 protocol, then the NetApp system will fall back to using SMB. If a session had been established over SMB 2.0 and later this option is disabled, existing sessions will not be terminated; the NetApp system will continue to use SMB 2.0 for the existing sessions, but no new sessions will attempt to use SMB 2.0. The default value is off.	off	Yes	Yes
7.3.1	cifs.smb2.durable_ handle.enable	This option is to enable/disable the durable handle functionality for SMB 2.0 clients. If this option is enabled, the open files from a client are preserved when the client gets disconnected from the NetApp system. These open files can be reclaimed when the client reconnects to the NetApp system. The default value is on.	on	Yes	Yes

From ONTAP Release	ONTAP Security Option Name	Description	Value	Applicability													
				Physical	vFiler												
7.3.1	cifs.smb2.durable_handle.timeout	This option is to configure the duration in seconds for which the NetApp system will preserve the durable handle after a temporary network failure. The default value is 16 minutes .	16m	Yes	Yes												
7.3.1	cifs.smb2.enable	This option is to enable/disable SMB 2.0 in Data ONTAP. When this option is disabled, the NetApp system will not accept any new SMB 2.0 sessions, but the existing sessions will not be terminated. The default value is off. Starting with ONTAP 8.1 this setting also configures support for SMB 2.1 as well as SMB 2.0 (a setflag workaround is required to disable SMB 2.1 independently of SMB 2.0 as part of bug 798842 for ONTAP 8.1.4P1 - See the following /etc/rc file section. For ONTAP 8.1.4P4 SMB 2.1 is disabled by default and there is no need to manually disable it in the /etc/rc file). Enabling SMB 2.0 is only advised if the SMB 2.0 permitted use requirements detailed in section 11.4.1.5.1 are met.	See notes in description field	Yes	Yes												
8.1.4P4	cifs.smb2.signing.max_threads	This option determines the number of processor cores which can be used concurrently for the handling and processing of signed SMB2 transactions. This is a hidden parameter introduced in ONTAP 8.1.4P4 and above. Follow the procedure documented in section 11.4.1.5.3 when configuring this value. Note when planning to change this value CIFS services on all vfilers (including vFiler0) must be stopped on the system before the setting is changed and can only be restarted after the value has been changed. The default value is 0 when upgrading from a version of ONTAP prior to 8.1.4. When upgrading from 8.1.4P1 the system will take the preferred value for the hardware platform. However the procedure documented in section 11.4.1.5.3 should still be followed in order to maintain a consistent experience. Starting with ONTAP 8.1.4P5 this setting will default to the correct value for the hardware platform when cifs.smb2.signing.multiprocessing is enabled. The recommended value for this setting varies according to FAS hardware platform and is only required to be configured on systems where SBM 2.0 is enabled on one or more vFilers. Consult the table below for the recommended setting for your particular system: <table><tr><th>Hardware Platform(s)</th><th>cifs.smb2.signing.max_threads value</th></tr><tr><td>FAS2040</td><td>0 (disabled)</td></tr><tr><td>FAS3140, FAS3210</td><td>1</td></tr><tr><td>FAS3160, FAS3170, FAS3220, FAS3240, FAS3270</td><td>3</td></tr><tr><td>FAS3250, FAS6210, FAS6220, FAS6240, FAS6250</td><td>7</td></tr><tr><td>FAS6280, FAS6290</td><td>11</td></tr></table>	Hardware Platform(s)	cifs.smb2.signing.max_threads value	FAS2040	0 (disabled)	FAS3140, FAS3210	1	FAS3160, FAS3170, FAS3220, FAS3240, FAS3270	3	FAS3250, FAS6210, FAS6220, FAS6240, FAS6250	7	FAS6280, FAS6290	11	See notes in description field	Yes	No
Hardware Platform(s)	cifs.smb2.signing.max_threads value																
FAS2040	0 (disabled)																
FAS3140, FAS3210	1																
FAS3160, FAS3170, FAS3220, FAS3240, FAS3270	3																
FAS3250, FAS6210, FAS6220, FAS6240, FAS6250	7																
FAS6280, FAS6290	11																

From ONTAP Release	ONTAP Security Option Name	Description	Value	Applicability	
8.1.4P5	cifs.smb2.signing. multiprocessing	This option determines whether the ability to spread the SMB2 signing processing workload over multiple processors is enabled or not. Enabling this feature automatically configures the appropriate value for the cifs.smb2.signing.max_threads setting. Follow the procedure documented in section 11.4.1.5.3 when configuring this value. The default value is 0 when upgrading from a version of ONTAP prior to 8.1.4 and is only required to be configured on systems where SBM 2.0 is enabled on one or more vFilers. Note when planning to change this value CIFS services on all vfilers (including vFiler0) must be stopped on the system before the setting is changed and can only be restarted after the value has been changed.	enabled (1)	Yes	No
8.1.1	cifs.smb2_1.branch_ cache.enable	Enables or disables BranchCache on the storage controller. The default value is off .	off	Yes	Yes
7.3	cifs.tcp_window_size	The TCP window size controls the number of TCP messages that can be transmitted between the storage system and the CIFS client at one time. The default value is 17520. The adjacent value is optimized the window size and provides a significant increase in performance for larger file transfers balanced against the performance impact if packet loss exists on the network. This option is 'hidden' in Data ONTAP meaning it must be queried by name to retrieve the value it is configured to and will not show up in a generic 'options' command to list all current settings.	65535	Yes	Yes
6.4	cifs.wins_servers	Defines the WINS servers used by the vFiler or physical Filer. The default value is unconfigured .	2 IP addresses (until ONTAP 8.1.4) "" (after WINS retirement)	Yes	Yes
7.2	disk.auto_assign	This option specifies if disks will be auto assigned on systems with software disk ownership. The default value is on .	off	Yes	No
7.1	disk.maint_center. enable	This option dictates the status of the maintenance center component of ONTAP. The purpose of the Maintenance Center is to improve storage reliability by reducing the number of unnecessary disk returns to NetApp due to transient errors. The default value is on .	on	Yes	No
7.2	ems.autosuppress. enable	The EMS autosuppression feature implements a dynamic suppression algorithm to filter out the bulk occurrences of only the most offensive event messages, leaving the remainder of the log intact and readable. The default value is on .	off	Yes	No
8.1	flexscale.rewarm	Determines whether the WAFL external cache preserves the cache in a Flash Cache module during a graceful shutdown. The default value is on .	on	Yes	No

From ONTAP Release	ONTAP Security Option Name	Description	Value	Applicability	
				Physical	vFiler
8.0.2	interface.blocked.mgmt_data_traffic	Blocks or allows data traffic on the management interface, e0M. The default value varies depending on whether the system was upgraded from ONTAP versions prior to 8.0.2, in which case the default is off, and when installing a new system on 8.0.2 or above the default value is on.	on	Yes	No
6.5	ip.fastpath.enable	If the option is on, the filer will attempt to use MAC address and interface caching ("Fastpath") so as to try to send back responses to incoming network traffic using the same interface as the incoming traffic and (in some cases) the destination MAC address equal to the source MAC address of the incoming data. This allows for automatic load-balancing between multiple interfaces of a trunk and between multiple filer interfaces on the same subnet. However having this option enabled has in some cases proven to cause communication problems between Filers and Citi AnyCast DNS infrastructure preventing ONTAP from participating in Active Directory. The default value is on.	off	Yes	No
8.1	ip.tcp.abc.enable	Appropriate Byte Counting (ABC) is an algorithm implemented in the TCP stack of Data ONTAP that enables you to improve TCP performance and security. The default value is on.	off	Yes	No
8.1	ip.tcp.rfc3390.enable	Enables or disables the use of RFC 3390 to increase the initial window used by TCP connections. The default value is on.	on	Yes	No
8.1.3	ldap.fast_timeout.enable	Controls the timeout behavior of LDAP client operations. The default value is on.	on	Yes	Yes
7.3	ldap.timeout	Specifies the timeout used for LDAP searches. This is the period (in seconds), after which an LDAP search request is timed out on the LDAP server, if incomplete. Changes to this setting are effective immediately and require implementation at both the physical and vFiler level. The default value is 20.	2	Yes	Yes
8.1.3	ldap.retry_delay	Sets the time, in seconds, to wait before trying to connect to LDAP servers after a failed attempt. The default value is 120	12	Yes	Yes
8.1	nfs.always.deny.truncate	Controls whether NFSv2 and NFSv3 clients can truncate files in UNIX qtrees that are also opened from a CIFS client with DENY write permissions. The default value is on.	on	Yes	Yes
7.2	nfs.export.exportfs_comment_on_delete	This option controls the deletion behaviour for exportfs -z. It controls whether entries are removed or commented from the /etc/exports file. The default value is true and entries are commented out. To remove entries on deletion set it to false. The default value for this option is on.	off	Yes	Yes
7.2	nfs.export.harvest.timeout	This option sets the idle expiration time for entries in the export access cache. The timer resets every time the export is accessed from the host. The default value for this option is 1800.	1800	Yes	Yes
7.2	nfs.export.neg.timeout	This option sets the refresh time for entries which were denied access in the export access cache. The default value for this option is 3600.	3600	Yes	Yes
7.2	nfs.export.pos.timeout	This option set the refresh time for entries granted access in the export access cache. The default value for this option is 36000.	36000	Yes	Yes

From ONTAP Release	ONTAP Security Option Name	Description	Value	Applicability	
				Physical	vFiler
7.3	nfs.ifc.rcv.high	This option sets the receive side TCP flow control high water mark for the NFS protocol. This value needs to be set to a value larger than nfs.tcp.recvwindow size or else the filer will go into TCP flow control if the client sends more data than nfs.ifc.rcv.high value because of the large receive window size. The value of nfs.ifc.rcv.high is proportionately related to the value of nfs.tcp.recvwindow size and should be 1.5 times the value. Although ONTAP should set this value automatically, due to bug 568415 storage administrators are required to set this manually.	98910	Yes	No
7.3	nfs.ifc.rcv.low	This option sets the receive side TCP flow control low water mark for the NFS protocol. This value is used to determine when the server exits flow control, once flow control becomes activated. The value of nfs.ifc.rcv.low is proportionately related to the value of nfs.tcp.recvwindow size. Unlike nfs.ifc.rcv.high, nfs.ifc.rcv.low is not changed when nfs.tcp.recvwindow size is adjusted.	33170	Yes	No
7.3	nfs.ifc.xmt.high	NFS goes into transmit flow control when the send window is full and the number of outstanding requests increases beyond nfs.ifc.xmt.high. At that time NFS will stop reading from the TCP input window. The default value for this option is set to 16 . Its maximum limit is 64. The value of this setting is persistent across reboots.	16	Yes	No
7.3	nfs.ifc.xmt.low	NFS comes out of flow control when the number of outstanding requests goes below nfs.ifc.xmt.low. The default value for this option is set to 8 . Its minimum value is 0. The value of this setting is persistent across reboots.	8	Yes	No
7.1	nfs.netgroup.strict	Determines whether netgroups should be prefixed with the '@' symbol. Strictly defining netgroups improves name resolution performance. The default value for this option is off .	on	Yes	Yes
8.1.1	nfs.nfs_rootonly	Controls all NFS requests from nonreserved ports. The default value for this option is off .	off	Yes	Yes
7.2	nfs.ntacl_display_permissive_perms	In ONTAP 7.2.1 and above the display permissions to NFSv3 clients on files with NT ACLs are based on the minimum access granted by NT ACL to any user and are therefore closer to the actual permissions. Some applications which check the rights granted by display permissions before attempting to access the files may start behaving differently. To prevent any unexpected issues set this option to on. The default value for this option is off .	on	Yes	Yes
7.3	nfs.rpcsec.ctx.high	The default value is zero . If set to a value other than zero it sets a high-water mark on the number of stateful RPCSEC_GSS (see RFC2203) authentication contexts (today, only Kerberos V5 produces stateful authentication state in NFS). If it is zero, then no explicit high-water mark is set.	0	Yes	No
7.3	nfs.rpcsec.ctx.idle	The default value is 360 seconds. This is the amount of time, in seconds, an RPCSEC_GSS context (see the description for the nfs.rpcsec.ctx.high option) will be permitted to be unused before it is deleted.	360	Yes	No
8.1.2	nfs.rpcsec.trace	Enables EMS messages for troubleshooting Kerberos issues. The default value for this option is off .	Off	Yes	Yes
7.3	nfs.tcp.recvwindow size	Determines the number of TCP messages that can be transmitted between the Filer and an NFS client at one time. As of ONTAP 7.3 the default value for this option is 65940 .	65940	Yes	Yes

From ONTAP Release	ONTAP Security Option Name	Description	Value	Applicability	
				Physical	vFiler
8.1.1	nfs.v2.enable	Enables or disables NFSv2. The default value for this option is off . Prior to version 8.1.1 there was no way to disable NFSv2 within ONTAP. Before changing this setting to off in the Citi build an environment analysis is required to confirm no NFSv2 traffic currently exists on the storage system.	on	Yes	Yes
8.1	nfs.v3.snapshot. active.fsid.enable	Specifies the file system ID handling behavior for NFSv3. The default value for this option is on .	on	Yes	Yes
7.3.3	nfs.v4.enable	When enabled, the NFS server supports NFS version 4. NFS version 4 support is only over the TCP protocol. Valid values for this option are on (enabled) or off (disabled). The default value for this option is off .	off (if NFSv4 is not required) on (if NFSv4 is required)	Yes	Yes
7.3.3	nfs.v4.acl.enable	When enabled, ACLs are supported for NFS version 4. The ACL option controls setting and getting NFSv4 ACLs. It does not control enforcement of these ACLs for access checking. This feature is not supported over NFS versions 2 and 3. The default value for this option is off . Enabling this feature will have no effect unless the setting nfs.v4.enable is also enabled. Due to BURT 458144 if setting nfs.v4.acl.enable on a vfiler this option must also be set to "on" on the physical filer as well for the feature to function correctly if the system is running ONTAP 8.0.1. When running 8.0.2P4 or above and implementing NFSv4 this feature only needs to be enabled in the affected vFiler, and not on the physical hosting system (unless, of course, you plan to implement NFSv4 at the physical level too).	on	Yes	Yes
8.1.1	nfs.v4.id.allow_ numerics	Allows or disallows NFSv4 user and group IDs as strings. The default value for this option is off .	off	Yes	Yes
7.3.3	nfs.v4.id domain	This option controls the domain portion of the string form of user and group names as defined in the NFS version 4 protocol. The standard for the NFSv4 ID domain name is to use the string "nfsdomain", this value is to be used for all Citi locations. The same value will be set on all Citi OS platforms looking to connect using NFSv4.	nfsdomain	Yes	Yes
7.3.3	nfs. v4.read_delegation	Read delegations allow NFS version 4 clients to do read operations locally without contacting the server. These include open for read, read locks and file read operations. Both the server and client must support read delegations for this feature to work. When enabled, read delegations are supported for NFS version 4. This feature is not supported for NFS versions 2 and 3. The default value for this option is off.	on	Yes	Yes

From ONTAP Release	ONTAP Technical Option Name	Description	Value	Applicability	
				Physical	vFiler
7.3.3	nfs.v4.write_delegation	Write delegations allow NFS version 4 clients to do write operations locally without contacting the server. These include open for write, write locks and writing to files. Both the server and client must support write delegations for this feature to work. When enabled, write delegations are supported for NFS version 4. This feature is not supported over NFS versions 2 and 3. Valid values for this option are on (enabled) or off (disabled). The default value for this option is off .	on	Yes	Yes
7.3.3	nfs.v4.acl_preserve	The default Data ONTAP behavior is to drop the ACL on a file or directory when a chmod, with mode bits (apart from SUID/SGID/STVX), is performed on that file or directory. This hidden option, provided in ONTAP 7.3.3, to changes this behavior so that the ACL is no longer dropped. The default value for this is off . When running 8.0.2P4 or above and implementing NFSv4 this feature only needs to be enabled in the affected vFiler, and not on the physical hosting system (unless, of course, you plan to implement NFSv4 at the physical level too). Enabling this feature will have no effect unless the setting nfs.v4.enable is also enabled. Due to bug 438464 this option should never be enabled on systems running ONTAP 8.0.1 in order to avoid a data corruption issue. With ONTAP 8.0.2 or later this bug is resolved.	off (8.0.1) on (8.0.2)	Yes	Yes
8.1	nfs.v4.snapshot.active.fsid.enable	Specifies the file system ID handling behavior for NFSv4. The default value for this option is off .	off	Yes	Yes
8.1.1	nfs.vstorage.enable	Enables or disables VMware vStorage over NFS (VAAI-NFS). The default value for this option is off .	off	Yes	Yes
6.5	raid.reconstruct.wafliron.enable	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 bypassing 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.	on	Yes	No
6.5	raid.scrub.duration	The duration of the automatic disk scrubbing operation (in minutes). If scrubbing does not finish in this defined period, Data ONTAP records where it stops. The default value for this option is 360 .	600	Yes	No
6.4	snmp.enable	Enables the SNMP server on the filer. Valid values for this option are on or off. The default value for this option is on.	on	Yes	No
7.3	stats.archive.enable	Enables or disables the Performance Archiver feature within ONTAP. Valid values for this option are on or off. The default value for this option is on. This is a hidden option within ONTAP whose value will only be displayed if the option name is queried specifically.	on	Yes	No
6.5	timed.enable	Enables the time daemon (timed) synchronization to an external source. Valid values for this option are on or off. The default value for this option is on. In a clustered environment it is important that both nodes on the cluster have the same setting for this option.	on	Yes	No

From ONTAP Release	ONTAP Technical Option Name	Description	Value	Applicability	
				Physical	vFiler
6.5	timed.log	This setting is obsolete as of ONTAP 8.0 or above. See NetApp Knowledge Base article 1010058 for more details. Enables whether the time changes initiated by the time daemon should be logged to the console. Valid values for this option are on or off. The default value for this option is off. Enabling this option allows for the system time state of the Filer to be recorded in the message log for use in troubleshooting.	n/a	-	-
6.5	timed.max_skew	This setting is obsolete as of ONTAP 8.0 or above. See NetApp Data ONTAP release notes for details. Specifies the maximum amount of skew between the time reported by the time server and the filer's time that ONTAP will allow before making a synchronization attempt. If the difference in the time reported by the server and the filer's time is greater than this value, the filer will not synchronize to the time reported by the time server. The default value for this option is 30m.	n/a	-	-
6.5	timed.min_skew	This setting is obsolete as of ONTAP 8.0 or above. See NetApp Data ONTAP release notes for details. Specifies the minimum amount of skew between the time reported by the time server and the filer's time that ONTAP will allow before making a synchronization attempt. Note as of ONTAP 7.2 this value defaults to 0 and must not be changed.	n/a	-	-
6.5	timed.proto	Specifies the protocol used to synchronize time. Valid values for this option are rdate, ntp or rtc. The default value for this option is ntp. Only the ntp protocol is supported within Citi	ntp	Yes	No
6.5	timed.sched	This setting is obsolete as of ONTAP 8.0 or above. See NetApp Data ONTAP release notes for details. Specifies the timed synchronization schedule. There are several pre-defined schedules: hourly, multihourly, and daily. The value 1h represents a synchronization event every hour.	n/a	-	-
6.5	timed.servers	Specifies up to five time servers used by the time daemon. Time servers are contacted in the order specified; if a server can't be contacted, the time daemon tries the next one in the list. The default value for this option is the null string. Refer to Appendix C of Network Document http://ipns.ny.smb.com/documents/Product-Development/CA-NTP/GNAST/NTP_design_RFP.doc to determine the entries for your Filer depending on its site location.	see notes	Yes	No
6.4	wafil.default_security_style	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 NTFS only filer, in which case the default is 'ntfs'.	ntfs	Yes	No
8.0.1	wafil.trunc.throttle.hipri.enable	This is an advanced option, hidden by default, which should be set to mitigate the impact of large deletes within the system. The default value for this option is off. See Customer Support Bulletin 0803-03 and bug article 90314 for more details.	off	Yes	No

From ONTAP Release	ONTAP Technical Option Name	Description	Value	Applicability	
				Physical	vFiler
7.3.3	waf1.trunc.throttle. min.size	This is an advanced option, hidden by default, which should be set to mitigate the impact of large deletes within the system. The recommended settings are platform specific: FAS250, FAS270, FAS920, FAS2020, FAS2050, and FAS3020: 510 ALL OTHER PLATFORMS: 1530 See Customer Support Bulletin 0803-03 and bug article 90314 for more details.	See description	Yes	No
8.0.1	waf1.trunc.throttle. slice.enable	This is an advanced option, hidden by default, which should be set to mitigate the impact of large deletes within the system. . The default value for this option is off . See Customer Support Bulletin 0803-03 and bug article 90314 for more details.	off	Yes	No
7.3.3P5	waf1.trunc.throttle. system.max	This is an advanced option, hidden by default, which can be set to mitigate the impact of large deletes within the system. The recommended settings are platform specific: FAS250, FAS270, FAS920, FAS2020, FAS2050, and FAS3020: 10 ALL OTHER PLATFORMS: 30 See Customer Support Bulletin 0803-03 and bug article 90314 for more details.	See description	Yes	No
7.3.3	waf1.trunc.throttle. vol.max	This is an advanced option, hidden by default, which can be set to mitigate the impact of large deletes within the system. The recommended settings are platform specific: FAS250, FAS270, FAS920, FAS2020, FAS2050, and FAS3020: 10 ALL OTHER PLATFORMS: 30 See Customer Support Bulletin 0803-03 and bug article 90314 for more details.	See description	Yes	No

There are a growing number of settings that are defined in the registry. The purpose of the following table is to capture these registry specific settings that are defined outside the scope of the traditional 'options' commands.

