



Technical Report

Operations Manager and Provisioning and Protection Manager Sizing Guide

Nagendra Krishnappa, Shiva Raja, NetApp
July 2010 | TR-3440

TABLE OF CONTENTS

1. PURPOSE	3
2. TERMINOLOGY	3
3. INTRODUCTION.....	3
4. PERFORMANCE EXPECTATIONS	4
5. SIZING FACTORS.....	4
6. TEST SUMMARY	5
6.1 DATAFABRIC MANAGER SERVER 4.0.....	5
6.2 DATAFABRIC MANAGER SERVER 3.8.....	13
7. DATABASE GROWTH	16
8. RECOMMENDED OPTIONS.....	17
9. CONCLUSIONS.....	18

1. PURPOSE

This document provides the necessary information to allow Operations Manager, Provisioning Manager, and Protection Manager administrators to choose the correct system for hosting the DataFabric® Manager server.

Note: The suggestions in this report are guidelines only; many environmental factors might influence the choice. This report discusses some of those factors.

2. TERMINOLOGY

Operations Manager, Performance Advisor, Protection Manager, and Provisioning Manager are software products that are delivered through software licensing to enable the capabilities of three installable components:

- DataFabric Manager Server and its repository (Sybase database) provide a central data storage and agent coordination to Operations Manager, Provisioning Manager, and Protection Manager. This component is installed on a server (Windows® and Linux®). Hence Operations Manager, Performance Advisor, Provisioning Manager, and Protection Manager have a common repository: DataFabric Manager.
- NetApp® Management Console provides a rich graphical user interface (GUI) for Performance Advisor, Protection Manager, and Provisioning Manager. This component can be installed on one or more workstations or PCs and is not required to be running for the server to continue operation. It is designed with easy and intuitive workflows/wizards for a simplified user experience.
- NetApp Host Agent is an agent that can be installed on one or more open systems hosts to provide the DataFabric Manager server with additional monitoring capabilities for SAN and FSRM features. Also, when it is installed on the same host as the Open Systems SnapVault® agent, It provides management features for Protection Manager.

Operations Manager is a client server application that delivers comprehensive monitoring and management of NetApp storage systems from a single console with alerts, reports, performance tools, and configuration tools to keep your storage infrastructure in line with business requirements for maximum availability and efficiency.

Performance Advisor is licensed as part of Operations Manager. It offers a rich GUI in the NetApp Management Console that delivers a single pane of glass for comprehensive performance management of NetApp storage systems and MultiStore®.

Provisioning Manager is a separate license that is added to the DataFabric Manager server to enable policy-based provisioning of NAS, SAN, or secondary storage. It also uses NetApp Management Console to provide a rich GUI, and the DataFabric Manager server to enable monitoring, alerting, and active management of the storage entities.

Protection Manager is a separate license that is added to the DataFabric Manager server. This license enables Protection Manager DR and Provisioning Manager functionality. Protection Manager revolutionizes data protection setup by automating discovery, provisioning, and reporting for data protection operations based on Snapshot™ through NetApp Management Console. It can configure and control all Snapshot, SnapMirror®, SnapVault, and Open Systems SnapVault (OSSV) operations.

NetApp Management Console (NMC) is a Java application that can be downloaded from the Setup menu of the Operations Manager GUI. It houses the Performance Advisor, Provisioning Manager, and Protection Manager UI.

3. INTRODUCTION

In growing NetApp environments, it's important to make sure that as the infrastructure grows, the DataFabric Manager server can scale to meet the increased needs. Also, the NetApp Management Console has been

introduced to extend the features from solely monitoring an environment to actively managing it as well. When more management features are used, more resources are needed to support those activities. Specific resource requirements are outlined in the installation guide for each release of Operations Manager (using DataFabric Manager server). These are minimum configurations. How many appliances; volumes; relationships between SnapMirror, SnapVault, and Open Systems SnapVault; and so on can be added determined by the size of the configuration. To determine the impact of monitoring and managing all these objects, this report looks at some of the factors that affect product performance.

