



CATE Storage and Backup Engineering
Citigroup Architecture & Technology Engineering

NetApp OnCommand Unified Manager 5.2

CSBE-11001-SG

Solutions Guide

Version 2.9

©Citigroup Inc. 2014, All Rights Reserved.

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

1 Version Control

Version	Date	Author	Reviewed	Comments
0.1	18-APR-2011	Richard Sopp	Stephen Hopkins	Peer Review
1.0	06-MAY-2011	Richard Sopp	CTI NAS Ops	Initial production release
1.1	11-MAY-2011	Richard Sopp	n/a	Change VM CPU count
1.2	27-MAY-2011	Richard Sopp	n/a	Add Perl TPM package
1.3	24-JUNE-2011	Richard Sopp	n/a	Updated cluster monitor intervals
1.4	08-FEB-2012	Tan Jackson	Armstrong, Tim Sopp, Richard	Peer Review
1.5	23-APR-2012	Tan Jackson	Global Stakeholders	Released For Testing
1.6	14-MAY-2012	Richard Sopp	n/a	Production release
2.0	03-JUN-2012	Tan Jackson	Sopp, Richard	Peer Review for OC 5.1
2.1	09-JUL-2013	Tan Jackson	n/a	RFC
2.2	30-JUL-2013	Tan Jackson	n/a	RFP
2.2.1	05-AUG-2013	Richard Sopp	n/a	Update link for dfm_cfg2
2.2.2	10-JAN-2014	Tan Jackson	Richard Sopp	Updated NetApp OnCommand Unified Manager 5.1 configuration script
2.2.3	14-JAN-2014	Tan Jackson	RFT	Minor changes
2.2.4	03-MAR-2014	Tan Jackson	N/A	RFP (Updated) Release
2.2.5	25-MAR-2014	Tan Jackson	Richard Sopp	Draft version for Unified Manager 5.2
2.3	16-APR-2014	Tan Jackson	RFT	Release for Testing
2.4	9-May-2014	Tan Jackson	RFP	Release for Production
2.5	18-AUG-2014	Tan Jackson	Richard Sopp	Added NMC 3.3
2.6	23-SEPT-2014	Tan Jackson	N/A	RFP
2.7	7-OCT-2014	Tan Jackson	Richard Sopp	Updated with OCUM5.2R1P1 details Peer Review

2.8	12-NOV-2014	Tan Jackson	n/a	RFT
2.9	19-NOV-2014	Tan Jackson	n/a	RFP

2 Legal Notice

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

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

3 Contact

Please address any comments or questions regarding the information contained in this document via email to 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 and Backup 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 Last Reviewed
Asia Pacific	Jill Christine Nuqui	1.5
EMEA	Tomasz Grala	2.3
Japan	NA	NA
Latin America / Mexico	NA	NA
North America	Shahid Hussain	1.5

Table 1 - Regional stakeholder review

Table of Contents

1	Version Control.....	2
2	Legal Notice.....	4
3	Contact.....	4
4	Latest Versions.....	4
5	Regional Reviewers.....	5
6	Introduction	7
7	Assumptions, Conventions, Definitions and References.....	9
7.1	Assumptions	9
7.2	Issuing commands.....	9
7.3	Documentation conventions.....	9
7.4	Definitions and Abbreviations	9
7.5	References.....	10
8	Installation	11
8.1	Getting started.....	11
8.2	Platform Requirements	13
8.3	Installation Procedure.....	16
9	Best Practices.....	20
9.1	Implementation Overview	20
9.2	UM Application Protection	25
9.3	UM Maintenance.....	52
9.4	UM Migration.....	53
9.5	UM Capacity Planning.....	56
9.6	UM Application Options.....	58
9.7	UM Storage Device Management.....	59
9.8	UM Reporting.....	66
9.9	UM Storage Configuration Management.....	86
9.10	UM SnapMirror (Replication) Monitoring.....	117
9.11	UM SnapVault (Backup) Monitoring.....	120
9.12	UM Events and Alarms	121
9.13	UM Autosupport	127
9.14	Performance Monitoring and NMC Performance Advisor	128
9.15	UM Script Plug-in Library.....	135
9.16	UM Qtree Quota Management.....	136
9.17	UM Command Aliasing.....	139
9.18	Run Command function	140
9.19	Chargeback.....	140
10	Security.....	141
10.1	Vulnerability and Threat Management.....	141
10.2	Known Issues	141
10.3	Web server SSL configuration	142
10.4	Required IP ports	146
10.5	Functional IDs.....	147
10.6	Access Rights to UM Management Station	148
10.7	UM Security Model Overview.....	148
10.8	UM Security Model Implementation.....	152
10.9	UM Entitlement Reporting.....	163
10.10	UM Filer and vFiler Password Management.....	164
10.11	UM Log Files.....	168
11	Appendices	173
11.1	Unified Manager Server Settings (Options) Configuration Matrix for Small to Medium Environments	173
11.2	Unified Manager Additional Server Settings Configuration Matrix	186
11.3	NetApp Management Console (Client) Settings Configuration Matrix	187
11.4	Unified Manager Server Settings (Options) Configuration Addendum Matrix for Large to Very Large Environments.....	188
11.5	Unified Manager Event Listing.....	190
11.6	UM Database Objects and Attributes	191
11.7	NetApp OnCommand Unified Manager Console.....	196
12	Document Version Control	199
13	Index.....	202

6 Introduction

NetApp OnCommand Unified Manager¹ (UM) is the centralized monitoring, management and reporting console provided by NetApp for use with NetApp's NAS Filers running the ONTAP 7G and ONTAP 8.X / 7-mode operating system².

The purpose of this document is to provide a reference resource to storage administrators, SAs and other technology support teams involved with supporting or using NetApp NAS Filer solutions. This Solutions Guide is written with UM release 5.2 in mind and also focuses on the advanced performance monitoring features offered by the Performance Advisor client (part of the NetApp Management Console [NMC] client side portion of the UM client / server architecture).

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 reporting interpretation when using technologies such as thin provisioning, storage over subscription and de-duplication; 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 and tries to present information in terms of the different CITMP disciplines such as change management and capacity planning.

The Solutions Guide is split into four main sections:

- Chapter 8 – [Installation](#)
 - High level overview of required infrastructure components
 - Rules for determining specification of hosting hardware for UM application
 - High level overview of the steps involved in installing and configuring the UM application

For detailed instructions on the installation and upgrade process refer to the [CATE UM 5.2 Release Notes](#) document.
- Chapter 9 – [Best practices](#)
 - Implementation options and models

Refer to the [CATE UM 5.2 Release Notes](#) document for a list of unsupported / uncertified functionality

 - Application optimization and management guidelines
 - Technical deep dives on:
 - storage reporting
 - storage configuration management
 - event and performance monitoring
- Chapter 10 – [Security best practices](#)
 - User entitlements and Functional IDs
 - Network port requirements
 - Application log files
- Chapter 11 – [Reference appendices](#)
 - Standard build configuration checklist
 - Technical extended references

¹ NetApp OnCommand Unified Manager is formerly known as Operations Manager (OM) [also known as DFM (Data Fabric Manager) before OM]. The core server component of this monitoring and management solution is still referred to as DFM and most of the CLI commands still use the keyword 'DFM'. Unified Manager is actually the name of the web interface for OM/ DFM however CATE Storage Engineering have decided to adopt the term Unified Manager to describe the entire product.

² The Citi UM 5.2 TPM package cannot be used to monitor or manage Clustered ONTAP.

For the experienced administrator the recommended fast track approach for using this document is:

1. Chapter 8 – Installation overview
2. [CATE UM 5.2 Release Notes](#)
3. Chapter [11.1](#) – Configuration guide
4. Chapter [9.8](#) - Storage Reporting
5. Chapter [9.9](#) - Storage Configuration Management
6. Chapter [9.12](#) - Event handling and Performance Monitoring

This will provide the necessary information on changes and new features for an experienced administrator, however administrators of all capabilities and experience levels are recommended to use this reference document whenever required. Where appropriate this document can be used to assist in training storage administrators.

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.

7 Assumptions, Conventions, Definitions and References

7.1 Assumptions

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

7.2 Issuing commands

All commands listed in this document are to be executed on the UM hosting server or in the NetApp Management Console (NMC). Where neither location is specified it is assumed the command is being executed on the UM server, otherwise the NMC will be explicitly defined in each example. There are two methods by which commands can be issued on the UM hosting server: the command line, and the web user interface. Where necessary, the type of interface used will be indicated throughout this document.

7.3 Documentation conventions

The following conventions are employed within this document:

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

7.4 Definitions and Abbreviations

The following table contains terms and acronyms that are used throughout this document. The definitions listed here relate to terms used specifically with this application or generic terms that have a non-standard meaning when used with Unified Manager.

Term or Acronym	Explanation or Definition
Management Station	Server on which UM is installed
ManageOnTap	Management API for NetApp's ONTAP operating system
MSSA	Management Station System Administrator. The SA responsible for the underlying O/S and physical server (if applicable) on which Unified Manager is installed
NMC	NetApp Management Console. The Java thick client installed on remote workstations. Used for Performance Advisor
UM	Unified Manager. Formerly known as Operations Manager/ DataFabric Manager. NetApp's management console for use in the NAS infrastructure
UM SA	Unified Manager System Administrator. The individual or team responsible for managing the UM application. Usually the application administrator is considered to be the same as the NAS Administrators group
UM user	Unified Manager User. A team or individual using Unified Manager to monitor and manage the storage assigned to them. Typically UM user is a business aligned SA team
PCM	Process Control Manual
Storage Controller	A generic term recently adopted by NetApp to replace the term 'Filer'

Table 2 – Definition of terms

7.5 References

Item Name	Location / Value
Functional Account Technical Requirements (CTI Information Security Administration (ISA) organization)	https://globalconsumer.collaborationtools.consumer.citigroup.net/sites/CTIWebHost/WHGcompl/CAP/RECOMM_CAPS/0176BI.2-DAP/CTI_STANDARDS_FOR_EMERGENCY_DEVELOPER_ACCESS.docx
ISO 3166 Country Code List	http://www.iso.org/iso/country_codes
UM 5.2 Solutions Guide	http://catecollaboration.citigroup.net/domains/platstor/gdse/products/platforms/NASFiler/11001-SG.pdf
UM 5.2 Release Notes	https://catecollaboration.citigroup.net/domains/platstor/gdse/products/platforms/NASFiler/14053-RN.pdf
UM 5.2 Project Scope	https://catecollaboration.citigroup.net/domains/platstor/gdse/products/platforms/NASFiler/CSBE-14053-PS.pdf
UM 5.2 Database Storage Estimator	http://catecollaboration.citigroup.net/domains/platstor/gdse/products/platforms/NASFiler/OM_database_growth_estimator_v6.xls
NetApp UM 5.2 Installation and Upgrade Guide*	http://mysupport.netapp.com/documentation/docweb/index.html?productID=61832&language=en-US
NetApp UM 5.2 Admin Guide*	http://mysupport.netapp.com/documentation/docweb/index.html?productID=61832&language=en-US
NetApp UM 5.2 Performance Advisor Admin Guide*	http://mysupport.netapp.com/documentation/docweb/index.html?productID=61832&language=en-US
Configuring VirusScan Settings for Manager Servers via ePO	https://catecollaboration.citigroup.net/domains/security/isp_standards/hsestds/AntiVirus/McAfeeVirusScan/stdsdocs/VS88ePO46ConfigStandards1.9.doc

Table 3 - Document references

*Login required to <http://support.netapp.com>.

8 Installation

The purpose of this section is to describe the pre-requisites for deploying [UM](#) in a Citi NetApp NAS Filer environment, and provide a step-by-step process for new installations of UM.

The rest of this chapter discusses new deployments of UM within an environment. Storage Administrators requiring the procedure to upgrade from UM 5.2 should refer to section 8.2 of the [UM 5.2 Release Notes](#) for this process and use the rest of this chapter for reference purposes only.

8.1 Getting started

The purpose of this section is to assist storage administrators and project managers when planning to implement new instances of Unified Manager in a NetApp NAS Filer environment.

8.1.1 Hardware

This section details the hardware required when deploying [UM](#):

1. **Storage Controller (Filer)** – if NetApp NAS does not exist in the storage environment then there is no storage infrastructure for UM to monitor and manage.
2. **Windows server(s)** – required to host the UM application. Refer to section [8.2](#) for platform requirements and the UM release notes for supported platforms.
A contingency server is highly recommended for UM.
A separate Windows server is required for the NMC. Note that a dedicated server for the NMC is not required so it can be installed on another server in your environment already performing another function (as long as there are no conflicts in application requirements, however NetApp does not recommend installing the NMC on the same server as Unified Manager.
Refer to section [10.6](#) for the access rights required by the storage administration team to access the Windows Server.

8.1.2 Network Infrastructure

This section is included to ensure the reader is aware of the requirement for an IP-based network infrastructure to connect the aforementioned hardware components and software components detailed in the next section. It is recommended that the Filer and Windows Server both be equipped and provisioned with at least 1GBe connectivity.

For details on the IP ports used for communication between each hardware component refer to section [10.4](#).

8.1.3 Software

This section details the software required when deploying [UM](#). All packages listed below are the evaluation release versions:

Software	Platform	Type	ID	Repository	Status
UM 5.2R1P1	Windows 2008 Server (SP2 and R2) Windows 2012 Server (SP0 and R2SP0)	TPM	CITI_WIN_NetAppOnCommandUnifiedManager_52_00	tpm.nam.nsroot.net	Production Release
OM license keys ³	Windows 2008 Server (SP2 and R2) Windows 2012 Server (SP0 and R2SP0)	n/a	n/a	NetApp	n/a
ONTAP	NetApp NAS	n/a	7.3.6P2, 8.0.2/ 7-mode, 8.1.4Px/ 7-mode	NetApp	Available from NetApp Support Site
NMC 3.3	Windows 2008 Server (SP2 and R2) Windows 2012 Server (SP0 and R2SP0)	TPM	CITI_WIN_x64_NetAppMgmttCo_n_33_01	tpm.nam.nsroot.net	Production Release

Table 4 - Required software

For detailed installation instructions refer to section 8 of the [CATE Release Notes](#).

8.1.4 Configuration Parameters

This section details some additional configuration parameters required when deploying [UM](#). These items require some planning and research prior to implementation and so are included here.

1. Web server SSL key configuration parameters – refer to section [10.3](#) for further information.
2. Installation of Citi Certificate Authority (CA) signed certificate – refer to section [10.3](#) for further information.
3. Functional IDs – refer to section [10.5](#) for further information.

³ Each installation of UM requires separate license keys. Software license keys should be obtained from NetApp as part of the procurement process. For more details on license keys refer to section [8.3.1](#)

8.2 Platform Requirements

The purpose of this section is to outline the platform requirements for the system hosting the [UM](#) application. UM is certified within Citi for deployment on the Wintel server platform, for systems running Windows 2008 server (R2) and Windows 2012 server. If considering a medium or large installation (as defined in [Table 5](#) - UM Management Station Hardware Specification) ensure you have the Enterprise edition of Windows Server 2008/ 2012 installed.

In order to determine the hardware specification of the hosting server for UM, consult the table at this end of this section. Note that CATE Storage and Backup Engineering does **not support** UM on Linux, Windows 2000 or Windows 2003 server platforms.

Based on high level testing, and in conjunction with NetApp's documentation, CATE Storage and Backup Engineering recommends consulting the current hardware Bill of Materials on the [Platforms Operating Environment engineering group website](#) to determine the applicable server models compared to the hardware requirements listed in [Table 5](#) - UM Management Station Hardware Specification

Best practices dictate that the UM host server should rely on SAN to mount the application and its database. CATE Storage and Backup Engineering does **not support** the use of iSCSI LUNS for hosting UM. In order to help plan the storage requirements for the UM database and how much disk should be installed and allocated to the hosting server, consult [Table 7](#) at this end of this section.

The following table provides recommendation in terms of Management Station hardware scaling according to the size of the storage environment (devices) being monitored. The values shown below are bigger than those published in NetApp's general UM sizing document ([OnCommand Sizing Guide](#)) as the design specification takes into account known factors specific to the Citi storage environment.

Environment	Small	Medium	Large
Devices monitored	1 - 25	26 - 100	101 - 250 ⁴
OM Hosting server CPU core count ⁵	2	4	8
OM Hosting server memory size	4GB	4GB-8GB	8GB-16GB ⁶

Table 5 - UM Management Station Hardware Specification

⁴ Publicly NetApp does not support greater than 250 devices per OM server. CATE storage engineering is recommending an infrastructure review if the size of your environment is beyond 250 storage controllers.

⁵ CPU resources are listed in terms of cores and not processors due to the licensing scheme for the embedded Sybase component of UM. The database can leverage up to a maximum of four cores on the Windows platform. The other services in the UM application can leverage the remaining CPUs in the system.

⁶ The recommended starting size is 8GB of RAM. Memory is not usually the bottleneck with UM the move to 16GB should only be undertaken once the performance of the UM server has been reviewed by NetApp and the subsequent recommendation is to increase memory. Expanding the available memory to 16GB yields a diminishing performance improvement in the UM server.

The following table describes the recommended configuration when using a VMware Windows Guest to host an instance of Unified Manager with Performance Advisor enabled⁷.

Devices monitored	1 - 50
OM Hosting server CPU core count	2
OM Hosting server memory size	10GB

Table 6 - OM Management Station VMware Windows Guest Specification for Performance Advisor

The above tables are missing one vital piece of information – the amount of storage required by the management station to support UM. This requirement can be estimated using the following set of formulas. These calculations involve more than just dependencies on the number of monitored devices and hence why they are presented separate to the above table.

Unified Manager Core Database	
Space required	Current year = $((\#hosts * 2.25MB) + (\#vols * 0.098MB) - (\#aggrs * 0.27MB)) * 1.25 * 9$ Historical data = $((\#hosts * 2.25MB) + (\#vols * 0.098MB) - (\#aggrs * 0.27MB)) * 1.25$ per year
Example	Number of systems: 50 Number of volumes: 500 Number of aggregates: 120 Database size at end of current year: 1.42GB $((\#(50 * 2.25MB) + (500 * 0.098MB) - (120 * 0.27MB)) * 1.25) * 9$ At the end of the second year the expected size of the environment is: Number of systems: 75 Number of volumes: 800 Number of aggregates: 150 Database size at end of second year is the size of the historical data from year 1 plus the current year data for the end of year 2: 2.40GB $((\#(50 * 2.25MB) + (500 * 0.098MB) - (120 * 0.27MB)) * 1.25) + ((\#(75 * 2.25MB) + (800 * 0.098MB) - (150 * 0.27MB)) * 1.25) * 9$
Notes	Storage administrators are recommended to use the UM storage requirement estimator tool
Performance Advisor Data Files	
Space required ⁸	Varies according to the configuration of each Filer being monitored
Notes	The calculation is fairly complex and there is little to be gained in explaining it in this document. Storage administrators are recommended to use the UM storage requirement estimator tool

Table 7 - UM Management Station Hardware Storage Specification

In order to ascertain to total storage required for UM add the space required for the UM database and the space required for the Performance Advisor data files.

In order to simplify the calculation process CATE Storage and Backup Engineering has devised a spreadsheet template to estimate the storage requirement, available from http://catecollaboration.citigroup.net/domains/platstor/gdse/products/platforms/NASFiler/OM_database_growth_estimator_v6.xls.

In addition to the space calculations provided by version six of the growth estimation spreadsheet the following working state and database restore state scenarios need to be considered⁹.

⁷ If a VMware environment is not available for hosting then this specification can also be used for a physical server.

⁸ Assumes no custom views have been created. Custom views will increase storage requirements.

⁹ These will be added to a later version of the growth estimation spreadsheet tool.

During normal operation UM requires a minimum of 15% free space (as defined by the configuration parameter [monMinFreePercent](#)) on the disk partition where the UM database file is stored. Use `'dfm database get dbdir'` to see the partition and path location information of the UM database file.

During a restore operation a significant amount of additional disk space is required: equivalent in size to the `monitordb.db`, `monitordb.log` files and the script-plugins and `perfddata` directories. The rollback directories created during a database restore process may be created on a separate partition on the UM server depending on how you have configured and moved data files around on your UM server. So, for example, if you have moved your performance data directory to a different partition due to space constraints UM will by design still create the rollback area to protect against a failed restore on the same partition as where the UM application is installed (which may not be the same as the current performance data partition). For further details on the database restore process refer to section [9.2.1.6](#).

8.2.1 Hosting UM on VMware

Both NetApp and CATE support the installation of [UM](#) on VMware for systems running ESXi 5 or above. It should be noted that UM is a fairly CPU, memory and disk I/O intensive application that may exhibit a performance profile that will require a LVM (Large VM) type configuration, especially when considering this NAS monitoring VM will require SAN storage.

However when considering deploying Performance Advisor enabled instances of Unified Manager in your environment the recommendation is to use a Window VMware Guest where at all possible. This is because the Performance Advisor architecture discussed in section [9.1.1](#) is scaled to a smaller size population of Filers per UM server and is using less functionality from the complete UM application thereby exhibiting less load on the underlying hosting platform. Consult [Table 6](#) for the VMware Windows Guest specification for a Performance Advisor enabled instance of UM.



NetApp OnCommand Unified Manager is used to monitor NAS. Therefore, if provisioning this application on a virtual machine, the storage for this VM must be on SAN to ensure that circular dependencies on infrastructure components do not exist.

8.3 Installation Procedure

The purpose of this section is to present an overview of the steps involved when installing [UM](#). A detailed set of installation instructions for new deployments can be found in section 8.1 of the [UM 5.2 Release Notes](#). For instructions on how to perform an upgrade from UM 5.2P2, refer to section 8.2 of the [UM 5.2 Release Notes](#).

UM has been packaged for deployment via Tivoli Provisioning Manager (TPM).

The installation of this package requires a server reboot, and this needs to be considered as part of your deployment plan.

Once the installation and configuration is completed, you can access the application interface using the following URL format:

For access to OnCommand Console: https://your_full_host_name/start.html

For access to Operations Manager: https://your_full_host_name/dfm/report/

For example if you have installed the package on the server nasdfm01.nam.nsroot.net, the URL would be:

For access to OnCommand Console: <https://nasdfm01.nam.nsroot.net/start.html>

For access to Operations Manager: <https://nasdfm01.nam.nsroot.net/dfm/report/>

Please refer to [Appendix](#) for more details on OnCommand Console.

8.3.1 License Keys

[UM](#) supports the use of difference license keys to enable to use of certain licensed features within the product. Based on the features currently certified and supported by CATE Storage and Backup Engineering, the following license keys are required as part of the installation process for a new system:

1. OnCommand Unified Manager core license

Note that the Unified Manager license includes the ability to run Performance Advisor component in the [NMC](#).

These license keys should be obtained from NetApp via your local technology procurement process and form part of the formal Volume Pricing Agreement (VPA) between NetApp and Citi. License keys vary per OM instance so it is not possible to use a global license key for all installations of OM within Citi.



The core UM license key is inputted during the TPM installation process (see section 8 of the [CATE UM Release Notes](#)). There is no requirement to update the license keys in UM if you are upgrading from an earlier version of the product. License keys only need to be installed if you are building a new instance of UM. The process to install license keys is defined below.

License keys can be installed in one of two ways: a) via the CLI on the [management station](#) or b) via the web UI. Both methods are briefly presented here.

- a) Installing license keys via the CLI
 - i. **RDP** (Terminal server) to management station and login to Windows
 - ii. On the Windows CLI type the '**dfm license install <license_key>**' command (replacing the term <license_key> with the actual key value provided by the Vendor) and press **ENTER**
 - iii. Confirm the licensed feature is installed with the '**dfm license list**' command. In the case of the Business Continuity option this will be listed as 'bizcon'.
- b) Installing license keys via the GUI
 - i. Point your web browser at your UM server's URL and **login** to the UM UI
 - ii. From UM Dashboard, select '**File**' menu and select '**Operations Manager**'
 - iii. From the '**CONTROL CENTER**' tab, select the '**SETUP**' menu and choose the '**OPTIONS**' menu item
 - iv. From the '**EDIT OPTIONS**' list in the left pane, select the '**LICENSED FEATURES**' option
 - v. In the '**NEW LICENSE KEY**' text box input the license key provided by the Vendor
 - vi. Click the '**UPDATE**' button
 - vii. Once the screen has refreshed the '**BUSINESS CONTINUANCE OPTION**' and the license expiry date should appear in the list



When concurrently monitoring a NetApp storage system with more than one UM server (for example when using the primary reporting and alerting approach with a second instance for Performance Advisor as discussed in section [9.1.1](#)) ensure that each UM server has a different server license key installed. If you try to monitor the same storage system concurrently with two instances of UM both using the same server license key you may obtain spurious data from the storage system or in worse case panic the storage system.



Please use only DFM Server licenses serial numbers starting with 1-50. Serial numbers starting with 1-51 **must not** be inputted when running dfm_inst.ps1 script

8.3.2 Installation Path

The installation procedures specify to install to C:\Program Files\NetApp\ as the default installation directory instead of D:\apps_sr\NetApp\. The application binaries will be installed on the C partition and the configuration script reconfigures the UM server to use D:\Program Files\NetApp\ as the repository for the applications data files.

8.3.3 Installation Performance

The TPM installation has been tested in the CATE Lab using a VMware session on an ESX4i cluster. The virtual machine was provisioned with two vCPU and 4GB of RAM. The installation took approximately 20 minutes.

8.3.4 Default Database Access



During the installation of UM (as opposed to upgrading an older version), the account used by the SA to install the application will be given global (full control) privileges over the database. The installation process cannot remove this account because it is not possible to remove an account while authenticated to the application with that same account. **Storage administrators must delete the install administrator account from UM prior to commissioning UM for production use.**

8.3.5 Configuring Web Server SSL Certificates

The next key step in preparing the Unified Manager application for use is to configure the SSL certificates for the embedded Web Server.

Securing of the web server is required prior to commissioning the application for use in order to adhere to CISS policy on securing authentication data over the network. Refer to section [10.3](#) for the procedure on configuring the SSL certificates for the web server. Certificate creation requires some planning and research to ensure the correct parameters are defined.

With the release of OM 3.8 and above CATE Storage and Backup Engineering is updating the web server SSL configuration procedures in line with updates to the CISS whereby storage administrators will be required to obtain a Certificate Authority signed version of the web server key. Detailed procedures for this process can be found in section [10.3](#).



It is recommended that the SSL configuration parameters are determined and prepared prior to installing and configuring Unified Manager.

8.3.6 Citi Recommended UM Configuration

The next step is configuring the UM application options. UM has over 100 configurable options which fall into four broad categories:

- Application security
- Application behavior / application performance tuning
- Regional specific settings (email server and distribution lists)
- General settings

All of the application options that are relevant to Citi are listed in section [11.2](#). In the appendix the default settings and the recommended setting for each option is listed according to how you are intending to use UM. For further information please consult section [9.6](#).

Starting with UM 5.1 CATE Storage and Backup Engineering is providing a configuration script `dfm_cfg.ps1` in the UM install folder which will automate the configuration of the application and optimize according to the size of your environment and the features you plan to use.

Further information can be found in section [11](#)

8.3.7 Storage Controller Discovery

The next stage is to add storage devices to UM so it can begin monitoring and reporting on these systems. A step by step guide on adding a storage device to UM, including best practices to reduce device communication and connectivity issues, is presented in section [9.7.1](#).

8.3.8 Storage Management Functional ID

The final step in this installation overview chapter is to create a Functional ID on each storage device and configure UM to use this functional ID for all management communication. Refer to step 6 in section [9.7.1](#) for the procedure to configure UM to use this functional ID for storage management communication. For additional information regarding the rights, privileges and requirements of this functional ID, refer to section [10.5](#) and using the example account listed as 'dfmproxy'.

This concludes the high level overview of installing and configuring UM to manage NetApp NAS storage devices. Use the sections referenced in this chapter for a deeper dive on the processes discussed here.

9 Best Practices

The purpose of this chapter is to provide a technical reference and training guide for Storage Administrators using Unified Manager (UM) to report on and manage their NetApp NAS Filer environment. For administrators new to UM it is recommended that you review this entire document. Experienced administrators should use this document as a technical reference. This chapter focuses on many aspects of UM including:

- Technical overview
- Protecting Unified Manager (application backup / replication / server backup)
- Unified Manager maintenance, migration, capacity planning and configuration best practices
- Adding NetApp NAS Filers to UM
- NetApp NAS Storage reporting, replication management, disk-based backup (snapvault) management, and centralized Filer configuration management
- Storage monitoring with performance advisor, events and alarms
- Customizations available through the script plug-in library and command aliasing

For a detailed discussion of security related material regarding UM please refer to section [10](#).

To begin this chapter let's start with an overview of how Unified Manager fits into a NetApp NAS Filer storage environment.

9.1 Implementation Overview

The purpose of this section is to provide a high-level overview of how Unified Manager interacts with a NetApp NAS File environment. The diagram overleaf illustrates the network implementation for [UM](#). Clients (depicted on the left of the diagram) mainly access the application via the web interface. The production UM hosting server polls the storage infrastructure using a combination of SNMP, HTTPS (tunneling the ManageONTAP API), and SSH protocols to assess the health, status and performance of each storage appliance, using predefined parameters in terms of polling frequency and thresholds.

Use of a contingency server is recommended to provide a 'hot-standby' storing a copy of the UM database generated by an offline full application backup (see section [9.2.1](#)). Should a critical failure occur on the production [management station](#), it is then possible to restore the database from the backup copy on the contingency management station.

This diagram has been simplified for demonstration purposes only. The main requirement is that there is at least a production and contingency instance of UM per NAS environment. An environment is usually considered as a collection of systems supported by a single support group. The diagram **does not imply** that you need one UM server per production site; nor does it imply that the UM server needs to be installed at the same site as your production systems (although installing at the same site as the majority of your production systems would be beneficial from a network perspective); nor does it imply that you can only use UM in environments that only have two sites. The diagram is provided simply to demonstrate how the UM application interacts with the NAS Filer devices.




Figure 1 - UM physical implementation

9.1.1 Recommended architecture with UM and Performance Advisor in medium environments

The purpose of this section is to present some of the challenges involved in continuing to use UM in the traditional implementation “one production, one contingency” instance per geographic region considering the increased functionality added to the product with modules such as Performance Advisor.

Medium size environments may face some scale issues when considering trying to monitor and manage all storage devices with a single instance of UM. NetApp currently lacks robust tooling that would allow multiple instances of UM to be combined for monitoring, management and reporting purposes into a single screen.

Essentially what differentiates a UM server from a Performance Advisor server is the UM option setting that enables the Performance Advisor module. All of the other UM settings essentially remain the same however an UM server with Performance Advisor enabled is recommended to be used purely for performance monitoring only and not for other UM management, alerting and reporting functionality.

To assist in strategic data center consolidation activities the recommended configuration for a Performance Advisor server is described in [Table 6](#).

The following recommendation should be used in medium environments that contain more than 100 devices to be monitored but less than 200:

- A central event monitoring, management and reporting UM management station
This is the “traditional” UM implementation where all storage devices are managed by a single UM instance. It allows all events, alerts and reports to be handled by a single instance. It simplifies the reporting and management aspects of the environment as all the relevant information is contained within a single database.
- A contingency instance of the central management station
For obvious reasons it is recommended that this contingency management station be installed at a different location to the primary (production) instance.
- One instance of UM per strategic data center (usually at a location that houses many storage devices). This instance of UM is specifically configured to disable most of the event management and reporting functionality, thereby focusing on the collection and recording of performance advisor data for storage systems at that particular site.

There are a number of benefits to this approach:

1. Isolating performance advisor data capture to devices at the local site reduces the impact of WAN latencies on the data gathering process that can impact the data sampling reliability.
2. Splitting the storage requirements for performance advisor into smaller subsets reduces the storage required by the main UM server and will also shorten the time it takes to perform a database backup on the main UM server (while performance advisor data is not stored in the database itself, the performance advisor data files are included as part of the database backup process).
3. Performance advisor dedicated smaller instances of UM has lower requirements in terms of hosting hardware, therefore making the host systems candidates for virtualization with VMware.

There are also a number of minor disadvantages to this approach that should be considered as part of any implementation:

1. The NetApp Management Console can only connect to one UM server at a time. Therefore if you wanted to interactively monitor a particular application (such as VMware on NFS) across all sites, you would need a separate console session per site. This would only affect real time viewing of the performance data.
2. More instances of UM to manage. Increases management overhead when performing RCSI activities and checking UM database backups have completed successfully. When adding a new storage device to a site you would be required to add it to the central UM instance as well as the site based instance capturing the performance advisor data.



Refer to section 8.3.1 for special licensing considerations when deploying a dedicated performance advisor environment in an infrastructure where you already have Unified Manager deployed or are intending to concurrently monitor the same NetApp storage system with more than one UM server.

This approach is demonstrated in the diagram on the following page.




Figure 2 - Performance Advisor Architectural Reference (medium environments)

For environments that have a significant number of distributed remote sites where it is not viable or practical to implement a performance advisor server at these sites, determine the network topology back to the strategic data centers in the region and place your Performance Advisor enabled UM instance within the data center. This is demonstrated above in the relationship between sites C and D.

9.1.2 Recommended architecture with UM and Performance Advisor in large environments

Large environments of 200 NetApp Filers or more will face significant scale out issues when considering trying to monitor and manage all storage devices with a single instance of UM. NetApp currently lacks robust tooling that would allow multiple instances of UM to be combined for monitoring, management and reporting purposes into a single screen.

Essentially the roles are redistributed according to the geographical and functional distributed model presented in the [NetApp Distributed OM Server Strategy document](#). It introduces the concept of a regional master reporting server whose primary function is to provide a single reporting repository for the regional environment beyond the NetApp publically supported limit of 250 storage controllers in a single Unified Manager instance. In this architecture many of the monitoring features are dialed down or disabled.

The Regional Master Reporting Server concept is combined with enhanced performance advisor enabled site specific instances, as used in the medium sized environment. The difference in large environments is these distributed instances take on the proactive monitoring and alerting functionality that was traditional the scope of the regional master server. The configuration parameters required to achieve this architecture are automated in the CATE Unified Manager configuration script and are detailed in section [11.4](#).




Figure 3 - Geographical and Functional Architectural Reference for Performance Advisor (large environments)

This concludes the discussion on recommended deployment methodologies for Unified Manager and Performance Advisor in large environments. Additional recommendations may be added to this section as and when they are identified and deemed viable in the Citi environment.

9.2 UM Application Protection

The section discusses the steps involved in protecting Unified Manager through the use of database backup technologies, management station backup, management station disaster recovery and graceful shutdown procedures.

9.2.1 UM Database Backup

This section discusses what is backed-up as part of a database backup, which UM functionality is and is not available during a database backup and describes the steps involved in configuring a regular scheduled backup as well as the procedure for performing an immediate database backup.

The database backup process captures the following data on the Management Server:

- Database
- Custom reports
- Script plug-ins
- Performance data

The database backup process **does not** capture the following data on the Management Server:

- Configuration plug-in files

Restoration of configuration plug-in files can be achieved via a different method. Refer to section [9.9.14](#) for this procedure.

You can still use UM and it will continue to monitor your environment while the backup is running. Operations that require a write to the database will be queued in a redo log until the locks on the database are released. Once the backup completes the redo log will replay until all database requests have caught up. All application services (DFMhttp, DFMwebui, DFMSybase, DFMEvent, DFMMonitor, DFMScheduler and DFMServer) maintain a 'running state' during a database backup. In earlier versions of Operations Manager (pre 3.5) these services were stopped / paused during a database backup.



However, data collection and view modifications of Performance Advisor come to a halt during the backup process. If you are using Performance Advisor to monitor a Filer which traditionally exhibits poor performance overnight, you should configure the database backup to happen at a time that does not conflict with the high workload period on the Filer you are monitoring. For all other systems that are being monitored by Performance Advisor you should correlate any gaps in data collection with the UM event log for database backup (look for the 'Info' event called 'Database Backup Succeeded').

A typical database backup can take between 15 and 90 minutes depending on size of the database and the load on the [Management Station](#). Backups are written to \data\dfm_temp_dest_dir and then the final backup file is copied to \data.

9.2.1.1 Database backup retention policy

OM 3.7 introduced the ability to define a retention policy in terms of how many database backups are kept on the [Management Station](#). This policy simplifies the house keeping of the Management Station and prevents the system from inadvertently running out of disk space. The default and Citi recommended values for this policy are defined in the [backupRetentionCount](#) UM option (click on the link to go to section [11.1](#) for the details).

When using the backup retention count in UM, when the threshold is reached the oldest database backup file is deleted before the new backup process begins. Ensure you forecast for enough space on the [Management Station](#) to accommodate the amount of database backup files defined by the retention policy. UM offers some basic space forecasting tools to assist with this task to determine how many

backups can be kept on the Management Station. Refer to [Figure 8](#) for the forecasting information via the GUI and section [9.2.1.3](#) when using the CLI.



The retention policy applies to both manual and scheduled backups. Therefore if your retention policy is set to five and five manual backups are initiated, all the scheduled database backups will be deleted. If you want to keep a certain database backup the recommended course of action is to copy the ndb archive file to a different location on the [Management Station](#).

To configure the database backup retention in the GUI use the menu bar: CONTROL CENTER / SETUP / DATABASE BACKUP. Use the following steps to setup the retention policy:

1. Click on the Backup Retention Count value (shown below by the value 'No Limit')

Figure 4 - Database backup retention policy

2. Enter the value of the 'Database Backup Retention Count'
3. Click on the 'Update' button

Figure 5 - Database backup options screen

4. A confirmation message will be received when the backup retention policy has been applied




Figure 6 - Database Retention Policy Applied Confirmation Screen

This completes the process to define the UM database backup retention policy via the GUI.

To configure the same database backup retention policy using the CLI on the [Management Station](#), use the 'dfm option set' command:

```
C:\>dfm option set backupRetentionCount=5
Changed database backup retention count to 5.
```

9.2.1.2 Scheduled Database Backup via GUI

To configure a database backup schedule in the GUI use the menu bar: CONTROL CENTER / SETUP / DATABASE BACKUP. Use the following steps to setup a database backup schedule:

1. Select backup type '**archive**' radio button to create a compressed form database backup file. CATE Storage and Backup Engineering does not support the use of 'snapshot' backup type as it requires the database to be stored on iSCSI or FiberChannel NetApp storage, which is an unsupported configuration in Citi.
2. Select '**enable schedule**' check box (this leverages the DFM scheduler service not the native Windows Scheduler service).
3. Select the '**daily**' radio button.
4. Enter a **time** (in hh:mm format between the two text boxes) to define the time at which the backup will start.



CATE Storage and Backup Engineering does not recommend scheduling the backup of the UM database during core business hours due to the increased load placed on the server during the backup process. If you are running Performance Advisor in your environment [Performance Advisor will be unavailable until the database backup completes.](#)

5. Click on the '**update schedule**' button to submit your backup scheduling request.

Figure 7 - Database Backup Screen

6. Click on the 'update' button to confirm your backup scheduling request.

Figure 8 - Database Backup Schedule Confirmation Screen

7. A confirmation message will be received when the backup schedule has been applied.

Figure 9 - Database Backup Schedule Applied Confirmation Screen

9.2.1.3 Scheduled Database Backup via CLI

To configure the same database backup schedule using the CLI on the [Management Station](#), use the 'dfm backup schedule' command:

```
C:\>dfm backup schedule set -t archive -D 02:00
Backups will be created daily at 02:00
Available space = 26.3 GB
Estimated space required = 1.58 GB
Can hold around 16 daily backups of 1.58 GB each.
```



CATE Storage and Backup Engineering does not recommend scheduling the backup of the UM database during core business hours due to the increased load placed on the server during the backup process. If you are running Performance Advisor in your environment [Performance Advisor will be unavailable until the database backup completes](#).

9.2.1.4 Manual (Ad-hoc) Database Backup via GUI

Outside of the scheduled database backups refer to section [9.3](#) for situations where a manual database backup is recommended.



CATE Storage and Backup Engineering does not recommend performing a manual backup of the UM database during core business hours due to the increased load placed on the server during the backup process. If you are running Performance Advisor in your environment [Performance Advisor will be unavailable until the database backup completes](#).

To configure a manual database backup in the GUI use the menu bar: CONTROL CENTER / SETUP / DATABASE BACKUP. Use the following steps to setup a database backup schedule:

1. Click the '**backup now**' button to submit the manual backup job.

Figure 10 - Database Backup Screen

2. Click the '**Run Backup**' button to start the manual backup operation.

Figure 11 - Database Backup Confirmation Screen

3. A confirmation message will be received when the backup has begun.




Figure 12 - Database Manual Backup Start Confirmation Screen

9.2.1.5 Manual (Ad-hoc) Database Backup via CLI

To configure the same database manual backup using the CLI on the [Management Station](#), use the 'dfm backup create' command.



It is not recommend performing a manual backup of the UM database during core business hours due to the increased load placed on the server during the backup process. If you are running Performance Advisor in your environment [Performance Advisor will be unavailable until the database backup completes](#).

```
C:\>dfm backup create -t archive
Disabling Performance Advisor.
Adaptive Server Anywhere Backup Utility Version 9.0.2.3396
(94682 of 94682 pages, 100% complete)
(10627 of 10627 pages, 100% complete)
Database backup completed
Adaptive Server Anywhere Stop Engine Utility Version 9.0.2.3396
Adaptive Server Anywhere Stop Engine Utility Version 9.0.2.3396
Waiting for existing writes to stop, which could take a few minutes.
Please be patient ...
Creating the archive ...
Verifying the archive ...
Enabling Performance Advisor.
Completed backup
D:/apps_srv/NetApp/DataFabric/DFM/data/dfm_backup_2008-07-28_19-57-32.ndb.
```

9.2.1.6 Performing a Database Restore

Like any backup technology, detailed instructions and best practices regarding database backup are worthless without the procedures to restore a backup when required. A database restore will also overwrite the performance data repository of Performance Advisor.



During a database restore operation none of the functions of Unified Manager are available.

- The web server and CLI will report the database is unavailable.
- No scripts will execute and no performance data will be collected from the Filers.
- Performance advisor and other features in the [NMC](#) will be unavailable.

The restoration function of the UM database is only available using the CLI on the [Management Station](#). Use the '**dfm backup list**' command to determine the backup files available for restore and then the '**dfm backup restore**' command:

```
C:\>dfm backup list
```

Backup Name	Backup Size	Backup Time
dfm_backup_2008-07-30_03-00-01.ndb	345 MB	30 Jul 03:13
dfm_backup_2008-07-29_03-00-04.ndb	344 MB	29 Jul 18:40



A database restore operations requires a significant amount of space. If there is not enough space on the UM server the restore operation will rollback the database to the point just prior to when the restore was requested.

Depending on how UM is configured in terms of which disk partition application data is stored while determine where, if any, disk space resource contention will occur during a restore procedure. The following list explains the high level process of a restore operation:

- Rollback directory of the current UM Performance Advisor data
 - Created in the UM application directory (D:\Program Files\NetApp\DataFabric Manager\DFM, although your location may vary for legacy reasons) and called tmp.perf.rstr
 - Equivalent in size to the Performance Data directory defined by `dfm option get perfarchivedir`
- Rollback directory of the current UM Script-Plugin data
 - Created in the UM application directory (D:\Program Files\NetApp\DataFabric Manager\DFM, although your location may vary for legacy reasons) and called tmp.script.rstr
 - Equivalent in size to the Performance Data directory defined by `dfm option get scriptdir`
- Rollback versions of the current database and database log files
 - Created in the UM database directory defined by `dfm database get dbDir`
 - Equivalent in size to the monitordb.db and monitordb.log files prior to commencing the restore operation
 - Database rollback files called tmp.db.rstr and tmp.log.rstr

In the following example the backup from 2008-07-29_03-00-04 will be restored.