From ONTAP Release	Purpose	Settings
6.5	Location of system configuration information	Check with <code>registry get options.admin.config.list</code> If the outcome is not <code>/etc/registry.default,/etc/registry.local</code> then implement setting with the following command: <code>priv set advanced</code> <code>registry set options.admin.config.list /etc/registry.default,/etc/registry.local</code>
7.3	Prevent redundant API schema output validation information from being recorded in the message log	<code>priv set advanced</code> <code>registry set state.api.schema_output_validate.enable off</code>
8.0.1P5	Generates a single SNMP trap for each volume that breaches a utilization threshold and enables alerting when an aggregate breaches a utilization threshold	<code>priv set advanced</code> <code>registry set options.thresholds.fsReportEachFullVolumeChange on</code> <code>registry set options.thresholds.fsReportFullAggregates on</code> Note the aggregate level alerting requires both registry settings to be enabled (you cannot alert on aggregate utilization without enabling <code>fsReportEachFullVolumeChange</code> first)

Table 35 - Registry custom settings

With the introduction of support for the Service Processor and RLM remote management ports in Citi with ONTAP 8.1 and above, it is now nexessary to check the following configuration parameters on the SP or RLM (the port type will depend on the model of FAS hardware deployed).

From SP Firmware Release	From RLM Firmware Release	Purpose	Settings
1.3	4.0	Determines whether the remote support agent (RSA) on the remote management port is enabled or not. The default value for this setting is disabled.	Confirm the RSA is not configured by logging into the remote management port and running the following command. <pre>sp> rsa show Remote Support Agent is disabled.</pre> If the rsa show response indicates the RSA is enabled as follows then re-run the RSA setup process to disable it: <pre>Remote Support Agent is enabled.</pre> <pre>sp> rsa setup The Remote Support Agent improves your case resolution time and minimizes your manual support overhead. Would you like to enable Remote Support Agent? [yes]: no</pre>

Table 36 - Remote Management Port software configuration settings

There are certain configuration parameters within Data ONTAP which cannot be defined by an option command either due design considerations made by NetApp or the feature is new enough that a traditional 'option' configuration parameter is yet to be developed. As such these features can be configured using alternate commands. The purpose of the following table is to capture all settings which fall outside of the three mechanisms (ONTAP options, ONTAP registry and Remote Management Port) shown above.

Applies to	Location	Mechanism	Purpose	Settings
8.1.4P1 (not required for 8.1.4P4 or above)	/etc/rc	setflag	Disables SMB 2.1 to prevent occurrence of critical known issues in the SMB 2.1 oplock lease functionality. See bug 798842 and Support Bulletin: 7010114 This workaround is not required in 8.1.4P4 where SMB 2.1 is disabled by default	Append the following three lines to the /etc/rc file to disable feature at system boot: # Start workaround for bug 798842 - Disable SMB 2.1 priv set -q diag;setflag smb_enable_2_1 0;priv set # End workaround for bug 798842
8.1.4P1 (not required for 8.1.4P4 or above)	cli	setflag	Disables SMB 2.1 to prevent occurrence of critical known issues in the SMB 2.1 oplock lease functionality. See bug 798842 and Support Bulletin: 7010114 This workaround is not required in 8.1.4P4 where SMB 2.1 is disabled by default	Run the following commands to disable this feature in the running configuration: ontap> priv set -q diag;setflag smb_enable_2_1 0;priv set Then terminate and restart CIFS for every vFiler on the storage system including vFiler0

Table 37 - Additional configuration settings

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

14.4.1 Data ONTAP Install Procedure

1. Login to NetApp Support Website:
 - a. <http://support.netapp.com/NOW/cgi-bin/software/>
 - b. Go to end of Software list "To access specific version"
 - c. Select Data ONTAP and type version number (**8.1.4P1**)
 - 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 814P1_<x>_image.tgz or 814P1_<x>_image.zip files for deployment via a web server (where <X> represents the one character hardware platform identifier). Refer to section 8 of the [CATE ONTAP 8.1.4 Release Notes](#) for the correct software image for your platform.
 - i. Select the appropriate version for your hardware platform according to the previous step
 - j. Right-click on the appropriate 814P1_<x>_image file link and select "Save as" to download the software. Make a note of the location you are downloading the file to.
 - k. Copy the file to a web server / web site. Make a note of the website name and path for use in the next step. Note this could be the /etc/http directory on an existing Filer.
2. SSH or Remote Console to Filer CLI.
 - a. From the Filers console issue the software get command to retrieve the ONTAP kernel image from the web server in the previous set. For example:
`Filer> software get http://webserver/netapp/ontap/814P1_q_image.tgz`
 - b. Confirm the image was retrieved with the `software list` command
3. By using the Filer CLI session from the previous step.
 - a. Run the software update command to install the ONTAP code on the filer "software update <package_name> -d". Use the name of the package defined in 1h) to replace the variable <package_name> in this command.
 - b. If filer is a HA node; type "cf disable" command to disable cluster failover to partner node.
 - c. Type "download" command. Once completed, your filer will have the downloaded DOT in place.
 - d. Type "reboot" command to complete upgrade.
 - e. Repeat steps 2 – 3 to upgrade partner node.
 - f. Upon completion of both nodes, type "cf enable" from any node to re-enable cluster failover.

14.4.2 Disk Firmware Upgrade Procedure

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

```

SOE SSH Client - ssh -l root 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 19 - Disk firmware update

14.4.3 Shelf Firmware Upgrade Procedure

1. Download the latest shelf firmware pack from the NetApp Support site
2. Expand zip file to `\\filename\c$\etc\shelf_fw`
3. Set the Filers console to advanced mode '`priv set advanced`'
4. On the Filers console run '`storage download shelf`'
5. Download the latest IOM ACP firmware pack from the NOW site
6. Expand zip file to `\\filename\c$\etc\acpp_fw`
7. On the Filers console run '`storage download acp`'

The above steps are demonstrated in the following example:

```

11c2-fas3240-01> priv set advanced
11c2-fas3240-01> ls /etc/shelf_fw
.
..
AT-FCX.3800.SFW
AT-FCX.3800.SFW.FVF
ESH.9520.SFW
ESH.9520.SFW.FVF
ESH2.2002.SFW
ESH2.2002.SFW.FVF
ESH4.1400.SFW
ESH4.1400.SFW.FVF
IOM3.0121.SFW
IOM3.0121.SFW.FVF
LRC.1103.SFW
LRC.1103.SFW.FVF
SAS.0500.SFW
SAS.0500.SFW.FVF
11c2-fas3240-01*> logger "Copy new firmware to system volume"
11c2-fas3240-01*> Tue Apr 12 20:01:17 EDT [11c2-fas3240-01: logger.usr:notice]:
    Copy new firmware to system volume

```



```
11c2-fas3240-01*> ls /etc/shelf_fw
.
..
AT-FCX.3800.SFW
AT-FCX.3800.SFW.FVF
ESH.9520.SFW
ESH.9520.SFW.FVF
ESH2.2002.SFW
ESH2.2002.SFW.FVF
ESH4.1400.SFW
ESH4.1400.SFW.FVF
IOM3.0121.SFW
IOM3.0121.SFW.FVF
LRC.1103.SFW
LRC.1103.SFW.FVF
SAS.0500.SFW
SAS.0500.SFW.FVF
IOM3.0160.SFW
IOM3.0160.SFW.FVF
11c2-fas3240-01*> storage download shelf
```

Downloading disk shelf firmware may take up to 10 minutes,
but will NOT disrupt client access during that time.

```
Are you sure you want to continue with shelf firmware update? y
11c2-fas3240-01*> Tue Apr 12 20:01:48 EDT [11c2-fas3240-01:
sfu.ctrllerElmntsPerShelf:info]: [storage download shelf]: 2 ES controller elements
can be updated on 3a.shelf1.
Tue Apr 12 20:01:48 EDT [11c2-fas3240-01: sfu.ctrllerElmntsPerShelf:info]: [storage
download shelf]: 2 ES controller elements can be updated on 3c.shelf11.
Tue Apr 12 20:01:48 EDT [11c2-fas3240-01: sfu.downloadStarted:info]: Update of disk
shelf firmware started on 2 shelves.
Tue Apr 12 20:01:48 EDT [11c2-fas3240-01: sfu.downloadingController:info]: [storage
download shelf]: Downloading IOM3.0160.SFW on disk shelf controller module B on
3a.shelf1.
Tue Apr 12 20:01:48 EDT [11c2-fas3240-01: sfu.downloadingController:info]: [storage
download shelf]: Downloading IOM3.0160.SFW on disk shelf controller module A on
3c.shelf11.
Tue Apr 12 20:02:11 EDT [11c2-fas3240-01: sfu.rebootRequest:info]: Issuing a request
to reboot disk shelf 3a.shelf1 module B.
Tue Apr 12 20:02:11 EDT [11c2-fas3240-01: sfu.rebootRequest:info]: Issuing a request
to reboot disk shelf 3c.shelf11 module A.
Tue Apr 12 20:02:11 EDT [11c2-fas3240-01: sfu.adapterSuspendIO.ndu:info]: Suspending
SMP to SAS adapter 4b for 30 seconds while shelf firmware is updated.
Tue Apr 12 20:02:11 EDT [11c2-fas3240-01: sfu.adapterSuspendIO.ndu:info]: Suspending
SMP to SAS adapter 3a for 30 seconds while shelf firmware is updated.
Tue Apr 12 20:02:11 EDT [11c2-fas3240-01: sfu.adapterSuspendIO.ndu:info]: Suspending
SMP to SAS adapter 4a for 30 seconds while shelf firmware is updated.
Tue Apr 12 20:02:11 EDT [11c2-fas3240-01: sfu.adapterSuspendIO.ndu:info]: Suspending
SMP to SAS adapter 3c for 30 seconds while shelf firmware is updated.
Tue Apr 12 20:02:40 EDT [11c2-fas3240-01: sfu.downloadingController:info]: [storage
download shelf]: Downloading IOM3.0160.SFW on disk shelf controller module A on
3a.shelf1.
Tue Apr 12 20:02:40 EDT [11c2-fas3240-01: sfu.downloadingController:info]: [storage
download shelf]: Downloading IOM3.0160.SFW on disk shelf controller module B on
3c.shelf11.
Tue Apr 12 20:03:04 EDT [11c2-fas3240-01: sfu.rebootRequest:info]: Issuing a request
to reboot disk shelf 3a.shelf1 module A.
Tue Apr 12 20:03:04 EDT [11c2-fas3240-01: sfu.rebootRequest:info]: Issuing a request
to reboot disk shelf 3c.shelf11 module B.
Tue Apr 12 20:03:04 EDT [11c2-fas3240-01: sfu.adapterSuspendIO.ndu:info]: Suspending
SMP to SAS adapter 4b for 30 seconds while shelf firmware is updated.
Tue Apr 12 20:03:04 EDT [11c2-fas3240-01: sfu.adapterSuspendIO.ndu:info]: Suspending
SMP to SAS adapter 3a for 30 seconds while shelf firmware is updated.
Tue Apr 12 20:03:04 EDT [11c2-fas3240-01: sfu.adapterSuspendIO.ndu:info]: Suspending
SMP to SAS adapter 4a for 30 seconds while shelf firmware is updated.
```

```
Tue Apr 12 20:03:04 EDT [11c2-fas3240-01: sfu.adapterSuspendIO.ndu:info]: Suspending
SMP to SAS adapter 3c for 30 seconds while shelf firmware is updated.
Tue Apr 12 20:03:33 EDT [11c2-fas3240-01: sfu.downloadSuccess:info]: [storage
download shelf]: Firmware file IOM3.0160.SFW downloaded on 3c.shelf11.
Tue Apr 12 20:03:33 EDT [11c2-fas3240-01: sfu.downloadSuccess:info]: [storage
download shelf]: Firmware file IOM3.0160.SFW downloaded on 3a.shelf1.
Tue Apr 12 20:03:33 EDT [11c2-fas3240-01: sfu.downloadSummary:info]: Shelf firmware
updated on 2 shelves.
```

```
11c2-fas3240-01*> sasadmin expander_map
```

```
Expanders on channel 0a:
Cannot Complete operation on channel 0a; Status Not Available.
Expanders on channel 0b:
Cannot Complete operation on channel 0b; Status Not Available.
Expanders on channel 3a:
Level 1: WWN 50050cc102896b7f, ID 1, Serial Number ' SHU0954292G13EP', Product
'DS424IOM3', Rev '0160', Slot A
Expanders on channel 3b:
Cannot Complete operation on channel 3b; Status Not Available.
Expanders on channel 3c:
Level 1: WWN 50050cc1029fb07f, ID 11, Serial Number ' SHU0954293G17RX', Product
'DS424IOM3', Rev '0160', Slot B
Expanders on channel 3d:
Cannot Complete operation on channel 3d; Status Not Available.
Expanders on channel 4a:
Level 1: WWN 50050cc1029fb1bf, ID 11, Serial Number ' SHU0954293G17RX', Product
'DS424IOM3', Rev '0160', Slot A
Expanders on channel 4b:
Level 1: WWN 50050cc10289717f, ID 1, Serial Number ' SHU0954292G13EP', Product
'DS424IOM3', Rev '0160', Slot B
Expanders on channel 4c:
Cannot Complete operation on channel 4c; Status Not Available.
Expanders on channel 4d:
Cannot Complete operation on channel 4d; Status Not Available.
```

```
11c2-fas3240-01*> ls /etc/acpp_fw
```

```
.
..
ACP-IOM3.0120.AFW
ACP-IOM3.0120.AFW.FVF
ACP-IOM3.0131.AFW.FVF
ACP-IOM3.0131.AFW
```

```
11c2-fas3240-01*> storage show acp
```

```
Alternate Control Path: Enabled
Ethernet Interface: e0P
ACP Status: Active
ACP IP Address: 192.168.0.75
ACP Domain: 192.168.0.1
ACP Netmask: 255.255.252.0
ACP Connectivity Status: Full Connectivity
```

Shelf Module	Reset Cnt	IP Address	FW Version	Module Type	Status
3a.01.A	000	192.168.3.35	01.20	IOM3	active
3a.01.B	000	192.168.3.59	01.20	IOM3	active
3c.11.A	000	192.168.0.59	01.20	IOM3	active
3c.11.B	000	192.168.0.55	01.20	IOM3	active

```
11c2-fas3240-01*> storage download acp
```

Downloading ACP firmware can take several minutes and will not disrupt client access during that time. However, normal ACP recovery capabilities will not be available while the firmware upgrade is in progress.

```
Are you sure you want to continue with ACP processor firmware update? y
[storage download acp] firmware will be downloaded to 4 ACP modules.
Tue Apr 12 20:10:19 EDT [11c2-fas3240-01: acp.command.sent:info]: Sent firmware
download (image: ACP-IOM3.0131.AFW) command to 3c.11.B (192.168.0.55), (disk shelf
serial number: SHU0954293G17RX).
11c2-fas3240-01*> Tue Apr 12 20:10:19 EDT [11c2-fas3240-01:
acp.command.response:info]: Command firmware download to 3c.11.B (192.168.0.55) was
successful, (disk shelf serial number: SHU0954293G17RX).
Tue Apr 12 20:10:19 EDT [11c2-fas3240-01: acp.command.sent:info]: Sent firmware
download (image: ACP-IOM3.0131.AFW) command to 3c.11.A (192.168.0.59), (disk shelf
serial number: SHU0954293G17RX).
Tue Apr 12 20:10:19 EDT [11c2-fas3240-01: acp.command.response:info]: Command
firmware download to 3c.11.A (192.168.0.59) was successful, (disk shelf serial
number: SHU0954293G17RX).
Tue Apr 12 20:10:19 EDT [11c2-fas3240-01: acp.command.sent:info]: Sent firmware
download (image: ACP-IOM3.0131.AFW) command to 3a.1.A (192.168.3.35), (disk shelf
serial number: SHU0954292G13EP).
Tue Apr 12 20:10:19 EDT [11c2-fas3240-01: acp.command.response:info]: Command
firmware download to 3a.1.A (192.168.3.35) was successful, (disk shelf serial number:
SHU0954292G13EP).
Tue Apr 12 20:10:19 EDT [11c2-fas3240-01: acp.command.sent:info]: Sent firmware
download (image: ACP-IOM3.0131.AFW) command to 3a.1.B (192.168.3.59), (disk shelf
serial number: SHU0954292G13EP).
Tue Apr 12 20:10:19 EDT [11c2-fas3240-01: acp.command.response:info]: Command
firmware download to 3a.1.B (192.168.3.59) was successful, (disk shelf serial number:
SHU0954292G13EP).
```

```
storage show acp
```

```
Alternate Control Path: Enabled
Ethernet Interface:      e0P
ACP Status:             Active
ACP IP Address:         192.168.0.75
ACP Domain:             192.168.0.1
ACP Netmask:            255.255.252.0
ACP Connectivity Status: Full Connectivity
```

Shelf Module	Reset Cnt	IP Address	FW Version	Module Type	Status
3a.01.A (upgrading firmware)	000	192.168.3.35	01.20	IOM3	inactive
3a.01.B (upgrading firmware)	000	192.168.3.59	01.20	IOM3	inactive
3c.11.A (upgrading firmware)	000	192.168.0.59	01.20	IOM3	inactive
3c.11.B (upgrading firmware)	000	192.168.0.55	01.20	IOM3	inactive

```
11c2-fas3240-01*> Tue Apr 12 20:13:46 EDT [11c2-fas3240-01:
acp.upgrade.successful:info]: ACP module 3a.1.A (192.168.3.35) successfully upgraded
firmware, (disk shelf serial number: SHU0954292G13EP).
Tue Apr 12 20:13:48 EDT [11c2-fas3240-01: acp.upgrade.successful:info]: ACP module
3a.1.B (192.168.3.59) successfully upgraded firmware, (disk shelf serial number:
SHU0954292G13EP).
Tue Apr 12 20:13:49 EDT [11c2-fas3240-01: acp.upgrade.successful:info]: ACP module
3c.11.B (192.168.0.55) successfully upgraded firmware, (disk shelf serial number:
SHU0954293G17RX).
```

Tue Apr 12 20:14:02 EDT [11c2-fas3240-01: acp.upgrade.successful:info]: ACP module 3c.11.A (192.168.0.59) successfully upgraded firmware, (disk shelf serial number: SHU0954293G17RX).

storage show acp

Alternate Control Path: Enabled
Ethernet Interface: e0P
ACP Status: Active
ACP IP Address: 192.168.0.75
ACP Domain: 192.168.0.1
ACP Netmask: 255.255.252.0
ACP Connectivity Status: Full Connectivity

Shelf Module	Reset Cnt	IP Address	FW Version	Module Type	Status
-					
3a.01.A	000	192.168.3.35	01.31	IOM3	active
3a.01.B	000	192.168.3.59	01.31	IOM3	active
3c.11.A	000	192.168.0.59	01.31	IOM3	active
3c.11.B	000	192.168.0.55	01.31	IOM3	active

14.5 Appendix E: 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 - 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 38 - Tape SAN Infrastructure configuration steps

14.5.1 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.

14.5.1.1 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.

14.5.1.2 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.

14.5.1.3 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 39 - 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.

14.5.1.4 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 r st0a on the filer will be the first drive on the NetBackup console, and r st1a will be the second (and so on).

14.5.1.5 Lab performance benchmarks

The following performance statistics were recorded as part of the lab test using ONTAP 7.1:

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 40 - Tape SAN Lab test performance benchmarks

14.6 Appendix F: SOE VFiler Creation Example

14.6.1 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/vfxxp #</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/vfxxp # 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>filernam\c\$</code> directory, edit /etc/rc file, add the following line: <code>ifconfig interface alias xxx.xxx.xxx.yyy netmask 255.255.xxx.yyy up</code>	<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>

14.6.1.1 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>

14.6.1.2 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. LDAP should be configured after CIFS and NIS.

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.	<code>vfiler run <vfilername> setup</code>	<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	<code>vfiler run <vfilername> dns info</code>	<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>

14.6.1.3 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	<code>vfiler run <vfilername> cifs testdc</code>	<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	<code>vfiler run <vfilername> cifs shares</code>	<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	<code>vfiler run <vfilername> cifs shares - delete <sharename></code>	<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>

14.6.1.4 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.

14.6.1.5 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/PDolvw 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 /vfguen1_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 /vfguen1_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>


14.6.1.6 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>hostfiler1C\$</i> 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>targetfiler1C\$</i> share. Using your favorite editor: edit or create the <i>/etc/snapmirror.conf</i> file on the local filer. The snapmirror.conf file is used to keep the snapmirrored 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>

14.6.2 SOE Vfiler DR/COB Creation Example

14.6.2.1 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-ipSPACE
  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 00:00:00:00:00:76:68 (VIA Provider)
vif0: flags=1948043<UP,BROADCAST,RUNNING,MULTICAST,TCPCKSUM> 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)
```

14.6.2.2 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 LDAP namespace
- 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.