4. PERFORMANCE EXPECTATIONS

FACTORS AFFECTING DATAFABRIC MANAGER SERVER PERFORMANCE

Memory, CPU, and disk are the top factors affecting server performance.

HOW DOES DATAFABRIC MANAGER SERVER USE ITS RESOURCES?

- **Memory:** The DataFabric Manager server is a 32-bit application. With large databases this means that the database is not always loaded 100% in memory. 4GB to 8GB RAM should yield acceptable results.
- **LUN:** NetApp recommends installing the database onto NetApp LUNs for better performance, to reduce the RTO to less than an hour, and to give the flexibility to create more frequent Snapshot copies of the database.
Note: CIFS/NFS database store is not supported.
- **CPU:** Multiple core CPUs are standard in systems today, and DataFabric Manager Server can use many cores or CPUs. In particular, the Operations Manager version 3.8 Sybase database is limited to two cores on Linux and up to four cores on Windows. The Sybase database is limited to using two CPUs or cores, but the rest of the server's processes can use additional CPU resources. Much of the data that Operations Manager presents in its GUI views and reports (or using the server CLI) is calculated dynamically; therefore the biggest bottleneck to performance can be CPU. Because most of the performance of the server depends on the database, adding more CPUs might not necessarily improve performance of systems where database processing is dominant because of the database core limit. But, with Operations Manager 4.0, the limitation on CPU cores has been removed, and it is now bundled with new Sybase license that allows DB to use unlimited number of CPU cores by default. This is an attempt to remove bottlenecks that might arise due to CPU performance.
- **Disk:** The DataFabric Manager server can be operated off of local RAID (striped for performance) disk (although NetApp recommends using SAN storage). In most cases, the CPU is the bottleneck before the disk. However, if the database is sluggish and the CPU is not overwhelmed, then disk I/O is the next consideration for troubleshooting. As with any database application, spindles are your friends.
- **Network:** Today's networks are fast enough to carry out the tasks relegated to them by Operations Manager, Provisioning Manager, and Protection Manager. However, if your network is extremely busy and you use quality-of-service devices, then consideration (in the QoS device) of monitoring and management services is necessary.

5. SIZING FACTORS

- **Number of storage systems monitored:** This is the total number of FAS storage systems and NearStore® appliances being monitored by the DataFabric Manager server.
- **Monitored objects:** These are the objects that Operations Manager constantly manages and monitors: qtrees, volumes, aggregates, vFiler™ units, disks, quotas, and so on. All these add to the total count of objects that Operations Manager can manage and monitor. The more the number of objects, the more the data that needs to be collected. For example, quotas: the largest quantity we know of is ~100,000 quotas being monitored, which does complete during the one-day default interval for monitoring these relationships. If you reach this number, you should consider lengthening the interval beyond the one-day default.
- **Monitoring intervals:** NetApp recommends leaving all monitoring intervals at their factory defaults, at least initially. Shortening these intervals adversely affects the server performance. Conversely,

lengthening the intervals improves server performance. However, these measures should be taken only in extreme cases.

- **SnapVault and SnapMirror relationships:** Number of relationships should be considered when monitoring from Operations Manager or configuring and managing from Protection Manager. Each relationship has to be constantly managed and monitored for Snapshot schedules, data transfer schedules, backup jobs, lags, and conformance.
- **NetApp Host Agents:** NetApp Host Agents pose little burden on the DataFabric Manager server when used for host data monitoring purposes, such as SAN implementations. However, when they are used for the optional FSRM data classification, a load is imposed on both the storage systems and the DataFabric Manager server, and therefore it is recommended that FSRM jobs be run consecutively during nonproduction hours. For this reason it is not factored into server sizing.
- **Checkpoint frequency:** The database in the DataFabric Manager server has an internal checkpoint process to make sure that it is accurate and keeping up with monitoring intervals. This checkpointing frequency increases with higher utilization of the DataFabric Manager services. That is, if the services and server are all very busy and creating an extreme load, then the checkpointing increases as well, producing even more strain on resources. This is necessary for data integrity purposes. However, if you are concerned about performance on the server and you see an increase in checkpoint frequency, then your server might be becoming overburdened.