The command in the following example assumes the database backup file to restore exists in the /data directory of UM. Use the command in section [9.2.2.3](#) to determine the location of the data directory on your Management Station. If the file to restore is stored in a different location you can either copy the file to the data directory or specify the exact path as part of the restore command.

```
C:\>dfm backup restore dfm_backup_2008-07-29_03-00-04.ndb
Please do not interrupt this operation with CTRL-C.
155 File(s) copied
3 File(s) copied
Upgrading database to latest schema.
No database upgrade is necessary.
Redefining SQL Views ...
Restored from backup dfm_backup_2008-07-29_03-00-04.ndb.
Completed restore from dfm_backup_2008-07-29_03-00-04.ndb.
```

This completes the process to perform an UM database restore via the CLI.

The number of files restored will vary depending on the size of your environment, whether performance advisor is enabled and the amount of time that has elapsed between when the backup was originally taken and when the restore is taking place.

The following items are examples of the responses to [Management Station](#) requests made during a database restore operation:

1. If issuing a command from the Management Station CLI the following error is returned:

```
Error: The DataFabric Manager database is temporarily unavailable;
error code is -100.
Error: Please try again later.
```

2. If making a request to the Management Station web server the following error is returned:




Figure 13 - GUI error message during database restore

3. If attempting to connect to the Management Station via the [NMC](#), or during an existing NMC session, the NMC will return the following error:




Figure 14 - NMC error message during database restore

9.2.1.7 Checking Status of Database Backup

It is possible to check the status of a backup job from both the GUI and the CLI.

To check the status of a database backup in the GUI use the menu bar: CONTROL CENTER / SETUP / DATABASE BACKUP. Use the following steps to check the database backup status:

1. Confirm the 'status' shown in the database backup manager window.




Figure 15 - Confirming the backup job status

This completes the process to check the status of an UM database backup via the GUI.

To perform the same database backup status check using the CLI on the [Management Station](#), use the 'dfm backup status' command:

```
C:\>dfm backup status
Backup status: Running.
```

9.2.1.8 Aborting In-progress Backup

It is possible to abort a running database backup job from both the GUI and the CLI.

To abort a database backup in the GUI use the menu bar: CONTROL CENTER / SETUP / DATABASE BACKUP. Use the following steps to abort a database backup:

1. Click on the '**abort**' link shown in the database backup manager window.



Use the abort command with care. There is no confirmation pop-up dialogue box once you click on the abort link. Termination of the current running database backup job is immediate.




Figure 16 - Aborting the current database backup job

2. A confirmation message will be received when the backup has aborted.




Figure 17 - Backup abort confirmation message

This completes the process to abort the current UM database backup via the GUI.

Aborting a backup will remove the partially created backup archive file. However if at the start of the aborted backup the oldest backup file was deleted because the backup file count retention policy had been exceeded, the deleted backup file will not be recovered. The deletion process is permanent.

To perform the same database backup abort using the CLI on the [Management Station](#), use the 'dfm backup abort' command:

```
C:\>dfm backup abort
Database backup abort operation completed successfully.
```

9.2.1.9 Listing Database Backups

It is possible to obtain a list of all the database backups stored on the [Management Station](#) from both the GUI and the CLI.

To obtain a list of database backups in the GUI use the menu bar: CONTROL CENTER / SETUP / DATABASE BACKUP. Use the following steps to list the database backup files:

1. Click on the 'List Database Backups' link shown in the database backup manager window.




Figure 18 - Listing the current database backups

2. The list of database backups stored on the Management Station is then shown.




Figure 19 - Database backup list

This completes the process of listing the UM database backups via the GUI.

To perform the same database backup listing using the CLI on the [Management Station](#), use the 'dfm backup list' command:

```
C:\>dfm backup list
```

Backup Name	Backup Size	Backup Time
dfm_backup_2008-07-28_20-25-45.ndb	342 MB	28 Jul 20:35
dfm_backup_2008-07-28_19-57-32.ndb	342 MB	28 Jul 20:08
dfm_backup_2008-07-27_19-00-03.ndb	340 MB	27 Jul 19:09

9.2.1.10 Deleting Database Backups



Before reading this section, be sure you have read and implemented the UM Database Backup retention policy as discussed in section [9.2.1.1](#). With a suitable database backup retention policy implemented you should not have to manually delete any database backups. The procedure shown below is presented for reference purposes only.



Database backup files which are deleted by using either the GUI or CLI processes listed below, or those backups automatically deleted by a retention policy, are permanently deleted. As such they cannot be restored from the 'recycle bin' on the [Management Station](#).

It is possible to delete any or all of the database backups stored on the [Management Station](#) from both the GUI and the CLI.

To delete database backups in the GUI use the menu bar: CONTROL CENTER / SETUP / DATABASE BACKUP. Use the following steps to delete the database backup files:

1. Click on the 'List Database Backups' link shown in the database backup manager window.




Figure 20 - Listing the current database backups

2. Place a check mark next to the backups you want to delete, or alternately if you wish to select all the backups for deletion use the 'select all' link at the bottom of the screen.
3. Click the 'Delete Selected' button when you are ready to delete the chosen files.




Figure 21 - Backup deletion selection

- You will be prompted to confirm your decision to delete database backups. Click 'OK' to proceed.




Figure 22 - Backup deletion confirmation screen

- A confirmation message will be received when the backups have been deleted.




Figure 23 - Backup deletion acknowledgement screen

This completes the process of deleting the OM database backups via the GUI.

To perform the same database backup deletion using the CLI on the [Management Station](#), use the 'dfm backup list' command to determine the available backup files and then use the 'dfm backup delete' command to delete the backup files:

```
C:\>dfm backup list
Backup Name                               Backup Size Backup Time
-----
dfm_backup_2008-07-29_03-00-04.ndb        344 MB      29 Jul 03:14
dfm_backup_2008-07-28_19-57-32.ndb        342 MB      28 Jul 20:08

C:\>dfm backup delete dfm_backup_2008-07-29_03-00-04.ndb
dfm_backup_2008-07-28_19-57-32.ndb
Deleted backup dfm_backup_2008-07-29_03-00-04.ndb.
Deleted backup dfm_backup_2008-07-28_19-57-32.ndb.
```



Note that you can delete multiple database backup files from the command line by separating the name of each file with a ' ' (space) character. In order to successfully delete the database backups using the CLI you need to specify the full file name including the file type extension.

This completes the process of deleting the UM database backups via the CLI.

9.2.1.11 Protecting Configuration Management Files

The purpose of this section is to briefly offer some best practices regarding the restoration of Configuration Files that have been temporarily deleted from the database. Instead of restoring the entire UM database, ensure that both the Filer and vFiler configuration templates are exported from the UM server and stored on a file system different to the one on the OM server.

A missing configuration file can be restored by re-importing the configuration export file back into the database. Refer to section [9.9.8](#) for the procedures to export and import configuration files in UM.

This concludes the brief discussion on additional ways of backing up Configuration Management Files.

9.2.2 Unified Manager Server Backup

Unified Manager Server ([Management Station](#)) backup should be used in conjunction with the application database backup (discussed in section [9.2.1](#)) and application DR (discussed in section [9.2.3](#)) as part of a suite of tools and processes to protect the storage infrastructure management application.

- Database backup is there for immediate recovery in the event of an application issue assuming that the rest of the management infrastructure remains intact.
- Application DR via the transfer of the database backup to a remote system is there to provide a quick Recovery Time Objective (RTO) in the event of a management infrastructure failure.
- Server based backup provides the opportunity to restore the database via independent backup infrastructure. This allows for the restoration of the database backup when the Recovery Point Objective (RPO) is further back in time than the earliest database backup stored on the Management Station.

The next section discusses the justification for a server backup.

9.2.2.1 Why perform a server backup?

Using the database retention policy in UM (see section [9.2.1.1](#)) only a predefined amount of recent database backups will be stored on the Management Station. For longer RPOs restoration via a backup infrastructure is required, however it is important not to treat backup as an archive. The retention periods for Management Station backup discussed later in this section may appear extremely aggressive but before discussing the server backup retention period it is appropriate to present some of the background information from which the server backup retention policy was defined:

1. The database can grow to a significant size (20GB) and maintaining tens of instances of the database backup on the local disk of the Management Station is not the best use of primary storage. This is why the database retention policy ([BackupRetentionCount](#)) is relatively small.
2. Due to the transient and time sensitive nature of the data in many infrastructure management applications, the RPO is usually very recent especially when recovering the database could mean losing many samples of performance data. Therefore most restore operations will be serviced from the local database backups as opposed to an independent backup infrastructure.
3. There may be an occasion where a database problem with UM is not realized until after the database retention policy has deleted the local backup, so an alternate source is required from which to recover the database. It is difficult to anticipate how long it may take for a problem to be realized but a period of 30 days is assumed to be sufficient if the vendor cannot address the issue with the database by some other means of intervention (be it a software patch or some form of privileged command).

This section has discussed the high level justification for a server based backup. The purpose of the next section is to define the retention period for the server backup.

9.2.2.2 Recommended Server Backup Retention Period

CATE Storage and Backup Engineering recommend the retention periods outlined below when backing up the Unified Manager server ([Management Station](#)). Refer to the next section for information on which drives and directories on the Management Station require backing up.

- Monthly full back-up – retention period 1 (one) year
- Daily incremental – retention period 30 (thirty) days

9.2.2.3 Directories for inclusion in Server Backup Policy

The purpose of this section is to provide [MSSAs](#) with the necessary information to submit a request to implement a backup policy for the Unified Manager server. Only the application database backup directory needs to be backed up. In the event of a catastrophic failure of the Unified Manager server, the expectation is to invoke DR for the application. Refer to section [9.2.3](#) for details on how to invoke and failback from DR. Server backup is recommended in case you need to restore an old version of the database to recover information that is no longer available in the production database instance.

The critical data that requires backing up on the Unified Manager server ([Management Station](#)) can be found in the application data directory. If you are unsure of the location of the data directory on the Management Station use the 'dfm option list databaseBackupDir' command on the CLI of the Management Station to determine its location.

```
C:\>dfm option list databaseBackupDir
Option          Value
-----
databaseBackupDir D:\Program Files\NetApp\DataFabric\DFM\data
```

No other directories require backing up. The next section explains the process involved in restoring the database using the backup infrastructure.

9.2.2.4 Restoring Unified Manager application from Backup Infrastructure

The purpose of this section is to document the database restore process when restoring the database using the backup infrastructure.

The database backup file(s) contains the critical information needed to restore Unified Manager in the event of a fatal application error or the need to confirm the contents of the database at a given historical moment. The contents of the backup file are explained in section [9.2.1](#).

Before detailing the restoration process the following assumptions / pre-requisites should be noted:

1. The Windows server you are restoring to should be built with the same O/S family as the system from which the backup was originally taken (for example if the source system was Windows 2008, the restore system should be Windows 2008). Differences in the Windows service pack are not applicable.
2. The Windows server you are restoring to must be in the same Windows Domain as the system from which the backup was originally taken.
3. This Windows server is installed with the same or later TPM version of Unified Manager as the system from which the backup was originally taken.
4. The backup infrastructure client is installed on the Windows Server and the necessary permissions have been granted in the backup infrastructure if performing a redirected restore (restoring the backup to a different host from which the backup was originally taken)

That completes the list of pre-requisites and assumption about performing a database restore from tape. The following process details the steps necessary to perform the restore operation.

1. Stop UM services using the 'dfm service stop' command (refer to section [9.2.4](#) for details)
2. Restore the require database backup file from tape to the /data directory on the [Management Station](#)
3. Start UM services using the 'dfm service start' command (refer to section [9.2.4](#) for details)

4. Perform an UM database restore using either the CLI process (refer to section [9.2.1.6](#) for details)
5. Regenerate the UM web server SSL keys (refer to section [10.3](#) for details)
6. The restored version of Unified Manager is now ready for use.

This completes the process of restoring the Unified Manager database from tape.

9.2.3 Unified Manager Server Contingency (DR)

Due to the increasingly important role of Unified Manager in not only reporting but actively managing the NetApp NAS environment through automation and scripts, there is a greater need for DR than in earlier versions of UM. The contingency model uses a secondary (hot standby) server, to which the backup of the current production UM database is replicated. A basic file copy transfer process will suffice for this replication activity.

The sections within this chapter present the processes for invoking contingency for the [Management Station](#) and then failing back to production. It also presents some recommendations on an optimized contingency configuration for large environments that are distributed across multiple data center locations. The following diagram provides an overview of the contingency model for Unified Manager:




Figure 24 - Unified Manager Contingency Model Overview

The following pre-requisites are required in order for a successful implementation of application contingency for Unified Manager.

1. A second (DR) Windows Server is available at a different site to the production Windows Server hosting Unified Manager. The DR server should have the following attributes:
 - a. Built with the same O/S family as the production system (for example if the production system was Windows 2008, the DR system should be Windows 2008). Differences in the Windows service pack are not applicable.
 - b. Built with the same disk partitions as the production system. For example the recommended installation partition for UM on Windows is D:. Ensure the DR system has a D: partition.
 - c. Built with the same Citi TPM of Unified Manager as the production system.
 - d. A member of the same Windows Active Directory domain as the production system.
 - e. Shares the same common DNS infrastructure as the production system.

- f. Has IP connectivity to all the storage systems as the production system utilizing ports 443 for HTTPS tunneled API calls and 22 for SSH connectivity.
 - g. All UM application services are in a stopped state on the contingency server (refer to section [9.2.4](#) for the procedure to stop services)
2. There is an automated mechanism for copying database backups from the production to contingency servers (see section [9.2.3.6](#) for details).

9.2.3.1 DR server state during normal business operations

The purpose of this section is to briefly describe the recommended state of the contingency [Management Station](#) when the production Management Station is working as expect and DR invocation is not required.

During normal business operations the Unified Manager services on the contingency Management Station must be kept in a “stopped” state. This will prevent storage systems in the infrastructure from being over burdened by two Management Stations querying performance and status.

- Use the procedure listed in section [9.2.4](#) to list the status of the UM services
- If the services are shown in a running state use the other procedure listed in section [9.2.4](#) to stop the UM services running on the contingency Management Station

9.2.3.2 Invoking DR

This section details the process for invoking DR of the Unified Manager application:

1. If possible perform a database backup of the production system (using the either of the procedures outlined in sections [9.2.1.2](#) or [9.2.1.3](#)) and copy the database backup to the data directory on the contingency server (use the procedure outlined in section [9.2.2.3](#) to determine the location of the data directory). Note that invoking DR implies catastrophic infrastructure failure at the production site so this step may be a moot point and can be skipped if irrelevant.
2. Start the Unified Manager services on the contingency server using the service start procedure in section [9.2.4](#).
3. Restore the most recent version of the production database backup to the contingency server using the procedure detailed in section [9.2.1.6](#).
4. Generate a new set of Unified Manager Web Server SSL keys using the procedure described in section [10.3](#). These cannot be configured in advanced as they are overwritten with data from the production server. The keys need to be recreated to avoid host key and server name conflicts.
5. Update the value of the [localHostName](#) option within Unified Manager using the following command on the CLI of the contingency [Management Station](#). This will ensure any email alerts generated by the contingency system contain the correct host name using the Windows system variable %COMPUTERNAME%.

```
C:\>dfm option set localHostName=%COMPUTERNAME%
```

6. Your contingency Management Station is now ready for use.

9.2.3.3 Failing back to production

This section presents the steps required to revert Unified Manager to production once the need for DR has ceased. The assumptions and pre-requisites listed in section [9.2.3](#) also apply to the failback to production.

1. Perform a database backup of the DR system (using the either of the procedures outlined in sections [9.2.1.2](#) or [9.2.1.3](#)).
2. Copy the database backup file to the data directory on the production server (use the procedure outlined in section [9.2.2.3](#) to determine the location of the data directory).
3. Ensure all UM application services are in a started state on the production server (refer to section [9.2.4](#) for the procedure to list the status of services).
4. Restore the database backup file copied in step 2 on the production server using the procedure detailed in section [9.2.1.6](#).
5. Generate a new set of Unified Manager Web Server SSL keys using the procedure described in section [10.3](#).
6. Update the value of the [localHostName](#) option within Unified Manager using the following command on the CLI of the production [Management Station](#). This will ensure any email alerts generated by the production system contain the correct host name using the Windows system variable %COMPUTERNAME%.

```
C:\>dfm option set localHostName=%COMPUTERNAME%
```

7. Your production [Management Station](#) is now ready for use once more.
8. Once checkout of the production Management Station has been completed ensure the Unified Manager services on the contingency system are returned to a stopped state (refer to section [9.2.4](#) for the procedure to stop services).

9.2.3.4 Performing a Contingency Test of Unified Manager

This section details the processes involved in conducting a contingency test of Unified Manager. You should only perform this contingency test if required to do so by your Business Continuity representative. During this test, services on the production Management Station will be in a stopped state. Prior to conducting this test ensure the environment is configured according to the infrastructure pre-requisites defined in section [9.2.3](#).

1. Perform a database backup of the production system (using the either of the procedures outlined in sections [9.2.1.2](#) or [9.2.1.3](#)) and copy the database backup to the data directory on the contingency server (use the procedure outlined in section [9.2.2.3](#) to determine the location of the data directory).
2. Start the Unified Manager services on the contingency server using the service start procedure in section [9.2.4](#).
3. Restore the most recent version of the production database backup to the contingency server using the procedure detailed in section [9.2.1.6](#).
4. Generate a new set of Unified Manager Web Server SSL keys using the procedure described in section [10.3](#).

5. Update the value of the [localHostName](#) option within Unified Manager using the following command on the CLI of the destination [Management Station](#). This will ensure any email alerts generated by the destination system contain the correct host name using the Windows system variable %COMPUTERNAME%.

```
C:\>dfm option set localHostName=%COMPUTERNAME%
```

6. Confirm users can login to the UM GUI on the contingency Management Station and that the UM server can monitor a sample of the storage controllers in the environment. Once this is confirmed you can move on to step 7. Do not move on step 7 until access to the contingency Management Station has been confirmed.
7. Use the procedure in section [9.2.4](#) to stop all Unified Manager services on the production Management Station.
8. Perform the necessary testing on the contingency Management Station according to the test requirements specified by the Business Continuity representative. In lieu of any test points from the Business Continuity team the following points can be used as a base guide to check functionality in COB:-
 - a. Storage Administrator login successful to UM GUI.
 - b. Confirm storage controller connectivity from Management Station using the '**dfm host diag**' command for a small subset of Filers in your environment.
 - c. Using the [NMC](#) connect to the contingency Management Station and confirm monitoring statistics are available in Performance Advisor.
 - d. Ensure SnapMirror and SnapVault relationships are reported and up-to-date.
 - e. Run some sample reports. Confirm the reports run successfully and the content is accurate.
 - f. Run a configuration comparison report between a configuration template and a sample of storage controllers. Confirm the report completes without errors.
9. If any test points in Step 8 have **failed** use the procedure below to end the contingency test, otherwise if all points in Step 8 completed successfully use the procedure outlined in Step [10](#).
 - a. Do not copy the database instance running on DR back to production following a test that has failed on any test points. Troubleshoot the contingency Management Station as required.
 - b. Restart the Unified Manager services on the production server using the service start procedure in section [9.2.4](#).
 - c. Confirm users can login to the UM GUI on the production Management Station and that the UM server can monitor a sample of the storage controllers in the environment.
 - d. Once step 10c is confirmed, stop the UM services on the contingency server using the procedure in section [9.2.4](#).
 - e. Skip step 10 and proceed directly to step [11](#).

10. Only follow this step to failback to production if all points in Step 8 completed **successfully**. If any points in step 9 failed then follow the procedure outlined in Step 9.
 - a. Perform a database backup of the contingency system (using the either of the procedures outlined in sections [9.2.1.2](#) or [9.2.1.3](#)) and copy the database backup to the data directory on the production server (use the procedure outlined in section [9.2.2.3](#) to determine the location of the data directory).
 - b. Restart the Unified Manager services on the production server using the service start procedure in section [9.2.4](#).
 - c. Restore database backup from the contingency server to the production server using the procedure detailed in section [9.2.1.6](#).
 - d. Generate a new set of Unified Manager Web Server SSL keys on the production server using the procedure described in section [10.3](#).
 - e. Update the value of the [localHostName](#) option within Unified Manager using the following command on the CLI of the production [Management Station](#). This will ensure any email alerts generated by the production system contain the correct host name using the Windows system variable %COMPUTERNAME%.


```
C:\>dfm option set localHostName=%COMPUTERNAME%
```
 - f. Confirm users can login to the UM GUI on the production Management Station and that the UM server can monitor a sample of the storage controllers in the environment.
 - g. Once step 10f is confirmed, stop the UM services on the contingency server using the procedure in section [9.2.4](#).
11. This completes the procedure for performing a contingency test of Unified Manager.

9.2.3.5 Optimized DR configuration for large environments (3 site fan-in model)

CATE Storage and Backup Engineering has devised an optimized DR configuration for large environments that are present in three or more strategic data centers in a given geographic region. The purpose of this design is to minimize the requirement for infrastructure management servers (be they physical or virtual).

The principle assumption of this design is based on the probability that only one of two possible production strategic data centers in a given region has the potential to suffer a catastrophic failure. The likelihood of two strategic data centers in the same region concurrently suffering a catastrophic failure is hopefully extremely low. It is not possible to merge two database instances of UM into a single instance therefore concurrent support of two offline sites is not available in Unified Manager in the current release.

The basis of the model is in a three site environment, production sites A and B share a common contingency site C. Both [Management Stations](#) in Sites A and B replicate their UM database backups to the DR Management Station at Site C. The intention is to automate this replication using the processes discussed in the next section. The remainder of this section will be used to describe the different potential states of this 3 site infrastructure management DR model. The diagram below demonstrates the high level design of this model.



Figure 25 - Three site UM DR fan-in model

The following table represents a summary of the possible states of Unified Managers at each of the three sites in this model given a defined set of circumstances for each site. The last line in the table clearly demonstrates the fact the DR cannot be concurrently provided if both sites A and B were unavailable.

Site A (UM Production)	Site C (UM DR)	Site B (UM Production)
State: Up / OK	State: Offline / OK	State: Up / OK
State: Down / Fail	State: Online / DR for A	State: Up / OK
State: Up / OK	State: Online / DR for B	State: Down / Fail
State: Down / Fail	State: Choose DR for A or B	State: Down / Fail

Table 8 - Three site fan-in state summary

9.2.3.6 Automating the replication of database backups to contingency

This is a placeholder for future scripted functionality that will integrate both the database backup and backup archive copy to contingency functions into a single automated process.

9.2.4 Unified Manager Application Graceful Shutdown

It is very important to gracefully shutdown [UM](#) prior to any work being performed on the [Management Station](#) which may result in a reboot of the Management Station. A graceful application shutdown provides the opportunity for all writes to be completed to the database and the database service to be terminated in a controlled manner.



If the hosting server is shutdown or rebooted without gracefully terminating all Unified Manager services the [UM](#) database may become corrupt. It is important that the [MSSA](#) and [UM SA](#) groups both understand this requirement and update their system maintenance management level III (PCM) documentation accordingly.

Use the following procedure on the CLI of the [Management Station](#) to gracefully shutdown all [UM](#) services. Type '**dfm service stop**' and press [ENTER].



While the services are stopped all processes on the Unified Manager server are terminated and connectivity from the [NMC](#) will no longer function.

```
C:\>dfm service stop
Service: watchdog stopped.
Service: server stopped.
Service: scheduler stopped.
Service: monitor stopped.
Service: eventd stopped.
Service: http stopped.
Service: webui stopped.
Service: sql stopped.
```

You can use the '**dfm service list**' command on the CLI of the [Management Station](#) to confirm the services have stopped.

```
C:\>dfm service list
sql: not started
webui: not started
http: not started
eventd: not started
monitor: not started
scheduler: not started
server: not started
watchdog: not started
```

The [MSSA](#) is now free to perform the planned maintenance. If no reboot of the [Management Server](#) is required at the end of the maintenance period, then all UM services can be simply restarted using the '**dfm service start**' command on the CLI of the [Management Station](#):

```
C:\>dfm service start
Service: sql started.
Service: webui started.
Service: http started.
Service: eventd started.
Service: monitor started.
Service: scheduler started.
Service: server started.
Service: watchdog started.
```

This concludes the best practices for the graceful shutdown and restart of Unified Manager services.

9.2.5 Unified Manager Server Monitoring

For regions where Enterprise Management Systems support the use of tools to monitor the status of services on Windows servers, CATE Storage and Backup Engineering recommend the following [UM](#) application services be monitored to ensure they are running:

Service Name	Display Name	Description
DFMApache	DFM Apache	Apache/2.2.23 (Win64) mod_ssl/2.2.23 OpenSSL/1.0.1e
DFMEvent	DFM Event	DataFabric Manager event monitor
DFMMonitor	DFM Monitor	DataFabric Manager monitor
DFMScheduler	DFM Scheduler	DataFabric Manager job scheduler
DFMServer	DFM Server	DataFabric Manager Server
DFMSybase	DFM Sybase ASA	DataFabric Manager database server
DFMWatchDog	DFM Watchdog	DataFabric Manager self-monitor
DFMWebUI	DFM Web UI	Jetty Web Server for OnCommand console

Table 9 - Critical application services list recommended for monitoring

Storage Administrators are advised to check with their local monitoring teams and the SA group managing the operating system of the Management Station for the procedures to implement Windows service based monitoring.



All application services (DFMSybase, DFMEvent, DFMMonitor, DFMScheduler and DFMServer) maintain a 'running state' during a database backup. This is different from versions of Unified Manager prior to version 3.5 where these services were stopped / paused during a database backup.

In addition to service status, CATE Storage and Backup Engineering also recommend monitoring the free space on the D: partition (recommended installation partition for Unified Manager data files) on the Management Station. For regions where Enterprise Management Systems support the use of tools to monitor the amount of free space in a file system on Windows servers, CATE Storage and Backup Engineering recommends configuring an event policy that will generate an alert should less than 15% free space remain on the D: partition.

Please refer to section [9.5](#) for a discussion of other items that relate to Unified Manager capacity planning.

9.2.6 Data Lifecycle within the Unified Manager Database

With the exception of performance data, most of the information in the UM database should be considered transitory. Like any infrastructure management product its purpose is to record and report the state of the infrastructure at a certain time. Certain types of information, usually relating to performance and utilization, are indexed chronologically but most other information only reflects the current status of a given aspect of the storage infrastructure.

Like many other applications that incorporate a database, Unified Manager requires policies and procedures to ensure the database does not become bloated and unwieldy due to unnecessary data. Two important aspects of UM are how long it maintains performance and event data, and are discussed below:

1. Performance information per device will be maintained in the database for one year in a detailed state. Performance information older than a year will be compacted in the database and as such some granularity will be lost from the samples, aggregated and averaged out over a larger sample period. This will allow for historic trending analysis for periods greater than a year. If you need to retain a detailed historic record of performance data prior to the data compaction process executing, export the data using the **'Export'** function from the **'File'** menu in the NMC. This is the default behavior for OM and cannot be adjusted.
2. All deleted events will be purged from the UM database after 4 weeks as defined by the option [eventsPurgeInterval](#). This automatic purge is required to keep the database optimized in size and performance. If you need to retain a record of events beyond this duration as part of the Operational Level III processes for a point in time record, export the events from the database by navigating to the **'Events'** section of the **'Group Status'** screen, selecting **'History'** as the report type and exporting using the spreadsheet export function. Follow the link to [eventsPurgeInterval](#) for information on the default value for the event purge interval. The value of 4 weeks has been determined as a result of joint work conducted between CATE and NetApp given the size of the NAS environment within Citi. Also refer to the managing events and alarms best practices in section [9.12](#).

It is not the purpose of Unified Manager to be an indefinite archive of data regarding the storage infrastructure. Depending on business and operational requirements Storage Administrators may need to setup a regular set of reports for export from the database that can be used as a reference archive. If these reporting requirements are global in nature then please feedback via the contact details in section 3 so that these reports can be delivered as a global standard solution. An example would be Entitlement information. For Unified Manager entitlement information will be uploaded to EERS as part of a future CATE Storage and Backup Engineering project. EERS will form the historic repository of application entitlement data. There is no requirement to keep a chronological repository of entitlement data within UM itself.

9.2.7 ePO Policy on UM Server

Each Unified Manager server instance stores its data primarily in a database running locally and it is important that the virus scan settings for NetApp Unified Manager Servers is set to exclude monitordb.db, the database UM uses.

The following sections describe the importance and which policy should be applied.

9.2.7.1 Why should the above mentioned file be excluded for virus scan

NetApp Unified Manager uses Adaptive Server Anywhere (ASA) as its database management system (DBMS) which contributes to its good performance as it performs many reads and writes operations outside of the normal I/O operations. While performance is achieved, together with anti-virus scanning, this could cause potential performance degradation and/or corruption of database.

Anti-virus software might detect the reads and writes outside normal I/O operations as a potential problem and quarantine the file. This can be particularly hazardous if the .db or temporary files are quarantined, and it may cause corruption by interfering with the normal functions of the database. Degraded performance can also occur if the anti-virus software is checking all I/O operations performed by the database server and this could affect all programs on the system and not limited to the ASA DBMS.

9.2.7.2 Which policy to apply

Please refer to the exclusion policy outlined in the [Configuring VirusScan Settings for NetApp Operations Manager Servers via ePO](#) document available from [Host Security Engineering](#) (see the link in section 7.5, References, for document location)

9.2.7.3 Confirming policy implementation

In order to validate that the policy has taken affect, check the local anti-virus settings on the UM server. Use the process documented in section 'Checking Local Settings' in the [Configuring VirusScan Settings for NetApp Operations Manager Servers via ePO](#) document.

9.3 UM Maintenance

This section lists recommendations for various maintenance activities that must be conducted in order to ensure Unified Manager application availability and conformance with CISS policy.

1. Generate a new SSL web server certificate every two years. Use the procedure outlined in section [10.3](#) to do this.
2. Backup the Unified Manager database at least once every 24 hours (preferably outside of core business hours). Use the procedure outlined in section [9.2.1](#) to configure this.
3. Prior to any patching or maintenance of the Windows Server hosting Unified Manager ensure the following steps have been completed:
 - a. A manual backup of the Unified Manager Database (refer to section [9.2.1.4](#) for the procedure). This will provide a rollback point if the server maintenance was to cause an issue with the application database.
 - b. Stop all Unified Manager services (refer to section [9.2.4](#) for the procedure). This will allow all services to gracefully shutdown; most important of which is the sql service running the database. If the hosting server is shutdown or rebooted without gracefully terminating all Unified Manager services the application database may become corrupt.
4. Prior to installing any script plug-ins in [UM](#), NetApp recommends performing a manual backup of the UM Database (refer to section [9.2.1.4](#) for the procedure). This will provide a rollback point if the plug-in install process causes an issue with the UM database.
5. Prior to installing patch updates to [UM](#), NetApp recommends performing a manual backup of the UM Database (refer to section [9.2.1.4](#) for the procedure). This will provide a rollback point if the patch install process causes an issue with the UM database.
6. Prior to upgrading [UM](#) to a later release, NetApp recommends performing a manual backup of the UM Database (refer to section [9.2.1.4](#) for the procedure). This will provide a rollback point if the plug-in install process causes an issue with the UM database.

9.4 UM Migration

Technology Infrastructure is not a static environment due to activities such as technology refreshes and platform End-of-Life. As such there is a need to provide a documented process to migrate Unified Manager between different [Management Stations](#). This section details the steps involved in migrating Unified Manager from one physical server to another physical server, or (potentially) migrating from a physical server to a virtual server (VM) (see section [8.2.1](#) for details regarding hosting Unified Manager on VMware). Refer to section [8.2](#) regarding details on the supported hosting platforms for Unified Manager.

9.4.1 Server Migration (Unified Manager same version)

This section details the process of migrating Unified Manager between two [Management Stations](#) where the destination Management Station will run the same version of Unified Manager as the source Management Station. CATE Storage and Backup Engineering does not differentiate between changes in the Operating System version on the source and destination Management Stations because these differences are basically transparent to the day to day operation of the application. The difference in the source and destination application version is more significant.

To achieve this application migration, the following pre-requisites are required in order for a successful migration of Unified Manager between Management Stations:

1. The destination system (target for the migration) requires the following attributes:
 - a. Built with the same disk partitions as the source system and the partition sizes are the same or larger. For example the recommended installation partition for UM on Windows is C:. Ensure the destination system has a C: partition.
 - b. Built with the same Citi TPM of Unified Manager as the source system.
 - c. Configure Unified Manager using recommended settings in section [11.1](#)
 - d. A member of the same Windows Active Directory domain as the source system.
 - e. Shares the same common DNS infrastructure as the production system.
 - f. Has IP connectivity to all the storage systems as the source system utilizing ports 443 for HTTPS tunneled API calls and 22 for SSH connectivity.

The following steps detail the process for migrating Unified Manager between [Management Stations](#), when each management station is running the same version of Unified Manager:

1. Perform a database backup of the source system (using either of the procedures outlined in sections [9.2.1.2](#) or [9.2.1.3](#)) and copy the database backup to the data directory on the destination server (use the procedure outlined in section [9.2.2.3](#) to determine the location of the data directory).
2. Restore the most recent version of the source database backup to the destination server using the procedure detailed in section [9.2.1.6](#).
3. Generate a new set of Unified Manager Web Server SSL keys using the procedure described in section [10.3](#).
4. Update the value of the [localHostName](#) option within Unified Manager using the following command on the CLI of the destination [Management Station](#). This will ensure any email alerts generated by the destination system contain the correct host name using the Windows system variable %COMPUTERNAME%.

```
C:\>dfm option set localHostName=%COMPUTERNAME%
```

5. Your newly migrated instance of Unified Manager is now ready for use.

9.4.2 Server Migration (Unified Manager different version)

This section details the process of migrating Unified Manager between two [Management Stations](#) where the destination Management Station is running a different version of Unified Manager compared to the source Management Station. The main differences between the procedures listed in this section and the previous section are a change to pre-requisite point 1(b) and migration process step (2).



CATE Storage and Backup Engineering only supports upgrading Unified Manager using the database restore process if the difference in the two releases is n-1 or less. Therefore the oldest version of Unified Manager database backup that CATE Storage and Backup Engineering will support for restoration into UM 5.2 is a version UM 5.1 backup. If the version of Unified Manager currently used in your environment is older than UM 5.1 then CATE Storage and Backup Engineering recommends upgrading the database to version UM 5.1 before upgrading to UM 5.2. Any other database upgrade path is untested and therefore unsupported.

CATE Storage and Backup Engineering does not differentiate between changes in the Operating System version on the source and destination Management Stations because these differences are basically transparent to the day to day operation of the application.

To achieve this application migration, the following pre-requisites are required in order for a successful migration of Unified Manager between Management Stations:

1. The destination system (target for the migration) requires the following attributes:
 - Built with the same disk partitions as the source system and the partition sizes are the same or larger. For example the recommended installation partition for UM on Windows is C:. Ensure the destination system has a C: partition.
 - Built with the latest Citi TPM of Unified Manager version 5.2.
 - Configure Unified Manager using recommended settings in section [11.1](#)
 - A member of the same Windows Active Directory domain as the source system.
 - Shares the same common DNS infrastructure as the production system.
 - Has IP connectivity to all the storage systems as the source system utilizing ports 443 for HTTPS tunneled API calls and 22 for SSH connectivity.

The following steps detail the process for migrating Unified Manager between [Management Stations](#), when each management station is running a different version of Unified Manager:

1. Perform a database backup of the source system (using the either of the procedures outlined in sections [9.2.1.2](#) or [9.2.1.3](#)) and copy the database backup to the data directory on the destination server (use the procedure outlined in section [9.2.2.3](#) to determine the location of the data directory).
2. Restore the most recent version of the source database backup to the destination server using the procedure detailed in section [9.2.1.6](#). The output to this command will differ from that shown in section [9.2.1.6](#) as the database schema will be upgraded as part of the restoration process (the following example is taken from 3.8.1 so a 5.1 restore will look slightly different). An example of this output is shown below (the output shown below has been edited and certain lines have been removed to reduce the length of the output).

```
C:\>dfm backup restore dfm_backup_2008-08-04_03-00-01.ndb
Please do not interrupt this operation with CTRL-C.
1321 File(s) copied
290 File(s) copied
Upgrading database to latest schema.
```

Depending on the size of the monitored network and the age of the DataFabric Manager installation, this command may take between

many minutes and several hours to complete.
Please be patient...

```
Stopping services before upgrade...
Restarting the DataFabric Manager Database service in bulk update
mode...
Upgrading from 629 to 660
Upgrading database schema.
Adding host agent OSSV enabled field.
Adding key between events table and statusMap.
Adding application data for data set.
... ..
... ..
... ..
... ..
Removing defunct counter groups.
Initializing new counter groups.
Updating calc stats.
Updating counter groups.
Updating database caches.
Testing whether to create or update sample DP Policies.
Skipping update of sample backup schedules and throttles.
Skipping update of sample DP policies.
Testing whether to create or update sample DR Policies.
Creating or updating sample DR policies.
Updating event details.
Successfully upgraded to level 660.
Restarting database service in normal mode...
Restarting other services...
Restored from backup dfm_backup_2008-08-04_03-00-01.ndb.
Completed restore from dfm_backup_2008-08-04_03-00-01.ndb.
```

3. Generate a new set of Unified Manager Web Server SSL keys using the procedure described in section [10.3](#).
4. Update the value of the [localHostName](#) option within Unified Manager using the following command on the CLI of the destination [Management Station](#). This will ensure any email alerts generated by the destination system contain the correct host name using the Windows system variable %COMPUTERNAME%.

```
C:\>dfm option set localHostName=%COMPUTERNAME%
```

5. Your newly migrated and upgraded instance of Unified Manager is now ready for use.

9.5 UM Capacity Planning

The purpose of this section is to provide Storage Administrators with a list of key performance counters and other capacity related measurements in order to ensure the sustained successful operation of the UM application. They are included in this document to meet the requirements of the CITMP capacity planning discipline for the Unified Manager application.

The recommended capacity planning measurements that require checking are listed below, including a frequency at which these checks should be performed:

- **Management Station (UM hosting server)**
These performance counters should be reviewed in line with the frequency at which other capacity planning checks are conducted for your environment.
 - Memory utilization – collected by BMC patrol (or equivalent)
 - Disk utilization – collected by BMC patrol (or equivalent)
 - Network utilization – collected by BMC patrol (or equivalent)
 - CPU utilization – collected by BMC patrol (or equivalent)
 - It is important to note on multi-core hosting servers the database service (DFM Sybase ASA) can leverage up to four processing cores
 - All other application services (DFM Apache, DFM Event, DFM Monitor, DFM Scheduler, DFM Server and DFM Watchdog) are shared between the remaining processing cores (if any) on the system, otherwise they are split over the available cores.
- **UM application licenses**
These thresholds are monitored automatically as part of the UM application. If the thresholds are breached then an event is generated so that periodic checking of these values is not necessarily required.
 - **Age**
Check the expiry date of the licensed features within UM either by navigating to the Licensed Features option page in the GUI (CONTROL CENTER / SETUP / OPTIONS / LICENSED FEATURES) or running `'dfm license list'` on the [management station](#) command line. This checking process is automated through the use of the application option [licenseExpireWarningThreshold](#). In the Citi standard configuration this is set to generate an event in UM 90 days prior to the license expiring. This should provide enough time for new licenses to be ordered, delivered, and installed under change control.
 - **Capacity**
Check the ability of the UM server to monitor the number of storage devices assigned to it either by navigating to the About page in the GUI (CONTROL CENTER / HELP / ABOUT) or running `'dfm about'` from the management station command line. This checking process is automated through the use of the application option [nodesRemainingWarningThreshold](#). In the Citi standard configuration this is set to generate an event in UM when there are only 20 free node licenses left. This should provide enough time for new licenses to be ordered, delivered, and installed under change control.
- **UM SSL web server certificates**
It is recommended that the expiry date of the embedded web server SSL certificates be checked on a quarterly basis.
 - **Age**
Check the expiry date of the SSL certificates on the embedded web server using the following CLI command on the management station: `'dfm ssl server detail'` and look for the row entitled 'Validity' within the first ten rows of the output. Ensure that the 'Not After' date is greater than 3 months time otherwise raise a change control to regenerate the web server SSL certificates

- UM database backup

It is recommended that the database backup status be checked monthly in terms of database archive size and the time to complete the backup process.

- Size

Browse the DFM data directory on the UM management station (%INSTALLDIR%/DFM/data/) and check the size of the active database filer (monitor.db) and the backup files (*.ndb) to ascertain the rate of growth in the database and the ability of the database backup utility to compression the backup archive.

- Time to complete

Compare the time the database backup is scheduled to start with the triggered time on the Database Backup Succeeded informational event to provide the time the database backup took to complete. If performance advisor is enabled, ensure the backup is not taking so long as to significantly impact the ability to run performance advisor because performance advisor is disabled as part of the database backup operation (see section [9.14](#)).

To determine the time at which the database backup is scheduled to start navigate to the Database Backup management screen in the GUI (CONTROL CENTER / SETUP / DATABASE BACKUP). To determine the amount of time the backup took to complete click on the most recent database backup events in the event listing section of the Database Backup management screen. On the event details page look at the time stamp for the 'triggered' field and compare the difference to the scheduled start time to derive the elapsed time.

- Total number of active events with a warning or worse severity

It is recommended that the total number of open events with a severity of warning or worse be checked on a monthly basis.

- If the number of open warning or worse events is greater than 1000 it is recommended to conduct a review of your environment to address the cause of these events on the storage devices to order for the events to be cleared from the database (see section [9.12.2](#)).

- Total number of active events with an informational severity

It is recommended that the total number of open events with a severity of informational be checked on a monthly basis and any events no longer required be deleted so that OM can eventually purge them from the database (see section [9.12.2](#) for further information).

This concludes the discussion of the various application capacity planning metrics.

9.6 UM Application Options

The purpose of this section is to provide an overview of the application configuration parameters or options used by Unified Manager. There are over 100 configurable parameters that are accessible via the CONTROL CENTER / SETUP menu or through using the 'dfm option set' CLI command set. These options can be broadly categorized into four distinct purposes:

- Application security
- Application behavior / application performance tuning
- Regional specific settings (email server, distribution lists, team specific options)
- General settings

All of the application options that are relevant to Citi are listed in section [11.2](#). In the appendix the default settings and the recommended setting for each option is listed according to how you are intending to use UM. The options are presented in a manner that designates security best practices (mandatory) and technical best practices (recommended).

CATE Storage and Backup Engineering is providing a configuration script dfm_cfg.ps1 in the UM install folder which will automate the configuration of the application and optimize according to the size of your environment and the features you plan to use. Further information can be found in section [11.7](#).

This concludes the discussion of Unified Manager application configuration parameters.

9.7 UM Storage Device Management

Having discussed the wider architectural and implementation topics surrounding the use of [UM](#) this section is the first in a series that concentrates on actually using the application itself. Having already covered application installation and configuration, this section provides the necessary procedures to add your storage devices into the database so that they can be both monitored and managed, as well as best practices is modifying device records that already exist in UM and deleting devices from the database when they are retired.