14.6.2.3 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
```

14.6.2.4 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	<code>vfiler create vfilername -r /vfilerqtree</code>	<pre>filer2> vfiler create vfiler1-gue -r /vol/vol0/vfguen1 Thu Jun 5 16:17:43 EDT [vfiler1- gue@nbns_timer:info]: All CIFS name registrations complete for local server Thu Jun 5 16:17:51 EDT [vfiler1-gue@rc:info]: Lookup of _ldap._tcp.nyc._sites.NAM.NSR00T.NET failed with DNS server 146.128.201.218: Can't assign requested address. Thu Jun 5 16:17:58 EDT [vfiler1- gue@rc:warning]: Cannot connect to server. Socket connect error: Can't assign requested address Thu Jun 5 16:17:58 EDT [vfiler1-gue@rc:info]: CIFS - can't make TCP connection with \\NAMDCNY002 Thu Jun 5 16:18:13 EDT [vfiler1-gue@rc:error]: Can't connect socket to WINS server: Can't assign requested address Thu Jun 5 16:18:13 EDT [vfiler1-gue@rc:error]: Can't connect socket to WINS server: Can't assign requested address Thu Jun 5 16:18:16 EDT [vfiler1- gue@rc:warning]: CIFS - Could not find a domain controller. CIFS local server is running. Thu Jun 5 16:18:16 EDT [vfiler1- 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 [vfiler1- gue@nis_poll:info]: NIS server(s) not available Thu Jun 5 16:18:16 EDT [vfiler1- 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.smb.com vfiler1-gue running ipspace: default-ipspace 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 vfiler1-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 [vfiler1- gue@rc:notice]: Vfiler vfiler1-gue restored. filer2></pre>

Action	Example	Results
From COB/DR filer; Add IP address resource from COB/DR network to vfiler	<code>vfiler add vfilername -i ipaddress</code>	<pre>filer2> vfiler add vfiler1-gue -i 192.168.30.144 Thu Jun 5 16:22:49 EDT [vfiler1- 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 "vfiler1-gue" filer2></pre>
From COB/DR filer; Add the IP address as an alias to the COB.DR filer network interface	<code>ifconfig interface alias ipaddress netmask xxx.xxx.xxx.yyy</code>	<pre>filer2> 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.) filer2></pre>
From COB/DR filer; Check the vfiler status using the -r option. Ensure: vfiler is running IP address assigned is up		<pre>filer2> vfiler status -r vfiler0 running ipspace: default-ipospace IP address: 192.168.30.142 Path: / UUID: 00000000-0000-0000-0000-000000000000 Vfiler1-gue running ipspace: default-ipospace 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 filer2></pre>
From COB/DR filer; Check the vfiler's NIS info	<code>vfiler run vfilername nis info</code>	<pre>filer2> vfiler run vfiler1-gue nis info ===== vfiler1-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 filer2></pre>
From COB/DR filer; Check the vfiler's DNS info	<code>vfiler run vfilername dns info</code>	<pre>filer2> vfiler run vfiler1-gue dns info ===== vfiler1-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 filer2></pre>

Action	Example	Results
From User Workstation; Unmount the production vfiler exported <i>path</i>	<code>umount mountname</code>	<code>bash-2.03# umount /vfguen1_v1</code>
From User Workstation; Re-mount the COB/DR vfiler exported <i>path</i>	<code>mount vfilername:/path</code>	<code>bash-2.03# mount vfiler1-gue:/vol/gue_v1 /vfguen1_v1</code>
From User Workstation; Verify mount; list mount directory	<code>ls /mountpoint</code>	<code>bash-2.03# ls -ltr /vfguen1_v1</code> 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.

14.6.2.5 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	<code>vfiler stop vfilername</code>	<code>filer2> vfiler stop vfiler1-gue</code> vfiler1-gue stopped filer2>

Action	Example	Results
Reboot production filer	<code>boot</code>	<pre>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</pre>
From COB/DR filer; Resync snapmirror of vfiler resources.	<code>snapmirror resync path</code>	<pre>filer2> snapmirror resync gue_v1 The resync base snapshot will be: filer2 (0050385317)_gue_v1.112 These older snapshots have already been deleted from the source and will be deleted from the destination: filer2 (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 filer1:gue_v1 using filer2 (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 vfiler1- 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 vfiler1- gue: storage unit /vol/gue_v1 now online filer2> snapmirror resync /vol/vol0/vfguen1 The resync base snapshot will be: filer2 (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 filer1:/vol/vol0/vfguen1 : success Transfer started. Monitor progress with 'snapmirror status' or the snapmirror log. filer2></pre>

Action	Example	Results
From COB/DR filer; Verify snapmirror status.	<code>snapmirror status</code>	<pre>filer2> snapmirror status Snapmirror is on. Source Destination State Lag Status filer1:/vol/vol0/vfguen1 filer2:/vol/vol0/vfguen1 Snapmirrored 00:00:33 Idle filer1:gue_v1 filer2:gue_v1 Snapmirrored 00:01:34 Idle filer2></pre>
From production filer; Verify vfiler status	<code>vfiler status -r</code>	<pre>filer1> vfiler status -r vfiler0 running ipspace: default-ipospace 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 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 filer1></pre>
From User Workstation; Unmount the production vfiler exported <i>path</i>	<code>umount mountname</code>	<pre>bash-2.03# umount /vfguen1_v1</pre>
From User Workstation; Re-mount the COB/DR vfiler exported <i>path</i>	<code>mount vfilername:/path</code>	<pre>bash-2.03# mount vfiler1-gue:/vol/gue_v1 /vfguen1_v1</pre>

Action	Example	Results
From User Workstation; Verify mount; list mount directory	<code>ls /mountpoint</code>	<pre> 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# </pre>
From management workstation: Verify CIFS access to /etc directory	Windows explorer, map to network drive <code>vfilenamelc\$</code>	

14.7 Appendix G: Role Based Access Control Capabilities

14.7.1 Overview

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




Figure 20 - 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 grouped 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.

14.7.2 High-level Administrative User Concepts

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




Figure 21 - RBAC high level overview

14.7.2.1 Common Capabilities

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

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

14.7.2.2 System Administrator

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

a) File Storage Administrator

Create/manage qtrees, shares; run volume utilization reports

i) CIFS Administrator

Create/modify/delete CIFS shares

ii) NFS Administrator

Create/modify/delete NFS mount points

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

14.7.2.3 Account Administrator

Account administrator is the equivalent of the Infrastructure Security Administrator (ISA). The term 'account' is used in preference 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

14.7.3.2 File System Administrator




Figure 23 - File System Administrator Role

14.7.3.3 CIFS File System Administrator




Figure 24 - CIFS File System Administrator Role

14.7.3.4 NFS File System Administrator




Figure 25 - NFS File System Administrator Role

14.7.3.5 Storage Administrator




Figure 26 - Storage Administrator Role

14.7.3.6 Security Administrator




Figure 27 - Security Administrator Role

14.7.3.7 Performance Administrator




Figure 28 - Performance Administrator Role

14.7.3.8 Network Administrator




Figure 29 - Network Administrator Role

14.7.3.9 License Administrator




Figure 30 - License Administrator Role

14.7.3.10 Backup Administrator




Figure 31 - Backup Administrator Role

14.7.3.11 Account Administrator




Figure 32 - Account Administrator Role

14.7.3.12 Auditor / Monitor




Figure 33 - Auditor/Monitor Role

14.7.3.13 Application integration: EMC Vpr SRM




Figure 34 - EMC Vpr SRM Role

14.7.4 RBAC Deployment Script

The following scripts have been developed to automate the deployment of RBAC to NetApp NAS Filers.

- Download the script files (txt format) from the [CATE Storage and Backup Engineering website](#).
- Take the text file and save it to the /etc/ directory on the filer to which you wish to deploy RBAC
- Run the script file using the source command 'source /etc/<file_name>' (replace <file_name> with the name of the actual file)
- Check the system message log to confirm all groups and roles have been created successfully

14.7.4.1 RBAC for ONTAP 8.1.4 new installs on a 64-bit hardware platform

This script has been created to configure RBAC on NetApp Filers running 8.1.4 that are new system installations.

As ONTAP 8 is a 64-bit operating system, only 64-bit hardware is supported therefore there is no longer a requirement to provide 32-bit and 64-bit versions of the CATE RBAC scripts.

The script has been listed here for reference purposes only.

To avoid line wrapping issues do not copy the script from this document, download the files for the [CATE Storage and Backup Engineering website](#).

```
#
# Revision history
#   v1.0 - Initial release
#

logger "Start CATE Storage and Backup Engineering ONTAP 8.1.4 7-mode RBAC Script v.1.0 for new system installs"
logger "Start RBAC role creation"

priv set diag; useradmin role add rol_login_sys_admin -c "System Admin Login Role" -a login-console,login-ssh,login-http-admin,login-ndmp,login-sp
priv set diag; useradmin role add rol_cli_sys_1 -c "System Admin CLI System Role Set 1" -a cli-anvl*,cli-bmc*,cli-com*,cli-coredump*,cli-dbg*,cli-
download*,cli-dumpblock*,cli-dumpstack*,cli-ems*,cli-export_stats*,cli-filersio*,cli-fru_led*
priv set diag; useradmin role add rol_cli_sys_2 -c "System Admin CLI System Role Set 2" -a cli-gdb*,cli-getXXbyYY*,cli-halt*,cli-hammer*,cli-
hostname*,cli-httpstat*,cli-icbulk*,cli-ic*,cli-igroup*,cli-inodepath*,cli-iscsi*,cli-iswt*
priv set diag; useradmin role add rol_cli_sys_3 -c "System Admin CLI System Role Set 3" -a cli-kma_stats*,cli-kt*,cli-labelmaint*,cli-label*,cli-
log_fio*,cli-log*,cli-mbstat*,cli-memerr*,cli-mt*,cli-ndmpd*,cli-nm*,cli-ontapi*
priv set diag; useradmin role add rol_cli_sys_4 -c "System Admin CLI System Role Set 4" -a cli-panic*,cli-portset*,cli-printflag*,cli-prof*,cli-
raid_config*,cli-rastrace*,cli-rdate*,cli-reboot*,cli-registry*,cli-restore_backup*,cli-result*
priv set diag; useradmin role add rol_cli_sys_5 -c "System Admin CLI System Role Set 5" -a cli-revert_to*,cli-rm*,cli-rshkill*,cli-rshstat*,cli-
rsm*,cli-rtag*,cli-rtfo*,cli-savecore*,cli-sectrace*,cli-setflag*,cli-setup*,cli-show_bios_log*
priv set diag; useradmin role add rol_cli_sys_6 -c "System Admin CLI System Role Set 6" -a cli-show_faults*,cli-showfh4*,cli-showfh*,cli-sh*,cli-
signal*,cli-slist*,cli-sm_mon_old*,cli-sm_mon*,cli-sm_not*,cli-smb_hist*,cli-snmp*,cli-software*
priv set diag; useradmin role add rol_cli_sys_7 -c "System Admin CLI System Role Set 7" -a cli-source*,cli-stack*,cli-stty*,cli-sync*,cli-
syslog*,cli-tape_qual*,cli-test_lcd*,cli-timezone*,cli-time*,cli-toe*,cli-treecompare*,cli-treemigrate*
priv set diag; useradmin role add rol_cli_sys_8 -c "System Admin CLI System Role Set 8" -a cli-ttcp*,cli-ups*,cli-uptime*,cli-vtic*,cli-
waf_l_steal_stats*,cli-waf_l_susp*,cli-waf_l*,cli-vm_stat*,cli-waffinity_stats*,cli-xttcp*
priv set diag; useradmin role add rol_cli_sys_9 -c "System Admin CLI System Role Set 9" -a cli-acorn*,cli-autosupport*,cli-bringhome*,cli-
cxbgtool*,cli-ifswitch*,cli-iomem*,cli-mailbox*,cli-mviatest*,cli-nv8*
```

```

priv set diag; useradmin role add rol_cli_sys_10 -c "System Admin CLI System Role Set 10" -a cli-ct_phy_read_t3*,cli-ct_phy_read*,cli-
ct_phy_write_t3*,cli-ct_phy_write*,cli-ct_reg_read_t3*,cli-ct_reg_read*,cli-ct_reg_write_t3*,cli-ct_reg_write*
priv set diag; useradmin role add rol_cli_sys_11 -c "System Admin CLI System Role Set 11" -a cli-ct_tpi_read*,cli-
ct_tpi_write*,cli-du*,cli-rpc_stats*,cli-schedstat*,cli-setral*,cli-sldiag*,cli-smf*,cli-stsb*,cli-sum32*
priv set diag; useradmin role add rol_cli_sys_12 -c "System Admin CLI System Role Set 12" -a cli-spinhi_stats*,cli-spinnp_replay_stats*,cli-
spinnp_replay*,cli-systemshell*,cli-tcp_client*,cli-tcp_server*,cli-udp_client*,cli-udp_server*,cli-vmsservices*
priv set diag; useradmin role add rol_cli_sys_13 -c "System Admin CLI System Role Set 13" -a cli-bootargs*,cli-ct_dump_t3*,cli-ct_reset_t3*,cli-
system*,cli-vdom*,cli-vol_db*,cli-vstorage*,cli-wafl_cmd_restrictions*,cli-wafltop*
priv set diag; useradmin role add rol_cli_sys_file_1 -c "System Admin CLI File Role Set 1" -a cli-availtime*,cli-charmap*,cli-cifs*,cli-df*,cli-
exit*,cli-export_tbl_dump*,cli-exportfs*,cli-file*,cli-fpolicy*,cli-fsecurity*,cli-ftp*
priv set diag; useradmin role add rol_cli_sys_file_2 -c "System Admin CLI File Role Set 2" -a cli-lock*,cli-logout*,cli-ls*,cli-maxfiles*,cli-
mkfile*,cli-ndmcopy*,cli-nfs_hist*,cli-nfsstat*,cli-nfs*,cli-nis*,cli-partner*,cli-qtrees*,cli-quota*
priv set diag; useradmin role add rol_cli_sys_file_3 -c "System Admin CLI File Role Set 3" -a cli-clone*,cli-reallocate*,cli-ypcat*,cli-
ypgroup*,cli-ypmatch*,cli-ypwhich*
priv set diag; useradmin role add rol_cli_sys_file_nfs_1 -c "System Admin CLI NFS File Role 1" -a cli-df*,cli-exit*,cli-export_tbl_dump*,cli-
exportfs*,cli-lock*,cli-ls*,cli-maxfiles*,cli-mkfile*,cli-nfs_hist*,cli-nfsstat*,cli-nfs*,cli-nis*,cli-qtrees*
priv set diag; useradmin role add rol_cli_sys_file_nfs_2 -c "System Admin CLI NFS File Role 2" -a cli-quota*,cli-ypcat*,cli-ypgroup*,cli-
ypmatch*,cli-ypwhich*
priv set diag; useradmin role add rol_cli_sys_str_1 -c "System Admin CLI Storage Role Set 1" -a cli-/etc/rmt*,cli-aggr*,cli-availtime*,cli-
backup*,cli-blink_off*,cli-blink_on*,cli-bootfs*,cli-cf*,cli-charmap*,cli-cifs*,cli-config*,cli-date*,cli-dd*
priv set diag; useradmin role add rol_cli_sys_str_2 -c "System Admin CLI Storage Role Set 2" -a cli-df*,cli-disk_fw_update*,cli-disk_list*,cli-
disk_stat*,cli-disk*,cli-environment*,cli-envirom*,cli-exit*,cli-export_tbl_dump*,cli-exportfs*,cli-fcadmin*
priv set diag; useradmin role add rol_cli_sys_str_3 -c "System Admin CLI Storage Role Set 3" -a cli-fcnic*,cli-fcp*,cli-fcstat*,cli-file*,cli-
flexcache*,cli-fsecurity*,cli-led_off_all*,cli-led_off*,cli-led_on_all*,cli-led_on_off*,cli-led_on*
priv set diag; useradmin role add rol_cli_sys_str_4 -c "System Admin CLI Storage Role Set 4" -a cli-led_reset_all*,cli-led_test_one*,cli-
led_test*,cli-logout*,cli-ls*,cli-lun*,cli-mv*,cli-ndmcopy*,cli-nv*,cli-options*,cli-parityck*,cli-partner*
priv set diag; useradmin role add rol_cli_sys_str_5 -c "System Admin CLI Storage Role Set 5" -a cli-perf*,cli-ping6*,cli-ping*,cli-pktt*,cli-
priority*,cli-qtrees*,cli-quota*,cli-reallocate*,cli-restore*,cli-rmt*,cli-sasadmin*,cli-sasstat*,cli-sata*
priv set diag; useradmin role add rol_cli_sys_str_6 -c "System Admin CLI Storage Role Set 6" -a cli-scsi*,cli-sesdiag*,cli-shelfchk*,cli-sis*,cli-
snaplock*,cli-snapmirror*,cli-snapvault*,cli-snap*,cli-spares_zero*,cli-storage*,cli-stty*,cli-sysconfig*
priv set diag; useradmin role add rol_cli_sys_str_7 -c "System Admin CLI Storage Role Set 7" -a cli-acpadmin*,cli-fcmon*,cli-timezone*,cli-
vfiler*,cli-vol*,cli-vscan*,cli-wcc*,cli-wrfile*,cli-ypcat*,cli-ypgroup*,cli-ypmatch*,cli-ypwhich*
priv set diag; useradmin role add rol_cli_sys_sec -c "System Admin CLI Security Role" -a cli-ipsec*,cli-key_manager*,cli-keymgr*,cli-logout*,cli-
ls*,cli-options*,cli-partner*,cli-pktt*,cli-secureadmin*,cli-vscan*
priv set diag; useradmin role add rol_cli_sys_perf_1 -c "System Admin CLI Performance Role 1" -a cli-availtime*,cli-df*,cli-disk_stat*,cli-
environment*,cli-envirom*,cli-exit*,cli-fcstat*,cli-ftp*,cli-ifinfo*,cli-ifstat*
priv set diag; useradmin role add rol_cli_sys_perf_2 -c "System Admin CLI Performance Role 2" -a cli-mem_scrub_stats*,cli-mem_stats*,cli-
netstat*,cli-nfsstat*,cli-nis*,cli-partner*,cli-perf*,cli-pktt*,cli-ps*,cli-statit*,cli-stats*,cli-sysstat*
priv set diag; useradmin role add rol_cli_sys_net_1 -c "System Admin CLI Network Role 1" -a cli-arp*,cli-dns*,cli-exit*,cli-ifconfig*,cli-
ifinfo*,cli-ifstat*,cli-ipspace*,cli-nbtstat*,cli-ndmcopy*,cli-ndp*,cli-netdiag*,cli-netstat*,cli-orouted*
priv set diag; useradmin role add rol_cli_sys_net_2 -c "System Admin CLI Network Role 2" -a cli-partner*,cli-ping6*,cli-ping*,cli-pktt*,cli-
rlm*,cli-routed*,cli-route*,cli-sysconfig*,cli-traceroute6*,cli-traceroute*,cli-vif*,cli-vlan*
priv set diag; useradmin role add rol_cli_sys_net_3 -c "System Admin CLI Network Role 3" -a cli-cdpd*,cli-dcb*,cli-ifconfig_priv*,cli-ifgrp*,cli-
netmpstat*,cli-radius*,cli-rlm*,cli-route_priv*,cli-sp*
priv set diag; useradmin role add rol_cli_sys_lic -c "System Admin CLI License Role" -a cli-license*,cli-registry*
priv set diag; useradmin role add rol_cli_sys_bak -c "System Admin CLI Backup Role" -a cli-df*,cli-dump*,cli-ping6*,cli-ping*,cli-smtape*,cli-
sysconfig*
priv set diag; useradmin role add rol_cli_all -c "All Admin CLI Common Commands Role" -a cli-?,cli-echo*,cli-help*,cli-logger*,cli-man*,cli-
passwd*,cli-priv*,cli-rdfile*,cli-result*,cli-uptime*,cli-version*

```