To see how much memory and CPU each of the DataFabric Manager services is consuming over time, refer to the `dfmwatchdog.log` file in the log subdirectory of the server installation.

6. TEST SUMMARY

6.1 DATAFABRIC MANAGER SERVER 4.0

Extensive testing was performed in order to determine the optimum point to which an Operations Manager server could be loaded. The scope was to:

- Validate performance of Operations Manager, Protection Manager, and Provisioning Manager enabled on a single DataFabric Manager Server (without Performance Advisor enabled).
- Validate performance of Operations Manager along with Performance Advisor in scaled setup enabled on a single DataFabric Manager Server (without Protection Manager and Provisioning Manager enabled).

These tests provide a basic idea on the right number of storage systems or total number of objects that Operations Manager can handle with optimum performance without any failures. Also, a series of tests were performed and response times were gathered for the UI screens. An ample set of samples was considered to arrive at the averaged response times.

The following two sections document the test results of the above-mentioned validation exercise.

6.1.1. SETUP WITH OPERATIONS MANAGER, PROTECTION MANAGER, AND PROVISIONING MANAGER ENABLED (PERFORMANCE ADVISOR DISABLED)

One of the key requirements for a customer is to be able to install the entire DataFabric Manager stack on a single server. But, the concern or the question is will the server be able to handle or scale when all licenses are turned on: Operations Manager, Performance Advisor, Protection Manager, and Provisioning Manager. To help find an answer as to what is the maximum with which a DataFabric Manager server can be loaded, a series of tests was conducted to determine the optimum point. Tests revealed that, the server could not scale when Performance Advisor was also enabled. Hence, tests were performed with Performance Advisor disabled, which led to the following tested and proved numbers (to which a server with all licenses enabled can scale to without Performance Advisor). The following tables summarize the setup and the maximum object counts that a DataFabric Manager server can handle or scale.

Table 1) DataFabric Manager Server Setup

Operating System	RHEL AS release 4 (Nahant Update 8) 2.6.9-89.ELsmp x86_64
Hardware Configuration	IBM Blade Server HS21 based on 2 Intel® Dual Core 2.83 GHz CPU (Xeon, 5160) and 16GB
DataFabric Manager Database	Hosted on NetApp FC LUN

Table 2) Managed Setup

Objects	Count	Remarks
Storage Systems	41	Monitor 41 Data ONTAP® systems
Aggregates	150	~ 4 aggregates per node
Volumes	1,650	
Qtrees	5,800	
LUNs	2,300	
vFiler Units	80	2 vFiler units per node
Disks	480	
Host Users	180	
Operations Manager		
Host Agents	10	OSSV/SRM Hosts
Resource Groups	150	
Alarms		Configured for all Critical, Worse and Emergency events
SRM File Walk	1.5 million files	On 10 SRM Paths
Performance Advisor	Disabled	
Monitor Interval	Default	
DataFabric Manager Global Options	Default	
Events	~1,384,500	
Data Management		
Total Data sets	120	
SAN Data set (Mirror Policy)	16	Mirror transfers every 10 mins between 09:00 to 20:00
NAS Data set (Mirror Policy)	16	Mirror transfers every 10 mins between 08:00 to 19:00
SAN Data set (Backup Policy)	36	Backup once daily between 20:00 and 02:00
NAS Data set (Backup Policy)	36	Backup once daily between 20:00 and 02:00
Open System Data set	16	8 OSSV Hosts. 50 Directories/Host Backup once daily between 03:00 and 04:00
DP managed SnapVault relationships	2,120	Based on planned Backup Data sets

DP managed SnapMirror relationships	230	Based on planned Mirror Data sets
OSSV relationships	400	Based on planned OSSV Data sets
General		
DataFabric Manager Server Licenses	Operations Manager Protection Manager Provisioning Manager Disaster Recovery FSRM	
Clients	5	3 NMC and 2 Web UI will be continuously accessing the DataFabric Manager Server