9.7.1 Adding a Filer (Storage Controller)

Adding a Filer to UM is effectively a two stage process; where the first stage allows device monitoring and the second stage enables device management.

It's worthwhile providing a quick note about vFilers at this point. vFilers are automatically discovered and added to the UM database as long as the underlying physical storage system exists in the UM database. The UM application option that controls auto-discovery of vFilers is [discoverVfilers](#) and the option that controls the frequency of this operation is [vFilerMonInterval](#).

Before adding a Filer to the UM database ensure the following infrastructure prerequisites are met:

- The Filer is member of the same Windows Active Directory domain as the UM server (while it's possible to have the Filer in a different domain to the UM server – it's technically simpler if they are both in the same domain).
- Shares the same common DNS infrastructure as the UM server.
- The UM server has IP connectivity to all the storage system utilizing ports 443 for HTTPS tunneled API calls and 22 for SSH connectivity.

Use the following process to add a Filer to UM:

- On the main CONTROL CENTER tab, select MEMBER DETAILS and type the IP address or the Fully Qualified Domain Name of the Filer in the (New appliance) field.
- Click the 'Add' button to add the Filer to the UM database.




Figure 26 - Adding a Filer to UM

- The addition of the Filer to the database is confirmed with the following message (see [Figure 27](#)). At this point UM can perform some basic monitoring of the storage device but additional steps are required for UM to fully manage the device.

4. Click on the appliance fully qualified domain name, to go to the 'Appliance Details' screen.




Figure 27 - Filer Addition Confirmation Message

5. In the 'Appliance Tools' menu in the left pane of the OM GUI, click the 'Edit Settings' link.





Figure 28 - Edit Appliance Settings

6. In the 'Edit Appliance Settings' window configure each of the appliance connectivity options:
 - a. The '**Primary IP address**' will be pre-populated (this is the IP address associated with the host name provided in step 1).
 - b. '**Remote platform IP address**' is the IP address of the RLM card or BMC port if the system is installed with an RLM card or BMC port and it is configured.
 - c. '**Console Terminal Server address**' is the IP address of the terminal server and the associated port number on the terminal server that is connected to the console port of the Filer.
 - d. The '**login**' and '**password**' fields contain the credentials of the Functional ID used by UM to manage the storage device. Refer to section [10.5](#) for further information on this functional ID. Note that this account also needs to be created on the Filer (the authentication target). Refer to section 12.6.2.5 in the [ONTAP Solutions Guide](#) for information on configuring this account on the Filer.
 - e. The '**Login Protocol**', '**Hosts Equiv Enabled**', '**Administration Transport**', '**Preferred SNMP Version**' and '**Performance Advisor Transport**' options must be configured to '**Global Default Settings**'. These are configured by the procedures discussed in sections [8.3.6](#) and [9.6](#), and are detailed in the Appendix in section [11.1](#).
 - f. '**Owner Email**' and '**Owner Name**' is recommended to be defined as the team Distribution List and team name responsible for the system.
 - g. '**Resource Tag**' is optional but can be used to record the contents of any Citi inventory tracking label attached to the device.
 - h. Once finished click on the '**Update**' button.



Edit Appliance Settings 18 Aug 22:08

Edit nas3050gdse01

Primary IP Address	10.40.42.85
Remote Platform Management IP Address	10.40.35.43
Console Terminal Server Address (example: 192.0.32.50:6001)	10.40.32.252:2078
Login	dfmproxy
Password	*****
Login Protocol	(Global Default) ▼
Hosts.equiv Enabled	(Global Default) ▼
Administration Transport	(Global Default) ▼
Preferred SNMP Version	(Global/Network Default) ▼
Performance Advisor Transport	(Global Default) ▼
Owner Email	richard.sopp@citi.com
Owner Name	Richard Sopp
Resource Tag	

Update

Figure 29 - Configuring the connectivity options to the Filer

- Assuming all Filers in your environment use the same SNMP community string, these last few steps are only required when adding the first Filer to a new UM database. From the CONTROL CENTER TAB / SETUP MENU select the '**Network Credentials**' to configure the SNMP community string for the environment.




Figure 30 - Accessing the Network Credentials menu

8. Assuming all Filers in your environment use the same SNMP community string, click the 'Edit' hyperlink to modify the default community string.



Note the SNMP community string shown in the diagram is for illustrative purposes only. If you are unsure of the community string to use with your Filer contact your local Enterprise Monitoring support team for details.

Network Credentials 18 Aug 22:10

Add Network Credentials

IP Address

Network Mask

Preferred SNMP Version SNMPv1 ▼

SNMPv1 Settings

SNMP Community (If unspecified then SNMPv1 will be disabled for the network)

SNMPv3 Settings

Login (If unspecified then SNMPv3 will be disabled for the network)

Password (Mandatory when Login is specified)

Back Add

IP Address	Network Mask	Preferred SNMP Version	SNMP Community	Login	Edit	Delete
default		SNMPv1	nassnmp		edit	

Delete Selected

Figure 31 - Network Credentials Configuration Screen

This screen can also be used when divesting or merging business units in the support environment so that you can define custom community strings for the devices being divested from or merged into the current environment.

9. Ensure 'Preferred SNMP Version' is set to 'SNMPv1'¹⁰ and the 'SNMP Community' String is set to the standard recommended for your environment. Again, the entry in the screen shot below is included for illustrative purposes only.

Edit Network Credentials 18 Aug 23:20

Edit Network Credentials

IP Address default

Preferred SNMP Version SNMPv1 ▼

SNMPv1 Settings

SNMP Community (If unspecified then SNMPv1 will be disabled for the network) nassnmp

SNMPv3 Settings

Login (If unspecified then SNMPv3 will be disabled for the network)

Password (Mandatory when Login is specified)

Cancel Update

Figure 32 - Editing SNMP Network Credentials

¹⁰ Use of SNMPv3 is not supported within Citi at the time of writing this document.

10. The final steps relate to confirming connectivity between the UM server and the Filer recently added to the database. To do this click on the '**CONTROL CENTER**' tab and then the '**MEMBER DETAILS**' sub-tab and select appliances from the menu bar. Select the appliance you just added from the list in the same way as described in step 4. The click on the '**Diagnose Connectivity**' option in the '**Appliance Tools**' menu from left pane of the screen.




Figure 33 - Confirming UM to Filer Connectivity

11. The screen will take a few seconds to refresh while the tests are being conducted. Once the test is complete a screen similar to the following will be displayed.
- SNMPv3** is expected to fail as its not configured in UM or used in the Citi environment
 - RSH** is expected to fail as its disabled as part of the ONTAP Standard Build
 - SNMP Version in Use** should be listed as SNMPv1 and SNMP Community should be listed as the community string you defined in step 9 (the entry shown below is for illustrative purposes only).

Diagnostics - nas3050gdse02.namdev.nsrootdev.net 19 Aug 11:30

Network Connectivity	
IP Address	10.40.42.86
Network	(discovery disabled)
DNS Aliases	nas3050gdse02.namdev.nsrootdev.net
DNS Addresses	10.40.42.86
SNMP Version in Use	SNMPv1
SNMPv1	Passed (343 ms)
SNMP Community	nassnmp
SNMP sysName	nas3050gdse02.namdev.nsrootdev.net
SNMP sysObjectID	1.3.6.1.4.1.789.2.3 (Clustered Filer)
SNMP productId	0101173124
SNMPv3	Failed: Unable to initialize SNMP session: Timeout
SNMPv3 Username	dfmproxy (Edit SNMP Settings)
ICMP Echo	Passed (0 ms)
HTTP	Passed (0 ms)
NDMP Ping	Passed (port 10000, 0 ms)
NDMP Connect	Passed (1500 ms)
NDMP MD5 Passwd Check	Passed
RSH	timed out (Edit login and password)
SSH	Passed (375 ms)
RLM	Passed (1094 ms)
XML (https port 443)	Passed (1078 ms)

Appliance Details		
According to:	DataFabric Manager server	Host
Host Name	nas3050gdse02.namdev.nsrootdev.net	✓ nas3050gdse02.namdev.nsrootdev.net
System ID	0101173124	✓ 0101173124
Model	FAS3050	✓ FAS3050
Type	Clustered Storage System	✓ Clustered Storage System
OS Version	7.2.4P6	✓ 7.2.4P6
Revisions	300.7.2.1.19	✓ 300.7.2.1.19

Figure 34 - Connectivity diagnostics results screen

9.7.2 Implications of Modifying Filer System Information

This section presents management best practices to be considered when making significant changes to the storage environment managed via UM.

9.7.2.1 Changing the head (system ID / serial number), host name or IP address of an appliance

UM tracks the appliances using three pieces of information:

- Host name
- IP address
- System ID

Usually, if only one of the three items changes, the UM monitor discovers the change when it next polls system information from that appliance. The frequency of this update in UM is defined by the [sysInfoMonInterval](#) option. To get a more immediate update, you can request the monitor to refresh data from the appliance by invoking the Refresh Monitoring Samples tool from the Appliance Details page for that appliance, or running this command at the command line:

```
C:\> dfm host discover <old-hostname-or-IP-address>
```

If the IP address is changed, and the address is the one to which UM already communicated, it may be necessary to manually update the IP address on the Edit Appliance Settings page for the appliance, or by running this command, before refreshing the monitoring (as described above):

```
C:\> dfm host set <hostname> hostPrimaryAddress=new-address
```

If the IP address being changed is not the primary IP address, nothing needs to be done.

Depending on the situation, changing two of the three pieces of information at the same time may result in the system looking like a new appliance, and therefore the historical data would not be associated with this new appliance in UM. In some situations, the manual step of changing the "hostPrimaryAddress" described above can avoid this problem, or using this command:

```
C:\> dfm host rename <old-filer-name> <new-filer-name>
```

Changing all three pieces of information at the same will certainly result in the system looking like a new appliance in UM.

9.7.3 Deleting a Filer object from UM

This section details the process of deleting a Filer from the UM database.

Ensure any information you may require for historical reference is exported from the database before decommissioning a Filer and removing it from UM.

In order to gracefully remove a Filer and all its associated records from the database, UM needs to be one of the last steps in a system decommissioning process to avoid orphaned records and objects in the database. Here is a high level example of the steps involved in decommissioning a storage system and the checks that are recommended in UM to avoid any issues.

- Stop snapmirror replication and / or snapvault backups. Ensure UM updates to reflect this.
- Remove all snapmirror and snapvault relationships. Ensure UM updates to reflect this.

- Offline and destroy volumes. Ensure UM updates to reflect this.
- Offline and destroy aggregates. Ensure UM updates to reflect this.
- If monitoring a Filer with performance advisor, disable performance advisor for that object before deleting it from UM (refer to section [9.14.2.5](#) for this procedure).

If these steps are not followed there is a potential for an orphaned object to be associated when a new Filer is added to the database, especially if the new Filer has the same host name and a volume on the system has the same name as the orphaned object in the database.

Once the above steps have been completed use the following process to remove the Filer from the UM database.

1. Click on the 'Global' group in the left hand navigation pane
2. Select 'Member Details Reports' from the 'Current Group' menu in the left pane
3. Click on the 'Appliances' sub menu
4. Place a tick in the checkbox next to the Filer you wish to remove from UM
5. Click the 'Delete Selected' button




Figure 35 - Filer Deletion Process

6. Once completed a confirmation message will be displayed.



Figure 36 - Filer Deletion Confirmation Message

This concludes the Filer deletion process and the UM storage device management chapter.

9.8 UM Reporting

This section describes the way in which information is organized in the UM database; how the inter-relationship of objects and entities within the database determines what information is available in a report; and finally the classification of various pre-defined reports in terms of storage management (including new features such as thin provisioning and space management policies) as well as aligning certain reports to specific CiTMP domains.

9.8.1 Database Reporting Structure

The purpose of this section is to provide an introductory overview of the UM database. The UM database is split into different objects (referred to as base catalogs). A Filer and a Volume are examples of database objects. Each object has a set of attributes. Using the previous examples a Filer *has* vFilers, and a Volume *has* a rate of utilization. The number of vFilers and the rate of volume utilization are attributes.

When accessing pre-defined reports these can be found in the respective sub-menu items under the 'Group Status' and 'Member Details' tabs in the main 'CONTROL CENTER' window. Therefore volume reports can be found under the 'File Systems' menu and event related reports in the 'Events' menu, for example.

The interrelationship between each base catalog (object), that is relevant to how NetApp NAS is deployed within Citi, is shown in the diagram. For a full listing of the objects and their attributes refer to section [11.6](#).




Figure 37 - UM Database Objects

This concludes the brief introduction to the database reporting structure. The following sections focus on specific pre-defined reports and align them to the CiTMP disciplines used within CTI and CATE.

9.8.2 Storage Capacity Planning

The purpose of this section is to provide a reference list of reports that can be used for capacity planning purposes. It also provides an explanation for some of the more complex capacity and utilization rate calculations used within [UM](#).

The first item of note is that all capacity calculations in UM are based on binary math (KB to MB is a multiple of 1024).

The following table provides a summary of the commonly used pre-defined reports found in UM that relate to storage capacity planning tasks.

Entity	Report	Description
Aggregates	Aggregates, Capacity	Provides the used and allocated sizes of each aggregate on the Storage System. This report excludes aggregate level snapshot reserves (aggregate snapshots are not used within Citi). This report is the equivalent of a 'df -A' on the Filer.
Aggregates	Aggregate Space Availability	Provides similar information to the 'Aggregates, Capacity' report above but also includes the amount of free space left in the aggregate. This report excludes aggregate level snapshot reserves (aggregate snapshots are not used within Citi)
Aggregates	Aggregates, Committed Space	Provides information on the total committed capacity of each aggregate in order to determine the rate of oversubscription of physical resources that often result from a thin provisioning model. For further information refer to the 'Aggregate Committed Capacity' definition in Table 11 .
Aggregates	Aggregate Space Summary	Provides a report combining all the details contained in the 'Aggregate Space Availability' and 'Aggregates, Committed Space' report. This report excludes aggregate level snapshot reserves (aggregate snapshots are not used within Citi). This report provides the best capacity planning information for managing aggregates.
Appliances	Storage System Volume Capacity	Provides cumulative used and total space of all volumes on the Storage System. This report excludes snapshot utilization and allocation. If vFiles are implemented on your Storage System this report only contains volumes assigned to vFiler0, the physical filer. Use the File Systems / Volume Capacity report (below) for the utilization rates on all volumes on the Filer (not just those assigned to vFiler0).
File Systems	Volume Capacity	Provides used space within the active file system of the volume and the allocated size of the active file system in the volume. This report excludes snapshot utilization and allocation.
File Systems	Volume Space Availability	Provides similar information to the 'Volume Capacity' report above but also includes the amount of free space left in the volume. For this report thin provisioned volumes are treated the same as thick provisioned volumes, therefore the free space in the volume is paired with the volume and not the underlying pool of free space in the aggregate. This report excludes snapshot utilization and allocation.
File Systems	Volume Quota Committed Space	Provides the used space within a volume and compares it to the space committed in the volume through the use of qtree level quota policies, and expresses the committed space as a percentage of the total amount of space available in the active file system of the volume. This report excludes snapshot utilization and allocation. A committed space of 0 (zero) is reported when either 1) there are no qtrees on the volume or 2) there are qtrees on the volume but no qtree level quotas have been applied to the qtree. This report provides the best capacity planning information for managing oversubscription of volumes with the use of qtrees.

Entity	Report	Description
File Systems	Volume Autosize Policies	Provides the status of volume autosize (automatic space management) policies on the Filer (irrespective of whether the volumes are assigned to the physical Filer or a vFiler). The report indicates whether the autosize policy is active, the maximum defined size of the volume within the policy (irrespective of whether the policy is active or not), and the incremental size in which each volume will expand when the volume full threshold is triggered. This report provides the best capacity planning information when using space management policies.
File Systems	Volume Space Guarantees	Provides information as to whether each volume is thin provisioned (space guarantee = none) or thick provisioned (space guarantee = volume), as well as the amount of guaranteed space remaining in a volume (for thin volumes this is always 0 [zero]) and the current defined total size of the volume (this does not take into consideration the potential growth in size of a volume due to an autosize policy). This report excludes snaps hot utilization and allocation. This report provides the best capacity planning information when using thin provisioned volumes.
File Systems	Volume Space Savings – Deduplication and LUN/File Clones	Provides information on the storage savings achieved through ONTAP De-duplication and file level flex clones. The values returned by this report are equivalent to the ONTAP diag mode 'sis stat' command as opposed to the df -s command as the percentage saved value is adjusted to factor out snapshots which cannot be deduplicated as the blocks are locked.
File Systems	Qtree Capacity	Provides information on the used and allocation rates of qtrees. The allocation rate is defined as the 'Disk Space Limit' which is the equivalent of the 'Hard Limit' in the quota management screen in ONTAP. A used space of 0 (zero) is reported when there are qtrees on the volume but no qtree level quotas have been applied.
File Systems	Qtree Expandability	Provides information included in the 'Qtree Capacity' report (as above), and also includes the possible space that could be added to the qtree quota policy. This 'Possible Addition' is calculated by the amount of free space left on the volume containing the qtree.
File Systems	Qtree Space Availability	Provides similar information to the 'Qtree Capacity' report above but also includes the amount of free space left in the qtree.
vFiler	vFilers, Capacity	Provides the sum of used file system space for all volumes assigned as resources to a vFiler.
vFiler	vFilers Storage Objects Capacity	Provides a breakdown of the 'vFilers, Capacity' report, by listing the used file system space for each volume assigned to a vFiler.

Table 10 - Common Storage Capacity Planning Reports

In addition to the pre-defined reports available in UM, there are various summary screens such as 'Appliance Details' and 'Aggregate Details' that provide a consolidated view of the storage resources available on a Filer. Interpreting these consolidated statistics requires an understanding for the formulas used to derive the results. This interpretation process becomes more complex when advanced storage management features such as virtualization (flexible volumes), thin provisioning (space guarantees) and space management policies (vol autosize) are implemented.

Instead of providing a description of both the 'Appliance Details' and 'Aggregate Details' screens let's focus on the aggregate summary screen as the appliance summary screen is simply the summary of all the aggregates on the system. [Figure 38](#) on the following page provides an example of a 'Aggregate Details' screen that combines all the attributes of a complex storage implementation (thin provisioning, space management policies, volumes with no snap reserves, qtrees and qtree level quota policies).

The following table is presented in order to provide an insight as to how some of these complex summary values are calculated. Refer to [Figure 38](#) for the values listed in the following table. The items listed in the table are shown in the order they appear in the screen shot.

Entity	Value	Definition
Aggregate	Capacity Used	Equivalent of the 'Used' column in a 'df -A' command in ONTAP
Aggregate	Aggregate Size	Equivalent of the 'Total' column in a 'df -A' command in ONTAP
Aggregate	Total Committed Capacity	Sum of total allocated size of all volumes in the aggregate (includes both thin and thick provisioned volumes) <i>plus</i> sum of difference between current volume size and max volume size for all volumes configured with vol autosize (space management) policy on the aggregate (irrespective of whether the volumes are thin or thick provisioned) <i>plus</i> sum of snapshot reserves on all volumes on the aggregate <i>plus</i> size of the aggregate snapshot reserve (if defined).
Aggregate	Aggregate overcommitted / committed total	Sum of the total space of each thick provisioned volume on the aggregate <i>plus</i> sum of the used space of each thin provisioned volume on the aggregate <i>plus</i> sum of total snap reserve for each thick provisioned volume on the aggregate <i>plus</i> sum of used snap reserve for each thin provisioned volume on the aggregate <i>plus</i> sum of all qtree quotas on all thin provisioned volumes on the aggregate <i>plus</i> sum of difference between current volume size and max volume size for all volumes configured with vol autosize (space management) policy on the aggregate
Aggregate	Aggregate overcommitted / committed out of	Equivalent of the 'Total' column in a 'df -A' command in ONTAP
Aggregate	Aggregate overcommitted / using total	Sum of the total space of each thick provisioned volume on the aggregate <i>plus</i> sum of the used space of each thin provisioned volume on the aggregate <i>plus</i> sum of total snap reserve for each thick provisioned volume on the aggregate <i>plus</i> sum of used snap reserve for each thin provisioned volume on the aggregate <i>plus</i> sum of all qtree quotas on all thin provisioned volumes on the aggregate
Aggregate	Aggregate overcommitted / using out of	Equivalent of the 'Total' column in a 'df -A' command in ONTAP
Aggregate	Storage capacity / Volumes / Used space	Sum of the utilization value of the active file system (ie excluding snapshots) for all volumes on the aggregate (sum of used amount in 'df' command for all vols on the aggr)
Aggregate	Storage capacity / Volumes / Allocated space	Sum of the allocated value of the active file system (ie excluding snapshots) for all volumes on the aggregate (sum of total amount in 'df' command for all vols on the aggr)
Aggregate	Storage capacity / Qtrees / Used space	sum of all used space on qtree trees where a qtree quota policy has been applied (if qtree quota policy has not been applied to the qtree then used space is not calculated as the qtree has no space constraints or bounds)
Aggregate	Storage capacity / Qtrees / Allocated space	sum of all allocated space on qtree trees where a qtree quota policy has been applied (disk hard limit) + sum of difference between current volume size and max volume size for all volumes configured with vol autosize (space management) policy and containing an unbound qtree (a qtree without a qtree quota policy applied) on the aggregate + sum of the total allocated space of every volume on the aggregate containing an unbound qtree (a qtree without a qtree quota policy applied) and without a vol autosize (space management) policy applied

Entity	Value	Definition
Aggregate	Space Breakout / Aggregate Size	Equivalent of the 'Total' column in a 'df -A' command in ONTAP
Aggregate	Space Breakout / Snapshot Reserve Size	Equivalent of the 'Total' column in a 'df -A' command in ONTAP for the aggr/snapshot line entry
Aggregate	Space Breakout / Used Space	Equivalent of the 'Used' column in a 'df -A' command in ONTAP
Aggregate	Space Breakout / Available Space	Equivalent of the 'Available' column in a 'df -A' command in ONTAP
Aggregate	Space Breakout / Unused guaranteed space	sum of the available active file system space in all thick provisioned volumes on the aggregate <i>plus</i> sum of the available snapshot reserve space in all thick provisioned volumes on the aggregate

Table 11 - Consolidated Aggregate Utilization Definitions

The entries in the above table are presented in the following reference screenshot.

Aggreg...	Storage Sys...	Status	State	Type	Block Ty...	RAID Type	Used Capacity (GB)	Available Capacity (GB)	Committed Capacity (...)	Total Capacity (GB)
aggr2	d7-fas3140...	Error	Online	Aggre...	32 bit	RAID DP	5353.41	6458.66	20924.00	11812.07
aggr1...	d7-fas3140...	Normal	Online	Aggre...	32 bit	RAID DP	12.04	106.28	50.00	118.32
aggr1	d7-fas3140...	Error	Online	Aggre...	32 bit	RAID DP	1387.75	841.41	3515.00	2229.16
aggr2s	d7-fas3140...	Normal	Online	Aggre...	32 bit	RAID DP	0.21	1300.13	10.00	1300.34
aggr0	d7-fas3140...	Normal	Online	Aggre...	32 bit	RAID DP	48.28	137.48	48.00	185.76
aggr0	11c2-fas32...	Error	Online	Aggre...	32 bit	RAID DP	446.23	21.33	443.75	467.56
aggr0	11c2-fas32...	Error	Online	Aggre...	32 bit	RAID DP	1351.05	64.58	1343.53	1415.63
aggr1	11c2-fas32...	Normal	Online	Aggre...	64 bit	RAID DP	1861.02	17957.82	1850.00	19818.84
aggr1	11c2-fas32...	Normal	Online	Aggre...	64 bit	RAID DP	52.04	6493.81	160.00	6545.85

Figure 38 - Aggregate Status Summary Screen

This concludes the presentation of common storage capacity reporting and consolidated summary aggregate reporting. The following two sections have been provided to emphasize the tooling around two specific capacity reporting topics.

9.8.2.1 Thin Provisioning

Thin provisioning is used to optimize the utilization of the physical pools of storage so that the unused space in each file system (volume) acts as a shared resource to any other thinly provisioned volumes, rather than reserving the unused space of the file system exclusively for its own use. NetApp uses the term 'volume guarantee' to describe thin provisioning.

It is not the purpose of this document to recommend best practices as to when an aggregate is considered overcommitted. By default an aggregate is considered overcommitted when the allocation of total volume space on the aggregate exceeds the size of the aggregate.

The following reports have been identified as useful for managing a thinly provisioned NAS environment.

- Aggregate space summary
- Volume space guarantees

Refer to [Table 10](#) for further information regarding these report types.

9.8.2.2 Qtree Reporting

Qtree utilization reporting in UM requires a qtree quota policy to be applied on each qtree. Without a quota policy the only upper limit to the size of a qtree is the size of the volume in which the qtree reside. As a volume can contain more than one qtree, the absence of a quota policy on the qtree makes utilization based calculations impossible as both UM and ONTAP have no way of determining how much of the total volume space has been allocated to a qtree.

Qtrees are depicted with the following icon. 

Those that have been assigned a quota policy will be shown with one of the status icons, such as . Those that have not been assigned a quota policy will be shown with a status of unknown .

9.8.3 Change Management

The purpose of this section is to provide a reference list of reports that can be used for change management purposes. The table on the following page provides a summary of the pre-defined events that can be accessed via the events report tab found in UM that relate to change manage tasks. Use the following steps to navigate to this information.

1. Click on the '**CONTROL CENTER**' tab
2. Select '**Group Status**' tab
3. Click the '**Events**' item
4. Select '**Information or Worse**' from the '**Report**' drop-down menu




Figure 39 - Informational Events

The table below lists change management related events stored in the UM database.

Remember these events will be recorded in the UM database with a timestamp that reflects the time the storage device was polled for this information, not the actual time at which a change occurred in the storage system. The change would have occurred in a window of time calculated by timestamp on the UM event and the monitoring (polling) interval defined in the UM configuration options (see section [11.1](#)). You can use the event information in the UM database to help determine a window of time in which a change occurred on the storage device to assist in narrowing the scope of a search when reviewing the Filer's own audit log. The Filer audit log is the golden source that will provide the actual timestamp of the change and the user ID that initiated the change.

Event Name	Object	Status	Severity
Aggregate	Filer	Deleted	Information
Aggregate	Filer	Offline	Error
Aggregate	Filer	Online	Normal
Aggregate	Filer	Restricted	Normal
Alarm	Management Station	Created	Information
Alarm	Management Station	Deleted	Information
Alarm	Management Station	Modified	Information
CFO settings	Filer	Disabled	Normal
CFO settings	Filer	Takeover Disabled	Normal
Configuration Changed	Management Station	Config File	Information
Configuration Changed	Management Station	Config Group	Information
Configuration Changed	Filer	Local Config Changed	Warning
Database	Management Station	Restored Failed	Error
Database	Management Station	Restored Succeeded	Information
Enclosures	Filer	Disappeared	Warning
Enclosures	Filer	Found	Normal
Enclosures	Filer	Inactive	Warning
Storage System Configuration Push	Filer	Some options could not be set	Warning
Storage System Configuration Push	Management Station	Configuration pushed successfully	Normal
Storage System Configuration	Filer	Changed locally	Warning
Storage System Configuration	Filer	Matches with pushed configuration	Normal

Event Name	Object	Status	Severity
HBA Port	Filer	Offline	Warning
HBA Port	Filer	Online	Normal
Host	Filer	Cold start	Information
Host	Management Station	Deleted	Information
Host	Management Station	Discovered	Information
Host	Filer	Down	Critical
Host	Management Station	Modified	Information
Host	Management Station	Name changed	Information
Host	Management Station	System ID changed	Information
Host Domain User	Filer	Domain user modified on host	Information
Host Roles	Filer	Role discovered on host	Information
Host Roles	Filer	Role deleted on host	Information
Host Roles	Filer	Role modified on host	Information
Host User	Filer	User discovered on host	Information
Host User	Filer	User deleted on host	Information
Host User	Filer	User modified on host	Information
Host Usergroup	Filer	Usergroup discovered on host	Information
Host Usergroup	Filer	Usergroup deleted on host	Information
Host Usergroup	Filer	Usergroup modified on host	Information
Interface status	Filer	Down	Error
Interface status	Filer	Up	Normal
Management station	Management Station	License not expired	Normal
Management station	Management Station	Node limit OK	Normal
NVRAM	Filer	Missing	Error
NVRAM	Filer	Replace	Error
NVRAM	Filer	Normal	Normal
Report schedules	Management Station	Report schedule enabled	Normal
Report schedules	Management Station	Report schedule disabled	Warning
Resource group	Management Station	Created	Information
Resource group	Management Station	Deleted	Information
Resource group	Management Station	Modified	Information
Script	Management Station	Script schedule enabled	Normal
Script	Management Station	Script schedule disabled	Warning
Snapmirror	Filer	Abort completed	Normal
Snapmirror	Filer	Abort failed	Error
Snapmirror	Filer	Break completed	Normal
Snapmirror	Filer	Break failed	Error
Snapmirror	Filer	Delete aborted	Warning
Snapmirror	Filer	Delete completed	Normal
Snapmirror	Filer	Delete failed	Error
Snapmirror	Filer	Deleted	Information
Snapmirror	Filer	Initialize aborted	Warning
Snapmirror	Filer	Initialize completed	Normal
Snapmirror	Filer	Initialize failed	Error
Snapmirror	Filer	Not scheduled	Normal
Snapmirror	Filer	Off	Normal
Snapmirror	Filer	Quiesce aborted	Warning
Snapmirror	Filer	Quiesce completed	Normal

Event Name	Object	Status	Severity
Snapmirror	Filer	Quiesce failed	Error
Snapmirror	Filer	Resume completed	Normal
Snapmirror	Filer	Resume failed	Error
Snapmirror	Filer	Resync aborted	Warning
Snapmirror	Filer	Resync completed	Normal
Snapmirror	Filer	Resync failed	Error
Snapmirror	Filer	Update aborted	Warning
Snapmirror	Filer	Update completed	Normal
Snapshot	Filer	Discovered	Normal
Snapshot	Filer	Deleted	Normal
Snapshot	Filer	Local snapshot aborted	Warning
Snapshot	Filer	Local snapshot completed	Normal
Snapshot	Filer	Local snapshot failed	Error
Snapvault	Filer	Backup aborted	Warning
Snapvault	Filer	Backup completed	Normal
Snapvault	Filer	Backup failed	Error
Snapvault	Filer	Relationship create aborted	Warning
Snapvault	Filer	Relationship create failed	Error
Snapvault	Filer	Relationship created	Information
Snapvault	Filer	Relationship delete aborted	Warning
Snapvault	Filer	Relationship delete failed	Error
Snapvault	Filer	Relationship deleted	Information
Snapvault	Filer	Relationship discovered	Information
Snapvault	Filer	Relationship modified	Information
Snapvault	Filer	Restore aborted	Warning
Snapvault	Filer	Restore completed	Normal
Snapvault	Filer	Restore failed	Error
vFiler	Filer	Deleted	Information
vFiler	Filer	Discovered	Information
vFiler	Filer	IP address added	Information
vFiler	Filer	IP address removed	Information
vFiler	Filer	Storage unit added	Information
vFiler	Filer	Storage unit removed	Information
vFiler	Filer	Renamed	Information
vFiler Configuration	Management station	vFiler configuration push errors	Warning
vFiler Configuration	Management station	vFiler configuration push OK	Normal
vFiler Configuration	Management station	vFiler configuration changed	Warning
Volume	Filer	Clone deleted	Information
Volume	Filer	Clone discovered	Information
Volume	Filer	Destroyed	Information
Volume	Filer	New snapshot	Normal
Volume	Filer	Offline	Warning
Volume	Filer	Restricted	Warning
Volume	Filer	Online	Normal
Volume	Filer	Scheduled snapshots disabled	Warning
Volume	Filer	Scheduled snapshots enabled	Normal
Volume	Filer	Snapshots schedule modified	Information

Table 12 – Change management related events

This concludes the presentation of events for use in a change management indexing process. These events indicate a time when UM determines a change had occurred on a storage device since the last polling event and can therefore provide a range of time in window to search for the event in the Filer's audit log. The purpose here is not to suggest UM is a change reconciliation management tool but to indicate ways in which UM can assist in the determining the location of an event within an ONTAP audit log.

9.8.4 Asset Management

The purpose of this section is to provide a reference list of reports that can be used for asset management purposes. The following table provides a summary of the commonly used pre-defined reports found in UM that relate to asset management reporting tasks.

Entity	Report	Description
Appliances	Appliance Releases	Provides the host name, hardware model, O/S and system board firmware versions
Appliances	Appliance Locations	Provides the host name, hardware model, primary IP address, console address, SNMP contact information field and SNMP location information filed
Appliances	All disk	Provides the host name, disk name, aggregate name, firmware revision, disk vendor, disk model, disk type, shelf, bay, plex and disk size

Table 13 - Common Asset Management Reports

The current release of UM does not include a method to collect information for or to report on disk shelves, their modules and the module firmware releases. This information must be gathered directly from ONTAP using the '`sysconfig -a`' command or could be gather using an UM run command against multiple systems (see section [9.18](#)).

9.8.5 Problem Management

The purpose of this section is to provide a reference list of reports that can be used for problem management purposes. The following table provides a summary of the commonly used pre-defined reports found in UM that relate to problem management reporting tasks.

Entity	Report	Description
Appliances	Appliances, Down	Provides the type of the appliance, it's status, hostname, the time at which OM stopped receiving a response from the appliance, and the current ping status of the device
Appliances	Appliance Environmentals	Provides the type of the appliance, it's status, hostname, model and the status of the temperature within the appliance, it's fans, PSUs and NVRAM Battery
Appliances	Appliances, Up	This is the inverse of the Appliances Down report. It differs slightly from appliances down as it provides the timestamp of the last successful ping to the appliance. Note that a pingable appliance does not guarantee its serving data.
Appliances	Appliances, Uptime	Provides the type of the appliance, it's status, hostname, model, O/S version and the uptime of the device as reported by the ONTAP operating system
Appliances	Storage System Performance Summary	Provides the type of the appliance, it's status, hostname, model, CPU Busy (%) [see notes elsewhere in this document regarding how CPU busy is calculated], Total Ops / Sec, Network Throughput MB / Sec and Disk Throughput KB / Sec.
Appliances	Broken Disks	Provides the appliance name, the number of broken disks on the system and the list of the disks that are reported as broken
Appliances	Disk Performance Summary	Provides the name of the storage system, the disk name, the aggregate name to which the disk is assigned, the disk type, RPM, Read Ops / Sec, Write Ops / Sec, CP Read Ops / Sec, Read Latency (millisec), Write Latency (millisec), CP Read Latency (millisec), Throughput (Blocks / Sec), Busy (%)
vFilers	vFilers, Up	This is the inverse of the vFilers Down report. It differs slightly from vFilers down as it provides the timestamp of the last successful ping to the appliance. Note that a pingable vFiler does not guarantee its serving data.
vFilers	vFilers, Down	Provides the type of the it's status, hostname, underlying physical storage system name and the time at which OM stopped receiving a response from the vFiler, and the current ping status of the device
vFilers	vFiler Performance Summary	Provides the name of the vFiler, the hostname of the underlying physical storage system, Total Ops / Sec, Network Throughput KB / Sec, Data Transfer KB / Sec
File Systems	CIFS Share Permissions	Provides the CIFS share name, share path and SMB share permissions (not to be confused with NTFS permissions) for CIFS shares
File Systems	CIFS Shares	Provides the CIFS share name and share path for CIFS shares
File Systems	NFS Exports	Provides a detail NFS export report including export policies, paths and root access for NFS exports
File Systems	Volume Space Availability	Provides the volume name, the name of the aggregate hosting the volume, the name of the physical or virtual Filer hosting the aggregate, space used per volume, space available per volume and total size of the volume
File Systems	Volume Space Breakout	Provides a graphical representation of the snap reserve, active file system utilization, snap reserve utilization and any space reservation made on the volume
File Systems	Volume Quota Committed Space	Provides the volume name, the name of the aggregate hosting the volume, the name of the physical or virtual Filer hosting the aggregate, a graphical representation of used and overcommitted space, space used within the volume, space committed from the volume by quotas, percentage committed and the total / maximum size of the volume

Entity	Report	Description
File Systems	Volume Expandability	Provides the volume name, the name of the aggregate hosting the volume, the name of the physical or virtual Filer hosting the aggregate, the space currently available in the volume, the amount of unallocated and uncommitted space from the aggregate that could be added to the volume and the total possible amount of space available for the volume (current volume size plus available uncommitted space from aggregate).
File Systems	Volume Snapshot Counts	Provides the volume name, the name of the aggregate hosting the volume, the name of the physical or virtual Filer hosting the aggregate and a count of the number of snapshots on each volume
File Systems	Volume Snapshot Details	A more detailed version of the volume snapshot counts report listing the snapshot name, the volume hosting the snapshot, the name of the aggregate hosting the volume, the name of the physical or virtual Filer hosting the aggregate, and the access time (create time) of the snapshot
File Systems	Volume Autosize Policies	Provides the volume name, the name of the physical or virtual Filer hosting the volume, the state of the space management policy, the maximum growth potential size of the volume and the incremental growth amount for the volume.
File Systems	Volume Performance Summary	Provides the volume name, the volume type (flexible or traditional), the name of the aggregate hosting the volume, the name of the physical or virtual Filer hosting the aggregate, Total volume Ops / sec, Volume Latency (millisec) and Data Throughput (Bytes per second)
File Systems	Qtree Space Availability	Provides the qtree name, the name of the physical or virtual Filer hosting the aggregate, the name of the volume hosting the qtree, the space used in the qtree (if a qtree quota has been applied), the space available in the qtree (again if a qtree quota has been applied), soft limit (if defined), disk space limit (hard limit) and the percentage of available space (if a qtree quota limit is defined)
File Systems	Qtree Performance Summary	Provides the qtree name, the physical or virtual Filer hosting the aggregate, the name of the volume hosting the qtree, CIFS Ops/Sec, NFS Ops/ Sec and internal Ops / Sec
Aggregates	Aggregate Space Availability	Provides the aggregate name, the hosting storage system, the amount of space used, the amount of space available, the total allocated size of the aggregate and the utilization rate (%)
Aggregates	Aggregates, Committed Space	Provides the aggregate name, the hosting storage system, and the bytes committed in terms of a count and percentile value
Aggregates	Aggregate Performance Summary	Provides the aggregate name, the hosting storage system, and the total Ops / Sec of each aggregate
Disaster Recovery	All Snapmirror Relationships	Provides the source system name and path (to either volume or qtree depending on VSM or QSM), the destination system name and path, the replication policy, a failover policy (if defined), the state of the replication relationship, the status of the replication job, the lag time of the data on the destination system from that of the source system and the time stamp of when this relationship information was last updated in the OM database.

Table 14 - Common Problem Management Reports

9.8.6 Custom reports

The purpose of this section is to provide an overview of the custom reports functionality in UM. Like many other reporting products storage administrators may encounter situations where not all the required fields for a certain reporting purpose are available in a single report. In order to provide report structure flexibility UM includes a custom reporting feature.

The following example describes the process of creating and editing custom reports.

1. Click on the '**CONTROL CENTER**' tab
2. Select '**CUSTOM**' from the '**REPORTS**' menu
3. You are then presented with the following custom report library management screen

The screenshot shows the 'Custom Reports' interface. At the top, it says 'Custom Reports' and '15 Sep 11:21'. Below this is a 'Create a Report' section. It contains several input fields: 'Name' (callout 4), 'Web Display Name (Optional)' (callout 5), 'Description (Optional)' (callout 6), 'Base Catalog' (callout 7), 'Display Tab' (callout 8), and 'Related Catalogs' (callout 9). The 'Related Catalogs' section shows a tree view with 'Snapshot' selected, and sub-items like 'SnapReserve', 'Volume', 'Aggregate', 'Filer', 'Appliance', and 'vFiler'. Below this is a 'Choose From available fields' section with 'Usage' and 'Performance' radio buttons. A list of fields is shown, including 'Snapshot Access Time', 'Snapshot Dependency', 'Snapshot Id', 'Snapshot Name', and 'Snapshot SnapReserve (A)'. A 'Reported Fields' section is also present, with 'Add >>' and '<< Remove' buttons. At the bottom, there are fields for 'Enter Field Name Displayed on Report' and 'Choose Formatting to Apply on Field'. A 'Create' button is at the bottom right.

Figure 40 - Custom report creation screen

4. Enter a '**NAME**' for the custom report. Ensure '_' (underscore) or '-' (hyphen) characters are used instead of spaces as the name filed does not support the use of the ' ' (space) character. All report names must be unique. It is recommended to include the acronym that represents your support group at the start of the report name. This it's easier to determine which reports have been created by your group.
5. Provide a '**WEB DISPLAY NAME**' that will allow for a greater descriptive name for the report. This name will be displayed in the GUI only. Keep this field short as you can use the next field to contain a descriptive text of the report.

6. Provide a '**DESCRIPTION**' that explains the purpose of the report. This field is useful when you have similar reports. Use this field to describe why this custom report is being created instead of using one of the pre-defined reports.
7. Select a '**BASE CATALOG**' for the most granular storage object you want to include in the report. Refer to [Figure 37](#) which describes how each of the catalogs relates to one another. You need to start with the smallest object you wish to include in the report. For example if you want to create a report including disk information then disk is the base catalog required. It is not possible to go backwards through the catalog so you could not include disks in a report that started with the 'appliance' base catalog for example.
8. Choose the appropriate '**DISPLAY TAB**' for the report. The display relates to the tabs on the CONTROL CENTER screen. Use display tab that best suits your report purpose. For example a snapshot based report would be best placed under the 'file system' tab whereas a disk report is best placed under the 'aggregate' tab and so on.
9. In the '**RELATED CATALOGS**' box, click on the radio button for each catalog type you wish to include in the report. Follow step 10 to choose the fields from each catalog and then return to this step if you wish to choose another related catalog to choose more fields.
10. The final step is '**FIELD SELECTION**' or '**COUNTER SELECTION**'. This process is best described in stages:
 - a. Select the '**USAGE**' or '**PERFORMANCE**' radio button depending on the counter type you wish to add
 - b. Select the counter to add from the left list window
 - c. If you want to change the field name as it will appear in the report, enter the alternate name in the '**ENTER FIELD NAME DISPLAYED ON REPORT**' text box
 - d. If applicable to the counter, choose the formatting for the counters using the '**CHOOSE FORMATTING TO APPLY ON FIELD**' drop down menu
 - e. With the counter selected, click the '**ADD**' button to add the field to the report
11. Repeat steps 9 and 10 until all the fields are added. Use the '**MOVE UP**' and '**MOVE DOWN**' buttons to change the order of the fields in the report if required.
12. Once completed click the '**CREATE**' button to save the report template.

Once a custom report template has been created, you can then edit it by clicking the '**EDIT**' hyperlink on the same line as the report name.

Web Display Name	Report Name	Description	Display Tab	Edit	Select
aggr	aggr		Aggregates	edit	☐
vol test	vol test		File Systems	edit	☐
Volume utilization	Volume utilization		File Systems	edit	☐
Select All - Unselect All					Delete

Figure 41 - Custom report editing

In order to delete a custom report place a check mark in the relevant '**SELECT**' check box, then click the '**DELETE**' button.

It is not possible to export a custom report definition from one UM server to another.