```

priv set diag; useradmin role add rol_api_sys_admin_1 -c "System Admin API Role 1" -a api-aggr*,api-cf*,api-cg*,api-cifs*,api-clock*,api-clone*,api-dfm*,api-disk*,api-ems*,api-fcport*,api-fcp*,api-fc*,api-file*,api-fpolicy*,api-ic*,api-igroup*
priv set diag; useradmin role add rol_api_sys_admin_2 -c "System Admin API Role 2" -a api-iscsi*,api-license*,api-lock*,api-lun*,api-nameservice*,api-net*,api-nfs*,api-options*,api-perf*,api-portset*,api-priority*,api-qtree*,api-quota*,api-reallocate*
priv set diag; useradmin role add rol_api_sys_admin_3 -c "System Admin API Role 3" -a api-rsh*,api-sis*,api-snaplock*,api-snapmirror*,api-snapshot*,api-snapvault*,api-snp*,api-software*,api-storage*,api-system*,api-vfiler*,api-volume*,api-wafl*
priv set diag; useradmin role add rol_api_sys_admin_4 -c "System Admin API Role 4" -a api-ipospace*
priv set diag; useradmin role add rol_sec_sys_admin -c "System Admin Security Role" -a security-priv-advanced,security-priv-diagnostic
priv set diag; useradmin role add rol_login_sys_file_cifs -c "System Admin CIFS File Login Role" -a login-http-admin
priv set diag; useradmin role add rol_api_sys_file_cifs -c "System Admin CIFS File API Role" -a api-cifs*,api-nameservice*,api-qtree*,api-quota*
priv set diag; useradmin role add rol_login_sys_file -c "System Admin File Login Role" -a login-console,login-ssh,login-http-admin,login-ndmp
priv set diag; useradmin role add rol_api_sys_file_1 -c "System Admin File API Role 1" -a api-aggr-get-filer-info,api-aggr-list-info,api-aggr-options-list-info,api-cifs*,api-file*,api-fpolicy*,api-iscsi-service-status,api-lock*,api-nameservice*
priv set diag; useradmin role add rol_api_sys_file_2 -c "System Admin File API Role 2" -a api-nfs*,api-options-get,api-perf*,api-portset*,api-qtree*,api-quota*,api-reallocate*,api-system-get-info,api-volume-get-filer-info
priv set diag; useradmin role add rol_api_sys_file_3 -c "System Admin File API Role 3" -a api-volume-get-language,api-volume-get-root-name,api-volume-list-info,api-volume-options-list-info,api-volume-size
priv set diag; useradmin role add rol_sec_sys_file -c "System Admin File Security Role" -a security-priv-advanced,security-priv-diagnostic
priv set diag; useradmin role add rol_login_sys_file_nfs -c "System Admin NFS File Login Role" -a login-ssh,login-http-admin
priv set diag; useradmin role add rol_api_sys_file_nfs -c "System Admin NFS File API Role" -a api-options-get,api-nameservice*,api-nfs*,api-qtree*,api-quota*,api-volume-get-root-name
priv set diag; useradmin role add rol_sec_sys_file_nfs -c "System Admin NFS File Security Role" -a security-priv-advanced,security-priv-diagnostic
priv set diag; useradmin role add rol_login_sys_str -c "System Admin Storage Login Role" -a login-console,login-ssh,login-http-admin
priv set diag; useradmin role add rol_api_sys_str_1 -c "System Admin Storage API Role 1" -a api-aggr*,api-cf*,api-cg*,api-cifs*,api-disk*,api-ems*,api-fcport*,api-fcp*,api-fc*,api-file*,api-iscsi-service-status,api-lun*,api-options*,api-perf*
priv set diag; useradmin role add rol_api_sys_str_2 -c "System Admin Storage API Role 2" -a api-priority*,api-qtree*,api-quota*,api-reallocate*,api-sis*,api-snapmirror*,api-snapshot*,api-snapvault*,api-storage*,api-system*,api-vfiler*,api-volume*
priv set diag; useradmin role add rol_api_sys_str_3 -c "System Admin Storage API Role 3" -a api-wafl*
priv set diag; useradmin role add rol_sec_sys_str -c "System Admin Storage Security Role" -a security-priv-advanced,security-priv-diagnostic
priv set diag; useradmin role add rol_login_sys_sec -c "System Admin Security Login Role" -a login-ssh,login-http-admin
priv set diag; useradmin role add rol_api_sys_sec -c "System Admin Security API Role" -a api-options*,api-rsh*
priv set diag; useradmin role add rol_sec_sys_sec -c "System Admin Security Security Role" -a security-priv-advanced
priv set diag; useradmin role add rol_login_sys_perf -c "System Admin Performance Login Role" -a login-console,login-ssh,login-http-admin
priv set diag; useradmin role add rol_api_sys_perf -c "System Admin Performance API Role" -a api-perf*,api-system*
priv set diag; useradmin role add rol_sec_sys_perf -c "System Admin Performance Security Role" -a security-priv-advanced
priv set diag; useradmin role add rol_login_sys_net -c "System Admin Network Login Role" -a login-console,login-ssh,login-http-admin
priv set diag; useradmin role add rol_api_sys_net -c "System Admin Network API Role" -a api-file-delete-file,api-file-write-file,api-ipospace*,api-iscsi-service-status,api-net*,api-options-get,api-options-set,api-volume-get-root-name
priv set diag; useradmin role add rol_sec_sys_net -c "System Admin Network Security Role" -a security-priv-advanced
priv set diag; useradmin role add rol_login_sys_lic -c "System Admin License Login Role" -a login-ssh,login-http-admin
priv set diag; useradmin role add rol_api_sys_lic -c "System Admin License API Role" -a api-iscsi-service-status,api-license*
priv set diag; useradmin role add rol_login_sys_bak -c "System Admin Backup Login Role" -a login-ssh,login-http-admin,login-ndmp
priv set diag; useradmin role add rol_login_acc_admin -c "Account Admin Login Role" -a login-ssh
priv set diag; useradmin role add rol_cli_acc_admin -c "Account Admin CLI Role" -a cli-useradmin*
priv set diag; useradmin role add rol_api_acc_admin -c "Account Admin API Role" -a api-useradmin*
priv set diag; useradmin role add rol_sec_acc_admin -c "Account Admin Security Role" -a security-priv-advanced,security-priv-diagnostic,security-complete-user-control,security-passwd-change-others
priv set diag; useradmin role add rol_login_audit -c "Audit Login Role" -a login-ssh,login-http-admin
priv set diag; useradmin role add rol_api_audit -c "Audit API Role" -a api-snmp-get*

```

```

logger "End RBAC role creation"
logger "Start RBAC group creation and built-in group modification"

priv set diag; useradmin group modify "Power Users" -r rol_login_sys_file_cifs,rol_api_sys_file_cifs
priv set diag; useradmin group add grp_sys_admins_1 -c "System Admin Group 1" -r
rol_login_sys_admin,rol_cli_sys_1,rol_cli_sys_2,rol_cli_sys_3,rol_cli_sys_4,rol_cli_sys_5,rol_cli_sys_6,rol_cli_sys_7,rol_cli_sys_8,rol_cli_sys_9
priv set diag; useradmin group add grp_sys_admins_2 -c "System Admin Group 2" -r rol_cli_sys_10,rol_cli_sys_11,rol_cli_sys_12,rol_cli_sys_13
priv set diag; useradmin group add grp_sys_admins_3 -c "System Admin Group 3" -r
rol_cli_sys_file_1,rol_cli_sys_file_2,rol_cli_sys_file_3,rol_cli_sys_file_nfs_1,rol_cli_sys_file_nfs_2,rol_cli_sys_str_1,rol_cli_sys_str_2,rol_cli_
sys_str_3
priv set diag; useradmin group add grp_sys_admins_4 -c "System Admin Group 4" -r
rol_cli_sys_str_4,rol_cli_sys_str_5,rol_cli_sys_str_6,rol_cli_sys_str_7,rol_cli_sys_sec,rol_cli_sys_perf_1,rol_cli_sys_perf_2,rol_cli_sys_net_1,rol_
_cli_sys_net_2
priv set diag; useradmin group add grp_sys_admins_5 -c "System Admin Group 5" -r
rol_cli_sys_lic,rol_cli_sys_bak,rol_cli_all,rol_api_sys_admin_1,rol_api_sys_admin_2,rol_api_sys_admin_3,rol_api_sys_admin_4,rol_sec_sys_admin,rol_c
li_sys_net_3
priv set diag; useradmin group add grp_sys_admins_file_1 -c "File System Admin Group 1" -r
rol_login_sys_file,rol_cli_sys_file_1,rol_cli_sys_file_2,rol_cli_sys_file_3,rol_cli_sys_file_nfs_1,rol_cli_sys_file_nfs_2,rol_cli_all
priv set diag; useradmin group add grp_sys_admins_file_2 -c "File System Admin Group 2" -r
rol_api_sys_file_1,rol_api_sys_file_2,rol_api_sys_file_3,rol_sec_sys_file
priv set diag; useradmin group add grp_sys_admins_file_nfs -c "NFS File System Admin Group" -r
rol_login_sys_file_nfs,rol_cli_sys_file_nfs_1,rol_cli_sys_file_nfs_2,rol_cli_all,rol_api_sys_file_nfs,rol_sec_sys_file_nfs
priv set diag; useradmin group add grp_sys_admins_str_1 -c "Storage System Admin Group 1" -r
rol_login_sys_str,rol_cli_sys_str_1,rol_cli_sys_str_2,rol_cli_sys_str_3,rol_cli_sys_str_4,rol_cli_sys_str_5,rol_cli_sys_str_6,rol_cli_sys_str_7
priv set diag; useradmin group add grp_sys_admins_str_2 -c "Storage System Admin Group 2" -r
rol_cli_all,rol_api_sys_str_1,rol_api_sys_str_2,rol_api_sys_str_3,rol_sec_sys_str
priv set diag; useradmin group add grp_sys_admins_sec -c "Security System Admin Group" -r
rol_login_sys_sec,rol_cli_sys_sec,rol_cli_all,rol_api_sys_sec,rol_sec_sys_sec
priv set diag; useradmin group add grp_sys_admins_perf -c "Performance System Admin Group" -r
rol_login_sys_perf,rol_cli_sys_perf_1,rol_cli_sys_perf_2,rol_cli_all,rol_api_sys_perf,rol_sec_sys_perf
priv set diag; useradmin group add grp_sys_admins_net -c "Network System Admin Group" -r
rol_login_sys_net,rol_cli_sys_net_1,rol_cli_sys_net_2,rol_cli_sys_net_3,rol_cli_all,rol_api_sys_net,rol_sec_sys_net
priv set diag; useradmin group add grp_sys_admins_lic -c "License System Admin Group" -r
rol_login_sys_lic,rol_cli_sys_lic,rol_cli_all,rol_api_sys_lic
priv set diag; useradmin group add grp_sys_admins_bak -c "Backup System Admin Group" -r rol_login_sys_bak,rol_cli_sys_bak,rol_cli_all
priv set diag; useradmin group add grp_acc_admins -c "Account Admin Group" -r
rol_login_acc_admin,rol_cli_acc_admin,rol_cli_all,rol_api_acc_admin,rol_sec_acc_admin
priv set diag; useradmin group add grp_sys_audit -c "System Audit Group" -r rol_login_audit,rol_cli_all,rol_api_audit

logger "End RBAC group creation and built-in group modification"
logger "End CATE Storage and Backup Engineering ONTAP 8.1.4 7-mode RBAC Script v.1.0 for new system installs"

priv set admin

```

14.7.5 RBAC Upgrade Script

The following scripts have been developed to automate the upgrade of RBAC policies on NetApp NAS Filers when upgrading the Data ONTAP version.

- Download the script files (txt format) from the [CATE Storage and Backup Engineering website](#).
- Take the text file and save it to the /etc/ directory on the filer to which you wish to deploy RBAC
- Run the script file using the source command 'source /etc/<file_name>' (replace <file_name> with the name of the actual file)
- Check the system message log to confirm all groups and roles have been created successfully

14.7.5.1 RBAC for ONTAP 8.1.4 upgrades from ONTAP 8.0.2 on a 64-bit hardware platform

This script has been created to configure RBAC on NetApp Filers running 8.1.4 which are upgraded from ONTAP 8.0.2.

As ONTAP 8 is a 64-bit operating system, only 64-bit hardware is supported therefore there is no longer a requirement to provide 32-bit and 64-bit versions of the CATE RBAC scripts.

The script has been listed here for reference purposes only.

To avoid line wrapping issues do not copy the script from this document, download the files for the [CATE Storage and Backup Engineering website](#).

```
#
# Revision history
#   v1.0 - Initial release
#

logger "Start CATE Storage and Backup Engineering ONTAP 8.1.4 RBAC Script v.1.0 for 64-bit platform upgrade from ONTAP 8.0.2"
logger "Start RBAC role modification"

priv set diag; useradmin role modify rol_cli_sys_1 -c "System Admin CLI System Role Set 1" -a cli-anvl*,cli-bmc*,cli-com*,cli-coredump*,cli-
dbg*,cli-download*,cli-dumpblock*,cli-dumpstack*,cli-ems*,cli-export_stats*,cli-filersio*,cli-fru_led*
priv set diag; useradmin role modify rol_cli_sys_2 -c "System Admin CLI System Role Set 2" -a cli-gdb*,cli-getXXbyYY*,cli-halt*,cli-hammer*,cli-
hostname*,cli-httpstat*,cli-icbulk*,cli-ic*,cli-igroup*,cli-inodepath*,cli-iscsi*,cli-iswt*
priv set diag; useradmin role modify rol_cli_sys_3 -c "System Admin CLI System Role Set 3" -a cli-kma_stats*,cli-kt*,cli-labelmaint*,cli-
label*,cli-log_fio*,cli-log*,cli-mbstat*,cli-memerr*,cli-mt*,cli-ndmpd*,cli-nm*,cli-ontapi*
priv set diag; useradmin role modify rol_cli_sys_4 -c "System Admin CLI System Role Set 4" -a cli-panic*,cli-portset*,cli-printflag*,cli-prof*,cli-
raid_config*,cli-rastrace*,cli-rdate*,cli-reboot*,cli-registry*,cli-restore_backup*,cli-result*
priv set diag; useradmin role modify rol_cli_sys_7 -c "System Admin CLI System Role Set 7" -a cli-source*,cli-stack*,cli-stty*,cli-sync*,cli-
syslog*,cli-tape_qual*,cli-test_lcd*,cli-timezone*,cli-time*,cli-toe*,cli-treecompare*,cli-treemigrate*
priv set diag; useradmin role modify rol_cli_sys_8 -c "System Admin CLI System Role Set 8" -a cli-ttcp*,cli-ups*,cli-uptime*,cli-vtic*,cli-
wafl_steal_stats*,cli-wafl_susp*,cli-wafl*,cli-vm_stat*,cli-waffinity_stats*,cli-xttcp*
priv set diag; useradmin role modify rol_cli_sys_9 -c "System Admin CLI System Role Set 9" -a cli-acorn*,cli-autosupport*,cli-bringhome*,cli-
cxgbtool*,cli-ifswitch*,cli-iomem*,cli-mailbox*,cli-mviatest*,cli-nv8*
priv set diag; useradmin role modify rol_cli_sys_11 -c "System Admin CLI System Role Set 11" -a cli-ct_tpi_par*,cli-ct_tpi_read*,cli-
ct_tpi_write*,cli-du*,cli-rpc_stats*,cli-schedstat*,cli-setral*,cli-sldiag*,cli-smf*,cli-stsb*,cli-sum32*
priv set diag; useradmin role modify rol_cli_sys_13 -c "System Admin CLI System Role Set 13" -a cli-bootargs*,cli-ct_dump_t3*,cli-ct_reset_t3*,cli-
system*,cli-vdom*,cli-vol_db*,cli-vstorage*,cli-wafl_cmd_restrictions*,cli-wafltop*
priv set diag; useradmin role modify rol_cli_sys_file_1 -c "System Admin CLI File Role Set 1" -a cli-availtime*,cli-charmap*,cli-cifs*,cli-df*,cli-
exit*,cli-export_tbl_dump*,cli-exportfs*,cli-file*,cli-fpolicy*,cli-fsecurity*,cli-ftp*
```

```
priv set diag; useradmin role modify rol_cli_sys_str_7 -c "System Admin CLI Storage Role Set 7" -a cli-acpadmin*,cli-fcmon*,cli-timezone*,cli-
vfiler*,cli-vol*,cli-vscan*,cli-wcc*,cli-wrfile*,cli-ypcat*,cli-ypgroup*,cli-ypmatch*,cli-ypwhich*
priv set diag; useradmin role modify rol_cli_sys_perf_1 -c "System Admin CLI Performance Role 1" -a cli-availtime*,cli-df*,cli-disk_stat*,cli-
environment*,cli-environ*,cli-exit*,cli-fcstat*,cli-ftp*,cli-ifinfo*,cli-ifstat*
priv set diag; useradmin role modify rol_cli_sys_net_3 -c "System Admin CLI Network Role 3" -a cli-cdpd*,cli-dcb*,cli-ifconfig_priv*,cli-
ifgrp*,cli-netmpstat*,cli-radius*,cli-rlm*,cli-route_priv*,cli-sp*

logger "End RBAC role modification"
logger "End CATE Storage and Backup Engineering ONTAP 8.1.4 RBAC Script v.1.0 for 64-bit platform upgrade from ONTAP 8.0.2"

priv set admin
```


14.7.5.2 RBAC for ONTAP 8.1.4 upgrades from ONTAP 8.0.1 on a 64-bit hardware platform

This script has been created to configure RBAC on NetApp Filers running 8.1.4 which are upgraded from ONTAP 8.0.1.

As ONTAP 8 is a 64-bit operating system, only 64-bit hardware is supported therefore there is no longer a requirement to provide 32-bit and 64-bit versions of the CATE RBAC scripts.

The script has been listed here for reference purposes only.

To avoid line wrapping issues do not copy the script from this document, download the files for the [CATE Storage and Backup Engineering website](#).