The above-mentioned object count is per DataFabric Manager server instance. If any of the above-mentioned objects are more in your currently managed environment, add additional DataFabric Manager server instances. For example, if you have more than 41 storage systems with object counts more than shown in the above table, you will need another DataFabric Manager server instance to manage the same. For example, if you have 120 storage systems with high object counts compared to above table, install 3 DataFabric Manager server instances and configure each one to manage 41 storage systems, each with object counts approximately equal to or less than what is shown as manageable above. The same holds true for the number of data sets and relationships too. If you have more than 120 data sets, then add another DataFabric Manager server instance to manage the additional data sets and relationships.

The flip side to this discussion would be: what if my data sets/relationships are much less but have more storage systems? For example, if you have 800 relationships with 200 data sets and 50 storage systems, you should be able to manage with a single DataFabric Manager server instance.

PERFORMANCE CHECK RESULTS

The below table captures the response time for the frequently used Operations Manager, Provisioning Manager, and Protection Manager workflows using GUI.

Table 3) Operations Manager

Report Names	Response Time (Sec)
Summary	7.93
Appliance	2.13
Events	5.89
File System	2.73
Aggregate	2.44
Volumes	2.76
Qtree	2.44

Table 4) Provisioning Manager and Protection Manager

Report Names	Response Time (Sec)
Loading of Provisioning Dashboards page	3.61
Load Protection Manager Dashboards page	4
Loading of Data sets Overview page	2.41
Measure the time taken to load Protection Policies	2.32
Measure the time taken to load Resource pool page	4.22
Loading of Hosts page	2.21
Loading of vFiler page	3.1
Measure the time taken to load OSSV Hosts page. Note the number of OSSV hosts	2.55
Loading the members of a Data set	3.12
Loading of Events page with at least 1,000 provisioning events	4.75
Loading of Jobs page with at least 200 provisioning jobs	5.99
Measure time required to generate the Operations Manager canned reports below. Mirror Data Transfer Weekly Report.	2.39
Measure time required to generate the Operations Manager canned reports below. Backup Data.	2.99

6.1.2 SETUP WITH OPERATIONS MANAGER AND PERFORMANCE ADVISOR ENABLED (PROTECTION MANAGER AND PROVISIONING MANAGER DISABLED)

This setup was used to determine the optimum point to which a DataFabric Manager can be used with Operations Manager and Performance Advisor. Several tests were performed to arrive at the below-mentioned numbers.

Table 5) DataFabric Manager Server Setup

Operating System	Microsoft® Windows 2003 Enterprise Edition SP2 (Build 3790) x64 base
Hardware Configuration	IBM Blade Server HS21 based on 2 Intel Dual Core 2.83 GHz CPU (Xeon, 5160) and 8GB
DataFabric Manager Database	Hosted on NetApp FC LUN

Table 6) Managed Setup

Objects	Count	Remarks
Storage Systems	250	250 Data ONTAP Systems (versions 7.2.5.1 or 7.3.1)
Aggregates	1,749	7 aggregates per node
Volumes	9,490	~ 35 volumes per node
Qtrees	50,400	~ 200 qtrees per node
LUNs	9,600	~ 40 LUNs per node
vFiler Units	500	2 vFiler units per node
Disks	14,000	56 disks per node
Host Users	2,000	
Snapshot copies	47,400	5 Snapshot copies per volume
Host Agents	20	
Resource Groups	250	
Alarms	20	

SRM File Walk	15 million files	On 10 SRM paths
Performance Thresholds	12	lun:avg_latency volume:san_latency system:cpu_busy processor:processor_busy system:avg_processor_busy nfsv3:nfsv3_read_latency nfsv3:nfsv3_write_latency nfsv3:nfsv3_ops system:load_total_mbps volume:avg_latency volume:write_latency volume:read_latency
DataFabric Manager Server Licenses	Core FSRM	All Global options set to default perfDataExportEnabled = Yes
Performance Advisor	Enabled	Performance data and QTree monitoring enabled on all Data ONTAP 7-Mode storage systems
Events	~ 580,800	

The above-mentioned object count is per DataFabric Manager server instance. If you have more than 250 storage systems, you need to install another DataFabric Manager server instance to manage the same. For example, if you have 300 storage systems, then install 2 DataFabric Manager server instances and configure each one to manage around 150 storage systems each. Now let's consider a different case where the number of storage systems is 300, but other object counts like aggregates, volumes, disks, and so on are very less compared to the above table. In such cases, one instance of DataFabric Manager should be able to scale. To summarize, if total object counts are less, then one server should scale; if more than what is tested as per above table, add another instance of DataFabric Manager to share the load.