9.8.7 Scheduling reports

The purpose of this section is to provide an overview of the scheduled reports functionality in UM. Report scheduling provides the following benefits to storage administrators:

- Automate distribution of reports
- Provide a point in time capture of your environment for historical reference

There are two stages involved in configuring scheduled reports:

1. Defining a schedule
2. Defining a report type and scope

Let's begin by defining a schedule policy as outlined in the procedure below:

1. Go to the '**CONTROL CENTER**' tab, select the '**REPORT**' menu, and select the '**SCHEDULE**' item.
2. Click on the '**SCHEDULES**' tab on the '**REPORT SCHEDULES**' screen.
3. Click on the '**ADD NEW SCHEDULE**' button to create a new schedule policy.




Figure 42 - Managing schedules

4. Create a name for the schedule in the '**SCHEDULE NAME**' field.
 - Create a name that is meaningful or descriptive for your implementation
 - Consider that in a large environment you may not want too many scripts running the same schedule as you may cause queuing and process spikes on the [Management Server](#). From a conservative starting point CATE storage engineering suggests having no more than 20 scripts assigned to the same schedule.
5. Define a '**SCHEDULE**' policy according to options shown in [Figure 43](#).
6. Click the '**ADD**' button to create the schedule.

Figure 43 - Creating a schedule

This concludes the process to create a schedule.

The following process discusses the steps involved in defining the contents of the report:

1. Go to the '**CONTROL CENTER**' tab, select the '**REPORT**' menu, and select the '**SCHEDULE**' item.
2. Click on the '**REPORT SCHEDULES**' tab on the '**REPORT SCHEDULES**' screen.
3. Click on the '**ADD NEW REPORT SCHEDULE**' button to create a new schedule policy.

Figure 44 - Managing report schedules

4. Create a name for the report schedule in the '**NAME**' field.
5. Select the Unified Manager component responsible for producing the report which relates to the tabs in the UM GUI: '**CONTROL CENTER**', or '**BACKUP**' or '**DISASTER RECOVERY**'.
6. Select the '**REPORT NAME**' that will be generated in this schedule.
7. Select the '**FORMAT**' of the report. The following format options are available:
 - csv
 - html
 - text
 - perl
 - xls
 - xml
 - PDF (available in the OnCommand Console UI only)

These reports will be contained within a zip file that will be emailed to the recipient defined in step 11.

8. Define the '**REPORT ARGUMENTS**' by selecting which group you would like to run the report against.
9. If graphing is available as part of the selected report, select the '**GRAPH TO SCHEDULE**' as part of this job.
10. Choose a '**SCHEDULE**' (as created in the previous process) which defines the frequency at which this report will run.
11. Define an '**EMAIL RECIPIENT(S)**' who will receive this report once generated.
When the report is sent, it will show the sending email address as the SMTP address defined by the UM option [alertFrom](#).
12. Click the '**ADD**' button to create the report schedule.

Reports - Add a Schedule 06 Oct 10:38

Specify the Report Schedule Parameters

Name:

Report Application: ☒ Control Center ☐ Backup ☐ Disaster Recovery

Report to Schedule:

- Qtrees, Comments
- Qtrees, Compact
- Qtrees, Deleted
- Snapshot Capacity**
- Space Reservation Thresholds
- Volume Autosize Policies
- Volume Capacity
- Volume Capacity Graph
- Volume Chargeback by Allocation, Last Month
- Volume Chargeback by Allocation, This Month
- Volume Chargeback by Usage, Last Month
- Volume Chargeback by Usage, This Month
- Volume Details
- Volume Expandability
- Volume First Snapshot Analysis

Format:

Show Deleted Objects:
(only for Global group) ☐ Yes ☒ No

Show All Fields:
(only in perl and xml format) ☐ Yes ☒ No

Report Arguments:

Graph to Schedule:

Schedule: [Manage Schedules](#)

Email Recipients
(specify email addresses as comma separated list)

Figure 45 - Creating a report schedule

The contents of the email generated by the scheduled reporting process will resemble the example shown below. Remember that the actual report will be attached to the email in the form of a zip file.

Generated report 'volumes-snapshot-capacity' on 'Global' at 06 Oct 2008 17:23.

Report Output Details

Output Id: 1
 Output Format: xls
 Report Schedule: Snapshot utilization
 Report Name: volumes-snapshot-capacity
 Report Application: Control Center
 Graph Name: none
 Report Arguments: Global
 Status: Success
 Run by: NAMDEV\rs01146
 Run Time: 06 Oct 2008 10:40

The attachment contains the output data.
 -- NetApp DataFabric Manager

Once a report has been generated as scheduled, there are two ways of accessing the results. The first way is to be the recipient of the emailed report. The second is to view the report using the saved reports tab as described below:

1. Go to the '**CONTROL CENTER**' tab, select the '**REPORT**' menu, and select the '**SCHEDULE**' item.
2. Click on the '**SAVED REPORTS**' tab on the '**REPORT SCHEDULES**' screen.
3. Click on the relevant entry in the '**OUTPUT ID**' column to show the details of the reporting job.




Figure 46 - Accessing saved reports

4. Click on the '**OUTPUT**' link in the '**STATUS**' field of the output details screen as shown in [Figure 47](#).

This concludes the process to access historical reports.




Figure 47 - Report output details

This concludes the presentation of reporting schedules.

9.8.8 UM Reporting Archive

The purpose of this section is to provide an overview of the advantages and disadvantage of defining a network drive as the reporting archive repository. Due to the limitations and complexity of defining a mapped network resource for this purpose, CATE Storage and Backup Engineering recommends maintaining the default setting for the `reportsArchiveDir` option (%INSTALLDIR%\reports).

9.8.8.1 Reporting Archive Advantages

By default the reporting archive is a directory in the UM application tree structure. Over time the number of reports in the archive will build up to a significant amount of disk space. Defining a different location for the report archive has the following advantages:

- Allows an archive class storage tier to be used to store archive data
- Reduces the I/O load on the UM server
- Reduces server backup size and time taken for backups

9.8.8.2 Reporting Archive Disadvantages or Limitations

While the use of a reporting archive will lighten the load on the UM server there are some limitations to be aware of:

- All reports are stored in a flat directory structure in the archive. If storing thousands of reports in the archive be mindful of the maxfile and maxdir limitations of file system you are using. For example if the archive is a UNC path to NAS CIFS share remember that the maximum number of files ONTAP supports at the root of a volume is 99,999¹¹. CATE Storage and Backup Engineering does not anticipate that many reports concurrently existing in the UM report archive.
- While it is possible for other applications and users to read the report data files directly by mounting to the archive directory (again assuming the archive is a CIFS share on a NetApp Filer) the automated naming convention of the report will not be intuitive to most users. Here is an example of the name of an Excel spreadsheet format report: 1031_3_2008_9_15_14_28_1.xls. At the time of writing this document it is unknown what '1031' represents. '3' represents the schedule report job ID; '2008_9_15' is the date the report was run and '14_28' represents the time of day. '1' is counter used to resolve cases where more than one report exists with the same name.
- As the report archive lies outside the UM database the scheduled report result files are not included in the database backup. A reporting archive is not considered critical to the day-to-day management processes of UM and therefore storage administrators should not consider replicating the reporting archive to DR storage.
- If you change the location of the report archive directory you have to manually copy the existing reports from the original archive directory location to the new archive directory.
- Requires the UM server to stay logged in to the Windows Desktop to maintain a mapped network drive to the reporting archive.

¹¹ This is larger than the root directory width limits permitted on other platforms.

9.9 UM Storage Configuration Management

This section describes the features and functionality available with the configuration management tooling in Unified Manager. Configuration management is designed to ease the administration burden of managing a larger NetApp NAS environment and is one of the more powerful (and underutilized) features of [UM](#). The following tasks are possible within Unified Manager 5.2 for both physical filer and vFiler management:

1. Build standard configurations
2. Pull configurations from existing filers / vFilers
3. Compare differences between a configuration template and live configuration of a filer / vFiler or group of filers / vFilers
4. Deploy configurations to new filers / vFilers
5. Automated policy enforcement and conformance deviation alerting

The remainder of this section discusses the infrastructure pre-requisites for configuration management, creating, modifying, copying and adapting templates; on demand configuration reporting, and finally automated configuration and management of your storage systems through the UM group tree and configuration policies.

It is the goal of the CATE Storage and Backup Engineering team to bring about improved Data Center Automation and increased operational efficiencies by providing CTI with standard templates based on NetApp NAS Filer SBCC, SB and Storage Engineering ONTAP Solutions Guide documentation.

9.9.1 Pre-requisites to Storage Configuration Management

In order to conduct remote configuration management of a Filer, certain prerequisites have to be met for the functionality to work.

1. Review section [9.7.1](#) step [6](#) to ensure your Filer is configured correctly in UM for device management and has passed the connectivity check process described in section [9.7.1](#) step [10](#).

A quick way to determine which Filers have not been configured for management is to go to the '**CONTROL CENTER**' tab, select the '**MANAGEMENT**' menu, select the '**STORAGE SYSTEM**' sub-menu and then select the '**PASSWORDS**' item. The storage systems listed in the box at the bottom of the screen are the systems that are not currently managed by UM. You can use this dialogue box to add the Functional ID discussed in section [9.7.1](#) to allow UM to manage the Filer.




Figure 48 - Unmanaged Filer List

2. The storage administrator logging into UM requires GlobalWrite and GlobalDelete privileges in UM in order to add and delete configuration files on a group.
3. Ensure the Filer has options.admin.config.list defined as follows (the default setting):
/etc/registry.default,/etc/registry.local

Use the following process to check this setting:

- a. Set the CLI on the Filer to advanced privilege mode
- b. Type the following command and press ENTER.
filer1> `registry get options.admin.config.list`
- c. If the response is the value defined above then this is OK, but if the value is undefined or empty then set the value using the following command (note the following command line has been wrapped):
filer1> `registry set options.admin.config.list /etc/registry.default,/etc/registry.local`

4. The appropriate ONTAP configuration template needs to be installed in UM. It is possible that the version of ONTAP installed on your Filer may be newer than the templates available when UM was released. CATE Storage and Backup Engineering will release new ONTAP configuration templates when certifying newer versions of ONTAP. This template release approach will apply to patch releases of ONTAP as well as full version releases.

- Use the following procedure to check the ONTAP configuration plug-in templates installed in UM:
- From the '**CONTROL CENTER**' tab, select the '**HELP**' menu and choose the '**About Operations Manager**' item.
 - Check through the output to ensure the configuration plug-in is installed for the version of ONTAP running on the Filers in your environment




Figure 49 - Determining Installed Configuration Plug-ins

9.9.2 Differences between Filer and Vfiler Configuration Management

UM offers two different types of configuration management, one for physical Filers and one for Virtual Filers. The configuration templates are designed such that you can only apply a physical configuration to a physical Filer and a vFiler configuration to a vFiler. There is no risk of inadvertently applying the wrong template type to the device.

A vFiler configuration template represents the subset of configuration options found on a vFiler when compared to the complete option listing found on a physical system, in the same way that the output from an options query on a vFiler is significantly shorter than the output of the options command on physical Filer.

To access the physical Filer configuration management tools go to the '**CONTROL CENTER**' tab / select the '**MANAGEMENT**' menu / the '**STORAGE SYSTEM**' item and then the appropriate sub-item (discussed in the following sections).




Figure 50 - Filer Configuration Menu

To access the vFiler configuration management tools go to the '**CONTROL CENTER**' tab / select the '**MANAGEMENT**' menu / the '**vFILER**' item and then the appropriate sub-item (discussed in the following sections).




Figure 51 - vFiler Configuration Menu

The only differences between Filer and vFiler configuration are the menu locations and the template options discussed above. For the remainder of the chapter the examples will use Filers only unless there is functionality or behavior that is specific to vFiler configuration. Unless stated otherwise the reader is safe to assume the same procedure applies to vFiler configuration management as does to Filer configuration management.

9.9.3 Creating a Configuration File

The purpose of this section is to outline the process used to create a configuration file. The first item to note is that it is not possible to create a configuration file manually using any of the tools in the UM GUI, and, more to the point, creating a configuration file manually would be a very time intensive process that would be exposed to human error. However you can copy and modify a configuration that has been pulled from a Filer (see sections [9.9.4](#) and [9.9.7](#)).

Use the following procedure to create a configuration file:

1. Select the '**Configuration Files**' option from the appropriate sub-menu in the '**Management**' menu.
2. Provide a descriptive name for the configuration file.
3. Select a Storage System (or vFiler – depending on menu selected) from which to pull the configuration list from.
4. Click the '**Create**' button.




Figure 52 - Pulling a configuration file

5. Once the configuration file is pulled the file will be listed on the screen.

9.9.4 Copying and Renaming a Configuration File

The purpose of this section is to outline the processes used when copying and renaming configuration files. Use the following procedure to copy or rename a configuration file:

1. Select the '**Configuration Files**' option from the appropriate sub-menu in the '**Management**' menu.
2. Click on the '**Edit**' hyperlink next to the configuration file you wish to copy or rename.




Figure 53 - Placing a configuration file into Edit Mode

3. From the '**Configuration File**' menu in the 'Configuration File' window select the '**Copy/Rename**' link.




Figure 54 - Selecting the Copy / Rename menu

4. Depending on whether you intend to copy or rename an existing configuration file, follow the appropriate action as outlined below:
 - a. Type the name of the configuration file you are copying to in the top text field and click the **'Submit'** button.
 - b. Type the new name of the configuration file you are renaming in the bottom text field and click the **'Submit'** button.




Figure 55 - Configuration File Copy and Rename Functions

5. Close the 'Configuration File' window and refresh the underlying UM display to see the changes in the configuration file list.

9.9.5 Upgrading / Reverting the ONTAP template linked to a Configuration File

The purpose of this section is to outline the processes used when upgrading and reverting the ONTAP version template used with the configuration files. With each version of ONTAP there are changes to the configurable options within the operating system. Upgrading or reverting the configuration file will ensure the defined options match those available in the defined version of ONTAP to ensure unsupported option commands are not sent to the Filer.



If you wish to change the ONTAP template linked to a configuration file it is recommended that you create a copy of the original configuration file first as the upgrade / reversion process is directly applied to the configuration file. Follow the procedure in the previous section to create a copy of the configuration file.

Use the following procedure to upgrade or revert the version of ONTAP linked to the template rules for a configuration file:

1. Select the **'Configuration Files'** option from the appropriate sub-menu in the **'Management'** menu.

2. Place a check mark in the '**Select**' box next to the configuration file you wish to upgrade or revert.
3. Select the target version of ONTAP you want the configuration file upgraded or reverted to.
4. Click on the '**Upgrade / Revert Selected**' button to implement the change.




Figure 56 - Configuration File ONTAP template upgrade or reversion

9.9.6 Viewing Configuration File Settings

The purpose of this section is to outline the processes used when viewing the settings defined in a configuration file. Use the following process to confirm settings that are defined in a configuration file:

1. Select the '**Configuration Files**' option from the appropriate sub-menu in the '**Management**' menu.
2. Click on the '**Edit**' hyperlink next to the configuration file you wish to view (see [Figure 53](#)).
3. From the '**Configuration File**' menu in the 'Configuration File' window select the '**View Settings**' link.




Figure 57 - Selecting the View Settings menu

4. The settings in the configuration file are displayed on the screen. Where the configuration file includes ONTAP system files you can click the '**View**' hyperlink to see the contents of the template file. The example configuration file on the next page defines the network services for a physical filer (DNS / NTP and time zone policies).




Figure 58 - Viewing the settings within a Configuration File

9.9.7 Creating a Partial Configuration File

The purpose of this section is to outline the process used when creating a partial configuration file. Partial configuration files can be created once a complete configuration has been pulled from a Filer (see section [9.9.3](#)). Partial files are much more manageable than a file with all configuration options included due to the large number of configurable parameters within ONTAP. Refer to section [9.9.13.7](#) to a more detailed discussion on why partial files are better. Use the following process to create a partial configuration file:

1. Select the **'Configuration Files'** option from the appropriate sub-menu in the **'Management'** menu.
2. Click on the **'Edit'** hyperlink next to the configuration file you wish to view (see [Figure 53](#)).
3. From the **'Configuration File'** menu in the 'Configuration File' window select the **'View Settings'** link.




Figure 59 - Selecting the Create Partial Configuration menu

4. In the '**Create Partial Configuration**' screen, place a check mark in the check-boxes for each option you want to include in the configuration policy. By default if you select a parent option, all sub options under the parent are included. By clicking on the [+] symbol you can deselect the sub options that you do not want to include in your policy.
5. Once you have chosen all the options you intend to include in the partial configuration policy click on the '**Update**' button.




Figure 60 - Selecting which options to include in a partial configuration

6. If you need to update the values for the options selected in the previous step, select the relevant parent option class from the '**Edit Option Groups**' list in the left pane. In this example we are going to choose '**options.cifs**'. You can also use this menu when editing a full configuration file.




Figure 61 - Selecting options for editing in a partial configuration

7. You can now update the option values for the items included in the partial configuration policy. In the following example the value of '**options.cifs.add_sid_history_to_credentials**' is set to '**on**'.
8. Once you have finished changing any values in your template that required editing, click the '**Update**' button to save the changes.
9. You can now close the 'Configuration File' window and return to the underlying OM screen.




Figure 62 - Editing option values in a partial configuration

9.9.8 Importing and Exporting Configuration Files

The purpose of this section is to outline the processes used when exporting and importing configuration files. Importing and exporting configuration files may seem redundant when considering these are included as part of the UM database backup but there are certain situations where the transfer of the configuration files between [Management Stations](#) is necessary. For example when you have more than one production Management Station in your environment and want to copy configuration files between systems. Use the following process to export all configuration files on a Management Station (all configuration files are saved into a single export file):

1. Select the '**Configuration Files**' option from the appropriate sub-menu in the '**Management**' menu.
2. Click on the '**To file on your computer**' radio button in the '**Export Configuration Files**' section of the screen.
3. Click the '**Export**' button.
4. The 'Save as...' dialog box is opened and all the configurations will be saved in a single export file. Choose a location and a name for the export file and click the '**Save**' button.




Figure 63 - Configuration export dialogue

This completes the procedure to export the configuration files from the Management Station.

The procedure to import configuration files to the Management Station is equally as straight-forward, but there are some behavioral characteristics worth mentioning:

1. Select the '**Configuration Files**' option from the appropriate sub-menu in the '**Management**' menu.
2. Click on the '**From file on your computer**' radio button in the '**Import Configuration Files**' section of the screen.
3. Click the '**Browse ...**' button to open the 'File Open ...' dialogue box. Select the configuration export file you would like to import to your Management Station. Click the 'Open' button on the 'File Open ...' dialogue box.
4. Click the 'Import' button to load the configuration files into the UM database.




Figure 64 - Configuration import dialogue

There are some important behavioral aspects of the configuration import process that should be noted:

1. A configuration file will not be fully imported if one of the configurations in the file already exists in the UM database. A configuration is considered the same if it possesses the same name as an existing configuration in the UM database, irrespective of whether the options in the configuration differ or not. Simply put the configuration will not be imported fully if the configuration name already exists. However the other configurations in the file will be loaded successfully into UM.
2. If the above error persists even if it appears that a configuration with a same name does not exist in the database it is possible that an orphaned object exists in the database. To rectify this situation follow the forcible configuration file deletion procedure in the next section.

9.9.9 Deleting Configuration Files

The purpose of this section is to outline the process used when deleting configuration files.



Before you attempt to delete a configuration file you need to check if the configuration policy has been deployed to any of groups in the UM group tree. As of release 4.0.1 OM does not perform a check to determine if a configuration policy has been applied to a group in UM. Deleting a configuration policy that is applied to group could create a condition where that policy is no longer enforced, thereby eliminating the benefits of centralized configuration management. Before deleting a configuration file, storage administrators are strongly advised to follow the manual checking process outlined below.

Use the following process to determine where configurations have been deployed in the UM group tree.

1. On the '**CONTROL CENTER**' tab, in the groups section of the left pane, click the '**Edit Groups**' link.




Figure 65 - Edit Groups

2. On the 'Edit Groups' screen look for groups that contain the group membership type of 'Configurations' (as shown in the diagram below) and click on the 'Configurations' hyperlink to show the list of configurations that are applied to that group. Repeat this process for all groups that contain configurations to ensure the configuration file you wish to delete is not actively applied to any groups in the group tree.

Full Group Name	Group Membership	Edit Move Copy Select
.Configurations	Empty	Edit Move Copy ☐
.Configurations/CATE Lab 10 Core	Configurations	Edit Move Copy ☐
.Configurations/CATE Lab 10 Core/CIFS NAMDEV	Configurations Appliances	Edit Move Copy ☐
.Configurations/CATE Lab 10 Core/CIFS SEED	Configurations	Edit Move Copy ☐
.Environment	Empty	Edit Move Copy ☐
.Environment/Production	Empty	Edit Move Copy ☐

Figure 66 - Group Membership List

This completes the check and review process for configurations that are applied to groups in UM.

Use the following process to delete a configuration file from UM.

1. Select the 'Configuration Files' option from the appropriate sub-menu in the 'Management' menu.
2. Place a check mark in the 'Select' box next to the configuration file you wish to delete from UM.
3. Click on the 'Delete Selected' button to implement the change.

Configuration Files					
Name	Version	Edit	Last Modified	Select	
CATE File CIFS General Policy	7.2.4	Edit	14 Aug 20:14	☐	
CATE File NFS General Policy	7.2.4	Edit	14 Aug 20:20	☐	
CATE Network DNS / NTP Policy Warren NAMDEV	7.2.4	Edit	14 Aug 19:56	☐	
CATE Security (CIFS Audit) Policy	7.2.4	Edit	16 Aug 00:35	☐	
CATE Security (System) Policy	7.2.4	Edit	15 Aug 18:04	☐	
CATE System Alerting & Monitoring Policy	7.2.4	Edit	14 Aug 19:32	☐	
CATE System Storage (Disk, RAID, WAFL) Policy	7.2.4	Edit	14 Aug	☒	
- Choose Version - Upgrade/Revert Selected					Select All - Unselect All
					Delete Selected

Figure 67 - Configuration File Deletion

This completes the process to delete a configuration file from UM.

On occasions where the configuration file cannot be deleted, or you are trying to import a configuration file but UM is reporting a configuration file of that name already exists and it clearly does not, then use the following process to force a deletion of the configuration file in question.

1. To rectify this situation logon to the [Management Station](#) and start a CLI session. Enter the following command to forcibly removed the orphaned configuration file entry from the database:

```
C:\>dfm config delete -f "<config_name>"
```

Where "<config_name>" is the name of the configuration you are trying to delete from the UM database.



If a configuration file is accidentally deleted it best to restore the configuration file by re-importing the configuration back into the database from an export file. A full database restore would involve overwriting a considerable amount of data in the database however a configuration file restore will simply restore the missing configuration file. Refer to the processes in section [9.9.8](#) regarding export and import of configuration files and section [9.2.1.11](#) where the backup via export of configuration files is recommended.

9.9.10 ONTAP Configuration Plug-ins

Each version of ONTAP released by NetApp is accompanied by an Unified Manager configuration plug-in file. This plug-in contains the template of available options within a given release of ONTAP. It's this template that ensures only the legal and valid configuration options are pushed to a Filer.

For more information about configuration plug-ins refer to section [9.9.14](#).

9.9.11 On Demand Configuration Comparison between templates and storage systems

The purpose of this section is to describe the procedure for conducting an ad-hoc / on-demand comparison of differences between a configuration file and the current configuration on a Filer / vFiler or group of Filers / vFilers.



When performing on demand comparisons of a large environment, divide the number of Filers or vFilers to check into groups of around 20 devices instead of running the comparison against the entire estate. This approach is recommended because the length of the results URL becomes too long to be handled via the UM Web Server.

Use the following process to perform an on demand comparison of configuration files (policies) to Filers / vFilers managed by UM. The following example is written in terms of a Filer but applies to vFilers as well:

1. Select the '**Compare Configurations**' option from the appropriate sub-menu in the '**Management**' menu.
2. Choose the configuration file (template) from the '**Template Storage System Configuration Filer**' drop down menu.
3. Choose either to run the comparison against a collection of Filers by selecting a group (optionally you can choose all Filers in the child groups by checking the '**Include Subgroups**' checkbox), or select individual Filers from the list on the right of the screen.
4. Once you have made your choices click on the '**Compare**' button.




Figure 68 - Storage System Configuration Comparison

5. The job submitted in the previous step will appear in the jobs table at the bottom of the screen. Refresh the browser window occasionally until the '**Status**' field reads '**Completed**'. Then click on the hyperlink that represents the job number to see the results.

Status of Compare Storage System Configuration Jobs							
Job ID	Template Config	Status	Submitted	Total	Done	Failed	Delete
20	CATE S... Alerting & Monitoring Policy v1.0(970)	Pending	21 Aug 17:20	2	0	0	
19	CATE File CIFS General Policy v1.0(972)	Completed	21 Aug 17:13	2	2	0	
Select All - Unselect All							Delete

Figure 69 - Storage System Configuration Jobs Queue

6. Two examples of the results screen are shown below. The first one is an example of a situation where no differences were detected, and the second an example of a setting on the Filer conflicting with the options as defined in the configuration template.

- a. Click on the page icon to the right of '**View Consolidated Report**' to download an Excel spreadsheet containing the template values and listing the values present on the storage systems.

Compare Storage System Configuration - Results			21 Aug 17:21
Template Configuration File: CATE System Alerting & Monitoring Policy v1.0 (7.2.4)			View 6a Download Report
Appliance	Job Status	Results	
nas3050qdse02.namdev.nsrootdev.net	done	No differences	
nas3050qdse01.namdev.nsrootdev.net	done	No differences	

Figure 70 - Storage System Successful Configuration Comparison

- b. Click on the page icon to the right of '**View Consolidated Report**' to download an Excel spreadsheet containing the template values and listing the values present on the storage systems.
- c. Click on the '**View Differences**' link to view the exceptions on a particular storage system.

Compare Storage System Configuration - Results			21 Aug 17:30
Template Configuration File: CATE Security (System) Policy v1.0 (7.2.4)			View 6b Download Report
Appliance	Job Status	Results	
nas3050qdse02.namdev.nsrootdev.net		No differences	
nas3050qdse01.namdev.nsrootdev.net	6c	View Differences	

Figure 71 - Storage System Differences Configuration Comparison

- d. Change the '**View**' drop down menu to show only mismatching options or report on all options for the given storage system.

Appliance Vs Template Configuration			21 Aug 17:30
View 6d Mismatching Options			
Option	Template File (CATE Security (System) Policy v1.0)	nas3050qdse01.namdev.nsrootdev.net	
options.httpd.enable	off	on	

Figure 72 - Storage System Configuration Comparison Differences Listing

9.9.12 On Demand Configuration Comparison between different configuration templates

The purpose of this section is to describe the procedure for conducting an ad-hoc / on-demand comparison of differences between two or more configuration files loaded on the UM server. This process applies whether the configuration files contain a full or partial configuration. If comparing partial configuration files only the entries that are common to both files are evaluated.



When performing on demand comparisons of a large environment, divide the number of configuration files into groups of around 10 files instead of running the comparison against the entire estate of templates. This approach is recommended because the length of the results URL becomes too long to be handled via the UM Web Server.

Use the following process to perform an on demand comparison of configuration files (policies) registered with UM. The following example is written in terms of a Filer configuration files but applies to vFilers as well:

1. Select the '**Compare Configuration Files**' option from the appropriate sub-menu in the '**Management**' menu.
2. Choose the configuration file (template) from the '**Template Storage System Configuration Filer**' drop down menu which you wish to compare to other templates.
3. Choose either to run the comparison against one Configuration File or a selection of Configuration Files.
4. Once you have made your choices click on the '**Compare**' button.




Figure 73 - Storage System Configuration File Comparison

5. When the comparison is completed the results are immediately presented. The example below demonstrates the situation where differences are found between configuration files.

- a. Click on the page icon to the right of '**View Consolidated Report**' to download an Excel spreadsheet containing the template values and listing the values present on the storage systems.
- b. Click on the '**View Differences**' link to view the exceptions on a particular storage system.




Figure 74 - Storage System Differences Configuration Files Comparison

- c. Change the '**View**' drop down menu to show only mismatching options or report on all options for the given storage system.




Figure 75 - Storage System Configuration File Comparison Differences Listing

9.9.13 Management Automation with Configuration template deployment

The purpose of this section is to describe the procedure for an efficient and effective adoption of centralized NAS configuration management by leveraging the configuration management tools in UM. Many regional bespoke processes have been created to ensure the Operating System options within ONTAP are configured in accordance with the NetApp NAS Filer standard build. Each process differs by region and there is not global consistency in the test points that are being checked. Transitioning to the UM Configuration Management toolset will allow Citi Storage Operations teams to follow one global standard process and reduce the develop burden of updating bespoke processes in each region.

The goal of this process is for CATE Storage and Backup Engineering to produce configuration template files that reflect the configuration best practices identified in the NetApp NAS Filer product Standard Build. The first release of these templates is expected with the release of ONTAP 7.3.3.



Configuration policy management via UM should be planned and implemented carefully to avoid accidentally changing the options within ONTAP on a Filer / vFiler or groups of Filers / vFilers. Leverage the best practices presented in this section to simplify centralized management and configuration conformance of your NAS infrastructure and avoid increasing risk during the transition from existing bespoke processes to a global standard process.

The following list represents a high level implementation overview of configuration management automation using Unified Manager.

1. Create a group tree structure in UM for configuration management that has the flexibility and granularity for your NetApp NAS environment (see section [9.9.13.1](#)).
2. Ensure the pre-requisites for configuration management have been implemented (see section [9.9.1](#)).
3. Assign Filers and vFilers to the groups created in step 1 (see section [9.9.13.2](#)).
4. Perform an on-demand comparison report between each configuration template and the Filers / vFiler in each group. Record exceptions and take appropriate corrective action (see section [9.9.13.3](#)).
5. Apply policies to group tree structure (see section [9.9.13.4](#)).
6. Configure policy overrides where necessary (see section [9.9.13.6](#)).
7. Long term configuration management and BAU best practices (see section [9.9.13.7](#)).

This concludes the presentation of the high level overview of configuration management automation using Operations Manager. The following sections discuss each of the high level items in detail.

9.9.13.1 Group tree structure

The purpose of this section is to describe the best practices to use when designing and implementing the group tree (container) structure in UM. It is not the purpose of this section to present the step to create groups (nodes) and move groups within the tree hierarchy. For these procedures refer to section [10.8](#).

The goal of this section is not to provide a rigid design structure for the group tree in UM but to provide Operational Owners with a set of infrastructure management logic that will allow you to create a structure that best suits your environment. This could be a site specific approach, or a line of business approach, but will most likely involve a hybrid, multi-level approach to configuration management.

When it comes to implementing configuration management policies it is clear that not all Filer configurations are the same. Some differ due to location because to locale specific network infrastructure and some differ due to use cases (NFS only environments compared with CIFS only or a mixed environment). To avoid duplicate and redundant policies, and the challenges of managing multiple policies that could conflict by configuring the same option in ONTAP, a structured sub-set policy approach is required instead of a one-policy does it all approach. This rationale is discussed further in section [9.9.13.7](#) but it's important to introduce it here as it will impact the design choices in how each group in the tree relates to each other.

The following example leverages the implementation used in the CATE lab, which only represents a subset of a true production environment, but should provide the building blocks necessary to implement your own group structure for your environment.




Figure 76 - Group hierarchy example for configuration management

Here are some behavioral principles around group structure as it applies to configuration management:

1. Configuration policies applied to parent groups are inherited by child groups
2. Configuration policies inherited by a group are applied before the policies assigned explicitly to that group
3. When multiple configuration policies (subsets of a whole policy) are applied to the same group, the policies are applied in the order in which each configuration file is displayed in the configuration resource group list. For example in the following list the '**CATE File CIFS General Policy**' is applied first and the '**CATE System Alerting & Monitoring Policy**' is applied last.

Configuration Files				
Name	Version	Edit	Last Modified	Select
CATE File CIFS General Policy	7.2.4	Edit	14 Aug 20:14	☐
CATE File NFS General Policy	7.2.4	Edit	14 Aug 20:20	☐
CATE Network DNS / NTP Policy Warren NAMDEV	7.2.4	Edit	14 Aug 19:56	☐
CATE Security (CIFS Audit) Policy	7.2.4	Edit	16 Aug 00:35	☐
CATE Security (System) Policy	7.2.4	Edit	15 Aug 18:04	☐
CATE System Alerting & Monitoring Policy	7.2.4	Edit	14 Aug 19:32	☐

Figure 77 - Configuration files order of implementation

4. If you move a storage device (Filer or vFiler) between different groups in the tree then the override or exclusion policies for that device do not move with it to the new group. Always check whether override or exclusion rules exist for the storage device before moving it to a different configuration group. See section [9.9.13.6](#) for further details.

Now let's discuss the steps involved in implementing this design.

1. Create a group called '**Configurations**' or '**Configs**' as a direct sub group of the '**Global**' root group. Using this name allows the purpose of this branch to be distinguished from other groups, such as reporting or lines of business structures, within the hierarchy.
2. Next create a group under '**Configs**' that describes your entire environment. In this example it's called '**CATE Core**'. For your environment use a name that best describes it, either as a region or as a support organization. This group is used as the node where all policies common across your environment can be applied, such as general security policies that are consistent to every Filer.
3. Next create a group under the one created in step 2 that represents each physical site in your environment. In this example it's called '**Lab 10**'. For your environment use a naming scheme that's meaningful to your support organization. Repeat this step until all your sites are represented. This group is used as a node where all site specific policies such as DNS servers and NTP servers can be defined.
4. Depending on the scale and complexity of your environment this last step may be optional. Finally create a group under the one created in step 3 that allows for different types of NAS implementation at the site defined in the previous step. In this example we have two different A/D domains at the same site: '**NAMDEV**' and '**SEED**'. Each domain has a different DNS and domain configuration hence the two groups. For your environment you may wish to have one group for CIFS implementations and one for NFS implementations. Determine what works best for your environment and repeat this step to create groups under each site as many times as you feel you have unique configuration subsets in your environment.

This concludes the discussion on the group tree structure design best practices for configuration management. The next section discusses why separate group structures are not required for Filer and vFiler configuration management.

9.9.13.2 Appliance and vFiler group membership

This section describes why separate group structures are not required for Filer and vFiler configuration management. Filer and vFiler configuration policies are controlled separately in UM. UM implements separation of policy types by:

1. Providing a separate interface to control Filer and vFiler configurations (see highlighted designated areas in the diagram below).




Figure 78 - Separate Filer and vFiler Configuration Management

2. Only allowing Filer configurations to be applied to Filers.
3. Only allowing vFiler configurations to be applied to vFilers.
4. Preventing the accidental import of a vFiler configuration into the Filer configuration list.

Therefore it is permissible and preferred to assign both Filers and vFilers into the same group. This reduces the complexity of the group structure and makes the environment easier to manage.

9.9.13.3 Configuration template deployment quality assurance process

This section describes the process by which to conduct a quality control check prior to deploying a configuration policy to a group of Filers or vFilers. This process has been presented in a prior chapter. Refer to section [9.9.11](#) as to how to perform an on-demand comparison between a template policy and a group of Filers or vFilers by selecting one of the groups in the tree that contains the storage devices upon which you are attempting to implement configuration management.

If the on-demand comparison identifies any exceptions then address the exceptions appropriately either by updating the option to the required value or by adding a configuration exclusion rule. Refer to section [9.9.13.6](#) for the process to implement an exclusion rule.

9.9.13.4 Applying policies via the Group Tree Structure

This section describes the procedure used to deploy configuration policies to groups within the UM tree structure. The example in this section focuses on Storage System (Filer) configuration but the process for vFiler configuration is very similar (refer to the discussion in section [9.9.12](#)).



One important pre-requisite of storage system configuration management via groups is you cannot apply a policy to a group unless there are Filers / vFilers assigned to the group or any of its child groups. Refer to section [10.8](#) for the procedure on assigning storage devices to groups.

The following example provides the process for applying and distributing (pushing) configuration policies to Filers:

1. Select the '**CONTROL CENTER**' tab from the main UM screen.
2. Expand the '**CONFIG**' part of the '**GROUP TREE**' and navigate down the tree to the level at which you wish to apply your configuration policy. In this example a policy will be deployed to the '**CIFS NAMDEV**' group. Refer to section [9.9.13.1](#) for the discussion of the group tree layout for configuration management.
3. Click on the '**Edit Storage System Configuration**' link from the '**Current Group**' menu in the left hand pane to get to the '**Edit Group Storage System Configuration**' screen.
4. Select a policy you wish to add to the group using the drop-down menu from the '**Add Configuration Files**' table.
5. Click the '**Update**' button to apply the policy to the group.
6. Once the policy is applied it will appear in the '**Configuration Files**' table.

The screenshot shows the 'Edit Group Storage System Configuration' interface. On the left, the 'Groups' tree is expanded to 'CIFS NAMDEV' (callout 2). Below it, the 'Current Group' menu has 'Edit Storage System Configuration' selected (callout 3). The main area shows the configuration for the group. At the top, 'Add Configuration Files' has a dropdown menu with 'CATE File CIFS General Policy v1.0' selected (callout 4) and an 'Update' button (callout 5). Below this is a table of 'Configuration Files' with columns: Name, Version, Edit, Modification Time, Order, and Remove. The table lists several policies, including 'CATE File CIFS General Policy v1.0' and 'CATE Network DNS / NTP Policy Warren NAMDEV v1.0'. At the bottom, the 'Push Configuration' section has a dropdown menu with 'All appliances in this group' selected (callout 7) and a 'Push' button (callout 8).

Figure 79 - Group Storage System Configuration Management

One item of note regarding the above example diagram is the policy 'CATE Network DNS / NTP Policy Warren NAMDEV v1.0' is applied to the group 'CIFS NAMDEV', however the other policies are inherited from the parent group 'CATE Core'. This is indicated in the demarcation between the two sections 'Included from' and 'Defined in this group' sections of the 'Configuration Files' table.



Before proceeding to step 7 ensure you have conducted the policy implementation Quality Control review process as outlined in section [9.9.13.3](#).

7. Once the policies are applied to the group they need to be pushed to the storage devices. This can be done either on a per Filer basis or all Filers in the same group. From the '**Push Configuration**' section, use the drop-down menu to select either '**All appliances in the group**' or a particular **Filer name** option.
8. Click on the '**Push**' button to instigate the configuration deployment process.
9. The screen will refresh and the job will appear in the '**Status of Configuration Jobs**' table at the bottom of the screen. The initial state will show as pending while the job is executing.

Status of Configuration Jobs						
Job ID	Status	Submitted	Total	Done	Failed	Delete
22	Pending	22 Aug 12:44	2	0	0	☐
Select All - Unselect All						Delete

Figure 80 - Configuration Management Job Pending

10. Refresh the browser window after a minute (maybe longer depending on the number of Filers you are pushing the configuration to) and the '**Status**' field should now show as '**Completed**'.

Status of Configuration Jobs						
Job ID	Status	Submitted	Total	Done	Failed	Delete
22	Completed	22 Aug 12:44	2	2	0	☐
Select All - Unselect All						Delete

Figure 81 - Configuration Management Job Completed

11. Click on the 'Job ID' entry to see the per device summary of the job as shown in the diagram below.

Run Status22 Aug 12:46

Summary for job 22

Job ID	Submitted	Target	Total	Done	Failed	Pending
22	22 Aug 12:44	974	2	2	0	0
Command						
config push -t filer 974						

 nas3050qdse01.namdev.nsrootdev.net

Status	Started	Completed	Attempts	Login Protocol	Use RLM Port
done	22 Aug 12:44	22 Aug 12:44	1	https	No

 nas3050qdse02.namdev.nsrootdev.net

Status	Started	Completed	Attempts	Login Protocol	Use RLM Port
done	22 Aug 12:44	22 Aug 12:44	1	https	No

Figure 82 - Configuration Job Run Status

12. Having pushed the configuration policy you can then confirm the policy has been applied successfully by doing an on-demand comparison each configuration file in the policy to the settings on the Filers affected. Use the procedure in section [9.9.11](#) to perform this task.

All configuration management processes are logged on both the Filer and the [Management Station](#). Refer to section [9.9.13.7](#) for examples of the log entries created by this process.

Certain configuration deployment (push) jobs may take longer to complete if there are network problems between the [Management Station](#) and the Filer, or if there is an issues localized to the Filer. UM has a default retry behavior for configuration push jobs so that the same job is attempted five times if an issue is encountered and will wait five minutes between each attempt. Neither the retry nor the wait value can be changed. To force a configuration job push retry use the following CLI command on the [Management Station](#):

```
C:\>dfm config push -R <filer_name>
```

Where <filer_name> is the name of the device to which you are trying to push the configuration.

9.9.13.5 Configuration Policy Conformance Checking

Having defined and deployed the configuration policy now is a good opportunity to explain how UM enforces the deployed policy.

Periodically UM will check the configuration running on each storage system against the configuration applied to the group in which the Filer resides in UM. This interval is determined by the option [ccMonInterval](#). The CATE Storage and Backup Engineering recommendation is to define this value as once a day which should be frequent enough to sufficiently check the environment without placing excessive workload on the [Management Station](#). The default setting from NetApp is once every four hours but this is considered too aggressive for large environments.

Where a configuration (ONTAP option) has been changed locally on a Filer and the defined value is now different to the value contained in the configuration policy a 'Storage Systems Configuration: Changed Locally' warning event will be generated in UM. This event will be generated which the configuration conformance checking process is executed (according to the frequency defined by the [ccMonInterval](#) option).

The event entry is shown in the screenshot below. Unlike other events, configuration conformance events have an additional 'FIX' button to enable a workflow process to revert the configuration deviation or accept it onto an override list (see section [9.9.13.6](#) for complete details on exclusions).

Follow the process outlined below when resolving a configuration conformance issue:

1. Select the '**STORAGE SYSTEM CONFIGURATION: CHANGED LOCALLY**' event from the event list
2. Click on the '**FIX**' button
3. You will then be presented with a screen showing the new value for the local option and the old value. The old value will be the value enforced by the configuration management policy. The new value is the option value defined locally on the storage device.

You can now take **ONE** of two options:

Click '**REJECT**' to revert the setting back to the value defined in the configuration policy
Or
Click '**ACCEPT**' to add the local setting to the override table for this option on this device




Figure 83 - Local Configuration Change Warning




Figure 84 - Resolving Local Configuration Change Conformance Issues

Any changes will be recorded in the appropriate audit logs (as discussed in section [9.9.13.7](#)).

9.9.13.6 Overrides and Exclusions

This section discusses the behavior overrides and exclusion rules with respect to Filer configuration management policies.



Exclusion and override rules should be used sparingly within the configurations groups as UM does not contain a reporting mechanism that will provide a recursive summary of all groups where overrides and exclusions have been implemented. This will make the task of removing the override or exclusion policy very time consuming when it is no longer required as the Storage Administrator would have to search every group that contains a '**configuration**' group membership type (see section [9.9.9](#) for the procedure to identify which groups have configuration policies applied).

Before working through an example it is worthwhile briefly comparing the difference between an override function and an exclusion function.

- **Override** is used when you need to replace the value of an option (or options) in a configuration policy with a different value when certain systems may have one-off deviations from a standard configuration (deviation)
- **Exclusion** is used when you do not need to check the value of an option and thereby exclude it from the list of options to enforce (push) to the device

The following example provides the process for implementing configuration policy override and / or exclusion rules:

1. From the '**Edit Group Storage System Configuration**' screen shown in [Figure 79](#), go to the '**Configure Overrides and Exclusions List**' section.
2. Select the device name, from the drop-down menu, to which you wish to apply to Override or Exclusion policy.
3. Click the '**Edit**' button.