```
#
# Revision history
#   v1.0 - Initial release
#

logger "Start CATE Storage and Backup Engineering ONTAP 8.1.4 RBAC Script v.1.0 for 64-bit platform upgrade from ONTAP 8.0.2"
logger "Start RBAC role modification"

priv set diag; useradmin role modify rol_cli_sys_1 -c "System Admin CLI System Role Set 1" -a cli-anvl*,cli-bmc*,cli-com*,cli-coredump*,cli-
dbg*,cli-download*,cli-dumpblock*,cli-dumpstack*,cli-ems*,cli-export_stats*,cli-filersio*,cli-fru_led*
priv set diag; useradmin role modify rol_cli_sys_2 -c "System Admin CLI System Role Set 2" -a cli-gdb*,cli-getXXbyYY*,cli-halt*,cli-hammer*,cli-
hostname*,cli-httpstat*,cli-icbulk*,cli-ic*,cli-igroup*,cli-inodepath*,cli-iscsi*,cli-iswt*
priv set diag; useradmin role modify rol_cli_sys_3 -c "System Admin CLI System Role Set 3" -a cli-kma_stats*,cli-kt*,cli-labelmaint*,cli-
label*,cli-log_fio*,cli-log*,cli-mbstat*,cli-memerr*,cli-mt*,cli-ndmpd*,cli-nm*,cli-ontapi*
priv set diag; useradmin role modify rol_cli_sys_4 -c "System Admin CLI System Role Set 4" -a cli-panic*,cli-portset*,cli-printflag*,cli-prof*,cli-
raid_config*,cli-rastrace*,cli-rdate*,cli-reboot*,cli-registry*,cli-restore_backup*,cli-result*
priv set diag; useradmin role modify rol_cli_sys_7 -c "System Admin CLI System Role Set 7" -a cli-source*,cli-stack*,cli-stty*,cli-sync*,cli-
syslog*,cli-tape_qual*,cli-test_lcd*,cli-timezone*,cli-time*,cli-toe*,cli-treecompare*,cli-treemigrate*
priv set diag; useradmin role modify rol_cli_sys_8 -c "System Admin CLI System Role Set 8" -a cli-ttcp*,cli-ups*,cli-uptime*,cli-vtic*,cli-
waf_l_steal_stats*,cli-waf_l_susp*,cli-waf_l*,cli-vm_stat*,cli-waffinity_stats*,cli-xttcp*
priv set diag; useradmin role modify rol_cli_sys_9 -c "System Admin CLI System Role Set 9" -a cli-acorn*,cli-autosupport*,cli-bringhome*,cli-
cxgbtool*,cli-ifswitch*,cli-iomem*,cli-mailbox*,cli-mviatest*,cli-nv8*
priv set diag; useradmin role modify rol_cli_sys_11 -c "System Admin CLI System Role Set 11" -a cli-ct_tpi_par*,cli-ct_tpi_read*,cli-
ct_tpi_write*,cli-du*,cli-rpc_stats*,cli-schedstat*,cli-setral*,cli-sldiag*,cli-smf*,cli-stsb*,cli-sum32*
priv set diag; useradmin role modify rol_cli_sys_13 -c "System Admin CLI System Role Set 13" -a cli-bootargs*,cli-ct_dump_t3*,cli-ct_reset_t3*,cli-
system*,cli-vdom*,cli-vol_db*,cli-vstorage*,cli-waf_l_cmd_restrictions*,cli-waf_l_top*
priv set diag; useradmin role modify rol_cli_sys_file_1 -c "System Admin CLI File Role Set 1" -a cli-availtime*,cli-charmap*,cli-cifs*,cli-df*,cli-
exit*,cli-export_tbl_dump*,cli-exportfs*,cli-file*,cli-fpolicy*,cli-fsecurity*,cli-ftp*
priv set diag; useradmin role modify rol_cli_sys_str_7 -c "System Admin CLI Storage Role Set 7" -a cli-acpadmin*,cli-fcmon*,cli-timezone*,cli-
vfiler*,cli-vol*,cli-vscan*,cli-wcc*,cli-wrfile*,cli-ypcat*,cli-ypgroup*,cli-ypmatch*,cli-ypwhich*
priv set diag; useradmin role modify rol_cli_sys_perf_1 -c "System Admin CLI Performance Role Set 1" -a cli-availtime*,cli-df*,cli-disk_stat*,cli-
environment*,cli-envIRON*,cli-exit*,cli-fcstat*,cli-ftp*,cli-ifinfo*,cli-ifstat*
priv set diag; useradmin role modify rol_cli_sys_net_3 -c "System Admin CLI Network Role Set 3" -a cli-cdpd*,cli-dcb*,cli-ifconfig_priv*,cli-
ifgrp*,cli-netmpstat*,cli-radius*,cli-rlm*,cli-route_priv*,cli-sp*

logger "End RBAC role modification"
logger "End CATE Storage and Backup Engineering ONTAP 8.1.4 RBAC Script v.1.0 for 64-bit platform upgrade from ONTAP 8.0.2"

priv set admin
```


14.7.5.3 RBAC for ONTAP 8.1.4 upgrades from ONTAP 7.3.6 on a 64-bit hardware platform

This script has been created to configure RBAC on NetApp Filers running 8.1.4 which are upgraded from ONTAP 7.3.6.

As ONTAP 8 is a 64-bit operating system, only 64-bit hardware is supported therefore there is no longer a requirement to provide 32-bit and 64-bit versions of the CATE RBAC scripts.

The script has been listed here for reference purposes only.

To avoid line wrapping issues do not copy the script from this document, download the files for the [CATE Storage and Backup Engineering website](#).

```
#
# Revision history
#   v1.0 - Initial release
#

logger "Start CATE Storage and Backup Engineering ONTAP 8.1.4 RBAC Script v.1.0 for 64-bit platform upgrade from ONTAP 7.3.6"
logger "Start RBAC role modification"

priv set diag; useradmin role modify rol_cli_sys_1 -c "System Admin CLI System Role Set 1" -a cli-anvl*,cli-bmc*,cli-com*,cli-coredump*,cli-
dbg*,cli-download*,cli-dumpblock*,cli-dumpstack*,cli-ems*,cli-export_stats*,cli-filersio*,cli-fru_led*
priv set diag; useradmin role modify rol_cli_sys_2 -c "System Admin CLI System Role Set 2" -a cli-gdb*,cli-getXXbyYY*,cli-halt*,cli-hammer*,cli-
hostname*,cli-httpstat*,cli-icbulk*,cli-ic*,cli-igroup*,cli-inodepath*,cli-iscsi*,cli-iswt*
priv set diag; useradmin role modify rol_cli_sys_3 -c "System Admin CLI System Role Set 3" -a cli-kma_stats*,cli-kt*,cli-labelmaint*,cli-
label*,cli-log_fio*,cli-log*,cli-mbstat*,cli-memerr*,cli-mt*,cli-ndmpd*,cli-nm*,cli-ontapi*
priv set diag; useradmin role modify rol_cli_sys_4 -c "System Admin CLI System Role Set 4" -a cli-panic*,cli-portset*,cli-printflag*,cli-prof*,cli-
raid_config*,cli-rastrace*,cli-rdate*,cli-reboot*,cli-registry*,cli-restore_backup*,cli-result*
priv set diag; useradmin role modify rol_cli_sys_6 -c "System Admin CLI System Role Set 6" -a cli-show_faults*,cli-showfh4*,cli-showfh*,cli-
sh*,cli-signal*,cli-slist*,cli-sm_mon_old*,cli-sm_mon*,cli-sm_not*,cli-smb_hist*,cli-snmp*,cli-software*
priv set diag; useradmin role modify rol_cli_sys_7 -c "System Admin CLI System Role Set 7" -a cli-source*,cli-stack*,cli-stty*,cli-sync*,cli-
syslog*,cli-tape_qual*,cli-test_lcd*,cli-timezone*,cli-time*,cli-toe*,cli-treecompare*,cli-treemigrate*
priv set diag; useradmin role modify rol_cli_sys_8 -c "System Admin CLI System Role Set 8" -a cli-ttcp*,cli-ups*,cli-uptime*,cli-vtic*,cli-
wafl_steal_stats*,cli-wafl_susp*,cli-wafl*,cli-vm_stat*,cli-waffinity_stats*,cli-xttcp*
priv set diag; useradmin role modify rol_cli_sys_9 -c "System Admin CLI System Role Set 9" -a cli-acorn*,cli-autosupport*,cli-bringhome*,cli-
cxgbtool*,cli-ifswitch*,cli-iomem*,cli-mailbox*,cli-mviatest*,cli-nv8*
priv set diag; useradmin role modify rol_cli_sys_10 -c "System Admin CLI System Role Set 10" -a cli-ct_phy_read_t3*,cli-ct_phy_read*,cli-
ct_phy_write_t3*,cli-ct_phy_write*,cli-ct_reg_read_t3*,cli-ct_reg_read*,cli-ct_reg_write_t3*,cli-ct_reg_write*
priv set diag; useradmin role add rol_cli_sys_11 -c "System Admin CLI System Role Set 11" -a cli-ct_tpi_par*,cli-ct_tpi_read*,cli-
ct_tpi_write*,cli-du*,cli-rpc_stats*,cli-schedstat*,cli-setral*,cli-sldiag*,cli-smf*,cli-stsb*,cli-sum32*
priv set diag; useradmin role add rol_cli_sys_12 -c "System Admin CLI System Role Set 12" -a cli-spinhi_stats*,cli-spinnp_replay_stats*,cli-
spinnp_replay*,cli-systemshell*,cli-tcp_client*,cli-tcp_server*,cli-udp_client*,cli-udp_server*,cli-vmsservices*
priv set diag; useradmin role add rol_cli_sys_13 -c "System Admin CLI System Role Set 13" -a cli-bootargs*,cli-ct_dump_t3*,cli-ct_reset_t3*,cli-
system*,cli-vdom*,cli-vol_db*,cli-vstorage*,cli-wafl_cmd_restrictions*,cli-wafltop*
priv set diag; useradmin role modify rol_cli_sys_file_1 -c "System Admin CLI File Role Set 1" -a cli-availtime*,cli-charmap*,cli-cifs*,cli-df*,cli-
exit*,cli-export_tbl_dump*,cli-exportfs*,cli-file*,cli-fpolicy*,cli-fsecurity*,cli-ftp*
priv set diag; useradmin role modify rol_cli_sys_file_3 -c "System Admin CLI File Role Set 3" -a cli-clone*,cli-reallocate*,cli-ypcat*,cli-
ypgroup*,cli-ypmatch*,cli-ypwhich*
priv set diag; useradmin role modify rol_cli_sys_str_7 -c "System Admin CLI Storage Role Set 7" -a cli-acpadmin*,cli-fcmon*,cli-timezone*,cli-
vfiler*,cli-vol*,cli-vscan*,cli-wcc*,cli-wrfile*,cli-ypcat*,cli-ypgroup*,cli-ypmatch*,cli-ypwhich*
priv set diag; useradmin role modify rol_cli_sys_sec -c "System Admin CLI Security Role" -a cli-ipsec*,cli-key_manager*,cli-keymgr*,cli-
logout*,cli-ls*,cli-options*,cli-partner*,cli-pktt*,cli-secureadmin*,cli-vscan*
```

```

priv set diag; useradmin role modify rol_cli_sys_perf_1 -c "System Admin CLI Performance Role 1" -a cli-availtime*,cli-df*,cli-disk_stat*,cli-
environment*,cli-envIRON*,cli-exit*,cli-fcstat*,cli-ftp*,cli-ifinfo*,cli-ifstat*
priv set diag; useradmin role modify rol_cli_sys_net_2 -c "System Admin CLI Network Role 2" -a cli-partner*,cli-ping6*,cli-ping*,cli-pktt*,cli-
rlm*,cli-routed*,cli-route*,cli-sysconfig*,cli-traceroute6*,cli-traceroute*,cli-vif*,cli-vlan*
priv set diag; useradmin role modify rol_cli_sys_net_3 -c "System Admin CLI Network Role 3" -a cli-cdpd*,cli-dcb*,cli-ifconfig_priv*,cli-
ifgrp*,cli-netmpstat*,cli-radius*,cli-rlm*,cli-route_priv*,cli-sp*
priv set diag; useradmin role modify rol_cli_sys_bak -c "System Admin CLI Backup Role" -a cli-df*,cli-dump*,cli-ping6*,cli-ping*,cli-smtape*,cli-
sysconfig*
priv set diag; useradmin group modify grp_sys_admins_1 -c "System Admin Group 1" -r
rol_login_sys_admin,rol_cli_sys_1,rol_cli_sys_2,rol_cli_sys_3,rol_cli_sys_4,rol_cli_sys_5,rol_cli_sys_6,rol_cli_sys_7,rol_cli_sys_8,rol_cli_sys_9
priv set diag; useradmin group modify grp_sys_admins_2 -c "System Admin Group 2" -r rol_cli_sys_10,rol_cli_sys_11,rol_cli_sys_12,rol_cli_sys_13
priv set diag; useradmin group modify grp_sys_admins_3 -c "System Admin Group 3" -r
rol_cli_sys_file_1,rol_cli_sys_file_2,rol_cli_sys_file_3,rol_cli_sys_file_nfs_1,rol_cli_sys_file_nfs_2,rol_cli_sys_str_1,rol_cli_sys_str_2,rol_cli_
sys_str_3
priv set diag; useradmin group modify grp_sys_admins_4 -c "System Admin Group 4" -r
rol_cli_sys_str_4,rol_cli_sys_str_5,rol_cli_sys_str_6,rol_cli_sys_str_7,rol_cli_sys_sec,rol_cli_sys_perf_1,rol_cli_sys_perf_2,rol_cli_sys_net_1,rol_
_cli_sys_net_2
priv set diag; useradmin group add grp_sys_admins_5 -c "System Admin Group 5" -r
rol_cli_sys_lic,rol_cli_sys_bak,rol_cli_all,rol_api_sys_admin_1,rol_api_sys_admin_2,rol_api_sys_admin_3,rol_api_sys_admin_4,rol_sec_sys_admin,rol_c
li_sys_net_3

logger "End RBAC role modification"
logger "End CATE Storage and Backup Engineering ONTAP 8.1.4 RBAC Script v.1.0 for 64-bit platform upgrade from ONTAP 7.3.6"

priv set admin

```

14.7.6 RBAC Downgrade / Revert Script

The following scripts have been developed to automate the reversion of RBAC policies on NetApp NAS Filers when downgrading the Data ONTAP version.

- Download the script files (txt format) from the [CATE Storage and Backup Engineering website](#).
- Take the text file and save it to the /etc/ directory on the filer to which you wish to deploy RBAC
- Run the script file using the source command 'source /etc/<file_name>' (replace <file_name> with the name of the actual file)
- Check the system message log to confirm all groups and roles have been created successfully

14.7.6.1 RBAC for ONTAP 8.0.2 revert from ONTAP 8.1.4 on a 64-bit hardware platform

This script has been created to configure RBAC on NetApp Filers which have recently been downgraded to 8.0.2 as a backout to an unsuccessful ONTAP 8.1.4 upgrade. This script is only required to be run if the ONTAP 8.1.4 RBAC policy has been applied to the storage system.

As ONTAP 8 is a 64-bit operating system, only 64-bit hardware is supported therefore there is no longer a requirement to provide 32-bit and 64-bit versions of the CATE RBAC scripts. The script has been listed here for reference purposes only.

To avoid line wrapping issues do not copy the script from this document, download the files for the [CATE Storage and Backup Engineering website](#).

```
#
# Revision history
#   v1.0 - Initial release
#
#   This script is only required to be run if the ONTAP 8.1.4 RBAC policy has been applied to the storage system.
#

logger "Start CATE Storage and Backup Engineering ONTAP 8.0.2 RBAC Script v.1.0 for 64-bit platform revert from ONTAP 8.1.4"
logger "Start RBAC role modification"

priv set diag; useradmin role modify rol_cli_sys_1 -c "System Admin CLI System Role Set 1" -a cli-anvl*,cli-bmc*,cli-com*,cli-dbg*,cli-
download*,cli-dumpblock*,cli-dumpstack*,cli-ems*,cli-export_stats*,cli-filersio*
priv set diag; useradmin role modify rol_cli_sys_2 -c "System Admin CLI System Role Set 2" -a cli-gdb*,cli-getXXbyYY*,cli-halt*,cli-hammer*,cli-
hostname*,cli-httpstat*,cli-icbulk*,cli-ic*,cli-igroup*,cli-inodepath*,cli-iscsi*,cli-isnap*,cli-iswt*
priv set diag; useradmin role modify rol_cli_sys_3 -c "System Admin CLI System Role Set 3" -a cli-java*,cli-kma_stats*,cli-kt*,cli-labelmaint*,cli-
label*,cli-log_fio*,cli-log*,cli-mbstat*,cli-memerr*,cli-mt*,cli-ndmpd*,cli-nm*,cli-ontapi*
priv set diag; useradmin role modify rol_cli_sys_4 -c "System Admin CLI System Role Set 4" -a cli-panic*,cli-portset*,cli-printflag*,cli-prof*,cli-
raid_config*,cli-rastrace*,cli-rdate*,cli-read_serno*,cli-reboot*,cli-registry*,cli-result*
priv set diag; useradmin role modify rol_cli_sys_7 -c "System Admin CLI System Role Set 7" -a cli-source*,cli-stack*,cli-stty*,cli-sync*,cli-
syslog*,cli-tape_qual*,cli-test_lcd*,cli-timezone*,cli-time*,cli-toe*,cli-treecompare*
priv set diag; useradmin role modify rol_cli_sys_8 -c "System Admin CLI System Role Set 8" -a cli-ttcp*,cli-ups*,cli-uptime*,cli-vtic*,cli-
wafl_steal_stats*,cli-wafl_susp*,cli-wafl*,cli-write_serno*,cli-vm_stat*,cli-waffinity_stats*,cli-xttcp*
priv set diag; useradmin role modify rol_cli_sys_9 -c "System Admin CLI System Role Set 9" -a cli-acorn*,cli-autosupport*,cli-bringhome*,cli-
cxgbtool*,cli-ide_savecore*,cli-ifswitch*,cli-incpy_cmd*,cli-iomem*,cli-mailbox*,cli-mviatest*,cli-nv8*
priv set diag; useradmin role modify rol_cli_sys_11 -c "System Admin CLI System Role Set 11" -a cli-ct_tpi_par*,cli-ct_tpi_read*,cli-
ct_tpi_write*,cli-du*,cli-rc_loop_check*,cli-rc_loop*,cli-repl_aggr*,cli-rpc_stats*,cli-schedstat*,cli-setral*,cli-sldiag*
priv set diag; useradmin role modify rol_cli_sys_13 -c "System Admin CLI System Role Set 13" -a cli-bootargs*,cli-ct_dbg_fatal_err_t3*,cli-
vol_db*,cli-vstorage*,cli-wafl_cmd_restrictions*,cli-wafltop*
```

```

priv set diag; useradmin role modify rol_cli_sys_file_1 -c "System Admin CLI File Role Set 1" -a cli-availtime*,cli-charmap*,cli-cifs*,cli-df*,cli-
exit*,cli-export_tbl_dump*,cli-exportfs*,cli-filestats*,cli-file*,cli-fpolicy*,cli-fsecurity*,cli-ftp*
priv set diag; useradmin role modify rol_cli_sys_str_7 -c "System Admin CLI Storage Role Set 7" -a cli-acpadmin*,cli-timezone*,cli-vfiler*,cli-
vol*,cli-vscan*,cli-wcc*,cli-wrfile*,cli-ypcat*,cli-ypgroup*,cli-ypmatch*,cli-ypwhich*
priv set diag; useradmin role modify rol_cli_sys_perf_1 -c "System Admin CLI Performance Role 1" -a cli-availtime*,cli-df*,cli-disk_stat*,cli-
environment*,cli-envIRON*,cli-exit*,cli-fcstat*,cli-filestats*,cli-ftp*,cli-ifinfo*,cli-ifstat*
priv set diag; useradmin role modify rol_cli_sys_net_3 -c "System Admin CLI Network Role 3" -a cli-cdpd*,cli-dcb*,cli-ifgrp*,cli-radius*,cli-
rlm*,cli-sp*

logger "End RBAC role modification"
logger "End CATE Storage and Backup Engineering ONTAP 8.0.2 RBAC Script v.1.0 for 64-bit platform revert from ONTAP 8.1.4"

priv set admin

```

14.7.6.2 RBAC for ONTAP 8.0.1 revert from ONTAP 8.1.4 on a 64-bit hardware platform

This script has been created to configure RBAC on NetApp Filers which have recently been downgraded to 8.0.1 as a backout to an unsuccessful ONTAP 8.1.4 upgrade. This script is only required to be run if the ONTAP 8.1.4 RBAC policy has been applied to the storage system.

As ONTAP 8 is a 64-bit operating system, only 64-bit hardware is supported therefore there is no longer a requirement to provide 32-bit and 64-bit versions of the CATE RBAC scripts. The script has been listed here for reference purposes only.

To avoid line wrapping issues do not copy the script from this document, download the files for the [CATE Storage and Backup Engineering website](#).