PERFORMANCE CHECK RESULTS

Along with testing how much Operations Manager or Performance Advisor can scale, certain response times for the Web UI as well as NMC were recorded. These are reports when navigated from the Global Summary page. The response times are as shown below.

Table 7) Performance results

Report Names	Response Time (Sec)
Summary	7.58
Appliance	1.15
Events	11.47
File System	2.73
Aggregate	1.47
Volumes	1.66
Qtree	3.2
Performance Advisor	
Dashboard page	19

Hierarchy with logical objects	8.48
Hierarchy with physical objects	8.89

6.1.3 SETUP WITH PROTECTION MANAGER ENABLED (PROVISIONING MANAGER AND OPERATIONS MANAGER ENABLED. PERFORMANCE ADVISOR DISABLED)

The main focus of these tests was to measure the ability of Protection Manager to drive backups. With that in mind, the test script used the following steps.

- Installed DataFabric Manager and added primary and secondary storage systems. Initially we added only a handful of storage systems in DataFabric Manager so we could focus on the performance of Protection Manager itself.
- Created primary volumes and qtrees on the specified primary aggregates. The number of volumes and qtrees depended on whether medium or large load configuration was being tested.
- Selected one of the data protection policies for testing and created data sets with that policy.
- Created a single resource pool and added all the secondary aggregates to it. Attached the resource pool to secondary nodes of all the data sets.
- Added primary volumes to data sets. Volumes were equally distributed among the data sets. The data sets were in suspended mode.
- Started conforming five data sets at a time. This is done by first resuming five data sets and then resuming more as any of the data sets reach conformance.
- Waited for all the data sets to conform.
- Populated primaries with data using the Data ONTAP mkfile command. The amount of data depended on the whether a medium or large load was being tested.
- Started backups for all the data sets simultaneously by using the Protection Manager scheduler.
- Waited for the backup jobs to finish and measured data throughput of backups

The test configurations used was to test the scalability of Protection Manager alone. The primary use of this server is for Protection Manager only. The server also has Operations Manager and Provisioning Manager licenses enabled, but toned down for discovery and monitoring purposes only.

CONFIGURATION 1

Table 8) DataFabric Manager Server setup

Operating System	Redhat Linux AS 4
Hardware Configuration	2 Intel quad Core Xeon 5440 at 2.83 GHz CPU with 16GB RAM
DataFabric Manager Database	Hosted on NetApp FC LUN

Table 9) Managed Setup

	Number of Data Sets	Number of Primary Storage Systems	Number of Secondary Storage Systems	Number of Primary Volumes	Number of Primary Qtrees	Protocol
Load	600	4	2	1,200	----	Volume SnapMirror
Load	400	4	2	800	1,600	SnapVault

The above configuration was tested with the following policies:

- Mirror Policy
- Backup policy with qtree SnapMirror
- Backup policy with SnapVault

CONFIGURATION 2: STRESS TEST

Table 10) DataFabric Manager Server Hardware Configuration

Operating System	Redhat Linux ES 5
Hardware Configuration	2 Intel quad Core Xeon 5440 at 2.83 GHz CPU, with 16GB RAM
DataFabric Manager Database	Hosted on NetApp FC LUN

Table 11) Managed Setup

	Number of Data Sets	Number of Primary Storage Systems	Number of Secondary Storage Systems	Number of Primary Volumes	Number of Primary Qtrees	Protocol
Load	500	4	1 or 2	5,000	----	SnapVault