Figure 85 - Configuration Override and Exclusion List

4. The '**Appliance Configuration Summary**' screen will show if any configuration overrides or exclusions have been applied to this device.

Appliance Configuration Summary22 Aug 12:48

Appliance Details			
Appliance Name	nas3050gdse01.namdev.nsrootdev.net		
OS Version	7.2.4P6		
Group	.Configurations/CATE Core/Lab 10/CIFS NAMDEV		
Overridden Options for nas3050gdse01.namdev.nsrootdev.net			
Options	Overriden Value	Group Value	Config File
No options overridden for nas3050gdse01.namdev.nsrootdev.net			
Excluded Options for nas3050gdse01.namdev.nsrootdev.net			
Options			
No options excluded for nas3050gdse01.namdev.nsrootdev.net			

Figure 86 - Appliance Configuration Summary

- From the 'Edit Options Groups' menu in the left pane, select the high level option type from the list. In the example below 'options.ssh' was selected.

Operations Manager NetApp

Edit Configuration pushed for Appliance: nas3050gdse01.namdev.nsrootdev.net OS Version: 7.2.4P6

Appliance

Summary

Preview Configuration

Edit Option Groups

- options.nfs_kerberos
- options.pcnfsd
- options.raid
- options.rsh
- options.secureadmin
- options.security
- options.snmp
- options.ssh**
- options.ssh1
- options.ssh2

Edit Option Group - options.ssh

Option	Group Value	Overridden Value	Use Group Value	Override	Exclude	Config File
options.ssh.access	*		☒	☐	☐	CATE Security (System) Policy v1.1
options.ssh.debug.enable	off	on	☐	☒	☐	CATE Security (System) Policy v1.1
options.ssh.idle.timeout	600		☒	☐	☐	CATE Security (System) Policy v1.1
options.ssh.passwd_auth.enable	on		☒	☐	☐	CATE Security (System) Policy v1.1
options.ssh.port	22		☒	☐	☐	CATE Security (System) Policy v1.1
options.ssh.pubkey_auth.enable	on		☒	☐	☐	CATE Security (System) Policy v1.1

Apply settings to nas3050gdse01.namdev.nsrootdev.net Update

Figure 87 - Selecting an option to override

- In the example below the default 'Off' value for 'debug.enable' is overridden with 'On'. Type the new value in the 'Overridden Value' field and click the 'Override' radio button. Note: if the option was going to be excluded rather than overridden just click on the 'Exclude' radio button; there is no need to enter an overridden value.

Edit Option Group - options.ssh 22 Aug 12:56

Updated settings for filer nas3050gdse01.namdev.nsrootdev.net.

Option	Group Value	Overridden Value	Use Group Value	Override	Exclude	Config File
options.ssh.access	*		☒	☐	☐	CATE Security (System) Policy v1.1
options.ssh.debug.enable	off	on	☐	☒	☐	CATE Security (System) Policy v1.1
options.ssh.idle.timeout	600		☒	☐	☐	CATE Security (System) Policy v1.1
options.ssh.passwd_auth.enable	on		☒	☐	☐	CATE Security (System) Policy v1.1
options.ssh.port	22		☒	☐	☐	CATE Security (System) Policy v1.1
options.ssh.pubkey_auth.enable	on		☒	☐	☐	CATE Security (System) Policy v1.1

Apply settings to: nas3050gdse01.namdev.nsrootdev.net

Figure 88 - Configuration Policy Override

- To apply the override or exclude policy settings, select the appropriate storage system or group from the **'Apply Settings to'** drop down menu. The choice of candidate systems will reflect the storage systems contained in the group to which this configuration is applied, or the name of the group.
- Click the **'Update'** button.
- A confirmation message will be displayed at the top of the screen.
- Click on the **'Summary'** link in the left pane to return to the summary view of the recently applied Override option.

Appliance Configuration Summary 22 Aug 12:57

Appliance Details

Appliance Name	nas3050gdse01.namdev.nsrootdev.net
OS Version	7.2.4P6
Group	.Configurations/CATE Core/Lab 10/CIFS NAMDEV

Overridden Options for nas3050gdse01.namdev.nsrootdev.net

Options	Overridden Value	Group Value	Config File
☐ options.ssh.debug.enable	on	off	CATE Security (System) Policy v1.1

[Select All](#) - [Unselect All](#) nas3050gdse01.namdev.nsrootdev.net

Excluded Options for nas3050gdse01.namdev.nsrootdev.net

Options

No options excluded for nas3050gdse01.namdev.nsrootdev.net

Figure 89 - Configuration Policy Override Summary

Additionally, use the above screen to clear policy overrides or exclusions by selecting the checkbox next to the option, selecting the system(s) from which you wish to remove the override and then clicking the **'Clear Overrides'** button. This same process applies to clearing exclusions.

9.9.13.7 Template Management, Logging and Change Control

This final section of the chapter reinforces some of the key points discussed regarding configuration management and provides some best practice recommendations for long term management of the environment.

Template Management

The goal is for CATE to provide a standard set of configuration templates based on the Citi NetApp NAS Filer Standard Build. However there may be occasions in a production environment when configuration policies specific to a certain environment or implementation need to be deployed. The purpose of the first part of this section is to present some high level design goals when creating your own configuration management files.

When designing and preparing configurations specific to your environment (a policy that defines the autosupport options is a good example) implement configuration files that deal specifically with the options you are trying to manage. By breaking down an overall configuration management policy into constituent elements the task of updating certain configuration options can be simplified as the file does not contain all the ONTAP options (of which there are many).

This subset of options approach also allows for policies to be split into groupings depending on the amount of systems to which these configurations apply. For example a global policy such as administrator account password composition would apply to all filers in the environment compared to a policy that may, for instance, define the NTP and DNS server IP addresses for a given site (a subset of your Filer environment).

In summary, the two main factors in grouping configuration options into a configuration file are:

1. Options that perform a similar function (for example alerting and monitoring), and
2. Options that have a similar scope of deployment (entire environment or a subset of devices).

Here's an example of an approach to dividing options into appropriately sized and applicable partial configuration files (see section [9.9.7](#) for reference). **Again the table below is an example to illustrate the concept, not a recommendation or standard implementation.**

Policy Name	Scope	ONTAP Options
System Security	Global	Telnet, Console, SSH, RSH & password composition
Alerting	Global	Autosupport
Storage	Global	Disk, RIAD and WAFL
CIFS	Global	CIFS options (excluding audit)
NFS	Global	NFS
CIFS Domain	Global or Site	CIFS
CIFS Audit	Site or Line of Business	CIFS audit
Monitoring	Site Specific	SNMP
Network	Site Specific	DNS and NTP
NIS / LDAP domain	Site or Line of Business	LDAP or NIS

Table 15 - Example of grouping configuration options

Configuration Management Logging

When considering the change to operational processes caused by moving from a 'per device' to a centralized configuration management process one of the common concerns is regarding how the events generated by the new process are logged. The purpose of this section is to indicate which logs contain configuration management events and where these logs are located ([Management Station](#) or Filer).

On the Filer, a configuration management deployment event is recorded in both the system (message) log (/etc/messages) and the system audit log (/etc/log/auditlog). The example below contains an entry from each log (respectively) for a configuration push event. Note the log entries from the Filer have been word wrapped for presentation purposes; however the event record from the log file on the Management Station is a multi-line event.

Filer - /etc/messages:

```
Tue Aug 12 16:54:52 EDT [nas3050gdse01: api_mpool_06:info]: DataFabric Manager
10.40.43.172 (admin 'NAMDEV\rs01146') pushed configuration to
nas3050gdse01.namdev.nsrootdev.net(70) [Message id#1218574487].
```

Note that the message log entry includes the Windows user ID of the administrator (NAMDEV\rs01146) which initiated the configuration push. This ID was obtained when the administrator authenticated with UM prior to executing the configuration management task.

Filer - /etc/log/auditlog:

```
Tue Aug 12 16:54:52 EDT [nas3050gdse01: api_mpool_06:debug]: dfmproxy:API:in:
<?xml version='1.0' encoding='utf-8' ?> <!DOCTYPE netapp SYSTEM 'file:
/etc/netapp_filer.dtd'> <netapp xmlns="http://www.netapp.com/filer/admin"
version="1.0"><system-cli><priv>advanced</priv><args><arg>syslog</arg><arg>1
</arg><arg>6</arg><arg><arg>syslog</arg><arg>DataFabric Manager 10.40.43.172 (admin
&apos;NAMDEV\rs01146&apos;) pushed configuration to nas3050gdse01.namdev.
nsrootdev.net(70) [Message id#1218574487].</arg></args></system>
```

Note that the audit log entry initially includes the user ID of the functional account used to manage the Filer (see section [10.5](#)) but also includes the Windows user ID of the administrator (NAMDEV\rs01146) which initiated the configuration push as part of the XML format message string.

Management Station - \dfm\log\dfmconfig.log:

```
Aug 12 16:54:52 [dfmconfig: INFO]: A config push job (id=7, submitted by user
'NAMDEV\rs01146') pushed config to storage system
nas3050gdse01.namdev.nsrootdev.net(70) successfully [Message id#1218574487].
Change list:
options.autosupport.support.enable:
Old Value=on
New Value=off
options.autosupport.support.transport:
Old Value=smtp
New Value=https
Aug 12 16:54:52 [dfmconfig:ERROR]: Unable to free SQL environment.
```

Note that the log on the Management Station contains some important information that is useful in change control purposes:

- The User ID of the storage administrator that initiated the job (in this case 'NAMDEV\rs01146')
- The Job ID (in this case '7')
- The outcome of the configuration job (in this case 'success')
- The device to which the configuration was deployed (in this case 'nas3050gdse01')
- The name of each option changed, plus the pre and post change option values

The last line of the Management Station log "unable to free SQL environment" is caused by a race condition within the database as the logging process exits. The race condition does not affect the overall performance of the Unified Manager server and the log line entry is expected every time this process is executed.

Change Control

In addition to the audit logging on the Filer and the [Management Station](#), CiTMP requires additional change control management processes to log any changes in the configuration management rules on the Management Station.

These are recorded as part of the event management portion of the UM database which is discussed in detail in section [9.12](#). For the purposes of relevance in this section only the events relating change control of configuration management events are listed below.

Both events are defined with a severity of informational, are part of the '**Configuration Changed**' event type and are classed as '**Config Group**' events.

The first event appears when an existing configuration file is changed or updated. The event entry indicates an item had changed but does not list the option in the configuration that has changed. Therefore when considering the deployment of a new or updated configuration file always use the quality control process outlined in section [9.9.13.3](#).




Figure 90 - Configuration File settings change event

This second event appears when a configuration file is added to a group. A similar event is generated when a configuration file is removed from a group.




Figure 91 - Configuration File add event

If a storage administrator changes the contents of a deployed configuration the event depicted in [Figure 90](#) is created. If a storage administrator changes the local configuration of a Filer then the events depicted in section [9.9.13.5](#) are created.

This concludes the advanced configuration management section as well as the UM storage configuration management best practices chapter as a whole.

9.9.14 ONTAP Configuration Plug-ins and UM Database Backup

The purpose of this section is to briefly present a known limitation in the UM database backup where ONTAP configuration plug-in files are not included as part of the database backup file.

The intended goal is to make configuration plug-ins files available to Storage Operations teams each time a new version of ONTAP is certified in Citi. Until this process is put in place, there is a potential for configuration plug-ins to be missing on a contingency UM server when restoring a production database as the plug-ins are not included as part of the backup. Use the following process to install a missing configuration plug-in on your UM server:

1. Configuration plug-ins are obtained from NetApp's Support website (support.netapp.com) and are listed as part of the release notes download page prior to downloading a copy of ONTAP.
2. Copy the configuration plug-in zip file to a location on your [Management Station](#).
3. Logon to the console of your Management Station and on the CLI navigate to the same directory where you saved the zip file from the previous step.
4. Run the following command to install the configuration plug-in:

```
C:\>dfm plugin add <plugin name>
```

Where "<plugin name>" represents the name of the plugin zip file.

9.10 UM SnapMirror (Replication) Monitoring

The purpose of this section is to present an overview of the snapmirror monitoring available within Unified Manager. This chapter will focus on purely the monitoring part of the replication process and not the management aspect. The reason the centralized management processes of replication are not included in this document is that how [UM](#) handles replication management is changing. CATE Storage and Backup Engineering felt it may be unwise to adopt policies and procedures based around the current UM replication management features when the scope of change in the product is not clear at the time of writing this document. Therefore the remainder of this chapter will focus on the monitoring aspect of replication.

9.10.1 Pre-requisites

The purpose of this section is to list the necessary pre-requisites required for successful monitoring of snapmirror processes. These items are listed below and then discussed in detail directly afterward.

1. Defined NDMP credentials in UM to allow snapmirror reporting of the source and destination systems.
2. Successful communication from the UM server to both the source and destination systems in the snapmirror relationship.

9.10.1.1 License

In order to enable the Disaster Recovery Manager tab in the UM GUI and thereby permit use of the dfdrm executable, the Business Continuity Management license key must be installed on the UM server. Refer to section [8.3.1](#) for the procedure to install this license key once the initial UM server build is complete.

9.10.1.2 NDMP

UM uses the NDMP protocol to interrogate the Filer regarding the status of its snapmirror processes. In order to achieve this, a set of NDMP credentials needs to be defined in UM for each storage device being monitored. Rather than having to create a separate functional ID for snapmirror monitoring, CATE Storage and Backup Engineering recommends creating the NDMP password hash of the dfmproxy account created in ONTAP for storage device authentication (see section [10.5](#) for full detail on the dfmproxy account).

Follow the process outlined below to obtain the NDMP password hash of the dfmproxy functional ID on each Filer you are intending to monitor with UM.

1. Login to the **ONTAP CLI** on each Filer
2. Type the command '**ndmpd password dfmproxy**' to obtain the ndmp password hash of the account. (If you chose not to use an account named dfmproxy in your environment, replace the term dfmproxy above with the name of the account you wish to use). It is recommended you choose the same account you used when defining the host login credentials from OM for each storage device (see section [9.7.1](#)).
3. Select the hashed password output and copy it (into memory) for pasting into the next screen.

The following steps present the process of adding the newly acquired hashed password on the UM database.

4. **Login to the UM web UI**
5. Select the '**DISASTER RECOVERY**' tab
6. Click on the '**AUTHENTICATE STORAGE SYSTEMS**' link in the left pane
7. Select one or more of the Filers from the '**UNMANAGED STORAGE SYSTEMS**' list. Use CTRL to select more than one. You can input the NDMP password for more than storage device at once assuming the NDMP hashes from each machine are the same because the actual account password on each machine is the same.
8. Type the user name in the '**NDMP User**' field (in this example we are using the functional ID dfmproxy)
9. Paste the ndmp password hash obtained in step 3 into the '**NDMP Password**' field
10. Click the '**UPDATE**' button to apply the credentials




Figure 92 - Configuring OM with NDMP credentials

While the above process is concluded at step 10 it is worth briefly mentioning the process to un-configure these systems in UM for NDMP communication as the undo functionality is not available in the GUI and the back-out process is not necessarily intuitive. In order to back-out steps 4 thru 10 above use the following process:

1. RDP (Terminal Services) to the UM server
2. Launch the CLI
3. Type the command 'dfm host set <hostname> hostndmlogin="" ' to remove the NDMP user credentials, where <hostname> represents the name of the host you wish to remove UM NDMP management
4. In addition type the command 'dfm host set nas3050gdse01 hostndmppassword="" ' to remove the NDMP user password.

In addition use the above process if you wish to change the NDMP password hash stored in UM (except input a password rather than using "" to signify a null password). It is recommended that when changing an NDMP password hash do not delete the storage system from the OM Disaster Recovery tab and then re-add it with the new NDMP credentials. Instead use the CLI command to update the password accordingly.

9.10.1.3 Communication

The final pre-requisite is to ensure communication protocols between the UM server and both the source and destination systems in the snapmirror relationship are working as expected.

Follow the process in steps 10 and 11 of section [9.7.1](#) to confirm connectivity. Look for an output similar to the screen shot in step [11](#), especially the NDMP related tests, to ensure connectivity is successful.

This concludes the configuration and validation steps required to ensure snapmirror monitoring works as advertised. Once these steps are completed, all snapmirror relationships will be automatically discovered in the environment.

9.10.2 Delegation

The purpose of this section is to present items for consideration if storage administrators wish to delegate read access to the Disaster Recovery tab to business aligned support groups in order to check the status of the replication process. This access is achieved by a combination of resource and group assignments using the processes discussed in section [10.8](#).

The snapmirror status of a Filer is only available to business aligned support personnel under the Disaster Recovery tab in UM if the physical filer resource is assigned to the businesses container in the UM group tree view. This is because all snapmirror relationships are indexed against the physical storage system. This may be a viable approach in situations where a line of business has exclusive access to a Filer, but there are complications in a shared storage model.

When assigning a full physical filer to a group all events triggered by the physical system will be available for review and acknowledgement in the UM database by the business aligned support team. These events will be confusing for the support team who will mainly be interested in volume and file system level events. From a security and management perspective it is best practice to restrict the processing of physical storage system related events to storage administrators only.

In a shared storage model, assigning a physical Filer object to a group will confuse business aligned support personnel from all lines of business using that system because any event triggered by any file system will show in the UI of all the businesses. As such CATE Storage and Backup Engineering does not recommend providing business aligned support personnel access to the disaster recovery tab

In order to obtain high level information about the state of the replication process for a given volume, business aligned SAs can use the information in [Figure 93](#) from the file system summary screen. It shows the other file system (be it source or target) in the relationship and lists the timestamp of the last snapshot used as the point in time marker for the most recent snapmirror process.

In order to gain a summary view of all the replicated file systems supported by a business aligned team CATE Storage and Backup Engineering recommends creating a custom report using the SNAPMIRROR_RELATIONSHIP base catalog, adding the necessary fields and assigning to the report to the 'File Systems' tab. For more information on custom reports refer to section [9.8.6](#).




Figure 93 - Volume summary snapmirror status

9.10.3 Replication Events and Alarms

For event and alarm information specific to snapmirror replication refer to section [9.12.5](#) of the UM Events and Alarms section.

9.11 UM SnapVault (Backup) Monitoring

The purpose of this section is to present an overview of the snapvault monitoring available within Unified Manager. This chapter will focus on purely the monitoring part of the disk based backup process and not the management aspect.

Snapvault monitoring share similar prerequisites to snapmirror monitoring:

1. Defined NDMP credentials in UM to allow snapvault reporting of the primary and secondary systems.
2. Successful communication from the UM server to both the primary and secondary systems in the snapvault relationship.

However the current implementation of snapvault monitoring in UM has one significant limitation in terms of the discovery and monitoring of existing snapvault relationships.



UM 5.1 and above will only discover a snapvault relationship if the file system on the physical primary storage is a qtree snapvaulting to a qtree on the physical secondary storage. If the primary storage file system is a volume then UM cannot discover and monitor the relationship.

For monitoring from Physical Filers SnapVault to vFilers or vice versa, consult the following [Knowledge Base Article](#).

For further information on monitoring snapvault relationships consult the vendor [Admin Guide](#).

9.12 UM Events and Alarms

The purpose of this section is to present an overview of the best practices regarding event and alarm management in Unified Manager. This chapter initially provides an overview of events and then focuses on best practices surrounding event (problem) management. It will then move onto configuring alarms, including ways of simplifying alarm management, ending in a focused review of configuring replication (snapmirror) based alarms (a topic which has historically demonstrated the requirement for a deep dive review at Citi). Let's begin with a high level review of UM's event management functionality.

9.12.1 Event management overview

The monitoring schedules¹² within UM allow for the generation of predefined events. These events are classified as emergency, critical, error, warning and informational. A fully listing of events can be found in section 11.5 for reference purposes. UM events deal specifically with the health of the storage systems (problem management) as compared to general event management systems that can be used for other purposes such change control and auditing. Understanding this design specification of UM will assist in explaining why certain events are only listed as Informational event.

Events can be accessed globally from the '**EVENTS**' link in the UM application banner or on a host or file system basis from the relevant tabs in the '**MEMBER DETAILS REPORTS**'. The number of events listed in the bar corresponds to the scope of user access. An UM user with global access will see all events, whereas a user with only rights to view one group (subset) will only see the relevant events for that group. The events shown in a group will depend on the resource memberships assigned to the group (see section 10.8.2 for a complete overview of resource and groups assignments); thus if only file systems are assigned to a group then only file system related events will appear to users with RBAC access to the group. Events that relate to the underlying physical storage system will present themselves in the group to which the storage controllers are assigned.




Figure 94 - Event Toolbar

Events - <u>Current</u> , History				SNMP Traps (5)
Severity	Event	Condition	Triggered	Ack'd by
!	Storage System Configuration: Changed Locally		01 Oct 21:57	

Figure 95 - Storage object specific event example

Events are triggered when an error condition occurs, such as a disk failure, or when a predefined threshold has been breached, such as the amount of used space in a volume has exceeded 85%.

Once an event is triggered it is the cue for a Storage Administration to investigate the cause of the event and take the appropriate corrective action. There are recommended best practices in acknowledging and deleting events from the UM database and these are presented in the following section.

¹² [ccMonInterval](#), [cfMonInterval](#), [cpuMonInterval](#), [dfMonInterval](#), [diskMonInterval](#), [envMonInterval](#), [fsMonInterval](#), [hostRBACMonInterval](#), [ifMonInterval](#), [licenseMonInterval](#), [opsMonInterval](#), [pingMonInterval](#), [snapmirrorMonInterval](#), [snapshotMonInterval](#), [statusMonInterval](#), [sysInfoMonInterval](#), [vFilerMonInterval](#)

9.12.2 Event management best practices

The purpose of this section is to provide best practice procedures when managing events in UM. It is important to understand the implications of acknowledging and deleting events in the UM database, and the best way to communicate this is through the use of the following example.

When an event is created in UM it is a result of an error condition or a predefined threshold being breached. When the storage administrator begins to action the event the first step is to acknowledge the event using the 'ACKNOWLEDGE' button on the event management screen. This indicates to other administrators that you are aware of the problem, are actively working on it and are the problem owner.

Let's introduce an example here and use a volume over 90% full as our condition. Once the event appears on the UM console, one of the storage administrators in the team acknowledges the event to begin working the problem. If, after taking the appropriate action, the storage administrator is capable of reducing the utilization rate of the volume (either by asking the data owner to delete some data on their file system or by increasing the size of the volume) the event condition will automatically be set to a deleted state as the event condition no longer exists. Deleted events can be accessed and reviewed by using the 'History, Deleted Events' report.

The above example describes an ideal use case, however not every administration function will be a text book case. Therefore we need to understand what would happen if 1) the event was acknowledged and no administrative action was taken and 2) if the event was deleted without administrative action being taken on the storage device. Let's begin with the first situation.

If an event is acknowledged in UM but no corrective action is taken on the storage system then the event will continue to exist in UM. Allowing UM to operate continually with more than 1000 events of severity warning or worse will have a long term degradation effect on the performance of the database. Follow the review procedure suggested in section [9.5](#) to reduce the number of open events in the UM database).

In the second situation where an event is deleted in UM and no corrective action is taken on the storage system all subsequent events for that attribute of the monitor object (in our example this is the utilization rate on a certain file system) will be suppressed until the condition that generate the event is corrected on the storage system. Deleting the event without addressing the cause will have the following effects:

- The error condition on the storage object (file system, vfiler, storage controller, etc) will be cleared in the UM database. Unless a lower severity event also exists on the same storage object the event status on the object will turn green (effectively signifying an all clear / no errors condition).
- Creates a contradictory situation in UM whereby a pre-defined utilization threshold has been triggered and surpassed, but there is no longer a corresponding threshold breach event in the database. Upon closer examination an experienced storage administrator may notice the actual utilization error condition exists however because the error event has been cleared from the database the threshold breach is not flagged as an error condition for immediate attention.
- Event deletion without addressing the cause of the event will act as event suppression. Do not expect UM to generate a new event the next time it collects the various system metrics as the deleted event will effectively suppress all subsequent events until such a time as the condition that cause the event is actually addressed on the storage device.



The bottom line in event management best practices is you must either take corrective action on the storage device to correct the error condition, or if the event is threshold based consider increasing the size of the resource pool (file system, network bandwidth etc) or increasing the defined threshold for that object. Remember that customizing the threshold for many storage objects within the UM database will result in an increased administrative overhead.

Do not delete an event in the UM database without first taking the appropriate action to address the cause of the event on the storage device.

CATE Storage and Backup Engineering has implemented a policy as part of the UM standard build that will automatically purge deleted events from the UM database after a defined period of time. This time period is controlled by the option [eventsPurgeInterval](#). The goal of this setting is to reduce the number of redundant events within the UM database to maintain application performance. Event purge and clean-up related activities are recorded in the UM log file DFMEvent.log, located in the %INSTALLDIR%\DFM\log directory.

9.12.3 Configuring alarms

The purpose of this section is to present the processes involved in configuring alarms within Unified Manager. Alarms extend the functionality of events as they provide the notification framework to alert on the occurrence of an event or predefined set of events.

UM offers an alarm feature to notify storage administrators and SAs of issues with the storage environment. Alarms are a notification component of an event when a predefined threshold or error condition has been triggered. These alarms are dispatched via email and contain a link to the event record in the UM database. This allows notification of events to be targeted at specific support personnel without the need to be constantly viewing the UM console.

Storage administration teams should carefully consider the amount of alerts generated and possibly look to limit which events should generate an alarm; otherwise the effectiveness of the system will be compromised due to spamming.



Before configuring any alarms ensure security groups have been configured as discussed in sections [10.8.1](#) and [10.8.4](#).

Use the following process to configure alarms:

1. With the '**CONTROL CENTER**' tab highlighted, click on the '**SETUP**' menu and choose the '**ALARMS**' option.




Figure 96 - Alarm setup menu

2. On the '**ALARMS**' screen, click on the '**ADVANCED VERSION**' hyperlink next to the 'ADD' button. The advanced version of the alarm configuration screen provides the greatest flexibility when defining alarm policies.
3. Select the '**RESOURCE GROUP**' to which you wish to apply the alarm policy
4. Choose whether you wish to create a policy based on '**SEVERITY**' level or '**EVENT NAME**'

5. Define the '**RECIPIENTS**' of the alarm
 - a. For recipients who already have a login to UM provide the login ID of the user in the '**EMAIL RECIPIENTS – ADMINS**' field. Note the SMTP email address for the user needs to be defined (see section [10.8.4](#) for details).
 - b. For recipients who do not have or for some reason cannot have a login to OM then use the '**EMAIL RECIPIENTS – NON-ADMINS**' field to define the target SMTP address.
6. Define the '**OTHER OPTIONS**' of the alarm (if required)
 - a. To create an alarm policy that only runs during core business hours define the '**TIME FROM**' and '**TIME TO**' fields.
 - b. Use of the repeat alarm option is not recommended as the email tends to get treated as spam after a while and therefore defeats the purpose of the functionality.
 - c. Use the '**DISABLE**' alarm option to temporarily suspend an alarm notification.

Alarms 07 Oct 11:23

Add an Alarm (Advanced)

Conditions

Group:

- Group
 - .Configurations
 - .Environment
 - .Production**
 - Contingency
 - Development
 - Lab
 - Support
 - Businesses
 - Report Groups

Select Event By

☒ Severity Critical or Worse

☐ Event Name All

Event Class
(regular expression to select a class of events e.g. env\temp)

Recipients

Email Recipients - Admins
(specify full login names of administrators)

Email Recipients - Non-Admins
(specify email addresses)

Page Recipients - Admins
(specify full login names of administrators)

Page Recipients - Non-Admins
(specify pager addresses)

SNMP Trap Hosts

Other Options

Time From 00:00

Time To 23:59

Repeat Notify No

Repeat Interval 30

Disable No

[Simple Version](#) **Add**


Figure 97 - Alarm configuration

- Click the **'ADD'** button to create the alarm policy

9.12.4 Simplifying alarm management

The purpose of this section is to provide some useful suggestions to simplify alarm management in UM.

- Each storage object (appliance, file system, vFiler etc) has the capability of having an ownerEmail address assigned to it. If an event is triggered for that object and a valid alarm policy exists then the address defined in the ownerEmail field will also receive the alarm; irrespective of whether they are defined in the alarm recipient list or not. To simply alarm management decide whether to purely rely on the ownerEmail field for defining recipients of alarms or whether to use the defined list of recipients from the alarm policy. Using a mixture of both will complicate management of the alarm policies.
- Use the **'SHOW ALL COLUMNS'** view in the alarm management screen to determine which alarm has been temporarily disabled.




Event	Severity	Group	Recipients	Edit	Test	Delete
Volume Almost Full	All	Businesses/CATE/PSE/GDSE (Storage Engineering)/Elizabeth Abraham	✉ elizabeth.abraham@citi.com	edit	test	☐
All	Emergency Only	Global		edit	test	☐
Configuration Changed: Config File	All	.Configurations		edit	test	☐
Configuration Changed: Config Group	All	.Configurations		edit	test	☐
Configuration Changed: Local Config	All	.Configurations		edit	test	☐

[Show all columns](#)

[Delete Selected](#)

Figure 98 - Alarm policies list



Event	Severity	Event Class	Group	Recipients	Time From	Time To	Repeat Interval	Disable	Edit	Test	Delete
Volume Almost Full	All		Businesses/CATE/PSE/GDSE (Storage Engineering)/Elizabeth Abraham	✉ elizabeth.abraham@citi.com	00:00	23:59		No	edit	test	☐
All	Emergency Only		Global					Yes	edit	test	☐
Configuration Changed: Config File	All		.Configurations					No	edit	test	☐
Configuration Changed: Config Group	All		.Configurations					No	edit	test	☐
Configuration Changed: Local Config	All		.Configurations					No	edit	test	☐

[Show fewer columns](#)

[Delete Selected](#)

Figure 99 - Alarm policies advanced view

9.12.5 Replication alarms

The purpose of this section is to provide an overview of replication related alarm policies.

It is important to remember that the Disaster Recovery Manager component of the UM server does not generate events when lag thresholds of SnapMirror sources are crossed. Only lag thresholds of SnapMirror destinations are used for generating events.

Replication alarms can be defined at four different levels in UM:

- Global (root of the group tree)
- All snapmirror sources and destinations in a specific group
- All mirrors of a snapmirror source
- A specific snapmirror source or destination

Replication lag thresholds can be changed from the Edit Policy page (Disaster Recovery tab > Home link > SnapMirror relationship Source > Manage Replication Policies icon > Policy name). The following lag thresholds can be changed:

- **SnapMirror Lag Warning Threshold**
This option specifies the limit at which the SnapMirror destination contents are considered nearly obsolete. If this limit is exceeded, Disaster Recovery Manager generates a SnapMirror Nearly Out of Date event.
You can specify the lag warning threshold limit in weeks, days, hours, minutes, or seconds. For example, to specify 20 hours and 15 minutes, enter 20.25 hours or 20:15:00 as the value for this option.
- **SnapMirror Lag Error Threshold**
This option specifies the limit at which the SnapMirror destination contents are considered obsolete. If this limit is exceeded, Disaster Recovery Manager generates a SnapMirror Out of Date event.
You can specify the lag error threshold limit in weeks, days, hours, minutes, or seconds. For example, to specify 20 hours and 15 minutes, enter 20.25 hours or 20:15:00 as the value for this option.

This concludes the section on alarm and event management best practices.

9.13 UM Autosupport

The purpose of this section is to provide a brief overview of the Autosupport functionality within UM. Similar to the functionality of the same name in the ONTAP storage operating system, the Autosupport component of UM has the ability to send call home status information to NetApp. The options to configure Autosupport are found under the '**CONTROL CENTER / SETUP**' menu and then by selecting the '**AUTOSUPPORT**' option from the '**EDIT OPTIONS**' list in the left pane. Alternately they can be changed through the 'dfm option set' CLI command set.

The table below presents a summary of the configurable options relating to Autosupport. Clicking on the hyperlink in each row will take you to the relevant entry in the appendices that provides a description of the option, its function, default value and recommended value.

Option Name (GUI)	Option Name (CLI)
Autosupport Administrative Contact	autosupportAdminContact
Autosupport Content	autosupportContent
Autosupport Destination Email	autosupportDestinationEmail
Autosupport Enabled	autosupportEnabled
Autosupport Protocol	autosupportProtocol
Autosupport Retry Count	autosupportRetryCount
Autosupport Retry Delay	autosupportRetryDelay

Table 16- Autosupport Configuration Options

You may have noticed that there is no option for configuring a proxy server for https Autosupports. UM is configured with a transparent proxy functionality which essentially leverages the proxy configuration from Internet Explorer. For Autosupport messages to be posted by https you should ensure that the web proxy server settings are configured on your UM server using the Citi Internet Explorer TPM.

The Autosupport function in UM is used by NetApp to ensure the monitoring and management software is running as expected in the environment and can be used to diagnose issues when cases are raised. The Autosupport process runs on a weekly pre-defined schedule similar to that of ONTAP but is separate to the ONTAP autosupport process. The UM Autosupport function is not an ONTAP autosupport monitoring function.

The Autosupport functionality can be tested using the Windows CLI on the UM management station. Use the following command to test the Autosupport component via the default HTTPS protocol:

```
C:\>dfm autosupport send -p https -u support.netapp.com/asupprod/post/1.0/postAsup
```

This concludes the brief overview of the Autosupport functionality within UM.

9.14 Performance Monitoring and NMC Performance Advisor

The purpose of this section is to discuss some of the nuances of performance monitoring on the NetApp platform (especially as it relates to CPU load monitoring), and the new certified performance advisor client.

It is not the purpose of this section to provide performance data analytics or the ability to perform data interpretation as both these topics are well beyond the scope of this document. This chapter is more focused on using and configuring the appropriate tools. Let's begin by reviewing UM with regards to CPU performance data collection and presentation.



Ensure monitored Filers are running ONTAP 7.2 or above before implementing a Performance Advisor data collection regime. Refer to the known issues list in the [CATE UM Release Notes](#).

9.14.1 Filer CPU performance monitoring

The purpose of this section is to focus in on how CPU performance is captured and measured by Unified Manager. This information has been added in an attempt to address some of the historic confusion about the subject of reporting CPU performance.

UM polls each storage device at a defined frequency in order to obtain the status of the CPU(s) on each device. This frequency is determined by the option [cpuMonInterval](#). By default this is set to 5 minutes and the Citi implementation uses the default value. It is not recommended to change this option to a smaller value (thereby preventing performance issues on your UM server) and conversely not change the value to a large value (thereby averaging out any peaks in the sampling). Like all monitoring interval options in UM changing this frequency will change the interval for all systems. It is not possible to customize the sampling frequency for different storage systems.

Remembering the sampling frequency it is also important to note that the plotted CPU performance in Unified Manager will be the average utilization rate of the CPU over the monitoring interval. This averaging behavior is the same for all performance counters captured during sampling intervals.

The reporting of CPU performance is mainly dependent on the version of ONTAP running on the storage system. Many improvements have been delivered as part of ONTAP 7.2 and above. The most significant of the changes with 7.2 is RFE 173714 which improves the CPU usage summary available via SNMP and ONTAP Counter Manager (CM) ZAPs. Previously, prior to ONTAP 7.2, SNMP used a conservative algorithm for reporting CPU usage whereby any CPU cycles not available for any possible use were counted as busy. The new algorithm reports CPU cycles as used only if they cannot be utilized by other processes in the current workload. Thus the new CPU busy metric is a much more accurate statement of the fraction of available CPU resources that are in use. For most workloads this new metric will report a lower CPU usage than the previous algorithm.

The new CPU utilization algorithm is the following:

Use "average idle time per CPU" instead of "total system idle time" when computing the CPU usage. This simply involves dividing the value currently being used (total idle time) by the number of active CPUs in the system (average idle time). This calculation can be summarized as follows. Compute the maximum processing domain utilization and compute the average CPU utilization and essentially report the higher value.

Finally, RFE 204800 remains open and is pending inclusion in a later release of ONTAP. RFE 204800 relates to SNMP CPU information only being available for a single CPU on multi-CPU system. This means for tools such as Unified Manager and Performance Advisor that use SNMP to gather performance information, the utilization rates on each CPU on a multi-CPU system will have to be gathered and then aggregated manually to get the combined performance of all CPUs in the storage system. This means on multi-CPU systems you need to gather all processing domain information on each CPU / core.

9.14.2 Performance advisor

The purpose of this section is to provide an introduction to performance advisor, highlighting areas of special interest and detailing configuration specific to the Citi NetApp NAS Filer environment. It is not the intention of this section to quote the vendors documentation verbatim and storage administrators are strongly advised to read [NetApp's Performance Advisor Admin Guide](#) document to understand the concepts and general best practices involved in using the product.



Refer to the Known Issues chapter of the [CATE UM Release Notes](#) for issues when using Performance Advisor:

1. Changing the Performance Advisor `perfTopnchartinterval` setting might require you to restart NetApp Management Console
2. Client statistics data not collected for NFS and CIFS objects in storage systems running Data ONTAP 8.1 (7-mode)

The following diagram shows the architecture and interrelationship of the NetApp Management Console (NMC), UM server and the storage devices. The only Citi supported configuration is to install the NMC TPM (see section [8.1.3](#) for details) on another server rather than installing on a workstation. When using the TPM approach you should not install the NMC on the same server hosting UM.




Figure 100 - Performance Advisor architecture overview

Performance advisor runs as a component of the UM server and is activated using the [perfAdvisorEnabled](#) option. By default all NetApp devices registered in UM are polled for performance data. It is possible to disable performance advisor on a per device basis and this process is discussed later in this section. The device polling interval used by performance advisor is defined by the [perfAdvisorPollInterval](#) option. It is strongly recommended that the default polling setting of 5 minutes is maintained.

Performance advisor does not require any additional software licenses in ONTAP or UM – it is included as part of the core UM license (refer to section [8.3.1](#) for best practices and known issues when licensing a second instance of UM to run Performance Advisor in your environment). Login to NMC is achieved with the same credentials as you use with UM.

All performance data gathered by Performance Advisor is recorded in a proprietary flat file format external to the UM database. It is important to note that while the data is not part of the database, the data files are included as part of the database backup archive file.



Due to the inclusion of Performance Advisor data files in the [UM database backup process](#), performance advisor data collection is disabled during the database backup process.

9.14.2.1 Performance Advisor versus PerfStat

If you currently use NetApp's Perf Stat tool in your environment you are strongly advised to cease use of the tool once Unified Manager and Performance Advisor are deployed and configured in your environment. Performance advisor has the following benefits over Perf Stat:

- Ease of use
- Authentication built into Unified Manager (no need to configure SSH public keys)
- Graphical data readily available (no need to upload data file to NetApp for processing)
- Availability of historical performance data so that performance data is not missed when reviewing an unexpected event or problem

Perf Stat should only be used if specifically requested by NetApp support as part of a troubleshooting case.

9.14.2.2 NMC Software Package / Supported Platforms

CATE Storage and Backup Engineering has packaged the NMC as a TPM for Windows Server 2008 (see section [8.1.3](#) for details). It is recommended not to install NMC and the performance advisor client on the same server on which UM is installed so CATE Storage and Backup Engineering recommends installing the TPM on the dedicated terminal server used in your environment for storage management activities. Using the TPM version of the NMC offers the following advantages compared to the Vista client version that was originally released with the first internal release of OM 3.7:

1. Can be installed in a Data Center to improve communication performance with an UM server
2. With the trend in remote working, performance of the NMC over VPN is less than desirable
3. Vista has not been fully deployed within Citi and so certain teams will not have a valid platform to install the Vista OVCM package on

9.14.2.3 Performance Advisor Access

In versions of Unified Manager prior to 4.0.1 it has not been possible to delegate access to the performance data in Performance Advisor to an individual outside of the NAS Operations team due to issues with the Role Based Access Controls in the NetApp Management Console. These issues (BURTs [311560](#) and [311567](#)) have been fixed starting with Operations Manager 4.0 meaning that the security concerns no longer apply. However this should not directly imply that CTI NAS Operations should consider granting access to other teams. In the end this is an operational decision but consider that there is usually a considerable amount of training provided to technologists in performance analysis and granting access to the untrained eye may lead to more questions and confusion than any quantifiable benefit. The CATE recommendation is to limit access to Performance Advisor to individuals who are trained and instructed in reviewing performance data from NetApp Filers.

9.14.2.4 Diagnosing NMC connectivity issues

If you are having issues logging into the NMC, use the following steps to diagnose the issue:

1. **PING** the UM server you are trying to connect to with the NMC client to ensure you have basic network connectivity
2. Click the '**OPTIONS>>**' button on the NMC client, ensure the protocol is set to HTTPS and the port is set to the same value as [serverHTTPSPort](#) option on the OM server
3. On the UM server run the '`dfm service list`' command and ensure the **DFM SERVER** service is running
4. Ensure the following UM server options settings are set as defined in section [11.1](#): [serverHTTPSEnabled](#) and [serverHTTPSPort](#)
5. Ensure the user name and password are correct (these should be the same Active Directory credentials you use to login to UM)
6. Ensure your Active Directory account is not locked out

9.14.2.5 Disabling Performance Advisor on a per host basis

On occasion you may wish to disable the collection of performance statistics for a certain storage device. The easiest way to do this is in the NMC. All hosts managed by UM will be listed in this window.

1. Login to the NMC and click on the '**PERFORMANCE ADVISOR**' icon
2. Choose the '**SETUP**' option
3. Choose the '**Storage Systems**' option
4. On 'Set Up **Storage Systems**' panel, click on 'Edit'
5. Enabled (checked) or disable (uncheck) the '**DATA COLLECTION ENABLED**' field




Figure 101 - Changing the state of performance data collection per device

9.14.3 High level monitoring best practices

In addition to the best practices discussed in the [Performance Advisor Admin Guide](#) the overall best practice when using performance advisor for monitoring is to keep the layout and configuration of the tool as simple as possible. The tree structure in the NMC reflects the group structure as configured in the UM GUI (see section [10.8.1](#)). Consider the layout of your group structure in order to combine storage devices or even file systems with similar performance characteristics. This will provide you with a frame work where you can apply thresholds and alerts to a group of devices, rather than configuring the monitoring policies on a per device basis which would be inefficient in large environments.

Storage administrators should refrain from having multiple active real time performance monitoring sessions. This will increase network traffic and load on the UM server as all transactions will pass from the storage appliances to the desktop client via the server.

In addition, storage administrators should not leave a real time or customized view console open on their workstation if they are not at their desk. Real time feeds are not required unless being observed in real time. The information can always be reviewed from the performance log files stored on the UM server.

For long term trending purposes refer to the data life cycle management best practices discussed in section [9.2.6](#).

The following subsections discuss the best practices for performance monitoring of the supported NAS use cases within Citi starting with VMware over NFS. Best practices for other use cases will be added over time.



When deploying VMware over NFS in your environment Performance Advisor is required to monitor the Storage Controllers in order to provide the appropriate storage layer performance insight in a virtualized environment.

9.14.3.1 VMware NFS NAS monitoring best practices

At a high level the monitoring required for all NetApp Filers can be split into two broad categories:

- Capacity
- Performance

The distinction between these two categories can be blurred sometimes especially when considering a virtualized environment when capacity sometimes needs to be measured in performance (for example “how much more I/O can my storage device provide”).