```
#
# Revision history
#   v1.0 - Initial release
#
#   This script is only required to be run if the ONTAP 8.1.4 RBAC policy has been applied to the storage system.
#

logger "Start CATE Storage and Backup Engineering ONTAP 8.0.1 RBAC Script v.1.0 for 64-bit platform revert from ONTAP 8.1.4"
logger "Start RBAC role modification"

priv set diag; useradmin role modify rol_cli_sys_1 -c "System Admin CLI System Role Set 1" -a cli-anvl*,cli-bmc*,cli-com*,cli-dbg*,cli-
download*,cli-dumpblock*,cli-dumpstack*,cli-ems*,cli-export_stats*,cli-filersio*
priv set diag; useradmin role modify rol_cli_sys_2 -c "System Admin CLI System Role Set 2" -a cli-gdb*,cli-getXXbyYY*,cli-halt*,cli-hammer*,cli-
hostname*,cli-httpstat*,cli-icbulk*,cli-ic*,cli-igroup*,cli-inodepath*,cli-iscsi*,cli-isnap*,cli-iswt*
priv set diag; useradmin role modify rol_cli_sys_3 -c "System Admin CLI System Role Set 3" -a cli-java*,cli-kma_stats*,cli-kt*,cli-labelmaint*,cli-
label*,cli-log_fio*,cli-log*,cli-mbstat*,cli-memerr*,cli-mt*,cli-ndmpd*,cli-nm*,cli-ontapi*
priv set diag; useradmin role modify rol_cli_sys_4 -c "System Admin CLI System Role Set 4" -a cli-panic*,cli-portset*,cli-printflag*,cli-prof*,cli-
raid_config*,cli-rastrace*,cli-rdate*,cli-read_serno*,cli-reboot*,cli-registry*,cli-result*
priv set diag; useradmin role modify rol_cli_sys_7 -c "System Admin CLI System Role Set 7" -a cli-source*,cli-stack*,cli-stty*,cli-sync*,cli-
syslog*,cli-tape_qual*,cli-test_lcd*,cli-timezone*,cli-time*,cli-toe*,cli-treecompare*
priv set diag; useradmin role modify rol_cli_sys_8 -c "System Admin CLI System Role Set 8" -a cli-ttcp*,cli-ups*,cli-uptime*,cli-vtic*,cli-
wafl_steal_stats*,cli-wafl_susp*,cli-wafl*,cli-watchdog*,cli-write_serno*,cli-vm_stat*,cli-waffinity_stats*,cli-xttcp*
priv set diag; useradmin role modify rol_cli_sys_9 -c "System Admin CLI System Role Set 9" -a cli-acorn*,cli-autosupport*,cli-bringhome*,cli-
cxgbtool*,cli-ide_savecore*,cli-ifswitch*,cli-incpy_cmd*,cli-iomem*,cli-mailbox*,cli-mviatest*,cli-nv8*
priv set diag; useradmin role modify rol_cli_sys_11 -c "System Admin CLI System Role Set 11" -a cli-ct_tpi_par*,cli-ct_tpi_read*,cli-
ct_tpi_write*,cli-du*,cli-rc_loop_check*,cli-rc_loop*,cli-repl_aggr*,cli-rpc_stats*,cli-schedstat*,cli-setral*,cli-sldiag*
priv set diag; useradmin role modify rol_cli_sys_13 -c "System Admin CLI System Role Set 13" -a cli-vol_db*,cli-vstorage*,cli-
wafl_cmd_restrictions*,cli-wafltop*
priv set diag; useradmin role modify rol_cli_sys_file_1 -c "System Admin CLI File Role Set 1" -a cli-availtime*,cli-charmap*,cli-cifs*,cli-df*,cli-
exit*,cli-export_tbl_dump*,cli-exportfs*,cli-filestats*,cli-file*,cli-fpolicy*,cli-fsecurity*,cli-ftp*
priv set diag; useradmin role modify rol_cli_sys_str_7 -c "System Admin CLI Storage Role Set 7" -a cli-acpadmin*,cli-timezone*,cli-vfiler*,cli-
vol*,cli-vscan*,cli-wcc*,cli-wrfile*,cli-ypcat*,cli-ypgroup*,cli-ypmatch*,cli-ypwhich*
priv set diag; useradmin role modify rol_cli_sys_perf_1 -c "System Admin CLI Performance Role 1" -a cli-availtime*,cli-df*,cli-disk_stat*,cli-
environment*,cli-envIRON*,cli-exit*,cli-fcstat*,cli-filestats*,cli-ftp*,cli-ifinfo*,cli-ifstat*
priv set diag; useradmin role modify rol_cli_sys_net_3 -c "System Admin CLI Network Role 3" -a cli-cdpd*,cli-dcb*,cli-ifgrp*,cli-radius*,cli-
rlm*,cli-sp*

logger "End RBAC role modification"
logger "End CATE Storage and Backup Engineering ONTAP 8.0.1 RBAC Script v.1.0 for 64-bit platform revert from ONTAP 8.1.4"
priv set admin
```

14.7.6.3 RBAC for ONTAP 7.3.6 revert from ONTAP 8.1.4 on a 64-bit hardware platform

This script has been created to configure RBAC on NetApp Filers which have recently been downgraded to 7.3.6 as a backout to an unsuccessful ONTAP 8.1.4 upgrade. This script is only required to be run if the ONTAP 8.1.4 RBAC policy has been applied to the storage system.

As ONTAP 8 is a 64-bit operating system, only 64-bit hardware is supported therefore there is no longer a requirement to provide 32-bit and 64-bit versions of the CATE RBAC scripts. The script has been listed here for reference purposes only.

To avoid line wrapping issues do not copy the script from this document, download the files for the [CATE Storage and Backup Engineering website](#).

```
#
# Revision history
#   v1.0 - Initial release
#

logger "Start CATE Storage and Backup Engineering ONTAP 7.3.6 RBAC Script v.1.0 for 64-bit platform revert from ONTAP 8.1.4"
logger "Start RBAC role modification"

priv set diag; useradmin group delete grp_sys_admins_5
priv set diag; useradmin group modify grp_sys_admins_4 -c "System Admin Group 4" -r
rol_cli_sys_lic,rol_cli_sys_bak,rol_cli_all,rol_api_sys_admin_1,rol_api_sys_admin_2,rol_api_sys_admin_3,rol_api_sys_admin_4,rol_sec_sys_admin,rol_c
li_sys_net_3
priv set diag; useradmin group modify grp_sys_admins_3 -c "System Admin Group 3" -r
rol_cli_sys_str_4,rol_cli_sys_str_5,rol_cli_sys_str_6,rol_cli_sys_str_7,rol_cli_sys_sec,rol_cli_sys_perf_1,rol_cli_sys_perf_2,rol_cli_sys_net_1,rol
_cli_sys_net_2
priv set diag; useradmin group modify grp_sys_admins_2 -c "System Admin Group 2" -r
rol_cli_sys_file_1,rol_cli_sys_file_2,rol_cli_sys_file_3,rol_cli_sys_file_nfs_1,rol_cli_sys_file_nfs_2,rol_cli_sys_str_1,rol_cli_sys_str_2,rol_cli
sys_str_3
priv set diag; useradmin group modify grp_sys_admins_1 -c "System Admin Group 1" -r
rol_login_sys_admin,rol_cli_sys_1,rol_cli_sys_2,rol_cli_sys_3,rol_cli_sys_4,rol_cli_sys_5,rol_cli_sys_6,rol_cli_sys_7,rol_cli_sys_8,rol_cli_sys_9,r
ol_cli_sys_10
priv set diag; useradmin role delete rol_cli_sys_11
priv set diag; useradmin role delete rol_cli_sys_12
priv set diag; useradmin role delete rol_cli_sys_13

priv set diag; useradmin role modify rol_cli_sys_1 -c "System Admin CLI System Role Set 1" -a cli-anvl*,cli-bmc*,cli-cmos_read*,cli-
cmos_write*,cli-com*,cli-dbg*,cli-download*,cli-dumpblock*,cli-dumpstack*,cli-ems*,cli-export_stats*,cli-filersio*
priv set diag; useradmin role modify rol_cli_sys_2 -c "System Admin CLI System Role Set 2" -a cli-gdb*,cli-getXXbyYY*,cli-halt*,cli-hammer*,cli-
hostname*,cli-httpstat*,cli-icbulk*,cli-ic*,cli-igroup*,cli-inodepath*,cli-iscsi*,cli-isnap*,cli-iswt*
priv set diag; useradmin role modify rol_cli_sys_3 -c "System Admin CLI System Role Set 3" -a cli-java*,cli-kma_stats*,cli-kt*,cli-labelmaint*,cli-
label*,cli-log_fio*,cli-log*,cli-mbstat*,cli-memerr*,cli-mt*,cli-ndmpd*,cli-nm*,cli-ontapi*
priv set diag; useradmin role modify rol_cli_sys_4 -c "System Admin CLI System Role Set 4" -a cli-outb*,cli-panic*,cli-portset*,cli-printflag*,cli-
prof*,cli-raid_config*,cli-rastrace*,cli-rdate*,cli-read_serno*,cli-reboot*,cli-registry*,cli-result*
priv set diag; useradmin role modify rol_cli_sys_5 -c "System Admin CLI System Role Set 5" -a cli-revert_to*,cli-rm*,cli-rshkill*,cli-rshstat*,cli-
rsm*,cli-rtag*,cli-rtfo*,cli-savecore*,cli-sectrace*,cli-setflag*,cli-setup*,cli-show_bios_log*
priv set diag; useradmin role modify rol_cli_sys_6 -c "System Admin CLI System Role Set 6" -a cli-show_faults*,cli-showfh4*,cli-showfh*,cli-
sh*,cli-signal*,cli-slist*,cli-sm_mon_old*,cli-sm_mon*,cli-sm_not*,cli-smb_hist*,cli-snet*,cli-snmp*,cli-software*
priv set diag; useradmin role modify rol_cli_sys_7 -c "System Admin CLI System Role Set 7" -a cli-source*,cli-stack*,cli-stty*,cli-sync*,cli-
syslog*,cli-tape_qual*,cli-test_lcd*,cli-timed*,cli-timezone*,cli-time*,cli-toe*,cli-trace*,cli-treecompare*
```

```

priv set diag; useradmin role modify rol_cli_sys_8 -c "System Admin CLI System Role Set 8" -a cli-ttcp*,cli-ups*,cli-uptime*,cli-vtic*,cli-
waf*_steal_stats*,cli-waf*_susp*,cli-waf*_*,cli-watchdog*,cli-write_serno*,cli-vm_stat*,cli-waffinity_stats*,cli-xttcp*
priv set diag; useradmin role modify rol_cli_sys_9 -c "System Admin CLI System Role Set 9" -a cli-acorn*,cli-ct_dbg_fatal_err_t3*,cli-
ct_phy_read_t3*,cli-ct_phy_read*,cli-ct_phy_write*,cli-fru_led*,cli-ifswitch*,cli-iomem*,cli-sftp*,cli-sldiag*
priv set diag; useradmin role modify rol_cli_sys_10 -c "System Admin CLI System Role Set 10" -a cli-ct_phy_write_t3*,cli-ct_reg_read_t3*,cli-
ct_reg_read*,cli-ct_reg_write_t3*,cli-ct_reg_write*,cli-ct_tpi_par*,cli-ct_tpi_read*,cli-ct_tpi_write*,cli-cxgbtool*

priv set diag; useradmin role modify rol_cli_sys_file_1 -c "System Admin CLI File Role Set 1" -a cli-availtime*,cli-charmap*,cli-cifs*,cli-df*,cli-
exit*,cli-export_tbl_dump*,cli-exportfs*,cli-filestats*,cli-file*,cli-fpolicy*,cli-fsecurity*,cli-ftp*
priv set diag; useradmin role modify rol_cli_sys_file_3 -c "System Admin CLI File Role Set 3" -a cli-reallocate*,cli-yocat*,cli-ypgroup*,cli-
ypmatch*,cli-ypwhich*,cli-clone*
priv set diag; useradmin role modify rol_cli_sys_str_7 -c "System Admin CLI Storage Role Set 7" -a cli-timezone*,cli-vfiler*,cli-vol*,cli-
vscan*,cli-wcc*,cli-wrfile*,cli-yocat*,cli-ypgroup*,cli-ypmatch*,cli-ypwhich*,cli-acpadmin*
priv set diag; useradmin role modify rol_cli_sys_sec -c "System Admin CLI Security Role" -a cli-ipsec*,cli-keymgr*,cli-logout*,cli-ls*,cli-
options*,cli-partner*,cli-pktt*,cli-secureadmin*,cli-vscan*
priv set diag; useradmin role modify rol_cli_sys_perf_1 -c "System Admin CLI Performance Role 1" -a cli-availtime*,cli-df*,cli-disk_stat*,cli-
environment*,cli-environ*,cli-exit*,cli-fcstat*,cli-filestats*,cli-ftp*,cli-ifinfo*,cli-ifstat*
priv set diag; useradmin role modify rol_cli_sys_net_2 -c "System Admin CLI Network Role 2" -a cli-partner*,cli-ping6*,cli-ping*,cli-pktt*,cli-
rlm*,cli-rmc*,cli-routed*,cli-route*,cli-sysconfig*,cli-traceroute6*,cli-traceroute*,cli-vif*,cli-vlan*
priv set diag; useradmin role modify rol_cli_sys_net_3 -c "System Admin CLI Network Role 3" -a cli-cdpd*,cli-cna*,cli-sp*
priv set diag; useradmin role modify rol_cli_sys_bak -c "System Admin CLI Backup Role" -a cli-df*,cli-dump*,cli-ping6*,cli-ping*,cli-sysconfig*

logger "End RBAC role modification"
logger "End CATE Storage and Backup Engineering ONTAP 7.3.6 RBAC Script v.1.0 for 64-bit platform revert from ONTAP 8.1.4"

priv set admin

```

14.7.7 Application Infrastructure Integration RBAC Scripts

The following scripts have been developed to automate the creation of RBAC roles for integration into various infrastructure applications.

- Download the script files (txt format) from the [CATE Storage and Backup Engineering website](#).
- Take the text file and save it to the /etc/ directory on the filer to which you wish to deploy RBAC
- Run the script file using the source command 'source /etc/<file_name>' (replace <file_name> with the name of the actual file)
- Check the system message log to confirm all groups and roles have been created successfully

14.7.7.1 RBAC for EMC ViPR Storage Resource Manager (SRM)

This script has been created to configure RBAC on NetApp Filers which are running ONTAP 7.3.6P2, ONTAP 8.0.x, ONTAP 8.1P2 and ONTAP 8.1.4Px.. The script has been listed here for reference purposes only.

To avoid line wrapping issues do not copy the script from this document, download the files for the [CATE Storage and Backup Engineering website](#).

```
#
# Revision history
#   v1.0 - Initial release
#   v1.1 - Corrected design error in v1.0 role structure

logger "Start CATE Storage and Backup Engineering EMC SRM ONTAP RBAC Script v.1.1"
logger "Start RBAC role creation"

useradmin role add rol_cli_app_emc_srm -c "EMC Storage Resource Management CLI Role" -a cli-?*,cli-aggr*,cli-cifs*,cli-df,cli-disk*,cli-echo*,cli-
environment,cli-fcp*,cli-ifconfig,cli-lun*,cli-maxfiles*,cli-nfs*,cli-qtree*,cli-quota*,cli-snap*,cli-stats*,cli-storage*,cli-sysconfig,cli-
version,cli-vol*
useradmin role add rol_login_app_emc_srm -c "EMC Storage Resource Management Login Role" -a login-ssh

logger "End RBAC role creation"
logger "Start RBAC group creation"

useradmin group add grp_app_emc_srm -c "EMC Storage Resource Management Group" -r rol_cli_app_emc_srm,rol_login_app_emc_srm

logger "End RBAC group creation"
logger "End CATE Storage and Backup Engineering EMC SRM ONTAP RBAC Script v.1.1"
```

Once the above script is completed, add the domain based functional ID used by EMC ViPR SRM to the grp_app_emc_srm group using the useradmin domainuser add command.

14.8 Appendix H Snapvault – further reading

14.8.1 Detailed Snapvault Snapshot Schedule Configuration

Normally, administrators would configure hourly, nightly and weekly snapshots schedules. Below is an example of how to configure a complete set, but you can select your own schedule. There is no requirement for a complete set and there is no interdependence among the schedules.

The most important aspect of the snapvault snap sched command is the '-x' option. This is the option that tells the volume whether it should merely create a snapshot, or update the data in all its qtrees before taking the snapshot. Only the snapvault secondary allows the use of the '-x' option, but not all secondary schedules require the use of the '-x' option. If one is updating the vault every hour, then the nightly snapshot merely needs to create a snapshot, not transfer the new data. In this manner, the secondary can create weekly snapshots, without retaining a weekly snapshot on the primary system. After the Saturday night "sv_nightly" transfer, the sv_weekly schedule could create a snapshot of that data.

14.8.1.1 Example 1: Hourly Snapshot Schedule

An hourly schedule is finest temporal granularity of snapshot schedule supported by SnapVault. Customers often want to retain a different number of snapshots on each system. The primary filers may need several hourly snapshots, but the secondary filer may want only days worth of hourly snapshots. Hourly schedules are primarily used just to keep the secondary as up to date as possible and increase the granularity of single file recovery. In the example below, the primary filers keep four hourly snapshots created at 8:00am, noon, 4:00pm, and 8:00pm.