The above configuration was tested with the following policy:

- Mirror Policy

Table 12) NMC (NetApp management console) hardware and software configuration

Operating System	Windows XP SP 2
Hardware Configuration	Pentium® 4 3.2Ghz, Pentium 4 3.2Ghz

CONFIGURATION

- Events table refresh was testing with 25,000 rows in the events table.
- Compound data source load test was tested the time of the first set of data returned from the server of loading 1,274 resource members.
- Volume list test was testing the start JAPI of volume listing against 497 volumes.
- Load backups were loading 78 backups.
- Load files were loading 4 files within one volume.
- Data set loading was testing the JAPI of data set listing 103 data sets.

PROTECTION MANAGER TEST RESULTS

Protection Manager 4.0 has proven to be far more scalable and responsive than its predecessor (3.8). We have increased (compared to the previous test-bed - only 400 data sets with 4,000 relationships) the number of data sets (600) and relationships (5,000) without compromising the DataFabric Manager server performance.

Table 13) Protection Manager Test Results

DataFabric Manager	Protection Manager 4.0 Conformance	
	Data Sets in HH:MM	Data Set per Minute
VSM	600 in 17:39	1.76
SnapVault	400 in 06:36	0.99

GUI TEST RESULTS (NMC)

The performance of NMC client when compared to its predecessor is more or less the same. As mentioned before, the performance of NMC was improved significantly in 3.8, and it is still maintained in 4.0. Below table records the time taken in seconds for commonly used screens.

Table 14) NMC Test results

Report Name	Protection Manager 4.0 (Seconds)
Events table refresh	1
Compound Data-source load	0.6
Volume List (API)	1
Enable Restore Button	1.8
Load Backups	1.7
Load Files	0.5
Data Set Load	1

6.1.4 SETUP WITH SELECTIVE LICENSES ENABLED

This section attempts to provide an idea of how to size environments that do not match the above provided configurations and managed setups. There can be two such scenarios:

1. Object counts less than tested: All environments having object counts less than the managed setup table that is provided in sections 6.2.1 and 6.2.2 can easily use a single instance of DataFabric Manager to manage their environments smoothly. For example, an environment with around 2,000 volumes, 1,000 aggregates, and 9,000 qtrees, which are much less than tested setup, a single server with Operations Manager and Performance Advisor can be used to manage the entire environment though there are 300 storage systems.
2. Object counts more than tested: These are environments with object counts beyond the counts provided in sections 6.2.1 and 6.2.2. As an example, let's consider an environment with 80 storage systems and the need is to enable only Operations Manager, Performance Advisor, and Protection Manager. For Protection Manager consider around 3,000 relationships in total. This kind of an environment will need two instances of DataFabric Manager server installed as shown below:
 - i. First instance can have only Operations Manager and Performance Advisor enabled: this server can be used only for the purpose of alerting, monitoring, reporting, and performance management.

- ii. Second instance can then be primarily made available for Protection Manager and Provisioning Manager: this server can be made exclusively available for data protection. Make sure that Operations Manager is configured only to discover/monitor storage systems: no active management. Disable Performance Advisor completely for better performance (refer the recommended options section).

If observed, what actually matters are the object counts that need to be considered for any sizing. All objects managed by Operations Manager and data sets/relationship counts managed by Protection Manager.

6.1.5 SETUP WITH CLUSTERED DATA ONTAP (DATA ONTAP 8)

Operations Manager 4.0 is the first enterprise-class tool that is available for managing Data ONTAP 8: both 7-Mode and Cluster-Mode. Below tables give what hardware/OS was used and how many objects Operations Manager was able to manage. This setup was tested with all licenses disabled other than Operations Manager.

Please check Data ONTAP 8.0 documentation for scalability of 8.0 storage systems. At the time of this release, we recommend to use up to 4 Data ONTAP 8.0 cluster systems in your environment.