For the purpose of this exercise it can be said that capacity monitoring policies are usually configured in the UM GUI and performance monitoring policies are configured using the Performance Advisor module of the NMC.

The next steps are based on the assumption that the environment is configured correctly for monitoring. Ensure the processes outlined in sections [8.3.1](#), [8.3.5](#), [8.3.6](#), [8.3.7](#) and [8.3.8](#) have been followed and use the settings listed in [Table 24](#).

9.14.3.1.1 Principle performance counters

There are four components in the NetApp hardware architecture that require performance monitoring. These are all monitored by default using the built in views within the Performance Advisor interface.

- Front end network
 - Utilization rate (mbps) of the physical ports aggregated into a single virtual interface
Performance advisor -> Physical view -> Select Storage Controller -> view type: Storage System Summary view -> Chart: Network Throughput
- CPU
 - Percentage busy time of the CPUs on the storage controller
Performance advisor -> Physical view -> Select Storage Controller -> view type: Storage System Summary view -> Chart: CPU utilization
- Cache
 - Back to Back consistency pointer count as an approximation of cache efficiency and an indicator of partial writes (unaligned VMs)
Performance advisor -> Physical view -> Select Storage Controller -> view type: Storage System - CP Details -> Chart: CP activity
- Disk
 - Disk busy percentage as an indicator of backend storage I/O capacity
Performance advisor -> Physical view -> Select Storage Controller -> Select Aggregate -> view type: Aggregate Summary view -> Chart: Top Disks by disk_busy
 - Volume read latency
Performance advisor -> Logical view -> Select Storage Controller -> Select Volume -> view type: Volume Latency view -> Chart: Overall latency by Op Type
 - Volume write latency
Performance advisor -> Logical view -> Select Storage Controller -> Select Volume -> view type: Volume Latency view -> Chart: Overall latency by Op Type
 - NFS read latency
Performance advisor -> Logical view -> Select Storage Controller -> view type: NFS Basic -> Chart: nfsv3_read_latency
 - NFS write latency
Performance advisor -> Logical view -> Select Storage Controller -> view type: NFS Basic -> Chart: nfsv3_write_latency
 - Aggregate IOPS
Performance advisor -> Physical view -> Select Storage Controller -> view type: Top Aggregates -> Chart: total_transfers (note this is a point in time histogram)

9.14.3.1.2 Recommended performance thresholds

Based on the built-in counters discussed in the previous section the following thresholds are recommended as part of this initial release of VMware performance monitoring best practices:

- System level
 - Network
Performance Advisor -> System:load_total_mbps > 300 (300 sec interval)
 - CPU
Unified Manager -> Host CPU Too Busy Threshold = 95%
Unified Manager -> Host CPU Busy Threshold Interval = 15 minutes
 - Cache
Performance Advisor -> waf:cp_count:back-to-back CP > 0 (300 sec interval)
 - Disk
Unified Manager -> aggregate full threshold = 90%
Unified Manager -> aggregate nearly full threshold = 80%
Performance Advisor -> disk:disk_busy > 60% (300 sec interval)
Performance Advisor -> volume:nfs_latency > 10ms (300 sec interval)
Performance Advisor -> volume:read_latency > 20000µs (300 sec interval)

9.14.3.1.3 Performance monitoring phased rollout

When deploying performance monitoring for the first time in an environment the temptation exists to enable all features and alerts from day one. This is not a recommended approach as you may quickly find yourself overwhelmed by alerts and messages. Use the following steps to achieve a successful deployment:

1. Implement thresholds
2. Allow at least a week to pass after implementing Performance Advisor so that a baseline data set can be gathered
3. Run reports against the performance data to determine times when thresholds were breached. A good starting place is the 'History, Performance Events' report in Unified Manager, found under the Control Center tab, Group Status tab and the Events menu.
4. Look for other factors in the environment that could have caused the performance threshold breach
5. If required over time adjust the thresholds for your environment against the baseline data that has been collected because not all environments will have baseline thresholds exactly the same
6. Once performance threshold breaches are deemed the exception rather than the rule in an environment implement alerting policies in UM to notify support personnel when a performance threshold has been breached. See section [9.12](#) for information on configuring alarms.

This concludes the introduction to Performance Advisor and also this section on performance monitoring.

9.15 UM Script Plug-in Library

The purpose of this section is to provide links to reference documentation for the script plug-in library component of UM. The script plug-in library and related information regarding Perl and code writing best practices in general are beyond the scope of this text.

This completes the brief discussion of the UM script plug-in library.

9.16 UM Qtree Quota Management

The purpose of this section is to present an overview of the qtree quota management best practices when using Unified Manager in the Citi NetApp NAS environment. UM is capable of reporting on quota utilization, and management of the quota policies on the Filer. CATE Storage and Backup Engineering recommends the use of qtrees to enforce quota policies as opposed to user and group based quotas due to the complexities involved in managing such policies.

In order to manage the quota policies on a per qtree basis navigate to the appropriate qtree in the using the following process:

1. With the '**CONTROL CENTER**' tab highlighted, click on the relevant '**GROUP**' in the group tree that contains the Filer hosting the qtree you wish to manage. In this example we are selecting Global / Businesses / CATE / PSE / GDSE group from the tree.
2. Select the '**MEMBER DETAILS**' tab
3. Select the '**FILE SYSTEMS**' topic under the Member Details tab
4. Click on the name (link) of the qtree you wish to manage (in this example we are choosing /ladp02/unix4)




Figure 102 - Qtree Quota Management Screen One

5. The '**QTREE DETAILS**' window displays the capacity used for the qtree and the quota limit applied to the qtree
6. The graphing functions on this page can be used to assist storage and business management groups in plotting data growth

7. To manage the quota policy on the Filer, click the **'WINDOW ICON'** next to the **'USER QUOTA EVENTS'** section of the **'QTREE DETAILS'** page. This link is a shortcut to the Filerview interface in ONTAP, which will require Storage Administrator authentication to login and commence management of the quota policies on the filer.

9.16.1 Quota alert configuration

UM offers the ability to customize thresholds and alerting requirements on a per qtree basis. Storage administrators should remember that while this level of granularity may appear appealing, it increases the complexity of the environment to be managed and so use of the global alerting thresholds should be encouraged where at all possible. See [qtreeNearlyFullThreshold](#) and [qtreeFullThreshold](#) with regards to the global qtree utilization thresholds.

The one useful feature of the **'EDIT QTREE QUOTA SETTINGS'** page is the ability to assign an email distribution list to a qtree so that alerts can be sent to a specific business or support group. In order to configure this setting, and also to customize (if absolutely required) the alerting threshold for a specific qtree, use the following process:

1. Starting at the page list point (5) from the previous section, click the **'EDIT QUOTA SETTINGS'** link in the bottom left pane of the window, under the **'QTREE TOOLS'** menu.




Figure 103 - Qtree Quota Management Screen Two

2. In the top section of the **'EDIT QTREE QUOTA SETTINGS'** page, you can enter the name and SMTP email address of the person or group of persons (distribution list) who should receive email alerts if any of the qtree related thresholds are triggered.
3. Once the details have been entered, click the **'UPDATE'** button for the changes to take effect.

4. In the bottom section of the '**EDIT QTREE QUOTA SETTINGS**' page you can define the custom thresholds for the qtree.



You must ensure that the '**THIS QTREE**' is selected in the '**APPLY SETTINGS TO**' drop down box otherwise you will overwrite the quota threshold values for all the qtrees in that leaf node of the UM group tree.

Figure 104 - Edit Qtree Settings Screen

5. Once the details have been entered, click the '**UPDATE**' button for the changes to take effect.

9.16.2 Quota reporting

As with all the other built-in reports in UM, if the built-in quota reports do not provide the level of detail required by storage and business management teams then utilize the custom reports feature in UM to create a standard report template that is suited to your reporting requirements. For more information on creating custom reports, please refer to section [9.8.6](#) of this document.

This concludes the presentation of qtree quota management best practices using UM.

9.17 UM Command Aliasing

The purpose of this section is to describe the command aliasing function available within the UM CLI.

UM features the ability to import command alias definition files. These command alias files are much simpler than the scripts discussed in section [9.15](#), and considerably quicker to construct and implement. However, the available commands in an alias file are limited to the commands available within the version of ONTAP that you wish run the alias file against.

In order to run an aliased command in UM against a Filer or group of Filers you must configure the host login and password for the filer in UM as discussed in section [9.7.1](#).

Before reading the rest of this section, refer to the “Creating and using command aliases” subsection of the “Using the command-line interface” section of the “Maintenance and Management” chapter in the vendor [Administrator Guide](#) for the procedure and required components to create a command alias definition file.

The *Read* alias type does not apply to the Citi UM implementation as there are no administrators defined in the application security model that possess the *GlobalRead* role. The Business Aligned SAs will have a Read role in the UM database, but this role is limited to their specific branch of the UM group tree; therefore they do not possess a true global read function. The Storage Administration team, on the other hand, will be able to see both Read and Write aliases.

If you intend to use command aliasing, CATE Storage and Backup Engineering recommends you follow these simple best practices:

- If the commands to be included in the alias file are consider reporting functions (for example: `qtree status`, `vfiler status -r` or `vscan`) then the alias file can be assigned the Read operator (r)
- If the commands to be included in the alias file are considered management functions (for example: `snap delete -a -f`, `snapmirror off`, or `vscan off`) then the alias file should be assigned the Write operator (w)
- You cannot pass arguments to an aliased command at the time of execution, therefore you cannot include commands that require arguments such as `vol offline` which need to understand which object they are performing the function against
- You are strongly recommended not to alias commands that are destructive or could seriously disrupt the filer serving environment (for example `halt`)

When aliased commands are recorded in the audit logs the actual commands within the alias are logged rather than the alias name itself to simplify the command reconciliation process.

This concludes the presentation of best practices when using the command aliasing function in UM.

9.18 Run Command function

The purpose of this section is to describe the best practices involved in using the run command feature in UM. The 'Run Command' feature is available in the left pane of the main '**CONTROL CENTER**' window of UM when dealing with nodes in the group tree view (see section [10.8.1](#)) or single storage devices.

The run command screen is shown below and allows a command to be executed against a single device or group of devices, depending on whether the run command was chosen from the '**CURRENT GROUP**' menu in the left pane or the '**APPLIANCE TOOLS**' menu (also in the left pane).



Figure 105 - Run Command dialog box

The Appliance Command field will appear irrespective of whether a single storage device or group of devices has been selected. The text box will allow any command to be input to be qualified and then run against the target devices using the functional ID credentials discussed in section [9.7.1](#).

As any command strings are allowed in this open text box, the text box does not have sufficient business logic implemented to understand when a command is being typed in that involves a password string. For example 'useradmin user add nn12345 -g administrators -P aBcD1234' would create a user ID on the target storage devices with the password specified in the -P argument. The password is masked in the log files on the storage devices, and is secure in transit over the network as the traffic is tunneled via SSH. However the data after the -P parameter will be stored in clear text in the UM application log file. This is the only instance within UM where this information could be recorded. As such the following best practice must be followed.



Storage administrators must not use commands in the 'run command' dialog box that contain passwords as parameters to the command being issued.

This completes the best practice advisory when using the run command feature of UM.

9.19 Chargeback

The purpose of this section is to briefly discuss the chargeback functionality in UM. UM has the ability to calculate basic chargeback rates based on file system utilization; however the certified chargeback tool for NAS and deployed in North America is [Intermin's FileCensus](#). FileCensus has not been deployed in all regions. Check with your Finance and Chargeback team on the approved NAS chargeback process for your region.

This concludes the brief discussion on UM chargeback functionality as it pertains to Citi.

10 Security

10.1 Vulnerability and Threat Management

The purpose of this section is to present the information about Unified Manager that is pertinent to Citi Vulnerability and Threat Management processes. The component architecture and the vulnerability communication processes are presented below.

10.1.1 Architecture overview

The purpose of this section is to present an overview of the software components used within Unified Manager.

UM 5.1 and above consists of two “off-the-shelf” components used in conjunction to produce a “black box” NetApp NAS management tool. The following software components are used:

- Apache Web Server
- Sybase iAnywhere

The following communication protocols are used:

- SNMP (version 1) (version 3 is also supported by UM but is not available for use within Citi)
- SSH (version 2)
- SSL / HTTPS

10.1.2 Vulnerabilities

The purpose of this section is to present an overview on the processes used with dealing with any possible vulnerabilities that may arise in Unified Manager.

NetApp manages the vulnerability tracking and notification process for the software components listed above. Multiple departments within NetApp are advised of issues from component vendors via telephone and email notification.

NetApp seeks to inform its customers of any vulnerability as soon as possible using product bulletins and field alerts. The weekly NetApp / Citigroup call is another conduit for communication of issues.

NetApp will release a patched version of UM as soon as the component vendor has provided a component patch. The patched version of UM will undergo QA testing before being released to customers.

10.2 Known Issues

There are no known security related issues with UM 5.1 and above.

10.3 Web server SSL configuration

The purpose of this section is to present the procedure to be used when configuring or reconfiguring the SSL keys for the Unified Manager integrated Apache Web server. This procedure is required in the following situations:

- Building a new Management Station (see section [8.3](#))
- Migrating the Unified Manager database to a new Management Station (see section [9.4](#))
- Failing Unified Manager over to the DR Management Stations (see section [9.2.3.2](#))
- Reverting UM Manager back to production after DR (see section [9.2.3.3](#))
- Upgrading from UM5.1 to UM5.2

In order to bring OM 3.8 and above in line with updates in CISS policy it is now required that the SSL certificate on the DFM web server must be signed with Certificate Authority (CA) certificate. CA signed certificates are an important part of CISS as they prevent 'man in the middle' attacks. The steps involved in this process are outlined below followed by a detailed walkthrough of the process.



This process is time sensitive. Ensure codes are used to activate certificates within 30 days or the certificates will expire and you will have to submit a new VSR.

1. Raise a Virtual Security Request (VSR) via MarketPlace
2. Record VSR results in a safe place for future reference
3. Generate initial DFM self signed certificate
4. Restart the DFM web server services
5. Generate a certificate signing request on the DFM server
6. Go to Digital Certificate Service website and generate the signed certificate
7. Install signed certificate on DFM web server
8. Restart the DFM web server services
9. Record SSL certificate and private key information in a safe place for future reference

Now let's walkthrough the process in detail of requesting a new certificate:

1. Raise a Virtual Security Request (VSR)

Infrastructure Security Administrators (ISA) require a VSR in order to create a Certificate Authority (CA) signed certificate request

 - a. Login to the MarketPlace system at [CMP - Entrust Device](#) using your SSO credentials
 - b. Under **Disclaimer**, select the appropriate answers to the questionnaire and click **Next**
 - c. Complete the Entrust Device v7 questionnaire as follows:
 - i. What action do you require?
Select '**NEW**' from the drop down menu
 - ii. What type of certificate do you require?
Select '**Web Server SSL**'
 - iii. From which Certificate Authority do you require the certificate?
Select '**PRODUCTION**' for production systems or '**DEVELOPMENT**' for test and lab systems
 - iv. Provide alternate owner's SOEID
*Note: This cannot be the same as the Primary Owner ID
 - v. Retrieve shopping for info
Select '**Look up**'
 - vi. Click "Look Up" for validation to ensure Primary Owner defined on the "Shopping for" field is not the same person as the Alternate Owner.
 - vii. Provide description of Device, hardware, or application
 - viii. Provide a DNS Name, Host Name or Common Name. This is the name/cn that will be displayed in the certificate.
Enter a description (for example: 'NetApp Operations Manager production in SWDC')

- ix. IP address - ONLY ONE IP ADDRESS PER REQUEST
 - x. Provide SubjectAltName - if applicable
Leave **'BLANK'**
 - xi. Additional Comments:
Leave **'BLANK'**
 - xii. Please provide CSI number if applicable
Leave **'BLANK'**
 - xiii. Click the **'ADD TO CART'** button
 - xiv. Submit the Shopping Cart for checkout.
2. Record CMP results in a safe place for future reference.
Once CMP has been completed you will receive the following pieces of information. Keep items B), C) and D) in a safe place for records management purposes.
- a. Procedure to retrieve signed certificate
 - b. Device DN
 - c. Request Reference Number
 - d. Request Authorization Code
3. Generate initial DFM self-signed certificate



Before performing this task you will need to gather some important configuration information required for the SSL Server Setup command from the table below:

Parameter	Value	Description	Reference
-z	2048	RSA key length (Bits)	n/a
-d	730	Certificate life span (days)	n/a
-c	<variable>	Two Character Country Code	ISO 3166 Country Code List
-s	<variable>	State / Province	n/a
-l	<variable>	Location / City	n/a
-o	<variable>	CompanyName	n/a
-u	<variable>	Organizational Unit Name	n/a
-n	<variable>	Request reference number of VSR for CA signed certificate request	You will need to raise a Virtual Security Request to obtain this reference number
-e	<variable>	Email address of Storage Admin team	n/a

Table 17 - Web Server SSL configuration parameters

From the CLI on the DFM server issue the `dfm ssl server setup` command to generate the initial SSL certificate.

The following example demonstrates the web server SSL configuration process, using actual values. Values for each implementation will differ according to location and support team. Note: In this example an existing set of SSL keys already existed on the system, hence the confirmation for removal prompts.

```
C:\>dfm ssl server setup -z 2048 -d 730 -c US -s "New Jersey" -l
"Warren" -o "Citigroup" -u "CATE / GDSE" -n 12345678 -e gdse@citi.com
A private key file already exists:
C:\Program Files (x86)\NetApp\DataFabric\DFM\conf\server.key
Remove (y/[N])? y
A certificate file already exists:
C:\Program Files (x86)\NetApp\DataFabric\DFM\conf\server.crt
Remove (y/[N])? y
Loading 'screen' into random state - done
```

```

Generating a 2048 bit RSA private key
.....+++++
.....+++++
writing new private key to C:\Program Files
(x86)NetApp\DataFabric\DFM\conf\server.key'
-----
The DataFabric Manager SSL server has been reinitialized.
You must now restart the http service:

    dfm service stop http
    dfm service start http

```

The `-n` parameter 12345678 represents VSR request reference number that was generated as part of step 2c).

4. Restart the DFM web server services and confirm service status

```

C:\>dfm service stop http
Service: http stopped.

C:\>dfm service start http
Service: http started.

C:\>dfm service list
sql: started
webui: started
http: started
eventd: started
monitor: started
scheduler: started
server: started
watchdog: started

```

5. Generate a certificate signing request on the DFM server

- a. From the command line on the DFM server issue the `dfm ssl server req` command
- b. Copy into the buffer the entire output from this command (including the BEGIN and END certificate lines). This information will be used in step 6e

6. Go to Digital Certificate Service (DCS) website and generate the signed certificate

- a. Using the information obtained in step 2a go to the DCS website (the URL may differ depending on whether you requested a production or development certificate)
- b. Select the **'Create Web Server Certificate from PKCS#10 Request'** option
- c. Enter the **'Request Reference Number'** from step 2c
- d. Enter the **'Authorization Number'** from step 2d
- e. From the **'Options'** drop down menu ensure the default value is selected **'displayed as PEM encoding of certificate in raw DER'**
- f. Paste the **'DFM SSL certificate signing request'** output from step 5b into the main text box
- g. Click the **'SUBMIT REQUEST'** button

7. Install signed certificate on DFM web server
 - a. Download the Web Server Signed Certificate from the DCS website by following the instructions provide in step 2a (this should be step 3 in the provided instructions)
 - b. Save the certificate file to the DFM server desktop
 - c. Open a command line on the DFM server and navigate to the desktop directory
 - d. Execute the `dfm ssl server import` command

```
C:\>dfm ssl server import servercert.cer
```

Where `servercert.bin` is the name of the file downloaded from the DCS website
8. Restart the DFM web server services and confirm service status

```
C:\>dfm service stop http
```

Service: http stopped.


```
C:\>dfm service start http
```

Service: http started.


```
C:\>dfm service list
```

sql: started
webui: started
http: started
eventd: started
monitor: started
scheduler: started
server: started
watchdog: started
9. Record SSL certificate and private key information in a safe place for future reference
 - a. Use the `dfm ssl server show` command to list the RSA private key and web server certificate details
 - b. Copy this information and save it in a safe place for future reference

This completes the process of generating new web server SSL certificates and signing them with a Certificate Authority key.

10.4 Required IP ports

The purpose of this section is to provide an overview of the network protocols and ports required for a successful implementation of Unified Manager to monitor and manage a NetApp Filer storage environment. The following information is especially important if there is a firewall between the Unified Manager Management Station and the Filer.

Port	Type	Protocol	Source	Target	Direction	Purpose
22	TCP	SSH	UM Server	Filer	Outbound	Secure cluster console, secure storage systems takeover and giveback, secure remote command execution on managed appliances, vFiler monitoring and management. This port / function is configured by the hostLoginProtocol option.
25	TCP	SMTP	UM Server	SMTP relay	Outbound	SMTP server port used when Unified Manager send e-mail for alarms and for autosupport notification when "autosupportProtocol" option is set to SMTP. This port / function is configured by the SMTPServerPort option.
161	UDP	SNMP	UM Server	Filer	Outbound	Monitoring of Filers. This port / function is not configurable.
443	TCP	HTTPS	UM Server	Filer	Outbound	SecureAdmin-based storage system management. This port / function is configured by the hostAdminTransport option.
443	TCP	HTTPS	Storage Administrator Web Browser	UM Server	Outbound	Secure Unified Manager access. This port / function is configured by the httpsPort option. Note that the Citi implementation does not follow the vendor's default value for this option.
8443	TCP	HTTPS	Storage Administrator workstation running NMC	UM Server	Outbound	Secure NetApp Management Console Access. This port / function is configured by the serverHTTPSPort option. Note that the Citi implementation does not follow the vendor's default value for this option.

Table 18 - Required IP ports

10.5 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. The table below provides information on the Functional IDs associated with UM according to new Functional Account technical requirements recently published by the CTI Information Security Administration (ISA) organization [here](#). Each account is listed, along with its purpose; whether the account is mandatory or optional for a given implementation of UM (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.

Account Name	dfmproxy
Purpose	Used by UM to authenticate to NetApp NAS Filers for the purpose of pulling performance and configuration information, as well as proxying management commands issued from the Management Station to the Filers over APIs, SSH and NDMP.
Required	Yes – this account is mandatory when deploying UM to manage and monitor a NetApp NAS environment
Scope	Local – this account is created locally on each Filer to be monitored and is then entered into the hostLogin and password credentials for each Filer registered in the UM database
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	srv_netapp_omsched
Purpose	Used to run the DFM scheduler service on each UM server
Required	Yes – this account is recommended when following the CATE approach of backing up the UM database and copying to a DR UM server, and will also be used as part of the forthcoming EERS entitlement process
Scope	Domain - Windows
Interactivity	Non-interactive – this is a system account used for backend processes
Privileged	Non-privileged – this system account requires “logon as service” rights to each Windows server running UM
Password Change Frequency	A regular password change is not required

Account Name	srv_netapp_omadmin
Purpose	Placeholder for the second production release of this document where super user access to Unified Manager is restricted to this account
Required	No – not in the first production release of this documentation
Scope	Domain - Windows
Interactivity	Interactive – this is an application super user account
Privileged	Privileged – this functional account has full control capabilities over the UM application, and requires local administrator rights to the Windows server
Password Change Frequency	A regular password change is required every 6 months

Table 19 - Functional and Privileged Account Matrix

10.6 Access Rights to UM Management Station

The purpose of this section is to present the permissions (entitlement assignment) required by storage administrators to the Windows server hosting the Unified Manager application (the Management Station). Local administrator rights are required on the Windows Server host UM for the storage administrators who will be managing UM. Local administrator rights are required to access the Windows Service Control Manager (SCM) as the UM services can be stopped and started using UM CLI commands that call the SCM. There are no default local security policies in Windows 2008 that allow for the right to stop and start services to be delegated to a user who is not a local administrator of the server. Local administrator rights will also allow the storage administrator to log onto the server via the “right to logon using terminal services”.

10.7 UM Security Model Overview

The purpose of this section is to provide an overview of the Citi application security model implemented in Unified Manager. The security model delegates all authentication services to Windows Active Directory services. Using a delegated authentication implementation simplifies the application security model where by the following security features are externalized from the application:

- Password composition policies
- Password ageing policies
- Account lockout policies
- Account management policies (activation/deactivation)
- Entitlement request and review processes

The Citi UM security model is based on the following key concepts:

- Authentication is only provided by Windows Active Directory services
- All individual user accounts should be removed from UM and authentication should be achieved via Windows Active Directory global group memberships

For the storage administration team the access control list / application entitlement for Unified Manager is defined below. The configuration outlined below is demonstrated in an example in section [10.8.3](#) of this document.

1. The storage administration team is defined as the group of administrators who have an administrator or system administrator login to a NetApp storage device, and have rights to mount the C\$ root volume of the Filer using Windows Domain Global Group membership permissions.
2. A domain local group should be created in the Windows Active Directory domain. This group will be granted Global Full Control rights to the UM database. Its members will include the Windows Domain Global Group mentioned in point 1 (windows group nesting) and will contain any ad-hoc personnel that require global full control of the UM application. Refer to [Figure 106](#) for further details.
3. A second domain local group should be created in the Windows Active Directory domain. This group will be granted Global Read Only rights to the UM database. Its members will include any ad-hoc personnel that require read only access to the UM application, and a Windows Domain Global Group with read only rights to the C\$ of each Filer (if read only access has been configured in your Filer environment). Refer to [Figure 107](#) for further details.




Figure 106 - UM Security Model for Storage Administrators (Full Control)




Figure 107 - UM Security Model for Storage Administrators (Read Only)

For each business aligned data administration (SA) team the access control list / application entitlement for Unified Manager is defined below. The configuration outlined below is demonstrated in an example in the next section of this document.

1. The business aligned data administration (SA) team is defined as the group of administrators who manage the data on the file systems of the Filer. In a shared storage (utility) model there may be more than one business group managing different file systems on the storage device. On the Filer each business aligned data administration group is allowed to create shares by a Windows Domain Global Group being a member of the local Power Users group.
2. A domain local group should be created in the Windows Active Directory domain. This group will be granted limited rights to the relevant groups (containers) to the UM database that represent the particular line of business. Its members will include the Windows Domain Global Group mentioned in point 1 (windows group nesting) and will contain any ad-hoc personnel that require the same access but are not members of the Domain Global Group. The limited UM rights for business aligned SA group consist of database read access, and event read and write access to the relevant container in the UM database. Refer to [Figure 108](#) for further details and the walk through example in the next section.




Figure 108 - UM Security Model for Business Aligned Data Administrators [SAs]

UM offers different levels of access controls depending on the context of the object within the management database. As such this classification system can support the Citi segregation of duty model where one team is responsible for the operating system, contingency and virtual machine provision (Vfiler / Multi-store), and another team is responsible for the day-to-day administration of business data.

Using the security model outlined above, the basic access control framework within UM will remain unchanged unless a new business aligned support group is created. This would require the structure of the File System groups within UM to be updated accordingly. The ongoing access controls are actually

externalized to the application layer and exist in the Windows environment. Global UM rights should always remain under the control of the storage management team; whereas business aligned data administrators should maintain and control the memberships of their global groups. Regional Wintel operational support groups define the naming conventions for Windows global groups.

It should also be noted that Windows authentication information is cached for 20 minutes using the UM application option [credCacheTTL](#). When an administrator authenticates to UM, the server response is cached. If the same administrator attempts to re-authenticate within 20 minutes, UM does not query the authentication server and instead references the cached response. If an administrator attempts to re-authenticate after 20 minutes, UM again queries the authentication server for verification. If required, you can flush the cached user authentication data using the following CLI commands on the UM server:

```
C:\> dfm user eject <username>
Removes the cache authentication data for a specific user <username>

C:\> dfm user flush
Removes the cache authentication data for all users in the UM database
```

10.8 UM Security Model Implementation

The purpose of this section is to provide a step by step guide to configuring the key concepts of the security model overview presented in the previous section. UM provides a granular security structure that allows a user or group of users to be assigned to a role, and the role assigned to a group of objects (where an object can be a storage host or file system, for example). In order to successfully implement the UM security model, the group structure will be discussed first, followed by the roles for each group (if applicable) and finally the users that are assigned to each role.

10.8.1 Administering groups

Let's begin by looking at the concept of groups in UM and understanding how they can be implemented. Groups can be one (or more) of the following four types (you can add one or more object types to a group):

- Appliance
- Aggregate
- Configuration
- File system (volume and qtree)

Groups exhibit the following behaviors:

1. **Adding first object to a group.** The first object added to an empty group will determine the group type. Therefore if a group is empty and then a file system (volume / qtree) is added, then the group automatically becomes a file system group.
2. **Removing last object from a group.** The last object removed from a group returns the group to an unassigned type. In this way a file system group could be made an aggregate group (for example).

In order to delete the last member from a group, select the correct tab for that group type in the main window. At this point you can place a check mark next to the item and then click the delete button (note only the instance of the object is deleted from this group – the original object remains unaffected). If the delete button is not available at the bottom of the screen, then the likelihood is the wrong tab is selected for that group type.

Alternatively, click the *edit groups* link above the group list in the user interface, locate the group from which you wish to delete the last object. Click on the *edit* link for that group. In the edit group settings screen, select the appropriate *Group Memberships* type by clicking on the link. Using the arrow buttons in the middle of the screen, remove items from the right hand column. Note there is no



“done” or “OK” buttons, so any changes performed with the arrow buttons are committed to the database without an undo option. This method requires a greater amount of navigation but does support the use of multiple object selection in the list columns by keeping the CTRL key depressed when selecting hosts.

3. **Aggregates.** In ONTAP implementations still using traditional volumes (as opposed to flexible volumes) these traditional volumes will also be listed as aggregates.
4. **Hierarchical groups.** Hierarchical or nested groups are supported in UM. You can move child objects between groups to reflect organization changes. If you move a storage object from one group to another, the storage object will inherit any properties of the new group. If you move a group to become a child object of another group then the settings associated with a given group will move with it although it may inherit some settings from the new parent group. Also remember a storage object (appliance, vFiler, file system etc) can concurrently be a member of more than one group.
5. **Security context aware (access based enumeration).** Users will only see the groups in the group list to which they have permissions. See the comparison diagram further down this page.

A best practice implementation should include the following configuration of UM groups:

- Categorization of Configurations into global, site specific and domain specific groups
- Categorization of Filers into Environment (Production, Contingency, Development and Lab systems)
- Categorization of Filers, VFilers or file systems and qtrees by Business (data administration group (SAs))
- Categorization of storage resources into report groups where the grouping of resources is not reflected in the group structures defined above

UM 5.1 has no granular display or ordering controls in terms of presenting the list of groups within the interface. The base level presentation provided organizes the groups into alphabetical order. In order to further organize the presentation of groups, the following naming convention is presented as a suggestion. The convention makes use of the ‘.’ (period) character to promote entries in the group list. This suggestion is demonstrated in the screen shot below.



This example demonstrates two aspects of group views in UM. The example on the left shows the suggested best practices for configuring the overall group structure within Unified Manager: by configuration policy, by environment, by business unit and by reporting groups (if applicable).

Contrast the tree view on the left with the tree view above to see how the group tree view supports access based enumeration. The tree on the left demonstrates global access to all groups in the tree, whereas the tree above demonstrates access to the VMware group only within the business tree. It is important to note that even with a restricted view of the tree, the navigation to the accessible group (in the case VMware) is displayed in the context of the root group in the tree (global). The red boxes in the left diagram highlight the path to the group shown in the diagram above.

Figure 109 - Group Tree Structure example

Best practices would recommend that an UM Business Aligned Database User Group be created for each different support group that administers data in NetApp NAS environment. This would allow for the lowest cost administrative model, where a Windows domain local group is mapped to an UM application group, which contains a collection of database objects (in this case file systems objects) that are administered by the same business aligned support group.

Only storage administrators will see all the groups in the UM group tree. Business aligned SAs will only see the branches of the tree from which they are entitled to read.

10.8.2 Assigning resources

The purpose of this section is to build upon the group structures presented in the previous section by assigning storage resources to the management and reporting groups. Once you have created your group tree structure, the next step is to assign the relevant resources to each business group so that when the business aligned support group logs into UM they see the appropriate and relevant storage objects assigned to that particular support group.

Use the following process to assign resources to a group:

1. Expand the group list using the [+] symbols in the tree
2. Locate the group to which you would like to add resources and click on the 'GROUP NAME' in the tree
3. Click the 'EDIT MEMBERSHIPS' link in the 'CURRENT GROUP' menu in the lower left pane of the screen
4. In the *Edit Group Membership* window, select the resource object type you wish to add to the group by selecting the desired entry for the 'CHOOSE FROM ALL AVAILABLE' drop down menu.




Figure 110 - Editing Group Memberships

There are various storage objects to choose from but the ones that will usually be assigned to a group in a Citi NetApp NAS implementation are:

- **Appliances** (useful when a single business has exclusive use of a storage controller)
- **Volumes and qtrees** (useful when a business has data distributed in many locations and no dedicated storage devices)
- **Aggregates** (useful in a VMware implementation where an aggregate is traditionally assigned to a specific ESX cluster)
- **vFilers** (useful when a single businesses data is served by a vFiler)

By using the tips in the above list choose the appropriate storage object (or objects) to assign to the group that best represents the storage object level at which the business data is segregated from other lines of business.

5. **'SELECT'** the object (or objects by using the **'CTRL'** key) from the list in the left pane that you wish to add to this group
6. Click the **'>>'** button to add the selected objects to the group. Note there is no "done" or "OK" buttons, so any changes performed with the arrow buttons are committed to the database without an undo option. If you accidentally add an object to the group, select the unwanted object from the right window and use the **'<<'** button to remove it from the group.
7. Repeat this process for each business group defined in section [10.8.1](#) to which you need to add business aligned storage resources.

10.8.3 Administering Roles

The purpose of this section is to build upon the group structures and resource assignments presented in the previous two sections by discussing the access control aspect of the groups and the resources they contain. When considering roles in UM, think of them as the equivalent of an access control list. In referring to the security model overview diagrams ([Figure 106](#), [Figure 107](#) and [Figure 108](#)) the concept of a role is represented by the access control list object in the top-right quadrant of the diagrams.

The remainder of this section will present a walk-through example of the three access control configurations presented in section [10.7](#). Let's begin with the two roles used by the Storage Administration team to that provides the full control of the UM database and a global read only role. The two roles are the easiest to present as they are available by default in UM:

- **GlobalFullControl** – Provides full control rights to all objects and functions in the UM database
- **GlobalRead** – Provides read to all objects and functions in the UM database

Use the following steps to create a business aligned support group role. Repeat the following process for each business support group:

1. Under the **'CONTROL CENTER'** tab, choose the **'SETUP'** menu and select the **'ROLES'** option
2. In the **'ROLE NAME'** field type the name of the role. The suggested naming convention for the role name is "Bus_<Business Acronym>_<Sub-business Acronym>_Admins". So for the example below the Role name is "Bus_CATE_POE_VMware_Admin".

3. In the '**DESCRIPTION**' field type the description of the role. The suggested naming convention for the role description is "Business Admins / Business Name / Sub-Business Name". So for the example below the description reads "Business Admins / CATE / POE / VMware".
4. Click on the '**ADD CAPABILITIES**' link to define custom capabilities for this role




Figure 111 - Creating a business aligned support group role (part 1)

5. In the '**CAPABILITIES WINDOW**', click on the [+] symbol to expand the '**GROUPS**' resource in the left hand pane of the window, and then expand the '**BUSINESSES**' group
6. **LOCATE** the business group container in the list to which you wish to assign this custom role (access control list). In this example we will be creating a business support group role for the CATE VMware engineering team




Figure 112 - Creating a business aligned support group role (part 2)

7. In the right pane ('**OPERATIONS**') of the '**CAPABILITIES WINDOW**', scroll down to the '**DFM.DATABASE**' operation and select the '**READ**' checkbox.
8. Continue to scroll down the '**OPERATIONS**' section to the '**DFM.EVENT**' operation, and select the '**READ**' and '**WRITE**' checkboxes.
9. Click the '**OK**' button to assign these operations (permissions) to the resource.




Figure 113 - Creating a business aligned support group role (part 3)

10. Back in the '**ADD ROLES**' window, the business aligned support role is not yet created. Click the '**OK**' button to complete the role creation. The capabilities defined in the previous step are shown in the red box in [Figure 114](#).




Figure 114 - Creating a business aligned support group role (part 4)

This completes our walk-through example creating a business aligned support role in UM and this section on administering roles. Repeat this process for each business aligned data administration team in your environment.

10.8.4 Administering Users

The purpose of this section is the culmination of implementing the UM security model by adding users to the roles discussed in the previous section. The remainder of this section will present a walk-through example of the three user types discussed in section 10.7 (Storage Administrator Global Full Control; Storage Administrator Global Read Only; and Business Aligned Support Group Read view with Event review and acknowledgement).

Use the following steps to add the storage administrators to the 'NAS Admin OM Database Admin group' as shown in [Figure 106](#):

1. Under the '**CONTROL CENTER**' tab, choose the '**SETUP**' menu and select the '**ADMINISTRATIVE USERS**' option
2. In the '**ADMINISTRATOR or WINDOWS GLOBAL NAME**' field type the name of the domain local group as defined in the Windows Active Directory domain that will be used for full administrative control of the OM database. In the example below we are using the group 'Per_GDSE_NetApp_OM_Admins'. Note you do not need to specify the domain in front of the group name but the group must exist in the same domain in which the OM server exists.
3. From the '**ROLES**' list select '**GLOBALFULLCONTROL**'
4. Click the '>>' button to assign the selected role to the user group.
5. Optionally you can add the SMTP email address of the support team (note this must be the SMTP address not the outlook friendly DL name).
6. Click the '**ADD**' button to grant the users in the Windows Domain Local Group Full Control access to the OM database.




Figure 115 - Adding a Storage Admin User Group (Full Access)

This completes the example of adding the Storage Administration user group with full access to the UM database. Use the following steps to add ad-hoc storage administrators to the 'NAS Admin OM Database user group' for a read only access as shown in [Figure 107](#):

1. Under the '**CONTROL CENTER**' tab, choose the '**SETUP**' menu and select the '**ADMINISTRATIVE USERS**' option
2. In the '**ADMINISTRATOR or WINDOWS GLOBAL NAME**' field type the name of the domain local group as defined in the Windows Active Directory domain that will be used for full administrative control of the UM database. In the example below we are using the group 'Per_GDSE_NetApp_OM_Users'. Note you do not need to specify the domain in front of the group name but the group must exist in the same domain in which the OM server exists.
3. From the '**ROLES**' list select '**GLOBALREAD**'
4. Click the '>>' button to assign the selected role to the user group.
5. Optionally you can add the SMTP email address of the support team (note this must be the SMTP address not the outlook friendly DL name).
6. Click the '**ADD**' button to grant the users in the Windows Domain Local Group Read-Only access to the UM database.




Figure 116 - Adding a Storage Admin User Group (Read Only Access)

This completes the example of adding the Storage Administration user group with read- only access to the UM database.

The final example will add a business aligned support group to their group specific role in the UM database.

Use the following steps to add a business aligned support group to their respective UM Database user group for a read only access to their storage resources and the ability to read and acknowledge events that are generated by their allocated storage resources. [Figure 108](#) is the design reference for the following process:

1. Under the '**CONTROL CENTER**' tab, choose the '**SETUP**' menu and select the '**ADMINISTRATIVE USERS**' option
2. In the '**ADMINISTRATOR or WINDOWS GLOBAL NAME**' field type the name of the domain local group as defined in the Windows Active Directory domain that will be used for full administrative control of the UM database. In the example below we are using the fictional group 'Per_CATE_VMwareENG'. Note you do not need to specify the domain in front of the group name but the group must exist in the same domain in which the UM server exists.
3. From the '**ROLES**' list select the relevant role that was created for this specific business aligned support group in section [10.8.3](#). In continuing from the example in the earlier section we are selecting the 'Bus_CATE_POE_VMware_Admins' role for this demonstration.
4. Click the '>>' button to assign the selected role to the user group.
5. Optionally you can add the SMTP email address of the support team (note this must be the SMTP address not the outlook friendly DL name).
6. Click the '**ADD**' button to grant the business align support users in the Windows Domain Local Group resource specific access to the UM database.





Figure 117 - Adding a Business Aligned Admin User Group

Using Windows Domain groups for UM application authentication and authorization purposes simplifies the overall management of the UM application. Entitlement review of group memberships is externalized to the Windows Active Directory domain, and changes to group memberships can be processed by the appropriate security administration team without making continual change to the user security of the UM application.

Using user administration process presented above a mapping has been achieved between the Windows Domain Local Group, and the UM Database user group. Continuing the business aligned vmware administrators example presented earlier in this section, the mapping of the domain group to the application group is demonstrated below.



Having presented all the constituent elements we can now pull the security model together by showing the complete lists of all the available UM application groups used in this example, and place them in context to the listing administrative users (Windows Domain groups).

The red arched arrow indicates the group in the tree to which the user group Per_CATE_VMwareENG has access to.

The third diagram acts as a refresher that all group listings is presented in terms of access based enumeration therefore users will only see the groups they have access to.

Administrator Name	Roles	Email	Pager	View	Edit	Delete
Everyone				view	edit	
NAMDEVlom37admuser1	Bus_CATE_POE...	richard.sopp@citi.com		view	edit	☐
NAMDEVPer_CATE_VMwareENG	Bus_CATE_POE...	virtual.infrastructure.engineering.qt@imcnam.smb.com		view	edit	☐
NAMDEVPer_GDSE_NetApp_OM_Admins	GlobalFullCo...	gdse@citi.com		view	edit	☐
NAMDEVPer_GDSE_NetApp_OM_Users	GlobalRead	gdse@citi.com		view	edit	☐




Figure 118 - Mapping users to resource groups and access based enumeration

This discussion completes this section on administering Unified Manager users.

10.9 UM Entitlement Reporting

The purpose of this section is to provide reporting procedures to allow storage administrators to perform entitlement reporting regarding users who access the UM application. A regular entitlement review process is recommended to reduce the risk of accidental unauthenticated access or legacy access to the UM application. A subsequent CATE project will automate entitlement reporting of Unified Manager into the EERS application. The reports presented below are for use until the automated procedure is available, and can also be used as an ad-hoc reporting function after the EERS process has been implemented.

A basic entitlement review process can be implemented using the following set of commands on the UM CLI.

The following command produces a mapping of roles to administrator accounts for all users of the UM server:

```
C:\>dfm report view role-admins
```

To determine the roles held by a particular administrator or administration (global) group provide the user name (including authenticating domain) as an argument to the report command:

```
C:\>dfm report view role-admins NAMDEV\rs01146
```

In this example, this user has GlobalFullControl rights to UM. The output is shown below:

Object ID	Administrator	Role
27	NAMDEV\rs01146	GlobalRead
27	NAMDEV\rs01146	GlobalProvisioning
27	NAMDEV\rs01146	GlobalResourceControl
27	NAMDEV\rs01146	GlobalSDFullControl
27	NAMDEV\rs01146	GlobalSDSnapshot
27	NAMDEV\rs01146	GlobalSDConfig
27	NAMDEV\rs01146	GlobalSDStorage
27	NAMDEV\rs01146	GlobalPerfManagement
27	NAMDEV\rs01146	GlobalDataProtection
27	NAMDEV\rs01146	GlobalDataSet
27	NAMDEV\rs01146	GlobalConfigManagement
27	NAMDEV\rs01146	GlobalExecute
27	NAMDEV\rs01146	GlobalEvent
27	NAMDEV\rs01146	GlobalDistribution
27	NAMDEV\rs01146	GlobalSRM
27	NAMDEV\rs01146	GlobalSAN
27	NAMDEV\rs01146	GlobalMirror
27	NAMDEV\rs01146	GlobalRestore
27	NAMDEV\rs01146	GlobalBackup
27	NAMDEV\rs01146	GlobalDelete
27	NAMDEV\rs01146	GlobalWrite
27	NAMDEV\rs01146	GlobalQuota
27	NAMDEV\rs01146	GlobalFullControl
27	NAMDEV\rs01146	GlobalFailover

Interpreting this output can be somewhat challenging as not only the assigned capability is listed (in this case GlobalFullControl) but all the inherited capabilities are listed (GlobalFullControl inherits all the other Global prefixed capabilities).