```
toaster> snapvault snap sched proj sv_hourly 5@8,12,16,20
fridge> snapvault snap sched proj sv_hourly 5@8,12,16,20
```

The snapshots will be created in volume proj. The prefix of the created snapshot names will be sv_hourly and will result in successive snapshots named sv_hourly.0, sv_hourly.1, etc. The snapshot name prefixes "hourly", "nightly" and "weekly" are reserved and cannot be used as SnapVault snapshot name prefixes.

Now configure the secondary filer to create a set of sv_hourly snapshots. The secondary filer will keep only one of these snapshots and it will only update this snapshot from the primary filers twice during the day. The secondary filer should stay up to date, but it does not need to replicate all the snapshots and it certainly doesn't need to save them all. It would not be a disaster to lose a couple of hourly snapshots. On the secondary filer enter this command:

```
safe> snapvault snap sched -x sv_proj sv_hourly 1@12,20
```

Use of the -x option instructs the secondary filer to transfer data from the primary filers into the secondary filer qtrees before creating the snapshot.

14.8.1.2 Example 2: Nightly Snapshot Schedule

In this example, the primary filers keep four nightly snapshots on every day of the week except Saturday and the secondary filer keeps eight nightly snapshots.

On the primary filers enter these commands:

```
toaster> snapvault snap sched proj sv_nightly 4@sun-fri
fridge> snapvault snap sched proj sv_nightly 4@sun-fri
```

On the secondary filer enter this command:

```
secondary> snapvault snap sched -x sv_proj sv_nightly 8@sun-fri
```

The default behavior is to create a snapshot every night at midnight so if we wanted to create a nightly snapshot every day of the week we would just say 4 instead of 4@sun-fri. However, since we are going to create a weekly snapshot on Saturday in the next section, we limit nightly snapshots to Sunday to Friday. To conserve snapshots, you might also consider limiting nightly snapshots to Monday to Friday.

14.8.1.3 Example 3: Weekly Snapshot Schedule

In this example, the primary filer keeps one weekly snapshot and the secondary filer keeps twelve. These snapshots are created on Saturday mornings at 8 am.

On the primary filers enter these commands:

```
toaster> snapvault snap sched proj sv_weekly 1@sat@8
fridge> snapvault snap sched proj sv_weekly 1@sat@8
```

On the secondary filer enter this command:

```
secondary> snapvault snap sched -x sv_proj sv_weekly 12@sat@8
```

Note how it is possible to specify any combination of days of the week and hours of the day.

Verifying and Displaying Snapshot Configuration

To verify that the snapshots are configured the “snapvault snap sched” command can be used with no arguments in the following way:

```
secondary> snapvault snap sched
xfer sv_proj sv_hourly 1@12,20
```

Monitoring ‘schedule’ Progress with snapvault status

The SnapVault schedule progress can be monitored using the “snapvault status” command.

```
safe> snapvault status -s
Snapvault secondary is ON.
Volume      Snapshot      Status      Schedule
-----
sv_proj      sv_hourly      Active      1@12,20
```

14.8.1.4 Schedule Syntax Details

The snapvault snapshot schedule specification syntax is of the form:

```
<count>[@<day_list>][@<hour_list>]
```

Where:

- *count* is the number of snapshots in the series to keep.
- *day_list* is the day or days of the week that the snapshot should be created. Days are indicated by abbreviated names: “mon”, “tue”, “wed”, “thu”, “fri”, “sat” and “sun”. The *day_list* can be an enumerated, comma-separated list such as “mon,wed,fri”, a range such as “mon-fri” or a combination of enumerated and ranges such as “sat,tue-thu”. The abbreviated day names are not case sensitive so, for example, “Mon” is simply interpreted as “mon”. If no *day_list* is given, the default is to create snapshots every day which is the equivalent of a default *day_list* of “mon-sun”.
- The *hour_list* specifies the hours in the day the snapshot should be created. Hours are numbers in 24-hour format and local time. Thus, 8:00am is simply “8” whereas 5:00pm is “17”. The “0” hour is first thing after midnight. That is, “sat@0” means Saturday morning at midnight. The list may be enumerated or a range just as for days. Thus, “8-17,23” is a valid *hour_list*. The default *hour_list* is “0” which specifies the creation of just one snapshot in the morning at midnight.

The *day_list* and *hour_list* may come in any order, but the *count* must come first. For example, "5@mon@20" is the same schedule as "5@20@mon". SnapVault will create a snapshot only when the current time matches a day in the *day_list* and an hour in the *hour_list*.

14.8.2 Procedure to recover from an "Out of Snapshots" condition

1. Ensure no snapvault transfers are currently occurring to the volume in question on nasweh903r

```
snapvault status
```

All should be in an idle state.

2. Turn off snapvault on nasweh903r

```
snapvault.enable off
```

3. Repeat step 1.

4. Ensure that all snapvaulted qtrees on nasweh903r are at the same base snapshot. "snapvault status -l", look at "Base Snapshot:" for each qtree that's in the destination volume, and make sure that all refer to the same snapshot.

```
snapvault status -l
```

The base snapshot will be something like below and the name will depend on the volume.

```
nasweh903r(0050408378)_earp_lv_vol1-base.257
```

5. Create a new vol of at least the same size as earp_lv_vol1 (or whatever the vol name is). Refer to section 0 on how to do this. Make sure to set the proper retention periods.

Ensure that "vol lang" is the same on new vol as it was on "earp_lv_vol1". If you execute "vol lang NEWVOL SOME_LANG" then you will have to reboot in order for the new language mappings to take effect.

6. On the SnapVault primary system (nasny22, nasny23 and nasrut107) identify the snapshots on the volume from which SnapVault was running which were created **after** the last successful SnapVault operation:

```
snap list /vol/vol1
```

Note that you should use the appropriate volume name.

7. Since there was plenty of space left on the Snapvault secondary volume, we decide to ndmcopy the data that was not SnapVaulted there. Check the space on the existing SnapLock/LockVault volumes on the SnapVault secondary;

```
nasweh903r: df /vol/earp_lv_vol1
```

Understand that these copies will not be incremental like typical SnapVault transfers, they will be full transfers of the data in the qtree in that snapshot as of that particular snapshot. This will occur for each snapshot we want to copy.

8. Ensure that NDMP is enabled on all systems involved:

```
options ndmpd.enable
ndmpd.enable on
```

If it is not on, turn it on

```
options ndmpd.enable on
```

9. Use ndmcopy to move the data that didn't get SnapVaulted. It is recommended that the destination use a naming format that is descriptive of the data that would have been SnapVaulted. For example, the qtree /vol/vol1/citi on nasny22 from 10/12/2004 would use ndmcopy as shown below:

nasny22:

```
ndmcopy -sa root:password -da root:password
nasny22:/vol/vol1/.snapshot/FLASHOTNAME/citi
nasneh903r:/vol/earp_lv_vol1/citi_sv_nasny22_vol1.20041012
```

where FLASHOTNAME(i.e. - sv_nasny22_vol1.4) is the name of the snapshot that did not get transferred on 10/12/2004 from the Snapvault primary, nasny22.

10. Repeat this for each snapshot that did not transfer for each of the volumes/qtrees on the appropriate volumes/filers.
11. If necessary, **FIX THE SNAPVAULT SNAP SCHEDULES! Create new schedule reflecting the new volumes names.**

Here is an example of how **NOT** to setup it up:

```
===== SNAPVAULT SNAP SCHED =====
create NEWVOL 0@-
xfer NEWVOL sv_nasny22_vol1 250@2 retention_period=default
xfer NEWVOL sv_nasny23_vol1 250@2 retention_period=default
xfer NEWVOL sv_nasrut107_sec_v6 250@2 retention_period=default
```

Correct way to setup schedules on the secondary:

```
snapvault snap sched -x NEWVOL sv_nightly 250@2
```

The important point here is to use the SAME snapshot name, such as:

```
sv_nightly
sv_nasny22> snapvault snap sched vol1 sv_nightly 1@2
sv_nasny23> snapvault snap sched vol1 sv_nightly 1@2
sv_nasrut107> snapvault snap sched sec_vol6 sv_nightly 1@2
```

12. Initialize the SnapVault relationships to the new SnapLock volume created in Step 5:

```
snapvault start -S nasny22:/vol/vol1/citi -L nasneh903r:/vol/earp_lv_vol2/citi
```

Repeat this for each qtree that needs to be LockVaulted.

13. All of the snapshots that had to be ndmcopy'd (i.e. -/vol/earp_lv_vol1/citi_sv_nasny22_vol1.20041012) will belong to the vfiler0 on the SnapVault Secondary, (i.e. - nasneh903r). Assign each to the appropriate vfiler. For example:

On nasneh903r:

```
vfiler move vfiler0 nasneh903v3earp
/vol/earp_lv_vol1/citi_sv_nasny22_vol1.20041012
```

14.8.3 Overview for System Administrators

14.8.3.1 Introduction

This information was originally presented in the Lockvault for System Administrators Technical Manual (GDSE-04013-TM) and has been included here as part of a document consolidation process.

Network Appliance 'LockVault' solution for file system archiving per SEC 17a-4.

Citi has an immediate need to implement several storage related solutions that utilize WORM compliant storage to meet compliance mandates for SEC 17a-4, NASD 3010, NASD 3110, HIPAA, Sarbanes-Oxley Act, Graham Leahy, and others. Compliant storage solutions are presently used primarily for Document Imaging and Transaction Log Archiving but projects for E-mail and Document Retention require regulatory compliant storage solutions.

Today, Citi businesses mostly use optical platters to store compliant data. This is largely a result of regulations that require certain financial documents to be stored on unalterable media, typically referred to as WORM, Write-Once-Read-Many. WORM storage is any kind of storage medium that can be written to only a single time, but can be read from any number of times. Typically this has been optical disk, where the surface is permanently etched using a laser in order to record information. The leading magnetic storage vendors have begun introducing magnetic storage solutions that are "WORM-Like". The record retention features of these solutions enable them to compete with optical as a WORM alternative. These magnetic storage solutions have a total cost of ownership (TCO) that is lower than optical. In addition, it is widely believed that while optical disks, such as CDs and DVDs, will continue to be popular in the entertainment industry, optical technology's future in the storage industry is nearing the end of its production life-cycle. As a result of these factors, Citi businesses have begun to consider replacing their optical WORM storage with magnetic storage that is "WORM-Like".

Our focus on determining a solution is based on how it meets the standards set forth by the recent SEC Interpretive Release on Rule 17a-4 as it is the most stringent so far with regards to magnetic media and WORM definition. The SEC Interpretive Release clarified the Commission's views on the standards for broker-dealer firms preserving electronic records and email on disk-based storage media. Specifically, the Interpretive Release affirms that broker-dealers can meet the non-rewriteable and non-erasable requirements of Rule 17a-4 through the employment of a storage solution that uses integrated hardware and software codes to prevent the overwriting, erasure or alteration of the records. The SEC Interpretive Release also states that storage solutions that only mitigate the risk a record will be overwritten or erased do not meet the requirements of Rule 17a-4. "Such systems - which may use software applications to protect electronic records, such as authentication and approval policies, passwords or other extrinsic security controls - do not maintain the records in a manner that is non-rewriteable and non-erasable."

The NetApp 'LockVault' solution provides for archiving of information in an immutable format with security and audit capabilities that meet regulatory compliance guidelines per the SEC and CATE Storage Engineering. This solution is designed to meet the requirements for file system archiving on a daily basis with a three-year retention period. The retention period for data stored on this solution can be extended past the three-year mark and we are presently working with NetApp on methods to migrate data off to WORM tape to support longer retention periods.

14.8.3.2 How it works

An application creates a daily dump of information that needs to be maintained per regulatory requirements as defined in project EARP. A script deposits this data on a production NetApp filer on a nightly basis. This same script removes deposited data that is more than two days old. Each night a scheduled SnapVault process will sweep up the deposited data and store it on a NetApp NearStore FAS device on a WORM volume with three-year retention. A SnapMirror process will asynchronously copy this information to a second NearStore FAS WORM volume located in another data center. A WORM log (NetApp Compliance Journal) will maintain a record of all file and directory changes that take place between each SnapVault. This WORM log will also be replicated to the second R200. We recommend two nights worth

of data in case to script does not complete its data transfer to the production filer before the SnapVault process starts.

This solution maintains two copies of the information on WORM compliant storage for a default three-year period. The retention period on the data and its associated WORM log entry can be extended past the three-year period. These backups are always available online for restores. The WORM log entries are also always available online for searches and verification.

14.8.3.3 Future Enhancements

We are working with NetApp and Veritas to provide for compliant migration to WORM tape to support extended retention periods – past 10 years. We are also working on methods to enhance regular backup, event driven transfers and direct dumps from applications to the primary 'LockVault' filer.

Ongoing enhancements in management, reporting, new features, added functionality and any other changes will be forwarded to application owners as they become available.

14.8.4 Information for Systems Administrators

14.8.4.1 Script Requirements

A script may be required to copy data to a production filer to be swept nightly via SnapVault. The script will be required to cleanup data that is more than two days old. If the application is writing directly to the production filer than the script will only need to remove data more than two days old. If the SnapVault process is pulling data from the applications production storage then a script will not be required.

It is advisable that the scripts check for the successful completion of the prior night's backup before deleting data. This can be accomplished by checking that the directory of the prior night's backup in the production and/or COB R200 exists.

14.8.4.2 Domain Information

CIFS access will be through a single Vfiler based on NT domain trust authentication. NFS access will require separate Vfilers in different domains. A NetApp Vfiler will be created for each domain that an application resides in. This will allow for multiple domain authentication for NIS and AD domains.

14.8.4.3 Functional ID

The directory for information dumping will be permissions for a functional ID that must be unique to each application. This same ID will have read-only access to the backups for restore.

14.8.4.4 Two day requirements

We require two days worth of information to be maintained in order to make sure that the SnapVault process copies data that may not have been written within the defined window for backup. This way, any data that is missed will be picked up the next day and the time windows can be adjusted as needed.

If sufficient space exists the two-day requirement can be extended. This should be determined on a case-by-case basis. CATE Storage Engineering recommends a minimum of two days and a maximum of seven days.

14.8.4.5 Data copy window

Based on the initial feedback, we are proposing a 1AM SnapVault sweep for the initial deployment. This may be adjusted on an as needed basis based on changing business and technical requirements. Application owners should notify Storage Integration if they will not be able to meet this timeframe requirement.

14.8.4.6 Information Restore

The backups will be available online and will be viewable to the functional ID used to create them. The filer name and path will be provided prior to going into production for both R200 systems. Naming convention for the daily directories will also be provided with examples. Both the Production and COB mounts for restores will be provided to the application owner's.

14.8.4.7 Mount Points for CIFS and Unix

\\<Filer Name>\<Share Name> for CIFS and \\<Filer Name>\vol\<Share Name> for UNIX.

Under the mount point you will see a ~snapshot (.snapshot for UNIX) directory. Below ~snapshot you will see the SnapVault snapshot directories with the following naming convention:

sv<job name>_nightly.YYYYMMDD_HHMMSS_<time zone>

Example: svpaulvol_nightly.20040623_020511_EDT

This would represent a SnapVault snapshot of 'paulvol' taken June 23rd, 2004 at 2:05:11 AM eastern standard time.

Below the snapshot directory will be a directory named after the qtree that was SnapVaulted.

Example: ~snapshot\ svpaulvol_nightly.20040623_020511_EDT\paulvol\...

14.8.4.8 Naming Conventions

NetApp Vfilers and Qtree mounts will be named based on the application and business requesting access. Storage Integration will assign names on an as needed basis. Storage Integration will maintain a list of all production systems, Vfilers and application mounts with all of the appropriate information needed to setup each share and backup.

14.8.4.8.1 Mainframe Requirements

Mainframe environments that require access to NetApp LockVault for storage will need to use an FTP gateway to store information on a NetApp production filer. Scripts for cleanup should run on the FTP gateway. The Gateway will mount the production NetApp filer so that FTP writes will go directly to the NFS mounted share. The FTP gateway will also mount the production and COB R200 mounts for restores via FTP directly from the mainframe environment. Application groups will need to install and manage these FTP gateways through their local UNIX SA groups.

14.8.4.9 Min, Max, Default retention settings

File retention is based on the last file access field of a file. This is modified automatically through the SnapVault process or by SnapLock when a file is marked read-only when saved by an application or modified via a script. This field can be manually modified through a touch command or similar command to extend the retention of an individual file. The retention timeframe cannot be reduced once a file is marked read-only.

The retention settings are set per volume. If an application owner requires a change, please note that the destination volume may change to accommodate this. Several applications may be assigned to the same volume based on similar retention requirements. If these requirements change the target may be moved to another volume that will meet the new retention limits.

If an application owner attempts to change the retention of a file manually and attempts to set the retention outside of the maximum setting – the file's retention setting will not be changed.

The min, max and default settings take effect at the time a command attempts to manipulate the setting. For example – an administrator extends the retention on a file that has been on the system for three years by adding 5 years to the retention period, which was originally set to three years on a volume with a max setting of seven years. This would result in a total retention of eight years but would be allowed since that actual retention setting on the file itself would never be longer than five years. Attempting to set the retention to eight years initially would fail since that would be more than the seven-year max limit on the volume.

The min, max and defaults can be changed on a volume at any time without affecting the retention set on existing files. Changing these limits on existing production volumes requires approval from all application owners using the volume.

14.8.4.10 Data Deletion

On records have reached their retention period they may be deleted. Records that have reached their retention period will not be deleted automatically. A report will be sent to application owners telling them what data has reached its retention period – the method is still being developed. It is the responsibility of the application owner to delete records that have reached their retention period or to extend the retention period. This should be done in accordance with the Citi records retention policy.

14.8.4.11 Alternate Infrastructure Implementations

In addition to the 'traditional' primary, secondary and tertiary system implementations discussed elsewhere within this document, NetApp and CATE Storage Engineering also support the following alternative infrastructure implementations. These alternate implementations should be considered on a case by case basis according to application storage and retention requirements.




Figure 36 - Dual application write to both systems




Figure 35 – Direct write to primary with NAS replication

14.8.4.12 NetApp Information for SnapLock Developers

14.8.4.12.1 Why can't I rename directories on a SnapLock volume?

Directories do not behave any differently than they would on regular volumes with one important exception. They cannot be renamed or moved once created. It is a requirement for regulatory compliance that WORM data not only be non-erasable and non-rewritable but must also be locked down in the same location at which it was created. In the case of our WORM implementation this means that the directory path to WORM files must be locked down and never change. If it were allowed to change then multiple ugly possibilities for problems present themselves.

The first and most obvious problem is that the front-end databases our application partners use to keep track of the location of WORM data on the volume would lose track of any content contained underneath the moved or renamed directory. This would give them some major headaches with the SEC as they are required by SEC 240.17a-4 to be able to generate an index of all archived content and verify the integrity of all the content on demand or face substantial penalties and fines.

The second more serious problem is the security hole it would introduce into our WORM model. If I wanted to modify a WORM file in a covert way I could simply move the directory containing the target file to an obscure place (or rename it as "...") or something tricky like that), recreate the original directory, copy all original files except the file I want to tamper with into this new directory, copy off the file I want to tamper with to a non-WORM volume to edit it and finally move the tampered file back into the new directory and commit the tampered file as WORM.

The applications are free to create and delete temporary directories on WORM volumes as long as they don't commit any WORM files inside those directories. Regular directory semantics dictate that you cannot delete a non-empty directory so if you create WORM files in a directory it will be permanently non-empty and as a result will become impossible to delete.

14.8.4.12.2 How do I commit a file to a WORM state on a SnapLock volume?

Committing a file to WORM state on a SnapLock volume simply involved transitioning its attributes from a writable state to a read-only state. The exact operation is dependent on the file protocol (CIFS, NFS etc.) and the client operating system but is always straight forward. This transition can easily be done manually, via scripts or programmatically. Some examples for different environments follow:

Unix shell environment:

```
chmod -w <file>
```

Unix POSIX programmatic environment:

```
if (chmod(path, 0444) < 0) {
    return ERROR;
}
```

Windows shell environment:

```
attrib +r <file>
```

Windows Win32 programmatic environment:

```
fattr = GetFileAttributes(PathName);
if (fattr == -1) {
    return ERROR;
}
fattr |= FILE_ATTRIBUTE_READONLY;
if (SetFileAttributes(PathName, fattr) == 0) {
    return ERROR;
}
```

It is very important to note that for the committal to WORM state to occur the file must experience a transition from a writable state to a read-only state. Files that are "born" read-only will not experience this transition and hence will not be committed to WORM state. Applications should always ensure that the file is initially writable on the SnapLock volume before making it read-only to commit to WORM state.

The cases in which files are born read-only are very operating system and client dependent. You should never rely upon or make behavioral assumptions about external utilities, commands or GUI drag and drop interfaces when copying read-only files onto a SnapLock volume. Depending on the implementation, these utilities may or may not cause the read-only file to be committed to WORM state on the SnapLock volume.

14.8.4.12.3 How can I tell if a file is WORM on a SnapLock volume?

Unfortunately this is not as easy as just checking whether the file is read-only because there is no guaranteed one-to-one mapping between the read-only attribute of a file and its WORM state. This is because files that were "born" read-only and never experienced a writable to read-only transition will not be WORM. There is no passive way to determine whether a file on a SnapLock volume is WORM.

The simplest active technique for determining whether a file is WORM on a SnapLock is to perform a harmless meta data update operation such as setting the access time of the file. Access for this operation will be denied with a read-only error if the file is WORM. Please use an access time earlier than the existing value since you do have the ability to increase retention periods, i.e. move the retention date (stored in the last access time field) into the future. However no version of ONTAP allows you to change the access time on a WORM committed file to an earlier time than its existing value – so this is a true test of a file's WORM state. For more info on how to set the access time – refer to the question "How do I set the retention period of a WORM file?"

14.8.4.13 How do symbolic links work on SnapLock volumes?

Because symbolic links do not have any permissions of their own they cannot be made read-only and hence cannot be committed to WORM state. This is not an issue for CIFS because it does not support symbolic links. The majority of the partner applications we are targeting are CIFS based so this is not an issue for them. For the partner applications which use NFS, none so far have any use or requirement for WORM symbolic links.

14.8.4.14 Record Retention Periods

14.8.4.14.1 What is a record retention period

The retention period of a record is the length of time for which a regulated customer must retain a non-erasable and non-rewritable copy of a record. The length of the retention period is dependent upon factors such as the applicable regulations and type of document as well as other external factors such as applicability to an investigation or ongoing litigation. It is not unusual to see retention dates of 4-6 years (SEC etc.) up to much longer periods of 30+ years (mortgages, patient records etc.).

From the perspective of the customer, records transition from being an asset to a liability at the moment they reach the end of their mandated retention period and they want them to be deleted as soon as possible. With traditional WORM media this destruction was a physical process in which retention was tracked at the granularity of the WORM media. Records with the same expiry date were grouped on the same media and when that expiry date was reached the unit of WORM media would be physically destroyed. This is not a very scalable or cost effective way to operate.

To support this retention requirement, SnapLock allows for the assignment of record retention periods using a mechanism which is, in almost every way, superior to traditional batched retention management. SnapLock supports assignment of retention dates via open file protocol interfaces (CIFS, NFS etc.) at the granularity of individual files. Co-existence of records with different expiry dates on the same volume is permitted and records can be individually deleted once their associated retention dates have been reached. Data ONTAP 7.1 release SnapLock also supports the extension of assigned retention dates further into the future. This achieved on a per volume level using the snapvault snap retain command (see section 11.6.4.3.4.5 for details on using this command).

When compared with retention solutions utilizing traditional optical WORM media (or WORM tape), SnapLock provides a huge advantage in terms of performance, scalability and storage consolidation.

14.8.4.14.2 How does SanpLock store retention periods?

In order to support assignment of retention periods via the open file protocols and provide a retro-active solution to customers running Data ONTAP releases prior to the support of retention periods, SnapLock has hijacked the access time field of the file meta data for storage of the file retention date.

Note: Keep in mind that ONTAP still uses a 32 bit UNIX clock count non-negative seconds since time zero (which coincides with the UNIX epoch on January 1st, 1970). This means retention dates can currently only be set up to the year 2038. Future releases of ONTAP will extend this capability with the introduction of a 64 bit UNIX clock.

Aside from its use to store retention dates, the access time is very rarely used by applications and is, in fact, often turned off via a volume option to improve the performance of the volume. Performance is improved because Data ONTAP does not have to update the file metadata every time files are accessed for non-modifying operations. In Data ONTAP releases from 6.5 onwards, the no_atime_update file system option of SnapLock volumes is forced to "on" to clearly indicate to users that the access time is not being updated by Data ONTAP.

14.8.4.14.3 How do I set the retention period of a WORM file?

Setting the retention date of a file on a WORM volume is as simple as setting the access time of the file to the desired date immediately prior to committing the file to WORM state. Please note that if you attempt to do this on a file after committing a file to WORM state – the file will already have been set to the default period. This is not as much of a concern as long as the default retention period on the volume is initialized to an appropriate value. Moreover, Data ONTAP does allow retention periods to be extended (although it does not allow the retention periods to be reduced).

From a Unix shell these simple commands will set a file retention date and commit it to WORM:

```
touch -a -t <time> <file>
chmod -w <file>
```

Here is a C code programmatic function for performing the same operation on Unix clients:

```
#include <sys/types.h>
#include <sys/stat.h>
#include <inttypes.h>
#include <utime.h>
/* SnapLock commit return codes */
#define SL_SUCCESS 0 /* Success */
#define SL_ERR_STAT -1 /* Failed to stat() file */
#define SL_ERR_BAD_PERIOD -2 /* Bad retention period */
#define SL_ERR_SET_PERIOD -3 /* Failed to set retention period */
#define SL_ERR_WORM_COMMIT -4 /* Failed to commit to WORM */
#define SL_RETENTION_EPOCH 1044057600 /* 1/Jan/2003 00:00:00 GMT */
/*
 * snaplock_commit_file
 *
 * Commit a file to WORM state with specified retention period.
 *
 * Args: path = pathname to file to commit to WORM
 * secs = retention period in seconds
 *
 * Returns: 0 if successful
 * -1 if error
 */
int
snaplock_commit_file(const char *path,
time_t secs)
{
    struct stat s;
    struct utimbuf utb;
    /* Query the last modify time of the file */
    if (stat(path, &s) < 0)
        return SL_ERR_STAT;
    /* Set the file retention date in the access time */
    utb.actime = secs;
    utb.modtime = s.st_mtime;
    if (utime(path, &utb) < 0)
        return SL_ERR_SET_PERIOD;
    /* Commit the file to WORM state by making it read-only */
    if (chmod(path, S_IRUSR|S_IRGRP|S_IROTH) < 0)
        return SL_ERR_WORM_COMMIT;
    return SL_SUCCESS;
}
```

On Windows clients there is no built-in shell command to set the access time similar to "touch" on UNIX systems but such utilities are available from various installed, the operation can be performed from a Windows shell thusly:

```
touch -a -t <time> <file>
attrib +r <file>
```

A similar programmatic function to the C code above can be implemented using the Window POSIX API or alternatively using the native Win32 API. Examples of this are the SnapLocker.exe utility and the touch210 utility. When using the Win32 API you should pay careful attention to the time formats to ensure that the correct encoding is produced because even though Win32 supports a time format with a larger range, Data ONTAP only supports the standard time_t 32-bit format.

14.8.4.14.4 What's the shortest retention period I can set?

The shortest supported retention period setting is 1 day. Similarly, the default retention period can be set in days, months or years – as in the example below –

`vol options volname snaplock_default_period 3y` (sets the default retention period to 3 years)

14.8.4.14.5 What's the longest retention period I can set?

The longest retention period you can set in Data ONTAP 7.1 is 30 years. CATE Storage Engineering does not recommend setting retention past 7 years based on existing disk technology. However, since you can extend retention periods – you will be able to extend the retention period well before any of these files expire.

14.8.4.14.6 How do I extend the retention period of a WORM file?

The ability to extend the assigned retention date of a WORM file is available in Data ONTAP release 6.5D1X4 and above.

To extend a retention date you simply set the access time to the desired date. More precisely, retention dates of WORM files can never be shortened for any reason.

The need for extending record retention dates usually stems from investigations by the regulatory body or litigation for which the records to be extended are considered pertinent evidence. Such extensions can be very open ended if the litigation becomes protracted.

14.8.4.14.7 How do I delete a WORM file once it has reached its retention date?

Deletion of WORM files which have passed their retention date (called expired files) is allowed in Data ONTAP release 6.5D1X4 and above.

When a WORM file reaches its retention date it does not become non-WORM but instead changes its "disposition" to allow for deletion. The only operations Data ONTAP will permit on an expired WORM file are changing of the file permission to deletable state and the delete operation itself. Any other kind of modifications to the file is not permitted.

A Unix shell example illustrating how to delete an expired WORM file is as follows:

```
chmod +w <worm_file>
rm -f <worm_file>
```

A Windows shell example illustrating how to delete an expired WORM file is as follows:

```
attrib -r <worm_file>
del <worm_file>
```

Programmatic techniques will follow the same pattern as the above examples.

14.8.4.14.8 Is there a default retention period for SnapLock volumes?

Default retention dates are implemented in the Data ONTAP 6.5D1X4 and above. Minimum and Maximum retention periods are also set on each volume.

14.8.4.14.9 How can I check the default retention period on a given volume?

You can check the default/min/max retention periods the same way you would set them, ie by using the `vol options` command.

Example -

```
filer1> vol options snaplock
nosnap=off, nosnapdir=off, minra=off, no_atime_update=on,
raidtype=raid4, raidsize=8, nvfail=off, snapmirrored=off,
resyncsnaptime=60, create_ucose=off, convert_ucose=off,
maxdirsize=31457, fs_size_fixed=off, snaplock_enterprise,
svo_enable=off, svo_checksum=off, svo_allow_rman=off,
svo_reject_errors=off, fractional_reserve=100,
snaplock_default_period=min, snaplock_minimum_period=0y,
snaplock_maximum_period=30y
```

14.8.4.14.10 Does the volume default retention period mean all files will expire at the same time?

No. The default retention period is applied to all files on the volume (unless the admin or application over-rides it for a particular file), but the starting point is when the particular file was committed to WORM status. For example if the default is 3 years, a file committed 2.5 years ago will be eligible for deletion in 6 months (0.5 years), but a file committed today will be locked for 3 years (starting from today).

14.8.4.15 Will SnapLock automatically delete expired WORM files?

No. SnapLock will never autonomously delete any files from the volume.

It is the responsibility of the applications to delete expired records and almost all applications do this already using daemon processes which periodically sweep through the volume deleting expired records. Apart from the obvious philosophical stance that Data ONTAP should never autonomously delete user data under any circumstances, the most practical argument against such behavior is that most applications manage and index their records through databases like Oracle, Microsoft SQL, or IBM DB2. Deleting data out from underneath these databases will cause problems. Application vendors have also expressed clearly to us that they would have a problem with NetApp if they decided to delete records autonomously.

14.8.4.16 What's to stop me resetting the system clock to delete files early?

The most obvious question regarding the security of the file retention dates is how NetApp guarantees that we have a secure time base to prevent users from simply maliciously setting the system clock forward 30 years in order to delete protected records. We cannot usurp and lock down the system clock because it is crucial for correct operation of the open file protocols (CIFS, NFS etc.) that this clock remains accurately synchronized to the time in the data center via NTP. External connections outside of the data center to trusted third party time sources are also a show-stopper for the vast majority of customers.

For Data ONTAP release 6.5D1X4 through 6.5D1X10 NetApp has implemented a secure software "Compliance Clock" that tracks the hardware clock of the system but reacts to changes with a bounded drift rate. This means that if the system clock is set forward the ComplianceClock will try to match it, but so slowly that it may take years to move forward by a few days. The ComplianceClock is initialized once from the system clock at activation time and is stored persistently on SnapLock volumes. This ComplianceClock is the time source that is used for making all time-based WORM file retention decisions. To set the ComplianceClock, the following command is required – `date -c initialize`

Example:

```
filer> snaplock clock initialize
```

```
*** WARNING: YOU ARE INITIALIZING THE SECURE COMPLIANCE CLOCK ***
You are about to initialize the secure Compliance Clock of this system to the
current value of the system clock. This procedure can be performed ONLY ONCE
on this system so you should ensure that the system time is set correctly
before proceeding. The current local system time is: Wed Feb 4 23:38:58 GMT
2004
```

```
Is the current local system time correct? y
```

```
Are you REALLY sure you want initialize the Compliance Clock? Y
```

```
Compliance Clock: Wed Feb 4 23:39:27 GMT 2004
```

To view the ComplianceClock, the following command is required – date -c

Example:

```
snaplock clock status
```

```
Compliance Clock: Wed Feb 4 23:42:39 GMT 2004
```

14.9 Appendix I NFSv4 ACLs – further reading

An ACL is a list of permissions associated with a file or directory and consisting of one or more Access Control Entries (ACEs). NFSv4 ACLs provide finer granularity than typical POSIX read/write/execute permissions and are similar to CIFS ACLs.

Below is an example ACL displayed using the Linux “nfs4_getfacl” command:

```
A::sh939001@nfsdomain:rwatT
A::sh939002@nfsdomain:rt
A:g:nasgrp9001@nfsdomain:rt
A::OWNER@:rwatTnNcCy
D::OWNER@:x
A:g:GROUP@:rtncy
D:g:GROUP@:waxTC
A::EVERYONE@:tc
D::EVERYONE@:rwaxTCy
```

In the example above:

- The user sh939001 has read, write, append, read attributes and write attributes granted to it.
- The user sh939002 only has read and, read attributes granted to it.
- All members of the group nasgrp9001 have read and, read attributes granted to them.
- The OWNER@, GROUP@ and EVERYBODY@ are special ACE entries that essentially reflect the attributes granted by the Unix permissions (owner, group and everybody respectively)
- In this case the owner has all attributes except “execute”
- The group has essentially the “read” type attributes.
- The group everybody can only read the attributes and ACL data.
- If an attribute has not be specifically granted then it is denied so the deny ACE entries in the example (those with a leading “D”) are essentially superfluous.

As can be seen each ACL has a series of ACL entries with the following format:

type:flags:principal:permissions

ACE Types:

- A Allow – Allows the principal to perform the actions listed in the permissions field.
- D Deny – Prevents the principal from performing the actions listed in the permissions field.

There are also Audit ACE types (U and L) but these are not supported by ONTAP.

ACE Flags:

There are three types of flag types and these are group, inheritance and administrative. As ONTAP does not support the Audit ACE types we will not discuss the administrative ACE flags are these are only relevant to Audit ACE entries.

Group Flags

- g group – Indicates that the principal is the name of a group as opposed to a username. This flags can be used on any ACE type.

Inherited Flags

- d directory-inherit. Only new directories will inherit this ACE.

- f file-inherit. Newly created files (and directories) will inherit this ace.
- n no-propagate-inherit. Only new directories will inherit this ACE but this ACE will have the inherit flag removed on new directories.
- i Inherit-only. This ACE does not grant permissions to the principal to this directory but is inherited by files and directories created in this directory. The "i" inherit flag is removed from the ACE when it is placed on new files but retained on new directories.

ACE principals

A principal is either a username (e.g. fred@nfsdomain, where "nfsdomain" is the NFSv4 domain id name and **must** be supplied) or a group name (e.g. nasgrp9001@nfsdomain). If the principal is a group then it will have the group flag "g" set in the ACE. There are also three special principals that are used to mirror the standard Unix permissions, these are OWNER@, GROUP@ and EVERYBODY@, note that these do not have the NFSv4 domain id name attached but that GROUP@ is treated as a group and hence requires the "g" group flag.

ACE permissions

There are many ACE permissions and some of these permissions have more than one meaning depending on whether they are applied to a file or a directory (where neither file or directory is specified the permission is valid on either). Below is a list of the permissions available:

- r File – read data. Directory – list directory contents
- w File – write data. Directory – create file in directory
- a File – append data. Directory – create subdirectory
- x File – execute. Directory – transit through directory (like "x" on Unix permissions)
- d Delete the file or directory
- D Directory only – delete file or sub-directory within the parent directory
- t read file or directory attributes
- T write file or directory attributes
- n read named attributes of file or directory
- N write named attributes of file or directory
- c read ACL (NFSv4 ACL)
- C write ACL (NFSv4 ACL)
- o change owner of the file or directory
- y allow synchronous reads and writes to server

Inheritance

If an inherited ACL is set on a directory (inheritance is only valid on a directory) then the Unix permissions as specified by the users umask settings are not set on the files or directories created in that directory on the inherited ACL is set. Hence if an inherited ACL is set on a directory and the Unix permissions are still required (the ACL may provide sufficient access) then these must also be created manually as inherited ACE entries. Below is an example of an inherited ACL that also sets the Unix permissions.

```
A:d:sh939002@nfsdomain:rt
A:d:sh939001@nfsdomain:rwatT
A:d:OWNER@:rwaDxtTnNcCy
A:dg:GROUP@:rxtncy
D:dg:GROUP@:waDTC
A:d:EVERYONE@:rwaDxtTnNcCy
A:fi:OWNER@:rwatTnNcCy
D:fi:OWNER@:x
A:fig:GROUP@:rtncy
D:fig:GROUP@:waxTC
A:fi:EVERYONE@:rtncy
D:fi:EVERYONE@:waxTC
```

In this example:

- The user sh939002 has read access to this directory and this ACE is inherited by any directories created in this directory. ("d" flag)
- The user sh939001 has read+write access to this directory and this ACE is inherited by any directories created in this directory. ("d" flag)
- The first set of OWNER@, GROUP@ and EVERYONE@ ACE entries provide all access to the owner and everybody with the group being given only read access. These ACEs both provide access to this directory and will also be inherited by any files or directories created. ("dg" and "d" flags)
- The second set of OWNER@, GROUP@ and EVERYONE@ ACE entries do not provide any access to directory itself but are inherited by all files created in this directory. ("fig" and "fi" flags)
- Note that directories and files get different OWNER@, GROUP@ and EVERYONE@ ACE permissions; which is appropriate as the permissions mean different things on files and ACLs. For example you want the execute permission on a directory but not necessarily on a file.

A file created in this directory would have the following ACE applied

```
A::OWNER@:rwatTnNcCy
D::OWNER@:x
A:g:GROUP@:rtncy
D:g:GROUP@:waxTC
A::EVERYONE@:rtncy
D::EVERYONE@:waxTC
```

A directory created in this directory would have the following ACE applied

```
A:d:sh939002@nfsdomain:rt
A:d:sh939001@nfsdomain:rwatT
A:d:OWNER@:rwaDxtTnNcCy
A:dg:GROUP@:rxtncy
D:dg:GROUP@:waDTC
A:d:EVERYONE@:rDxtTnNcCy
A:fi:OWNER@:rwatTnNcCy
D:fi:OWNER@:x
A:fig:GROUP@:rtncy
D:fig:GROUP@:waxTC
A:fi:EVERYONE@:rtncy
D:fi:EVERYONE@:waxTC
D::OWNER@:
D::EVERYONE@:wa
```

Deny ACES

When a principal attempts to access a file or directory with an ACL set then the ACL is evaluated ACE-by-ACE from the top ACE to the bottom ACE (as they are displayed) for each permission. Once a permission has either been allowed or denied once then this permission cannot be changed by subsequent ACE entries. For this reason the order of ACE entries is very important and will determine the actual permissions given to a principal on accessing the file or directory.

14.10 Appendix J OnCommand Systems Manager

Starting with ONTAP 8.1 NetApp has retired the built-in FilerView GUI and replaced it with an off array stateless management web server called System Manager. System Manager is a graphical management interface that allows you to manage storage systems and objects (such as disks, volumes, aggregates etc.) and performing these tasks over a web browser.

14.10.1 System Manager - Overview

System Manager enables storage administrator to perform the following tasks with ease:

- Configure and manage storage objects such as disks, aggregates, volumes, qtrees and quotas
- Configure CIFS/ NFS protocol and provisioning of file sharing
- Verify and configure network configuration settings in the storage systems
- Create and manage vFilers
- Setup and manage SnapMirror and SnapVault relationships
- Manage HA configurations and perform takeover and giveback operations

Within Citi System Manager is designed to be installed on designated management servers within the Data Center and as such is only available for deployment on the Windows Server platform.

14.10.2 System Manager - Installation

Please refer to System Manager 2.2 Release Notes for more information on the steps and pre-requisites for installation.


14.10.3 System Manager - Configuration

System Manager 2.2, unlike other NetApp products such as NetApp OnCommand Unified Manager 5.x, is designed to run as a stateless web server based application accessible by the web browser on the server upon which System Manager is installed. Therefore no specific settings are required to configure.

- Upon launching the System Manager from the shortcut on the server desktop (or Start menu) you will see an empty **Home** screen. Click on the **Add** button to add the storage system that you wish to manage. This is a one time operation to add each system you wish to manage.



- Partner node will be added automatically if the filers are configured as High-Availability pair as long as the credentials are the same when managing each node



- To start managing/ configuring the filer, select the filer and click **Login**