Table 15) DataFabric Manager Server Setup

Operating System	Microsoft Windows 2003 Enterprise Edition SP2 (Build 3790) x64 base
Hardware Configuration	IBM Blade Server HS21 based on 2 Intel Dual Core 2.83 GHz CPU (Xeon, 5160) and 8GB
DataFabric Manager Database	Hosted on NetApp FC LUN

Table 16) Managed Setup

Objects	Count
Controllers	4 (FAS3070s)
Vservers	42
Aggregates	18
Volumes	1,752
Ports	16
Logical Interfaces	49
Snapshot copies	17,350

6.2 DATAFABRIC MANAGER SERVER 3.8

As per the scalability results of DataFabric Manager 3.7, there were no issues with respect to the response time of GUI in Operations Manager. But, the response times were quite longer for NMC, especially with respect to Protection Manager. Hence, it was decided to focus the scalability testing more toward the response time of NMC.

6.2.1 SETUP WITH OPERATIONS MANAGER AND PERFORMANCE ADVISOR ENABLED

Table 17) DataFabric Manager Server hardware Setup

Operating System	Windows 2003 Server Enterprise Edition VM hosted on a VMware® ESX server
Hardware Configuration	IBM Blade Server HS21 based on two Intel Dual Core 2.83 GHz CPU (Xeon, 5160) with 8GB of RAM

DataFabric Manager Database	DataFabric Manager DB on 400GB NetApp FC LUN
Web Browser Client	Internet Explorer 6 SP 2

Table 18) Managed Setup

Objects	Count
Storage systems	250
Volumes	8,500
Qtrees	100,000
LUNs	10,000
Quotas	130,000
Aggregates	1,500
Disks	11,000
Users	1,600
vFiler units	500
Operations Manager	Count
Resource group	250
Host agents	20
Alarms	60
Jobs	12
SRM File walks	15 Million files
Snapshot copies	60,000
Performance Advisor	Enabled
Performance Thresholds	60
Monitoring interval and Options	Default

PERFORMANCE CHECK RESULTS

The response time of Operations Manager reports should be less than 15 seconds.

OPERATIONS MANAGER

The response time of Operations Manager Web UI in DataFabric Manager Server 3.8 is slower (but negligible) when compared to DataFabric Manager 3.7.

Table 19) Response time

Report Names	Response Time (Sec)
Summary	4.4
Appliance	3.3
Events	14.7
File System	6.8
Aggregate	2.8
Storage systems	2.5
Volumes	2.5
Qtree	6

6.2.2 SETUP WITH PROTECTION MANAGER ENABLED

The main focus of these tests was to measure the ability of Protection Manager to drive backups. With that in mind, the test script used the steps as described in Section 6.1.3.

CONFIGURATION 1

Table 20) DataFabric Manager Server Setup

Operating System	Redhat Linux ES 5
Hardware Configuration	2 Intel quad Core Xeon 5440 at 2.83 GHz CPU with 16GB RAM
DataFabric Manager Database	DataFabric Manager DB on NetApp FC LUN

Table 21) Managed Setup

	Number of Data Sets	Number of Primary Storage Systems	Number of Secondary Storage Systems	Number of Primary Volumes	Number of Primary Qtrees	Data Transfer Size
Load	400	4	1 or 2	800	4,000	1200GB

The above configuration was tested with the following policies:

- Backup policy with SnapVault
- Backup policy with qtree SnapMirror
- Mirror policy

CONFIGURATION 2

Table 22) DataFabric Manager Server Setup

Operating System	Redhat Linux ES 5
Hardware Configuration	1 Intel quad Core Xeon 5440 at 2.83 GHz with 16GB of RAM
DataFabric Manager Database	DataFabric Manager DB on NetApp FC LUN

Table 23) Managed setup

	Number of Data Sets	Number of Primary Storage Systems	Number of Secondary Storage Systems	Number of OSSV Relationships	Data Transfer Size
Load	50	400 to 1,000	1	400 to 2,000	2MB to 20MB

The above configuration was tested with the following policy:

- "Remote backups only" policy for Open Systems SnapVault

NETAPP MANAGEMENT CONSOLE (NMC) HARDWARE AND SOFTWARE CONFIGURATION

HARDWARE CONFIGURATION

- Virtual Machine
- Xeon 3.20GHz
- 1GB RAM

OPERATING SYSTEM

Windows XP SP 2

PROTECTION MANAGER TEST RESULTS

Protection Manager 3.8 has proven to be far more scalable and responsive than its predecessor (3.7). We have doubled (compared to the previous test-bed - only 200 data sets) the number of data sets (400) and relationships (4,000) without sacrificing the DataFabric Manager server performance.