The following command produces a mapping of administrator accounts to roles for all users of the UM server:

```
C:\>dfm report view admin-roles
```

To determine the users who are members of a particular role provide the role name as an argument to the report command:

```
C:\>dfm report view admin-roles GlobalFullControl
```

In this example, these users have GlobalFullControl rights to UM. The output is shown below:

Object ID	Role	Administrator
25	GlobalFullControl	NAMDEV\om37admgrp1
25	GlobalFullControl	NAMDEV\rs01146
25	GlobalFullControl	NAMDEV\sh93900

This output is easier to interpret but in order to make the output meaningful requires the person generating the report to know each role they want to report on. If each role is not explicitly requested in this report then the out is complex and it will list all inherited as well as explicitly defined roles.

10.10 UM Filer and vFiler Password Management

The purpose of this section is to provide an overview of the Filer and vFiler password management features included with Unified Manager. Before presenting the password management process it's worth spending some time discussing what the password management component is capable of doing as historically there has been a certain amount of confusion over this feature.

10.10.1 Features

This section briefly describes the features of the password management function in UM. The password management feature allows for an administrator to authenticate to a Filer or group of Filers (alternately a vFiler or group of vFilers) and change their local password on the selected device(s). This feature will not be required for storage administrators SOE IDS once they are integrated with Active Directory for ONTAP management authentication. At that time the only accounts that will require this feature will be functional IDs such as root on the Filer.

The step by step procedure for changing a password can be found in section [10.10.3](#), however at a high level the password management feature allows an administrator to provide their current credentials and new password to a single system or group of systems in order to change their password on the defined systems. The actual password change process is achieved via an API call and most importantly of all does not use the "host login" credentials defined in UM for the storage device (see section [9.7.1](#)), but rather the username and old password information provided by the administrator in the UM password management screen (as described in section [10.10.3](#)).

10.10.2 Limitations

This section briefly describes the limitations of the password management function in UM. Due to the historical confusion about this feature the information of most interest can be found when describing the product limitations.

- **UM needs to be aware that the account exists on the Filer.**
If an account has recently been created on a Filer, UM may fail the password change attempt if the time transpired between the account creation in ONTAP and the password change attempt is less than the monitoring sample period defined in the [hostRBACMonInterval](#) UM global setting. If UM does not have the specific user account for the specific storage system in the Host Users table of the database then the password change event will fail for that device even though the account does exist on the storage system. This type of behavior is expected in monitoring and management software that uses polling interval to pull data from devices.
- **Old user password needs to be known and the same across all devices.**
As the password management feature authenticates to each storage device defined in the password change process using the username and old password defined in the request, the following rules need to hold true for the password change process to complete:
 - a. The old password needs to be known for authentication purposes. If the old password is unknown or forgotten then the password change feature in ONTAP will need to be used for the Filer(s) in question.
 - b. The old password needs to be the same across all devices on which you are attempting the password change. If the old password is different on some devices then the password change process will fail for these devices.

10.10.3 Procedures

The basic processes for Filer and vFiler password management are the same and in this section we will be using Filer password management as the example. The only difference in the processes is navigating the menu to reach the password management functionality.

For Filer (Storage System) password management:

1. Under the '**CONTROL CENTER**' tab, choose the '**MANAGEMENT**' menu and select the '**STORAGE SYSTEMS**' option, and then the '**PASSWORDS**' sub option




Figure 119 - Filer password management menu

For vFiler password management:

1. Under the '**CONTROL CENTER**' tab, choose the '**MANAGEMENT**' menu and select the '**STORAGE SYSTEMS**' option, and then the '**PASSWORDS**' sub option




Figure 120 - vFiler password management menu

For the rest of the password management process we'll be using the Filer procedure as the example for the rest of this chapter.

2. In the '**USER NAME**' field enter the user ID for the account you want to perform a password change
3. In the '**OLD PASSWORD**' field enter the existing password for the account defined in step 2
4. In the '**NEW PASSWORD**' field enter the new password for the account defined in step 2. Ensure the new password meets the password complexity requirements configured in ONTAP.




Figure 121 - Filer password management screen

5. Confirm the new password in the '**CONFIRM NEW PASSWORD**' field.
6. Define the '**STORAGE SYSTEMS**' (or vFilers) on which you wish to change the account password. This can be done in one of two ways:
 - a. **Group list**
From the group tree view select the group containing the Filers (or vFilers) on which you wish to change the user password by placing a check mark in the appropriate box. Check or uncheck the 'APPLY TO SUBGROUPS' as required. Applying to subgroups allows for Filers (or vFilers) in child groups to be included in the password change process otherwise just the systems in the current group are updated.
 - b. **Storage Systems**
Choose the Filers (or vFilers) directly for the list. Use the CTRL key to enable multiple device selection.
7. Click the '**UPDATE**' button to start the password change job.

This concludes the process to manage user passwords on multiple storage devices, however for auditing purposes it is worthwhile taking the opportunity to review which UM and ONTAP logs record these password related events. The following example is taken from a password change event on a vFiler:

Component	Log File	Log File Entry
ONTAP	/etc/messages	Mon Aug 18 13:30:12 EDT [vfiler1@nas3050gdse01: passwd.changed:info]: passwd for user 'rs01146' changed.
ONTAP	/etc/log/auditlog	Mon Aug 18 13:30:11 EDT [nas3050gdse01: sshd_1:debug]: dfmproxy:START:ssh2 shell:10.40.43.155:3233:password Mon Aug 18 13:30:11 EDT [nas3050gdse01: rshd_0:debug]: dfmproxy:IN:ssh2 shell:SSH INPUT COMMAND is passwd Mon Aug 18 13:30:12 EDT [nas3050gdse01: rshd_0:debug]: dfmproxy:END:ssh2 shell:
UM	%installdir%/log/audit.log	Aug 18 13:29:43 [dfm:NOTIC]: NAMDEV\rs01146:WEB:in:10.98.128.52:https://nasmgmtgdse03.nam.nsroot.net/dfm/edit/passwords?type=vFiler: Aug 18 13:30:10 [dfm:NOTIC]: NAMDEV\rs01146:WEB:in:10.98.128.52:dfm host password set type=vFiler user-name=rs01146 old-passwd=***** new-passwd=***** group= objects=264 subgroup-dont-recurse=0 ApplySubGroups=:Started job : 15. Aug 18 13:30:11 [Unknown:NOTIC]: NAMDEV\rs01146:API:in:127.0.0.1::<host-user-modify-password> <host-name-or-id>264</host-name-or-id> <host-user-name-or-id>rs01146</host-user-name-or-id> <new-password>*****</new-password> <old-password>*****</old-password> </host-user-modify-password>
UM	%installdir%/log/run.log	Aug 18 13:30:10 [dfm:DEBUG]: rjob_submit: job_id=15 host_id=264 cmd=passwd timeout=7200

Table 20 - Password management log files

This concludes the presentation on the password management functionality in UM.

10.11 UM Log Files

The purpose of this section is to provide an overview of the log files used by the UM application. These log files are stored on the file system of the UM [management station](#).

The log files are accessible by browsing the file system of the management station. The log files are located in the %INSTALL_DIR%/log directory on the UM server (%INSTALL_DIR% represents the installation location of the UM software –typically C:\Program Files (x86)\NetApp\DataFabric\DFM).

A subset of log files is available via a hidden URL in the web GUI for UM. You are required to login to UM in order to review these logs via the web. To access the **subset of logs** available via the web:

1. Point your browser to <https://<hostname>/dfm/diag>.
(where <host_name> is the name of your management station)
2. Scroll down to the logs section
3. Click on the log file name hyperlink to read the contents of the log file in the web browser

The following table presents a list of all the log files maintained by UM, together with a statement of the logs purpose, and whether the log file is accessible via the web interface.

Log File	Via GUI	Purpose
Access.log	No	Apache web server log (similar to Microsoft's IIS log). Records IP address of remote web client, time stamp, the type of web activity (POST or GET), the local path and the HTTP standard used to access the file
Alert.log	No	Tracks all alarm events generated when thresholds are triggered. Records timestamp, sub-system generating the message, alarm id, eventId, alarm message, and prescribed action (SMTP or SNMP alert) defined for the given event on the given storage device
Audit.log	Yes	Tracks all changes made in the web user interface by a user. Records time stamp, sub-system generating the message, user ID, interface, client IP address and activity performed.
DFDRM.log	No	Tracks operations that relate to the Disaster Recovery component (snapmirror management) of UM. Records time stamp, sub-system generating the message, and the message string.
DFM.log	Yes	Tracks all instances where application services are installed; stopped or restarted (including services during database backups), and group cache rebuild events. Records time stamp, sub-system generating the message, and the message string. Consider this a general application message log.
DFMbackup.log	No	Contains the database backlog. Records the state of the database backup process at each iteration of the Sybase DB backup process.
DFMcmd.log	No	Contains the login activity of each administrator on the UM console, and the resulting output of any commands executed in the GUI or CLI on the UM hosting server. This log records time stamp, sub-system generating the message, and the output of the command.
DFMconfig.log	No	Tracks all configuration push job information. Records the timestamp of the action, sub-system generating the message, the user initiating the configuration push, the target storage system(s), the system options to be change by this configuration push, their old value and the new value implemented as part of the configuration push process.
DFMevent.log	Yes	Contains the status and output of the DFM Event service. Examples of entries in this file include service started and active port notification, process and thread alerting. This log records time stamp, sub-system generating the message, and the status of the service.

Log File	Via GUI	Purpose
DFMmonitor.log	Yes	Contains the status and output of the DFM Monitoring service. Examples of entries in this file include service started and instance when changes are made to monitored objects in the database. This log records time stamp, sub-system generating the message, and the status or change of an object in the service.
DFMscheduler.log	Yes	Contains the status and output of the DFM Scheduler service. Examples of entries in this file include service started; database connectivity problems and scheduled queryerrors. This log records time stamp, sub-system generating the message, and a descriptive comment on the error.
DFMserver.log	Yes	Contains the performance gathering information for UM, including XML queries and host connectivity error reports. This log records time stamp, sub-system generating the message, type of message descriptor, and a descriptive comment of the message.
DFMwatchdog.log	Yes	Contains the status of each component of the UM application. An entry can be found for each of the service components of the application. Effectively watchdog is checking for stalled processes and services within the UM application. This log records time stamp, sub-system generating the message, and a descriptive comment regarding the application component uptime and system utilization.
Discovery.log	Yes	Contains the host and storage appliance discovery activity of the UM application. This log should be very small in size (1K) as the discovery options are disabled in the Citigroup DFM configuration. This log records time stamp, sub-system generating the message, and a descriptive comment regarding the subnet being discovered and the progress of scanning IP addresses within that subnet.
Error.log	No	Contains warning and informational error messages from any element of the UM application. This log is the best to check for overall configuration errors. Specific storage appliance connectivity and monitoring errors should be diagnosed with the use of DFMMonitor.log. This log records time stamp, nature of the message, and a descriptive comment of the message.
Mon_Perf.log	No	Contains performance data for use by NetApp in diagnosing problems with DFM and monitoring of remote storage appliances. Due to the nature of the log file there is no requirement to further explain its contents. Mon_perf.log will rollover when it reaches a limit of 500K in size then append a sequential number to the log file, and increment all existing log files. A maximum of 6 performance logs are retained.
Pingmon.log	Yes	Contains the change of ping status results for devices monitored in DFM. Depending on the ping protocol defined in the application options, and change in response is recorded in this log – including up -> down, down -> up and unknown -> up transitions. This log records time stamp, sub-system generating the message, target IP address and host name, transition description, ping protocol and error reference.
Run.log	No	Contains the input and execution result for jobs submitted using the run command or that are scheduled and subsequently use the run command.
Sharemon.log	No	Contains changes in the state of the shares on the storage systems. Specifically those shares that exist in the database but are no longer present on the target storage system.
Smmon.log	No	Contains snapmirror related activities for all snapmirror transfers monitored via UM.
Svmon.log	No	Contains snapvault related activities for all snapvault transfers monitored via UM.
Sybase.log	Yes	Adaptive Server Anywhere Network Server log. Records time stamp, checkpoint, and performance data for the Sybase component of the DFM application.
Vfmon.log	Yes	Contains all vFiler related monitoring communication events.

Table 21 - Application log file list

Having presented all the different logs generated by UM it is worthwhile spending some time focusing on one of the more important logs; the audit log.

Audit logging is the process of logging every activity performed by UM for a later review. UM logs all the activities in the audit log file. System administrators view the audit log file to:

- Determine the recently changed configurations to understand why a problem is occurring.
- Determine when a given change to the configuration of the system was made.
- Determine who made the given change to the configuration of the system and ask them why the change was made.
- Identify attempts to subvert the security of the system.

Audit logging is enabled in the Citi standard build of UM using the option [auditLogEnabled](#) and is the default value out of the box.

The events that can be audited in UM and the information that is recorded about each event are described here:

- **Authentication events:** UM logs each authentication action that succeeded or failed. The user name and source address associated with the authentication attempt is also recorded.
- **Authorization events:** UM logs each authorization failure and the user name associated with it.
- **Command execution:** UM logs execution of each command. The complete command line (including options and arguments) is recorded in the audit log file. UM also logs the name of the user who executed the command; the failure status of the command, if any; and the type of request: Web or CLI.
- **API calls:** UM logs invocation of any API by using DataFabric Manager service. The complete details of the API call and the authenticated user's name on whose behalf the API was invoked are recorded in the audit log file. The API calls made from the CLI are changed only to perform audit logging.
- **Scheduled actions:** When the scheduler starts a job by invoking a CLI, UM logs the scheduled action and the user affiliated with it in the audit log file. In addition, a timestamp is recorded for each event. In the case of APIs, the IP address of the machine from which the requests are received is logged. In the case of CLI requests, the IP address is always that of the [management station](#).

Audit logs can grow up to 3MB in size. The size limit of 3MB is a predefined value that cannot be changed. Once this size threshold is reached the log file is rotated and appended with a number. Up to 6 audit log files can exist on a management station before the oldest log file is deleted. More than 6 audit logs can be maintained on a management station if the option [auditLogForever](#) is enabled; however CATE Storage and Backup Engineering strongly discourages the use of this option because if the file system utilization on the management station is left unchecked the audit logs will simply fill up the UM application partition. The recommended value for [auditLogEnabled](#) is 'Yes'.

The entries in the audit logs follow a predefined format:

```
<timestamp> [<application-name>:<priority>]:<username>:<protocol><label>:
<ip-address>:<intent>:<message>
```

The following is an example of an audit log entry:

```
Apr 11 00:04:19 [dfm:NOTIC]:NAMDEV\rs01146:LOG:action:::Added new
administrator "testu1":
Apr 20 11:57:27 [dfm:NOTIC]: NAMDEV\rs01146:API:in:10.72.1.2:Listing database
backups: <dfm-backup-verify-settings><dfm-backup-directory>opt<dfm-backup-
directory><dfm-backup-verify-settings>
```

In this example, the *<timestamp>* field is the local time and it is in the format of %b, %d, %H, %M, %S where:

- The field %b logs the abbreviated month name
- The field %d logs the day of month as a decimal number (01 - 31)
- The field %H logs the hour in 24-hour format (00 - 23)
- The field %M logs the minute as a decimal number (00 - 59)
- The field %S logs the second as a decimal number (00 - 59)

The *<application-name>* value denotes the application invoking the audit log facility. For example, it can be dfm/dfbm/dfdrm/dfcm if a CLI or Operations Manager request is being audit logged from the dfm/dfbm/dfsrm/dfcm executable.

In the case of APIs, *<application-name>* is the actual name of the application that called the API (for example, dfm and sdu). If the API was called by an external application other than UM and it does not pass the name of the application in the API input parameters, the field :: is not printed in the log message.

The message priority field *<priority>* can have one of the following values:

- EMERG: unusable system
- ALERT: action required
- CRIT: critical conditions
- ERROR: error conditions
- WARN: warning conditions
- NOTIC: normal but significant condition
- INFO: informational messages
- DEBUG: debug-level messages

The *<username>* field logs names of the users who invoked the CLIs and APIs.

The *<protocol>* field describes the source of the event being logged. The protocol label can have one of the following values:

- API: In the case of an API invocation
- Cmd: In the case of CLI commands
- Log: When an event is explicitly logged by the system
- Web: In the case of a Unified Manager request

The *<label>* field describes the type of the event being logged. The message label can have one of the following values:

- in: for input
- out: for output (for example, in the case of API calls)
- err: for error (for example, in the case of log in failures)
- action: for actions initiated by the user

The *<ip-address>* value for APIs is the IP address of the system from which the API is invoked. In the case of the CLI, it is the IP address of the [management station](#). For requests coming through Unified Manager (Web Interface), it is the IP address of the work station on which Unified Manager is installed or the IP address of the client accessing the UM web server.

The *<intent>* field describes the following information:

- If the protocol is API, it conveys the intention of the API
- If the protocol is cmd, it is the actual command used
- If the protocol is Web, it is the URL of the Web page
- If audit-log API is called, this file remains blank

The *<message>* field content depends on the value of *<protocol>*, as follows:

If *<protocol>* is API, the XML input to the API is logged, excluding the *<netapp>* element

If *<protocol>* is cmd, it contains output or an error message

If *<protocol>* is Web, it is empty

The following is an example of the differences in the message field dictated by the protocol type:

```
July 04 22:11:59 [dfm:NOTIC]:NETAPP\tom:CMD:in:127.0.0.1:dfm host password
set -p ***** jameel:Started job:2
```

```
July 06 13:27:15 [dfm:NOTIC]:NETAPP\tom:WEB:in:127.0.0.1:dfm user login
username=tom password=*****:Logged in as<B>NETAPP\tom<\B><BR>
```

```
July 06 14:42:55[dfm:NOTIC]:NETAPP\tom:API:in:127.0.0.1:Add a role to a
user:<rbac-admin-role-add> <role-name-or-id>4</role-name-orid><admin-name-or-
id>TOM-XP\dfmuser</admin-name-or-id></rbacadmin-role-add>
```

As a security measure, UM does not log any passwords in clear text in the audit log file; including those used in any password management process.

This completes the review of the UM application log files and the deep dive on the structure and context of the application audit log file.

11 Appendices

11.1 Unified Manager Server Settings (Options) Configuration Matrix for Small to Medium Environments

The purpose of this section is to define and describe the configuration options used when implementing Unified Manager server in Citi. The vendor's default settings for each option are described in the 'description' column.

The Citi recommended value for each option has been developed with the vendor based on their experience and understanding of the Citi environment.

As of OM 3.8 and above CATE Storage and Backup Engineering is using a simplified approach to configuration management and classification. Configuration options are split into two categories: mandatory (security) and recommended (technical specification). Optional settings to enable additional features, over and above the basic Unified Manager functionality, are listed separately.

It is not the purpose of the following table to define every option in the Unified Manager application. This following table is intended to be used as a template for the Standard Build Configuration Checklist.

Where the recommended value for an option has no value (i.e. leave it blank) it is listed in the table as '<undefined>'.

Table 22 - Unified Manager Mandatory (Security) Configuration Matrix

Unified Manager Mandatory Settings Configuration Matrix			
From OM Version	Option Name	Description	Value
3.5	agentHostTransport	Hostagent/ UM server communications protocol (global setting). The default value for this option is http . The use of host agents is not supported by CATE Storage and Backup Engineering.	https
3.5	auditLogEnabled	Enables audit logging of administrative events on the Unified Manager server. The default value for this option is Enabled . (Note changing this setting requires the DFM server service to be stopped and restarted).	Enabled
3.5	auditLogForever	Keep each 3MB Unified Manager audit log indefinitely and do not roll the logs. The default value for this option is No . (Note changing this setting requires the DFM server service to be stopped and restarted).	No
3.5	authUsePam	Determines whether PAM is used for Unified Manager authentication. The default value for this option is No . (Note changing this setting requires the <i>DFM server</i> service to be stopped and restarted). The use of PAM is not supported by CATE Storage and Backup Engineering.	No

Unified Manager Mandatory Settings Configuration Matrix			
From OM Version	Option Name	Description	Value
3.5	autosupportProtocol	Defines the transport protocol used to deliver the autosupport message to NetApp. The default value for this option is https .	HTTPS
3.5	credCacheTTL	Defines the length of time Unified Manager holds user credentials before performing a validity check with an external authentication service. The default value for this option is 20 minutes .	20 minutes (20m)
3.5	discoverAgents	Defines whether the automatic discovery of host agents is permitted. The default value for this option is Enabled .	Disabled
4.0	discoverClusters	Defines whether the discovery of c-mode clusters from an already discovered storage systems is permitted by the Unified Manager server (for example if one node of a cluster was added to Unified Manager the rest of the nodes in the c-mode system would automatically be discovered). The default value for this option is Enabled .	Enabled
3.5	discoverEnabled	Defines whether automatic discovery of Filers is permitted by the Unified Manager server. The default value for this option is Enabled .	Disabled
3.8	discoverHostInit Enabled	Defines whether automatic discover of Host Agents through host initiated discovery mechanism is permitted by the Unified Manager server. The default value for this option is Enabled .	Disabled
3.5	discoverHosts	Defines whether host agent discovery of Hosts is permitted by the Unified Manager server. The default value for this option is Enabled .	Disabled
3.5	discoverNetworks	Defines whether the discovery of network is permitted by the Unified Manager server. The default value for this option is Disabled .	Disabled
3.5	discoverTimeout	Defines the time interval after which the Unified Manager server considers a discovery activity to have failed. The default value for this option is 5 seconds .	5 seconds (5s)
3.5	discoverVfilers	Defines whether the discovery of Vfilers on already discovered storage systems is permitted by the Unified Manager server. The default value for this option is Enabled .	Enabled
3.8	hostAdminPort	The port on the host where the ZAPI requests are sent. This port can also be set for a particular host. It can be any number that is valid as an Ethernet port number and not already used on the host. The default value for this option is 80 .	443
3.5	hostAdminTransport	Defines the transport protocol used for streaming statistics and displaying the FilerView and Appliance Manager. The default value for this option is HTTP.	HTTPS
3.5	hostLoginProtocol	Defines the transport protocol that Unified Manager will use when executing commands on a remote Filer (global setting). The default value for this option is RSH .	SSH
3.5	httpEnabled	Defines whether the Unified Manager web server will accept requests over the HTTP protocol. The default value for this option is Yes . (Note changing this setting requires the <i>DFM HTTP</i> service to be stopped and restarted).	No
3.5	httpPort	Defines the port the Unified Manager web server uses for HTTP traffic. The default value for this option is 8080 . (Note changing this setting requires the <i>DFM HTTP</i> service to be stopped and restarted).	80
3.5	httpsEnabled	Defines whether the Unified Manager web server will accept requests over the HTTPS protocol. The default value for this option is No . (Note changing this setting requires the <i>DFM HTTP</i> service to be stopped and restarted).	Yes

Unified Manager Mandatory Settings Configuration Matrix			
From OM Version	Option Name	Description	Value
3.5	httpsPort	Defines the port the Unified Manager web server uses for HTTPS traffic. The default value for this option is 8443 . (Note changing this setting requires the DFM HTTP service to be stopped and restarted).	443
3.5	ldapEnabled	Defines whether LDAP authentication is enabled in Unified Manager. The default value for this option is No .	No
3.5	networkDiscoveryLimit	Defines the number of network hops Unified Manager server will send storage system discovery requests. The default value for this option is 15 .	0
3.5	perfAdvisorTransport	Defines the transport protocol used by the Performance Advisor component of the Unified Manager server when communicating with the remote storage systems. The default value for this option is httpOnly .	httpsOK
3.5	profileTTL	Defines the length of time that web cookies are valid. Web cookies contain identifying information about administrators using the Unified Manager; the cookies allow the DataFabric Manager server to log users in automatically when they return to the Unified Manager. If profile TTL is zero, the DataFabric Manager server uses only session cookies which expire when the user closes the browser. The default value for this option is 45 days .	0
3.5	serverHTTPEnabled	Defines whether HTTP protocol is enabled for NetApp Management Console / Unified Manager server communication. The default value for this option is Enabled . (Note changing this setting requires the <i>DFM HTTP</i> service to be stopped and restarted).	Disabled
3.5	serverHTTPPort	Defines the port used for NetApp Management Console / Unified Manager server communication using the HTTP protocol. The default value for this option is 8088 . (Note changing this setting requires the <i>DFM HTTP</i> service to be stopped and restarted).	8080
3.5	serverHTTPSEnabled	Defines whether HTTPS protocol is enabled for NetApp Management Console / Unified Manager server communication. The default value for this option is Disabled . (Note changing this setting requires the <i>DFM HTTP</i> service to be stopped and restarted).	Enabled
3.5	serverHTTPSPort	Defines the port used for NetApp Management Console / Unified Manager server communication using the HTTPS protocol. The default value for this option is 8488 . (Note changing this setting requires the <i>DFM HTTP</i> service to be stopped and restarted).	8443
3.5	SMTPServerPort	Defines the port Operations Manager uses to communicate with the SMTP servers. The default value for this option is 25 .	25
3.5	snmpTrapListener Enabled	Defines whether the SNMP trap listening component is enabled or not. The default value for this option is Yes .	No
3.5	snmpTrapListenerPort	Defines the port the Unified Manager server will use to receive SNMP traps. The default value for this option is 162 .	162
3.5	snmpTrapRcvdMaxPer Window	Defines the maximum number of SNMP traps an Unified Manager server will receive in a given time window before traps are ignored. The default value for this option is 250 .	250
3.5	snmpTrapRcvdWindow Size	Defines the size of the window (duration) in which Unified Manager server will receive before SNMP traps are ignored. The default value for this option is 5 minutes .	5 minutes (5m)

Unified Manager Mandatory Settings Configuration Matrix			
From OM Version	Option Name	Description	Value
3.5	useHostsEquiv	Defines whether authentication to a remote Storage System will be achieved through the hostEquivfile. The default value for this option is No .	No

Table 23 - Unified Manager Recommended (Technical Reference) Configuration Matrix

Unified Manager Recommended Settings Configuration Matrix			
From OM Version	Option Name	Description	Value
3.5	agentHostCIFSAccount	Host agent CIFS credentials for Storage Resource Management (File SRM license) (global setting). The default value for this option is blank . The use of host agents is not supported by CATE Storage and Backup Engineering.	<undefined>
3.5	agentHostPort	Host agent defined communication port (global setting). The default value for this option is blank . The use of host agents is not supported by CATE Storage and Backup Engineering.	<undefined>
3.5	agentMonInterval	Host agent monitoring interval. Defines how often Unified Manager server checks each host agent. The default value for this option is 2 minutes . The use of host agents is not supported by CATE Storage and Backup Engineering.	0 (Off)
3.5	aggrFullThreshold	The utilization percentage at which an aggregate is considered full (global setting). The default value for this option is 90%	90
3.5	aggrFullThreshold Interval	The aggregate full event is only generated if the aggregateFullThreshold condition persists for this period of time (global setting). The default value for this option is 0 seconds . A value of 0 seconds means the alert is generated immediately.	0 seconds (0s)
3.5	aggrNearlyFull Threshold	The utilization percentage at which an aggregate is considered nearly full (global setting). The default value for this option is 80%	80
3.5	aggrNearlyOver-committedThreshold	The utilization percentage at which an aggregate is considered nearly overcommitted (global setting). The default value for this option is 95%	95
3.8	aggrNearlyOver DeduplicatedThreshold	The value (as an integer percentage) that will be used to generate an “aggregate nearly over deduplicated” event for this aggregate. Over de-duplicated is consider to be the ratio of re-hydrated (un-deduplicated) to physical space in the aggregate, such that if you disabled deduplication on every volume in the aggregate and performed a sis undo, there would not be enough space in the aggregate for all the data. The default value for this option is 140% .	140
3.5	aggrOvercommitted Threshold	The utilization percentage at which an aggregate is considered overcommitted (global setting). The default value for this option is 100%	100
3.8	aggrOverDeduplicated Threshold	The value (as an integer percentage) that will be used to generate an “aggregate over deduplicated” event for this aggregate. Over de-duplicated is consider to be the ratio of re-hydrated (un-deduplicated) to physical space in the aggregate, such that if you disabled deduplication on every volume in the aggregate and performed a sis undo, there would not be enough space in the aggregate for all the data. The default value for this option is 150% .	150
3.5	aggrSnapshotFull Threshold	The utilization percentage of the Snapshot reserve on an aggregate that can be used before the system raises the Aggregate Snapshots Full event (global setting). The default value for this option is 90% . The use of aggregate snapshots is not recommended by CATE Storage and Backup Engineering.	90
3.5	aggrSnapshotNearly FullThreshold	The utilization percentage of the Snapshot reserve on an aggregate that can be used before the system raises the Aggregate Snapshots Nearly Full event (global setting). The default value for this option is 80% . The use of aggregate snapshots is not recommended by CATE Storage and Backup Engineering.	80

Unified Manager Recommended Settings Configuration Matrix			
From OM Version	Option Name	Description	Value
3.5	alertFrom	The SMTP address of the sender account used to issue email alerts. This setting will vary depending on the regional storage support team managing the Operations Manager installation. The default value for this option is blank .	<regional / support team specific setting>
4.0	autoClientStatEnabled	Configures automatic collection of per-client statistics when a set of thresholds is breached. The default value for this option is no	No
3.5	autosupportAdmin Contact	Contact details to be used in Unified Manager autosupport messages (similar to contact details in ONTAP autosupports). The default value for this option is blank . It is recommended that you also include a group line or contact phone number and externally accessible email distribution list in this field.	<regional / support team specific setting>
3.5	autosupportContent	Content level for autosupports. Can be set to complete or minimal (similar to content setting for ONTAP autosupports). The default value for this option is complete .	Complete
3.5	autosupportDestination Email	NetApp recipient address for Unified Manager autosupports. (DO NOT CHANGE THIS VALUE). The default value for this option is autosupport@netapp.com .	autosupport @ netapp.com
3.5	autosupportDestination URL	NetApp destination URL when posting Unified Manager autosupports via HTTP(S). (DO NOT CHANGE THIS VALUE). The default value for this option is support.netapp.com/asupprod/post/1.0/postAsup .	support.netapp.com/asupprod/post/1.0/postAsup
3.5	autosupportEnabled	Defines whether autosupport is enabled for this instance of Unified Manager. The default value for this option is Yes .	Yes
3.8	autosupportIncludeAll DiagInfo	Includes all diagnostic information in the autosupport. If set to yes the diagnostic information is included. The default value for this option is No .	No
3.8	autosupportIncludePerf	Include Performance Advisor autosupport data along with DataFabric Manager server autosupport message. If set to yes Performance Advisor autosupport data shall be included. The default value for this option is Yes .	Yes
3.8	autosupportIncludeProv	Include Provisioning Manager autosupport data along with DataFabric Manager server autosupport message. If set to Yes, Provisioning Manager autosupport data shall be included. The default value for this option is Yes .	Yes
3.5	autosupportMonInterval	Defines the interval at which Unified Manager checks the status of the autosupport component. The default value for this option is 2 minutes .	30 minutes (30m)
3.5	autosupportRetryCount	Defines the number of times Unified Manager will attempt to post the autosupport message to NetApp. The default value for this option is 4 .	4
3.5	autosupportRetryDelay	Defines the time Unified Manager will wait between autosupport posting attempts. The default value for this option is 15 minutes .	15 minutes (15m)
3.7	backupRetentionCount	Defines the number of database backups Unified Manager should keep before deleting old backups. The default value for this option is 0 (none) .	5

Unified Manager Recommended Settings Configuration Matrix			
From OM Version	Option Name	Description	Value
3.5	ccMonInterval	Defines the interval at which Unified Manager server checks the status of the configuration conformance of Filers whose configuration is managed via UM. The default value for this option is 4 hours .	1 day
3.5	cfMonInterval	Defines the interval at which Unified Manager server checks the status of storage clusters for failover events for Filers registered with the Operations Manager server. The default value for this option is 5 minutes .	5 minutes (5m)
3.5	chargebackDayOf Month	Defines the date of the month on which the billing cycle begins. The default value for this option is 1 (first) .	1
3.5	chargebackIncrement	Defines how the chargeback amount is defined (number of days in the cycle versus a flat rate). The default value for this option is daily .	Daily
3.5	chargebackRate	Defines the amount to charge for storage space used per GB per year. The default value for this option is blank .	<undefined>
4.0	clientStatCifsLatency	The percentage increase in CIFS latency to be used to automatically collect per-client statistics. The default value for this option is 30 .	30
4.0	clientStatCpuThreshold	Specifies the value of the system:cpu_busycounter in %, to be used to automatically collect per-client statistics. The default value for this option is 80 .	80
4.0	clientStatMinTotalOps Rate	Specifies the minimum number of total operations performed on a storage system per second, for per-client statistics to be automatically collected from it. The default value for this option is 500 .	500
4.0	clientStatNfsLatency	The percentage increase in NFS latency to be used to automatically collect per-client statistics. The default value for this option is 30 .	30
4.0	clientStatThreshold Period	The period of time over which the rate of change in latency or operations per second should be measured, specified in seconds. The default value for this option is 300 .	300
4.0	clientStatTotalOpsRate	The percentage increase in total input operations to be used to automatically per-client statistics. The default value for this option is 30 .	30
4.0	clusterMonInterval	Defines the interval at which Unified Manager server obtains the c-mode cluster information for c-mode clusters registered with the Unified Manager server. The default value for this option is 15 minutes .	15 minutes (15m)
3.5	cpuBusyThreshold Interval	The CPU busy event is only generated if the cpuTooBusyThreshold condition persists for this period of time (global setting). The default value for this option is 15 minutes .	15 minutes (15m)
3.5	cpuMonInterval	Defines the interval at which Unified Manager server obtains the CPU performance metrics for Filers registered with the Unified Manager server. The default value for this option is 5 minutes .	5 minutes (5m)
3.5	cpuTooBusyThreshold	Defines the highest level of activity the CPU can reach before a CPU Too Busy event is triggered. The default value for this option is 95% .	95
3.5	currentEventsCache Size	The event handler caches all events, not just abnormal events, in order to speed up notifications. By default, all current events are cached (that is, there is no events cache size limit by default). The default value for this option is blank .	<undefined>

Unified Manager Recommended Settings Configuration Matrix			
From OM Version	Option Name	Description	Value
3.7	databaseBackupDbeng WaitTime	Defines the number of seconds the database engine waits while a database backup is being performed. The default value for this option is 600 .	600
3.7	defReportLinesPerPage	Defines the default number of result lines returned when a report query is run on the Unified Manager server. The default value for this option is 20 .	20
3.7	dfmDataExportEnabled	Defines whether the export of data from within the Unified Manager database is permitted. The default value for this option is No .	Yes
3.5	dfMonInterval	Defines the interval at which Unified Manager server obtains the free disk space metrics for Filers registered with the Unified Manager server. The default value for this option is 30 minutes .	30 minutes (30m)
3.5	discoverInterval	Specifies the time period after which Unified Manager server scans for new storage systems and networks. The default value for this option is 15 minutes .	260 weeks (260w)
3.5	diskMonInterval	Defines the interval at which Unified Manager server checks the status of the disks for each storage system registered with the Unified Manager server. The default value for this option is 4 hours .	4 hours (4h)
3.5	envMonInterval	Defines the interval at which Unified Manager server obtains the Environmental metrics for Filers registered with the Unified Manager server. The default value for this option is 5 minutes .	5 minutes (5m)
3.5	eventsPurgeInterval	Defines the length of time Unified Manager maintains events in the database before they are automatically purged from the Unified Manager database. The default value for this option is 25.71 weeks .	4 weeks (4w)
3.7	favoriteMaxReports	Defines the maximum number of favorite reports to display in the Unified Manager user interface. The default value for this option is 25 .	25
3.5	filerConfigSaveLocal Changes	Defines whether local configuration changes on a remote Filer are saved. The default value for this option is Yes .	Yes
3.5	fsMonInterval	Defines the interval at which Unified Manager server obtains the File System metrics for Filers registered with the Unified Manager server. The default value for this option is 15 minutes .	15 minutes (15m)
3.5	groupTreeShowStatus	Defines whether the group node in the Group Tree View (left pane) will show the worse case status of a child object. The default value for this option is Enabled .	Enabled
3.5	growthRateSensitivity	Defines the sensitivity of the volume and qtree growth rate calculations (and the "Volume Growth Rate Abnormal" event). The valid range is greater than 0 and less than or equal to 5. "5" means not very sensitive to growth rate changes, low values mean very sensitive to growth rate changes. The default value for this option is 2 .	2
3.5	guiRefreshInterval	Defines the time interval before a client web browser refreshes the display. The default value for this option is 00:05:00 .	00:15:00
3.5	hbaportTooBusy Threshold	The highest level of activity the HBA Port Traffic can reach before a HBA Too Busy event is triggered. The default value for this option is 90% .	90
3.5	hostEnclosure DiscoveryEvents	Defines whether Unified Manager will generate Enclosure Discovery events (see bug 245671 for details). The default value for this option is Disabled .	DisabledFor Cluster

Unified Manager Recommended Settings Configuration Matrix			
From OM Version	Option Name	Description	Value
3.5	hostPingMethod	Defines the method used by Unified Manager server to ping the remote Filer. The default value for this option is echo_snmp .	echo_snmp
3.7	hostRBACMonInterval	The interval at which Unified Manager server obtains RBAC (user, group, role, domainuser) information for Filers registered with the Unified Manager server. The default value for this option is 1 day .	1 day (1d)
3.5	ifMonInterval	Defines the interval at which Unified Manager server checks the status of the interfaces for each storage system registered with the Unified Manager server. The default value for this option is 15 minutes .	15 minutes (15m)
3.5	licenseExpireWarningThreshold	Defines the number of days prior to an Unified Manager license expiry at which time a license expiration warning event is triggered. The default value for this option is 5 days .	90 days (90d)
3.5	licenseMonInterval	Defines the interval at which Unified Manager server checks the status/validity of the software licenses in the Operations Manager server. The default value for this option is 4 hours .	1 day (1d)
3.5	localHostName	Defines the local host name of the server hosting Unified Manager. The default value for this option is blank .	<regional/ support team specific setting>
3.5	lunMonInterval	Defines the interval at which Unified Manager server checks the status of LUNS on each storage system registered with the Unified Manager server. The default value for this option is 30 minutes . The use of LUNS on NetApp Storage is not supported by CATE Storage and Backup Engineering.	0 (Off)
3.7	maxReportLinesPerPage	Defines the default maximum number of lines on a report view page. The default value for this option is 1000 .	1000
3.5	monMinFreeBytes	Defines the number of bytes of free disk space required (in the directory containing the database) for the monitor to continue to update the database. The default value for this option is 10000 .	10000
3.5	monMinFreePercent	Defines the percentage of free disk space required (in the directory containing the database) for the monitor to continue to update the database. The default value for this option is 10% .	15
3.5	monSNMPRetries	Defines the number of times to retry requests before failing when using SNMP-based polling. The default value for this option is 4 .	4
3.5	monSNMPTimeout	Defines how long to wait for SNMP replies before retrying when using SNMP-based polling. The default value for this option is 5 .	5
3.7	ndmpDataUseAllInterfaces	Backup Manager honors the NDMP preferred interface setting on the primary storage system when performing SnapVault backup and on the secondary storage system when restoring. Override this behavior by setting ndmpDataUseAllInterfaces to "yes". The default value for this option is 0 .	0
3.5	nodesRemainingWarningThreshold	Defines the number of unallocated Unified Manager node licenses remaining prior to a node limit nearly reached warning event is triggered. The default value for this option is 1 .	20
3.5	opsMonInterval	Defines the interval at which Unified Manager server checks the count of operations on each storage system registered with the Unified Manager server. The default value for this option is 10 minutes .	10 minutes (10m)

Unified Manager Recommended Settings Configuration Matrix			
From OM Version	Option Name	Description	Value
4.0	perfMaxObjectInstancesInBarChart	The maximum number of object instances allowed in a bar chart. The default value is 20 .	20
4.0	perfSampleRate1	Specify the sample rates. This sample-rate is positive number, optionally followed by a time period suffix to indicate seconds, minutes, hours, days, or weeks. The default value for perfSampleRate1 is 1 minute	1 minute (1m)
4.0	perfSampleRate2	Specify the sample rates. This sample-rate is positive number, optionally followed by a time period suffix to indicate seconds, minutes, hours, days, or weeks. The default value for perfSampleRate1 is 5 minutes	5 minutes (5m)
4.0	perfSampleRate3	Specify the sample rates. This sample-rate is positive number, optionally followed by a time period suffix to indicate seconds, minutes, hours, days, or weeks. The default value for perfSampleRate1 is 15 minutes	15 minutes (15m)
4.0	perfSampleRate4	Specify the sample rates. This sample-rate is positive number, optionally followed by a time period suffix to indicate seconds, minutes, hours, days, or weeks. The default value for perfSampleRate1 is 30 minutes	30 minutes (30m)
3.5	pingMonInterval	Defines the interval at which Unified Manager server pings each storage system registered with the Operations Manager server. The default value for this option is 1 minute .	1 minute (1m)
3.5	pingMonRetryDelay	Defines the delay (in seconds) before Unified Manager attempts to retry a ping monitoring test to a Storage System. The default value for this option is 3 seconds .	3
3.5	pingMonTimeout	Defines the timeout threshold (failure to respond) for a ping test from the Unified Manager server to the Storage System. The default value for this option is 3 seconds .	3
3.5	qtreeFullThreshold	Defines the percentage at which a qtree is considered full (global setting). The default value for this option is 90%	90
3.5	qtreeFullThresholdInterval	Qtree full event is only generated if the qtreeFullThreshold condition persists for this period of time (global setting). The default value for this option is 0 seconds . A value of 0 seconds means the alert is generated immediately.	0 seconds (0s)
3.5	qtreeGrowthEventMinChangePct	Defines the minimum qtree percentage utilization change required to generate a qtree growth abnormal message. The default value for this option is 1%	1
3.5	qtreeMonInterval	Defines the interval at which Unified Manager server checks the status of qtrees on each storage system registered with the Unified Manager server. The default value for this option is 8 hours .	8 hours (8h)
3.5	qtreeNearlyFullThreshold	Defines the percentage at which a qtree is considered nearly full (global setting). The default value for this option is 80%	80
3.7	recentMaxReports	Defines the maximum number of recent reports displayed when recent reports list is queried. The default value for this option is 25.	25

Unified Manager Recommended Settings Configuration Matrix			
From OM Version	Option Name	Description	Value
3.5	SANHostMonInterval	Defines the interval at which Unified Manager server checks the status of SAN hosts registered with the Unified Manager server. The default value for this option is 5 minutes . The use of SAN Attached NetApp Storage is not supported by CATE Storage and Backup Engineering (except Tape SAN).	0 (Off)
3.5	SANHostMonSnapshot LUNs	Defines whether Unified Manager server will monitor Snapshot LUNs on SAN hosts. The default value for this option is Enabled . The use of SAN Attached NetApp Storage is not supported by CATE Storage and Backup Engineering (except Tape SAN).	Disabled
3.5	SMTPServerBackup	Defines the name or IP address of the backup (secondary) SMTP server which Unified Manager will use to send email alerts should the primary SMTP server be unavailable. The default value for this option is blank .	<regional / support team specific setting>
3.5	SMTPServerName	Defines the name or IP address of the SMTP server which Unified Manager will use to send email alerts. The default value for this option is blank .	<regional / support team specific setting>
3.5	snapmirrorMonInterval	Defines the interval at which Unified Manager server checks the status of snapmirror jobs on each storage system registered with the Unified Manager server. The default value for this option is 30 minutes .	30 minutes (30m)
3.7	snapshotDiscovery EventsEnabled	Defines whether Unified Manager server will generate an event for Snapshot Discovered and Snapshot Deleted occurrences. The default value for this option is No .	No
3.5	snapshotMonInterval	Defines the interval at which Unified Manager server checks the status of snapshots on each storage system registered with the Unified Manager server. The default value for this option is 30 minutes .	30 minutes (30m)
3.5	statusMonInterval	Defines the interval at which the Unified Manager server collects global status information on each storage system registered with the Unified Manager server. The default value for this option is 10 minutes .	10 minutes (10m)
3.5	sysInfoMonInterval	Defines the interval at which Unified Manager server checks the values of System Information on each storage system registered with the Unified Manager server. The default value for this option is 1 hour .	1 hour (1h)
3.5	userEnableAlerts	Defines whether user level quota alerts are enabled on the Unified Manager server. The default value for this option is Yes . The use of User Quotas on NetApp Storage is not supported by CATE Storage and Backup Engineering.	No
3.5	userFullThreshold	Defines the percentage at which a user quota is considered full (global setting). The default value for this option is 90% . The use of User Quotas on NetApp Storage is not supported by CATE Storage and Backup Engineering.	90

Unified Manager Recommended Settings Configuration Matrix			
From OM Version	Option Name	Description	Value
3.5	userNearlyFullThreshold	Defines the percentage at which a user quota is considered nearlyfull (global setting). The default value for this option is 80% . The use of User Quotas on NetApp Storage is not supported by CATE Storage and Backup Engineering.	80
3.5	userQuotaMonInterval	Defines the interval at which Unified Manager server checks the user level quotas on each storage system registered with the Unified Manager server. The default value for this option is 1 day . The use of User Quotas on NetApp Storage is not supported by CATE Storage and Backup Engineering.	0 (Off)
3.5	vFilerMonInterval	Defines the interval at which Unified Manager server checks the status and attributes of vFilers on each storage system registered with the Unified Manager server. The default value for this option is 1 hour .	1 hour (1h)
3.5	volFullThreshold	Defines the percentage at which a volume is considered full (global setting). The default value for this option is 90% .	90
3.5	volFullThresholdInterval	Volume full event is only generated if the volFullThreshold condition persists for this period of time (global setting). The default value for this option is 0 seconds . A value of 0 seconds means the alert is generated immediately.	0 seconds (0s)
3.5	volGrowthEventMinChangePct	Defines the minimum volume percentage utilization change required to generate a qtree growth abnormal message. The default value for this option is 1%	1
3.5	volNearlyFullThreshold	Defines the percentage at which a volume is considered nearlyfull (global setting). The default value for this option is 80% .	80
3.5	volNearlyNoFirstSnapThreshold	Defines the percentage at which a volume is considered to have consumed most of the free space for its space reservation—space that the volume needs when the first Snapshot copy is created. The default value for this option is 80% .	80
3.7	volNearlyOver-committedThreshold	Defines the value (as an integer percentage) of the fullness threshold that will be used to generate an "volume nearly overcommitted" event for volumes which do not have a per-volume volNearlyOvercommittedThreshold. The default value for this option is 95% .	95
3.8	volNearlyOverDeduplicatedThreshold	The value (as an integer percentage) that will be used to generate "volume nearly over deduplicated" event for volumes which do not have a per-volume value. Over de-duplicated is consider to be the ratio of re-hydrated (un-deduplicated) to logical space in the volume, such that if you disabled deduplication on the volume and performed a sis undo, there would not be enough space in the volume for all the data. The default value for this option is 140% .	140
3.5	volNoFirstSnapThreshold	The percentage at which a volume is considered to have consumed all of the free space for its space reservation, space that the volume needs when the first Snapshot copy is created. The default value for this option is 90% . This option applies to volumes that contain space-reserved files, no Snapshot copies, a fractional overwrite reserve set to greater than 0, and where the sum of the space reservations in the volume is greater than the free space available to the volume.	90

Unified Manager Recommended Settings Configuration Matrix			
From OM Version	Option Name	Description	Value
3.7	volOvercommitted Threshold	Defines the value (as an integer percentage) of the fullness threshold that will be used to generate an "volume overcommitted" event for volumes which do not have a per-volume volOvercommittedThreshold. A volume is overcommitted when the quota limits of the qtrees within the volume add up to a size greater than the size of the volume. The default value for this option is 100% .	100
3.8	volOverDeduplicated Threshold	The value (as an integer percentage) that will be used to generate "volume over deduplicated" event for volumes which do not have a per-volume value. Over de-duplicated is consider to be the ratio of re-hydrated (un-deduplicated) to logical space in the volume, such that if you disabled deduplication on the volume and performed a sis undo, there would not be enough space in the volume for all the data. The default value for this option is 150% .	150
3.5	volReserveDepleted Threshold	Defines the percentage at which a volume is considered to have consumed all of its reserved blocks. This option applies to volumes with LUNs, Snapshot copies, no free blocks, and a fractional overwrite reserve of less than 100%. A volume that has crossed this threshold is getting dangerously close to having write failures. The default value for this option is 90% .	90
3.5	volReserveNearly DepletedThreshold	Defines the percentage at which a volume is considered to have consumed most of its reserved blocks. This option applies to volumes with LUNs, Snapshot copies, no free blocks, and a fractional overwrite reserve of less than 100%. A volume that crosses this threshold is getting close to having write failures. The default value for this option is 80% .	80
3.5	volSnapshotCount Threshold	Defines the number of Snapshot copies which, if exceeded, is considered to be too many for the volume. The default value for this option is 250 .	250
3.5	volSnapshotFull Threshold	Defines the percentage at which the space reserved for taking volume Snapshot copies is considered full. The default value for this option is 90% .	90
3.5	volSnapshotTooOld Threshold	Defines the age of a Snapshot copy which, if exceeded, is considered too old for the volume. The Snapshot age can be specified in seconds, minutes, hours, days, or weeks. The default value for this option is 52 weeks .	52 weeks (52w)
4.0	vserverMonInterval	Defines the interval at which Unified Manager server checks the status and attributes of vServers on each c-mode cluster system registered with the Unified Manager server. The default value for this option is 1 hour .	1 hour (1h)

Table 24 - Unified Manager Recommended (Technical Reference) Configuration Matrix for Performance Advisor

Unified Manager Recommended Settings Configuration Matrix for Performance Advisor Enabled Systems			
From OM Version	Option Name	Description	Value
3.7	perfAdvisorDisabledCounterGroups	Defines how the counter groups are configured within Performance Advisor. The default value for this option is ' Qtree Basic '.	Qtree Basic
3.5	perfAdvisorEnabled	Defines whether the performance advisor component of the Operations Manager server is enabled. The default value for this option is Enabled .	Enabled
3.5	perfAdvisorMaxMonitorThreads	Defines the maximum number of threads to be used for collecting performance advisor data from storage systems. The default value for this option is 32 .	32
3.5	perfAdvisorPollInterval	Defines the interval at which performance advisor polls the storage systems it is monitoring (global setting). The default value for this option is 10 minutes .	5 minutes (5m)
3.7	perfAdvisorShowAllViews	Defines whether performance advisor will list all available console views. The default value for this option is Disabled .	Enabled
3.5	perfAdvisorShowDiagCounters	Defines whether diagnostic counters are available when creating custom views in the NetApp Management Console. The default value for this option is Disabled . (Note changing this setting requires the DFM server service to be stopped and restarted).	Enabled
3.7	perfAdvThreshViolationMonInterval	Defines how often to check each storage system for violations of Performance Advisor thresholds. The default value for this option is 15 minutes .	15 minutes (15m)
3.7	perfDataExportEnabled	Defines whether the export of data from within Performance Advisor is permitted. The default value for this option is No .	Yes

11.2 Unified Manager Additional Server Settings Configuration Matrix

The purpose of this section is to define and describe other configuration settings (outside of the `dfm options` command) used when implementing Unified Manager server in Citi. These settings do not have vendor default values for each command.

Unlike the previous section, the Citi recommended value for each setting will not depend on the way in which you are planning to implement Unified Manager. These settings are considered standard across all implementations.

Option Supported as of Version	Setting	Description	Recommended Configuration
3.5	dfm ssl server details	Information regarding the SSL configuration of the Unified Manager web server	Difference between "Not Before" and "Not After" values should be no more than 2 years "RSA Public Key" should be 1024-bit

11.3 NetApp Management Console (Client) Settings Configuration Matrix

The purpose of this section is to define and describe the configuration options used when implementing NetApp Management Console (client for Performance Advisor) for Unified Manager in Citi.

The vendor's default settings for each option are described in the 'description' column.

Unlike the Unified Manager server there is only one optimized (recommended) configuration for the NetApp Management Console (NMC). The settings listed below are shown for reference purposes only.

The options in the following table can be classified into two broad categories. For reference purposes these categories have been color coded.

Security

Technical best practice

Option Supported as of Version	Option Name	Description	Recommended Setting
OM 3.7 NMC 2.2	<ServerLogin> <protocol>	Defines which protocol is used for NetApp Management Console / Unified Manager server communication. The default value for this option is HTTP . (Note changing this setting requires the user to exit the <i>NMC</i> and log back in).	HTTPS
OM 3.7 NMC 2.2	<ServerLogin> <port>	Defines the port used for NetApp Management Console / Unified Manager server communication using the HTTPS protocol. The default value for this option is 8488 . (Note changing this setting requires the user to exit the <i>NMC</i> and log back in).	8443

Table 25 – NMC Client Recommended (Technical Reference) Configuration

11.4 Unified Manager Server Settings (Options) Configuration Addendum Matrix for Large to Very Large Environments

The purpose of this section is to define and describe the configuration options used when implementing Unified Manager server in Citi in a region that has a large or very large environment (tending towards or over 250 NetApp storage controllers). For environments that qualify the settings in these tables supersede those settings listed in other subsection of this chapter ([11](#)).

The Citi recommended value for each option has been developed with the vendor based on their experience and understanding of the Citi environment. The values have been incorporated into the configuration workflow when a large environment is specified during the interview phase of the workflow.

Table 26 -Very Large Environment Custom Settings for Regional Master Reporting Server

Unified Manager Recommended Settings Configuration Matrix for Large to Very Large Environments				
From OM Version	Option Name	NetApp Default	Small to Medium Environment Value	Large to Very Large Environment Value
3.7	autosupportMonInterval	2 mins	30 minutes	1 hour
4.0	backupDirMonInterval	8 hours	n/a	Off
3.7	ccMonInterval	4 hours	1 day	Off
3.7	cfMonInterval	5 minutes	5 minutes	12 hours
4.0	clusterMonInterval	15 minutes	n/a	12 hours
3.7	cpuMonInterval	5 minutes	5 minutes	1 hour
3.7	dfMonInterval	30 minutes	30 minutes	12 hours
3.7	diskMonInterval	4 hours	4 hours	24 hours
3.7	dsConformanceMonInterval	1 hour	n/a	off
3.7	dsDRMonInterval	15 minutes	n/a	off
3.7	dsProtectionMonInterval	15 minutes	n/a	off
3.7	envMonInterval	5 minutes	5 minutes	off
3.7	fsMonInterval	15 minutes	15 minutes	12 hours
3.7	hostRBACMonInterval	1 day	1 day	Off
3.7	iffMonInterval	15 minutes	15 minutes	1 hour
3.7	licenseMonInterval	4 hours	1 day	24 hours
3.7	lunMonInterval	30 minutes	Off	Off
4.0	ndmpMonInterval	30 minutes	30 minutes	Off
3.7	opsMonInterval	10 minutes	10 minutes	6 hours
4.0	perfAdvThreshViolationMonInterval	15 minutes	n/a	Off
3.7	perfAdvisorEnabled	Enabled	Off	Off
3.7	pingMonInterval	1 minute	1 minute	1 hour

Unified Manager Recommended Settings Configuration Matrix for Large to Very Large Environments (cont.)				
From OM Version	Option Name	NetApp Default	Small to Medium Environment Value	Large to Very Large Environment Value
3.7	qtreeMonInterval	8 hours	8 hours	24 hours
4.0	respoolSpaceMonInterval	1 hour	n/a	Off
3.7	SANHostMonInterval	5 minutes	Off	Off
4.0	shareMonInterval	1 hour	1 hour	12 hours
3.7	snapmirrorMonInterval	30 minutes	30 minutes	12 hours
3.7	snapshotMonInterval	30 minutes	30 minutes	12 hours
4.0	snapvaultMonInterval	30 minutes	n/a	Off
3.7	srmMonInterval	10 minutes	Off	Off
3.7	statusMonInterval	10 minutes	10 minutes	6 hour
3.7	sysInfoMonInterval	1 hour	1 hour	6 hour
3.7	userQuotaMonInterval	1 day	Off	2 days
3.7	vFilerMonInterval	1 hour	1 hour	12 hours
4.0	vserverMonInterval	1 hour	1 hour	Off

Table 27 -Very Large Environment Custom Settings for Site Specific Monitoring and Alerting Servers

Unified Manager Recommended Settings Configuration Matrix for Large to Very Large Environments				
From OM Version	Option Name	NetApp Default	Small to Medium Environment Value	Large to Very Large Environment Value
3.7	snapshotMonInterval	30 minutes	30 minutes	3 hours

The terms “regional master reporting server” and “site specific monitoring and alerting servers” are particular to the geographic and functionally distributed reference architecture discussed in section [9.1.2](#).

11.5 Unified Manager Event Listing

A complete list of events types and their associated severities can be found in the [NetApp Operations Manager Admin Guide](#) in the section 'List of Events and Severity Types'.

11.6 UM Database Objects and Attributes

The purpose of this section is to provide a reference list of the database objects and their attributes that are relevant to how NetApp NAS is deployed within Citi. Therefore not all the available objects in the database will be listed here.

Object (Base Catalog)	Attributes / Values	
Aggregate	Aggregate Name Aggregate Full Name Aggregate Id Aggregate Group Id Aggregate Storage System (A) Aggregate Comment Aggregate Type Aggregate RAID Type Aggregate State SnapLock Aggregate Status Aggregate Mirrored Aggregate Deleted When Aggregate Deleted By Aggregate Used Capacity Aggregate Used Capacity % Aggregate Full Threshold Aggregate Nearly Full Threshold	Aggregate Nearly Over Deduplicated Threshold Aggregate Over Deduplicated Threshold Aggregate Full Threshold Interval Aggregate Overcommitted Threshold Aggr. Nearly Overcommitted Thresh. Aggregate Total Space Aggregate Space Available Aggregate Space Available in % Aggregate Bytes Committed Aggregate Bytes Committed % Aggregate Daily Growth Rate Aggregate Daily Growth (%) Aggregate Days to Full Aggregate Snapshot Disabled Aggregate Snapshot Autodelete Aggregate Snap Reserve Total Aggregate Snap Reserve Used Aggregate Snap Reserve Used %
Appliance	Appliance Name Appliance Full Name Appliance Id Appliance Group Id Appliance Type Appliance Model Appliance Serial Number Appliance Uptime Appliance Ping Timestamp Appliance Status Appliance Deleted When Appliance Deleted By Appliance GUI Link Appliance Secure GUI Link Appliance Comment Appliance Down Timestamp Appliance Temperature	Appliance Fans Appliance Power Units Appliance NVRAM Battery Appliance IP Address Appliance Console Address Appliance Contact Appliance Locations Appliance OS Version Appliance Firmware Version Appliance Config Group Appliance Config Matches Appliance CPU Threshold Appliance CPU Threshold Interval Appliance CPU % Appliance Ping Status Appliance RLM Status Appliance Partner
Disk	Disk Name Disk RAID Group Id Disk Storage System (A) Disk Aggregate (A) Disk Size Disk Firmware Revision Disk Vendor Name Disk Model Disk Type Disk RPM Disk Shelf Disk Bay Disk Plex Disk Bytes per Sector Disk Checksum Compatibility Disk Host Adapter	Disk Id Disk Is Zeroed Disk Physical Blocks Disk Pool Disk Port Disk Port Name Disk Raid Group Disk Raid Type Disk Raid State Disk SecondaryName Disk SecondaryPort Disk Serial Number Disk Used Blocks Disk Used Space Disk Volume Disk Zeroing Percent

Object (Base Catalog)	Attributes / Values	
Event	Event Description Event Id Event Severity Event Time of Trigger Event Acknowledged By Event Acknowledged Event Source Id Event Source Event Source Type Event Application Type Event Agent (A) Event Storage System (A) Event Aggregate (A) Event Appliance (A)	Event Qtree (A) Event Volume (A) Event SnapReserve (A) Event Interface (A) Event vFiler (A) Event Primary Directory (A) Event Script Job (A) Event Time of Deletion Event Deleted By SNMP Trap Description SNMP Trap Condition SNMP Trap Severity SNMP Trap Timestamp
Filer	Storage System Name Storage System Full Name Storage System Id Common Appliance Properties (A) Storage System Used Capacity Storage System Total Capacity Storage System Used % Storage System CPU % Storage System Total Ops/Sec Storage System NFS Ops/Sec Storage System CIFS Ops/Sec Storage System HTTP Ops/Sec Storage System FCP Ops/Sec Storage System iSCSI Ops/Sec Storage System NFS protocol Storage System CIFS protocol Storage System SnapMirror protocol Storage System DAFS protocol Storage System Vfiler protocol Storage System FCP protocol Storage System iSCSI protocol	Storage System SnapVault Primary protocol Storage System SnapVault Secondary protocol Storage System Beginning of Period This Month Storage System End of Period This Month Storage System Days in cycle This Month Storage System Avg. Usage This Month Storage System Rate/GB Storage System Monthly Rate/GB Storage System Usage Charge This Month Storage System Avg. Allocation This Month Storage System Allocation Charge This Month Storage System Beginning of Period Last Month Storage System End of Period Last Month Storage System Days in cycle Last Month Storage System Avg. Usage Last Month Storage System Rate/GB Last Month Storage System Usage Charge Last Month Storage System Avg. Allocation Last Month Storage System Allocation Charge Last Month Storage System Failed Disk Count Storage System Failed Disks Number of vFilers Storage System Physical Storage
HostRole	HostRole Name HostRole Id HostRole Description	Storage System (A) vFiler (A) HostRole Capabilities
HostUsergroup	HostUsergroup Name HostUsergroup Id HostUsergroup Description	Storage System (A) vFiler (A) HostUsergroup Roles
HostUser	HostUserName HostUserId Storage System (A)	vFiler (A) HostUser Usergroups
HostLocalUser	HostLocalUser Name HostLocalUser Full Name HostLocalUser Id HostLocalUser Description Storage System (A)	vFiler (A) HostLocalUser Usergroups HostLocalUser Min Password Age HostLocalUser Max Password Age HostLocalUser Status
HostDomain User	HostDomainUser Name HostDomainUser Id HostDomainUser Domain Name HostDomainUser SID	Storage System (A) vFiler (A) HostDomainUser Usergroups

Object (Base Catalog)	Attributes / Values	
Interface	Interface Name Interface Full Name Interface Id Interface Appliance (A) Interface Address Interface Mask Interface Physical Address Interface Type Interface Input Octets/sec	Interface Output Octets/sec Interface Input Errors Interface Input Errors/sec Interface Output Errors Interface Output Errors/sec Interface MTU Interface Speed Interface Status
IP Address	IP Address IP Address Interface (A)	IP Address vFiler (A)
Qtree	Qtree Name Qtree FullName Qtree Id Qtree Group Id Qtree Volume (A) Qtree vFiler (A) Qtree Status Qtree Time of Deletion Qtree Deleted By Qtree Comment Qtree Used Capacity Qtree Used Capacity % Qtree Soft Limit Qtree Disk Space Limit Qtree Files Limit Qtree Files Used Qtree Files Used % Qtree Available Qtree Available % Qtree Daily Growth Rate Qtree Daily Growth % Qtree Days to Full Qtree Possible Addition Qtree Possible Available Qtree Period Beginning This Month Qtree Period End This Month	Qtree Days In Cycle This Month Qtree Avg. Usage This Month Qtree Annual Rate/GB Qtree Rate/GB This Month Qtree Usage Charge This Month Qtree Avg. Allocation This Month Qtree Allocation Charge This Month Qtree Period Beginning Last Month Qtree Period End Last Month Qtree Days In Cycle Last Month Qtree Avg. Usage Last Month Qtree Rate/GB Last Month Qtree Usage Charge Last Month Qtree Avg. Allocation Last Month Qtree Allocation Charge Last Month Qtree Full Threshold Qtree Nearly Full Threshold Qtree Full Threshold Interval Qtree Bytes Used % Is PrimaryDirectory Qtree Snapvaulted To Snapvaulted From Is SnapMirror Source Is SnapMirror Destination Mirror Of Mirrored To
Script	Script Name Script Id Script Version Script Provider Script Scheduled Script Description	Script Installed By Script Installed On Script Required Privileges Script Installed In Script Interpreter
Script Job	Job Name Job Id Job Schedule (A)	Job Status Job Start Time
Script Schedule	Schedule Name Schedule Id Schedule Script (A) Schedule Enabled Schedule Schedule Schedule Last Run	Schedule Last Run Result Schedule Created By Schedule Arguments Schedule Applied to Subgroups Schedule Group(s)
Snapmirror Relationship	Id Source Source Storage System Destination Destination Storage System	State Lag Status Refreshed

Object (Base Catalog)	Attributes / Values	
Snapreserve	SnapReserve Name SnapReserve Full Name SnapReserve Id SnapReserve Volume (A) SnapReserve Used Space SnapReserve Total Space SnapReserve Available Space	SnapReserve Used % SnapReserve Status SnapReserve Full Threshold SnapReserve DailyGrowth Rate SnapReserve DailyGrowth % SnapReserve Days to full
Snapshot	SnapshotName SnapshotId SnapshotSnapReserve (A)	SnapshotAccess Time SnapshotDependency
vFiler	vFiler Name vFiler Full Name vFiler Internal Name vFiler Id vFiler Group Id vFiler System Id vFiler Comment vFiler Deleted At vFiler Deleted By vFiler Ping Timestamp vFiler Down Timestamp vFiler Hosting Storage System (A) vFiler Type vFiler Status vFiler Ping Status vFiler PrimaryIP Address vFiler IP Space vFiler Config Group vFiler Config Matches vFiler CPU % vFiler CPU Threshold vFiler NFS Allowed	vFiler CIFS Allowed vFiler iSCSI Allowed vFiler RSH Allowed vFiler PrimaryStorage Unit vFiler Beginning of Period This Month vFiler End of Period This Month vFiler Days in cycle This Month vFiler Avg. Usage This Month vFiler Rate/GB vFiler Monthly Rate/GB vFiler Usage Charge This Month vFiler Avg. Allocation This Month vFiler Allocation Charge This Month vFiler Beginning of Period Last Month vFiler End of Period Last Month vFiler Days in cycle Last Month vFiler Avg. Usage Last Month vFiler Rate/GB Last Month vFiler Usage Charge Last Month vFiler Avg. Allocation Last Month vFiler Allocation Charge Last Month vFiler Used Capacity
Volume	Volume Name Volume Full Name Volume Id Volume Group Id Volume Aggregate (A) Volume vFiler (A) Volume Type Volume RAID Volume State SnapLock Minimum Retention Period SnapLock Default Retention Period SnapLock Maximum Retention Period Volume Status Volume Time of Deletion Volume Deleted By Volume Comment Volume Clone Parent Volume Clones Volume Total Capacity Volume Autosize State Volume Maximum Size Volume Increment Size Volume Used Capacity Volume Used % Volume Quota Committed Volume Quota Committed % Volume Available	Vol. Agg. Space After 1st Snapshot Volume Space After 1st Snapshot % Vol. Agg. Spc. After 1st Snapshot % Volume Space Guarantee Volume Unused Guaranteed Space Volume Unused Guaranteed Space % Vol. Agg. Unused Guaranteed Space % Volume Snapshot Count Volume Nearly No Snapshot Threshold Volume No Snap Threshold Volume Reserve Depleted Threshold Vol. Res. Nearly Depleted Thresh. Volume Beginning of Period This Month Volume End of Period This Month Volume Days in Cycle This Month Volume Avg. Usage This Month Volume Annual Rate/GB Volume Rate/GB This Month Volume Usage Charge This Month Volume Avg. Allocation This Month Volume Allocation Charge This Month Volume Beginning of Period Last Month Volume End of Period Last Month Volume Days in Cycle Last Month Volume Avg. Usage Last Month Volume Rate/GB Last Month Volume Usage Charge Last Month

Object (Base Catalog)	Attributes / Values	
Volume (Continued)	Volume Available % Volume Daily Growth Rate Volume Daily Growth % Volume Days to full Volume Possible Addition Volume Possible Available Space Volume Space Breakout Volume Nearly Over Deduplicated Threshold Volume Overwrite Rate Volume Overwrite Reserve Available Volume Overwrite Days to Full Volume Over Deduplicated Threshold Volume Reserved Files Total Size Volume Fractional Overwrite % Volume Reserved Space Used Volume Reserved Space Used % Volume Reserved Space Available Volume Reserved Space Total Volume First Snapshot Reserve Reqd Volume Space After 1st Snapshot	Volume Avg. Allocation Last Month Volume Allocation Charge Last Month Volume Full Threshold Volume Nearly Full Threshold Volume Full Threshold Interval Vol. Quota Overcommitted Threshold Vol. Quota Nearly Overcommitted Th. Volume – Deduplication or File/LUN clone Volume Inodes Used % Hourly Retention Nightly Retention Weekly Retention Scheduled Snapshots Snapshot Directory Visible Is Secondary Volume Is SnapMirror Source Is SnapMirror Destination Mirror Of Mirrored To

11.7 NetApp OnCommand Unified Manager Console

The purpose of this section is to provide a brief description of the tabs on the new OnCommand Console. NetApp OnCommand Console in the figure below shows statistics of managed filers and the various tabs which includes dashboard, events, storage, server, policies, datasets and reports.

The sections to follow will describe each tab in more details.




Figure 122 – NetApp OnCommand Console

11.7.1 OnCommand console dashboard

OnCommand console dashboard contains multiple panels that provide cumulative at-a-glance information about the storage and virtualization environment which includes the following:

- **Availability** – this panel list the availability of the controllers (filers) and vFilers that being monitored in the environment.
- **Events** – this panel provides information about the status of the objects and the management host by listing the top 5 events for 7-Mode filers generated by OnCommand Unified Manager (Operations Manager/ DataFabric Manager) server
- **Resource Pools** – this panel displays the top 5 resource pools based on the total space allocated to and the percentage of space used by each resource pool
- **Full soon storage** – this panel displays the top 5 aggregates and volumes based on the full storage capacity, growth rate, and the number of days remaining before the object reaches full storage capacity.
- **Fastest growing storage** – this panel displays the top 5 aggregates and volumes based on the growth rate
- **Dataset overall status** – this panel summarizes the number of datasets in overall Error, Warning or Normal status. The dataset status based on the combined status conditions for disaster

recovery, protection, conformance, space and resources. If one of the conditions displays Error as the status value, the overall dataset status is displayed as Error

- External relationship lags – this panel summarizes the relative percentages of external SnapVault, qtree SnapMirror and volume SnapMirror relationships with lag times of Normal, Warning and Error status
- Unprotected data – this panel summarizes the total number of storage objects (volumes and qtrees) that are not protected
- Get started – this panel provides a link to NetApp Software page in NetApp University Web site

11.7.2 OnCommand console events

Events are generated automatically when a predefined condition occurs. This panel summarizes all the events that had been generated and ignoring these could potentially cause performance and system unavailability.

In the events tab, alarms could be configured to send notification automatically when specific events or events of a specific severity occur.

11.7.3 OnCommand console storage

OnCommand console storage tab summarizes information about storage controllers, aggregates, disks, vFiler Units, volumes, qtrees, luns¹³, user or group quotas¹⁴, snapshot copies and deleted objects that are being monitored.

Objects are monitored by DataFabric Manager when it is added to it but object deleted would have different results.

In the OnCommand console, monitoring could be stopped by deleting it from the global group. When an object is deleted, the DataFabric Manager server stops collecting and reporting data about it. Data collection and reporting is resumed only when the object is added back to the OnCommand console database.

When storage object are deleted from any group other than global group, the object is only deleted from that group i.e. DataFabric Manager Server will continue to collect and report data about it. Deletion of object must be done from the global group if monitoring and collection of data need to be stopped.

11.7.4 OnCommand console server

OnCommand console server tab summarizes information that relates to VMWare but in Citi, as we do not deploy NetApp agents on ESX servers, this tab will be empty.

11.7.5 OnCommand console policies

The OnCommand console policies tab enables you to add, edit, copy, delete and list local protection policies. Listed local policies could be assigned to virtual object datasets to configure local protection of their virtual object.

11.7.6 OnCommand console datasets

Datasets is a set of virtual or physical containers that could be used to configure as a unit for the purpose of group protection or group provisioning operations. During dataset configuration, you could additionally configure or assign a local protection or remote protection arrangements and schedules that apply to all objects in that dataset.

¹³ Not a supported configuration in Citi

¹⁴ Not a supported configuration in Citi

11.7.7 OnCommand console reports

OnCommand console reports tab allows you to do tasks related to reports management. In this tab, you could schedule a report and send it to one or more users in various file formats.

11.7.8 OnCommand console alarms

In OnCommand console alarms tab, you would be able to view a list of alarms configured based on event, event security type and event class. In this tab, you could also edit, delete, test and enable or disable alarms.

12 Document Version Control

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

Version	Changes
0.1	Initial release taken from version 1.1 of the OM 3.8.1 Solutions Guide, changes include: - Updated document designation from 10005 to 11001 - Reset version numbering from 1.1 to 0.1 - Reset table and figure numbering scheme to avoid including chapter references and added 'table' and 'figure' prefixes back into the document - Changed document title from Operations Manager 3.8.1 to Operations Manager 4.0.1 Solutions Guide - Update hyperlinks to Release Notes documentation - Updated team name to CATE Storage and Backup Engineering - Updated Table 4 with the latest SPAC and TPM package names - Updated CPU core count from 2 to 4 and memory from 4GB to 10GB in Table 6 based on recommendations from NetApp Operations Manager Technical Marketing Engineer - Updated sections 8.3.6 and 9.6 with details of the scripted automatic configuration of OM
1.0	- Corrected inconsistencies in intra-document hyperlinks - Added new 4.0.1 settings missed in Table 22 and Table 23 - autoClientStatEnabled - clientStatCifsLatency - clientStatCPUthreshold - clientStatMinTotalOpsRate - clientStatNFSLatency - clientStatThresholdPeriod - clientStatTotalOpsRate - clusterMonInterval - discoverClusters - perfMaxObjectInstancesInBarChart - perfSampleRate1 - perfSampleRate2 - perfSampleRate3 - perfSampleRate4 - vserverMonInterval - Removed section 11.6 as the database upgrade console output was repetition of information available from the NetApp Support site. - Added a new section Error! Reference source not found. providing detailed walkthrough of installation automation - Added section 11.4 detailing optimized configuration parameters for large to very large environments, originally incorporated in CATE Operations Manager for Large Environments technical addendum - Resized previous large environment architectural reference to medium sized environment and added new architectural reference for very large environments as section 9.1.2 based on the CATE Operations Manager for Large Environments technical addendum - Updated Performance Advisor access recommendations in section 9.14.2.3
1.1	Decrease CPU core count requirement for 4 to 2 in Table 6 based outcome of support case 2002059998

Version	Changes
1.2	- Updated Table 4 with details of the TPM Perl Package for Windows 2008 server - Corrected link error in section 10.11 as to the location of OM log files - OM server disk space capacity requirements (section 8.2) updated to include free space during normal operation and additional space required for database restore operation - Added OM server disk space requirements for OM DB restore operations to section 9.2.1.6 - Added section Error! Reference source not found. detailing the console output produced during the upgrade procedure
1.3	Updated cfMonInterval and clusterMonInterval recommended values for the Regional Reporting Master Server in very large environments from 'off' to '12 hours' is allow reporting of active/active cluster partner names when generating reports. See Table 26 -Very Large Environment Custom Settings for Regional Master Reporting Server for updated details
1.4	- Added link to HSE document for Virus Scan for OM in section 7.5 - Added section 9.2.7, ePO Policy on OM Server
1.5	Minor changes following peer review
1.6	- Removed draft document designation - Added section 9.2.7.3
2.0	- Initial draft for NetApp OnCommand Unified Manager 5.1 - Removed installation automation appendices
2.1	- Updated CATE website URL in Section 4 - Included link for HSE document for AV Scanning in Table 3 - Business Continuance License requirement removed - Added Section 11.7 NetApp OnCommand Unified Manager Console - Updated various diagrams to reflect UM instead of OM - Added draft document designation
2.2	- RFP - Added software requirement under Section 8.1.3
2.2.1	Updated link to dfm_cfg2.ps1 in Section 8.1.3
2.2.2	- Removed link for dfm_cfg2.ps1 in Section 8.1.3 - Updated package name to CITI_WIN_2K8_NetAppUnifiedMgr_51_01 in Section 8.1.3
2.2.3	Fixed hyperlink in Section 8.3.6 and Section 9.6
2.2.4	RFP (Updated) Release
2.2.5	RFC for OnCommand Unified Manager 5.2 – Peer Review

2.3	• RFT • Edited the hyperlink in Section 11.5
2.4	• Changed the link to the FID document in Section 10.6 • Removed the statement “This step does not apply if you are upgrading from an earlier version of UM on the same server” in Section 8.3.5 Configuring Web Server SSL Certificates due to an upgrade bug in the executable provided by NetApp • Added a condition “Upgrading from UM5.1 to UM5.2” in Section 10.3 Web server SSL configuration
2.5	• Updated document with NMC3.3 details • RFT
2.6	• RFP
2.7	• Updated document with OCUM 5.2R1P1 details • Peer review
2.8	• RFT
2.9	• RFP

13 Index

Backup		Procedure	16
OM Database	25–39	Citi Configuration	18
Abort backup	35	Default DB Security	18
Checking status	34	Device Discovery	18
Defining a scheduled backup	27	Install Path	17
Deleting backup files	37	License Keys	16
Listing backup files	36	Management Functional ID	19
Manual (on demand) backup	30	Performance	17
Restoring	32	SSL Configuration	18
Retention (backupRetentionPolicy) ...	25, 178	Maintenance	
OM Server	40–42	OM Database	
Backup Retention	40	Data Lifecycle	50
Justification	40	General Best Practice	52
Optimization	41	Graceful Shutdown	48
Restoration	41	Server Migration	53, 54
Server Backup	40	Monitoring	
Snapvault	120	OM Server	
Capacity Planning	56	Services	49
Command Aliasing	140	NetApp Management Console	128
Configuration Management	86	NMC	See NetApp Management Console
Configuration File		Operations Manager Server Settings	58, 173–86,
Automated Management	102	173–86, 173–86, 173–86, 173–86, 173–86,	
Copying and Renaming	90	173–86, 173–86, 173–86, 173–86, 173–86	
Creation	90	agentHostCIFSAccount	177
DB Backup	116	agentHostPort	177
Deletion	96	agentHostTransport	173
Import / Export	95	agentMonInterval	177
On Demand Device Comparison	99	aggrFullThreshold	177
On Demand Template Comparison	101	aggrFullThresholdInterval	177
Partial File	93	aggrNearlyFullThreshold	177
Upgrade/Revert	91	aggrNearlyOvercommittedThreshold	177
View Settings	92	aggrNearlyOverDeduplicatedThreshold	177
Filers and vFilers	89	aggrOvercommittedThreshold	177
Prerequisites	87	aggrOverDeduplicatedThreshold	177
Contingency		aggrSnapshotFullThreshold	177
OM Server	42–48	aggrSnapshotNearlyFullThreshold	177
3 site fan-in	46	alertFrom	178
Automation	47, 51	auditLogEnabled	170, 173
DR server normal state	43	auditLogForever	170, 173
Failing back	44	authUsePam	173
Invoking	43	autosupportAdminContact	178
Overview	42	autosupportContent	178
Testing	44	autosupportDestinationEmail	178
SnapMirror	117	autosupportDestinationURL	178
Disaster Recovery	See Contingency	autosupportEnabled	178
DR	See Contingency	autosupportIncludeAllDiagInfo	178
Events and Alarms	121	autosupportIncludePerf	178
Functional ID	19, 60, 87, 117, 140, 147, 164	autosupportIncludeProv	178
Installation	11–19	autosupportMonInterval	178
Getting Started	11	autosupportProtocol	174
Configuration Parameters	12	autosupportRetryCount	178
Hardware	11	autosupportRetryDelay	178
Network	11	backupRetentionCount	25, 27, 40, 178
Software	12	ccMonInterval	108, 179
Platform Requirements	13	cfMonInterval	179
VMware	15	chargebackDayOfMonth	179

<i>chargebackIncrement</i>	179	<i>perfAdvisorShowDiagCounters</i>	186
<i>chargebackRate</i>	179	<i>perfAdvisorTransport</i>	175
<i>cpuBusyThresholdInterval</i>	179	<i>perfAdvThreshViolationMonInterval</i>	186
<i>cpuMonInterval</i>	128, 179	<i>perfDataExportEnabled</i>	186
<i>cpuTooBusyThreshold</i>	179	<i>pingMonInterval</i>	182
<i>credCacheTTL</i>	174	<i>pingMonRetryDelay</i>	182
<i>currentEventsCacheSize</i>	179	<i>pingMonTimeout</i>	182
<i>databaseBackupDbengWaitTime</i>	180	<i>profileTTL</i>	175
<i>defReportLinesPerPage</i>	180	<i>qtreeFullThreshold</i>	182
<i>dfmDataExportEnabled</i>	180	<i>qtreeFullThresholdInterval</i>	182
<i>dfMonInterval</i>	180	<i>qtreeGrowthEventMinChangePct</i>	182
<i>discoverAgents</i>	174	<i>qtreeMonInterval</i>	182
<i>discoverEnabled</i>	174	<i>qtreeNearlyFullThreshold</i>	182
<i>discoverHostInitEnabled</i>	174	<i>recentMaxReports</i>	182
<i>discoverHosts</i>	174	<i>SANHostMonInterval</i>	183
<i>discoverInterval</i>	180	<i>SANHostMonSnapshotLUNs</i>	183
<i>discoverNetworks</i>	174	<i>serverHTTPEnabled</i>	175
<i>discoverTimeout</i>	174	<i>serverHTTPPort</i>	175
<i>discoverVfilers</i>	59, 174	<i>serverHTTPSEnabled</i>	175
<i>diskMonInterval</i>	180	<i>serverHTTPSPort</i>	175
<i>envMonInterval</i>	180	<i>SMTPServerBackup</i>	183
<i>eventsPurgeInterval</i>	50, 123, 180	<i>SMTPServerName</i>	183
<i>favoriteMaxReports</i>	180	<i>SMTPServerPort</i>	175
<i>filerConfigSaveLocalChanges</i>	180	<i>snapmirrorMonInterval</i>	183
<i>fsMonInterval</i>	180	<i>snapshotDiscoveryEventsEnabled</i>	183
<i>groupTreeShowStatus</i>	180	<i>snapshotMonInterval</i>	183
<i>growthRateSensitivity</i>	180	<i>snmpTrapListenerEnabled</i>	175
<i>guiRefreshInterval</i>	180	<i>snmpTrapListenerPort</i>	175
<i>hbaportTooBusyThreshold</i>	180	<i>snmpTrapRcvdMaxPerWindow</i>	175
<i>hostAdminPort</i>	174	<i>snmpTrapRcvdWindowSize</i>	175
<i>hostAdminTransport</i>	174	<i>statusMonInterval</i>	183
<i>hostEnclosureDiscoveryEvents</i>	180	<i>sysInfoMonInterval</i>	64, 183
<i>hostLoginProtocol</i>	174	<i>useHostsEquiv</i>	176
<i>hostPingMethod</i>	181	<i>userEnableAlerts</i>	183
<i>hostRBACMonInterval</i>	165, 181	<i>userFullThreshold</i>	183
<i>httpEnabled</i>	174	<i>userNearlyFullThreshold</i>	184
<i>httpPort</i>	174	<i>userQuotaMonInterval</i>	184
<i>httpsEnabled</i>	174	<i>vFilerMonInterval</i>	59, 184
<i>httpsPort</i>	175	<i>volFullThreshold</i>	184
<i>ifMonInterval</i>	181	<i>volFullThresholdInterval</i>	184
<i>ldapEnabled</i>	175	<i>volGrowthEventMinChangePct</i>	184
<i>licenseExpireWarningThreshold</i>	181	<i>volNearlyFullThreshold</i>	184
<i>licenseMonInterval</i>	181	<i>volNearlyNoFirstSnapThreshold</i>	184
<i>localHostName</i>	43, 44, 45, 46, 53, 55, 181	<i>volNearlyOvercommittedThreshold</i>	184
<i>lunMonInterval</i>	181	<i>volNearlyOverDeduplicatedThreshold</i>	184
<i>maxReportLinesPerPage</i>	181	<i>volNoFirstSnapThreshold</i>	184
<i>monMinFreeBytes</i>	181	<i>volOvercommittedThreshold</i>	185
<i>monMinFreePercent</i>	181	<i>volOverDeduplicatedThreshold</i>	185
<i>monSNMPRetries</i>	181	<i>volReserveDepletedThreshold</i>	185
<i>monSNMPTimeout</i>	181	<i>volReserveNearlyDepletedThreshold</i>	185
<i>ndmpDataUseAllInterfaces</i>	181	<i>volSnapshotCountThreshold</i>	185
<i>networkDiscoveryLimit</i>	175	<i>volSnapshotFullThreshold</i>	185
<i>nodesRemainingWarningThreshold</i>	181	<i>volSnapshotTooOldThreshold</i>	185
<i>opsMonInterval</i>	181	<i>Performance Advisor</i>	128–34
<i>perfAdvisorDisabledCounterGroups</i>	186	<i>VMware NFS NAS</i>	133–34
<i>perfAdvisorEnabled</i>	186	<i>Reporting</i>	66–85
<i>perfAdvisorMaxMonitorThreads</i>	186	<i>Asset management</i>	75
<i>perfAdvisorPollInterval</i>	186	<i>Change management</i>	71
<i>perfAdvisorShowAllViews</i>	186	<i>Custom reports</i>	78

Database structure	66	SnapMirror	See Contingency: SnapMirror
Problem management.....	76	Snapvault.....	See Backup: Snapvault
Scheduling	80	Storage Device Management	59
Storage capacity planning	67	Adding Devices	59
Qtrees	71	Deleting Devices	64
Thin provisioning.....	71	Modifying Devices	64
Script Plug-in Library.....	135		