GUI TEST RESULTS (NMC)

- The performance on the Events page has significantly improved. It only takes 2 – 3 seconds to populate 6,000 plus events.
- The performance of most of the overview pages in NMC have improved. Previously, when a new event came in, the entire table of data would need to refresh; now only the table row associated with the event is refreshed. This significantly improves performance and memory usage, especially on virtual machines.
- Performance improved:
 - The time taken to enable the restore button from the data set overview page
 - Listing the backups in the restore wizard
 - Listing the files contained in a backup in the restore wizard. It now only takes a couple of seconds to populate 200 plus files.

7. DATABASE GROWTH

Space in the DataFabric Manager database is dominated by history tables. A rough breakdown is that 75% of the file is history tables and about 25% is free pages and overhead. The remaining tables, including qtree and quota histories, tend to only use a small percentage of the file.

From a regression analysis of history table contents, the following formula estimates the space consumed as a function of the number of hosts, aggregates, and volumes:

$$(\# \text{ hosts} * 2.25\text{MB}) + (\# \text{ vols} * 98\text{KB}) - (\# \text{ aggrs} * 270\text{KB}) = .\text{bytes/table}$$

For example, suppose that a customer has 250 hosts, 1,587 volumes, and 1,587 aggregates. The customer expects that the number of hosts will grow to 350. To accommodate this growth, the customer should plan for 2,200 volumes and 2,200 aggregates:

$$(350 * 2.25\text{MB}) + (2200 * 98\text{KB}) - (2200 * 270\text{KB}) = .409\text{Mbytes/table}$$

DataFabric Manager keeps two each of daily, weekly, monthly, and quarterly tables, plus an open-ended number of yearly tables. At the end of the year, there are nine filled history tables, and DataFabric Manager adds a new yearly table each year. If we add the 25% overhead (free pages and internal bookkeeping), each table uses about 509MB. Nine tables mean 4.58GB, plus a new 509MB yearly table each year.

The bottom line is that after a year, the DataFabric Manager database will be about 4.6GB, and it will grow by about 509MB/year.

If Performance Manager is used, additional space is consumed by the DataFabric Manager server. This space is not directly part of the DataFabric Manager database but is stored in a special directory.

NetApp recommends storing the database on a NetApp LUN and creating more frequent Snapshot copies, so that RTO is reduced to less than one hour. As the database grows with the use of Protection Manager and Provisioning Manager, the volumes can be easily resized, or thin provisioning can be used.

PERFORMANCE ADVISOR DATA REPOSITORY GROWTH

From the moment the DataFabric Manager server is started and a storage system is enabled for performance monitoring, it begins to collect data for the predefined views configured in Performance Advisor, such as the views used in its dashboard. This data collection for Performance Advisor is stored on the DataFabric Manager server but not in the database. Performance Advisor keeps a separate accounting of the data it collects.

For each counter that is part of a view, one set of data is collected. This collection of data can range anywhere from one-minute intervals (the default is one minute) down to real-time data collection. Each instance of Performance Advisor opened on desktop queries the data for a view from the same “view repository” as the others. This scheme reduces the overall amount of storage and resources necessary. For some views, the data is recycled every seven days. For other views, such as CPUs, the data is kept for a year and then recycled. For real-time views, the data is recycled each time the real-time view is disabled.

Amount of disk space needed per monitored system for Performance Advisor:

From the scaled configuration from section 6.3, the object count per host is derived as follows:

Disks per controller: 51
Aggregates per controller: 4
Volumes per controller: 33
LUNs per controller: 31
vFiler units per controller: 2

The following object counts are assumed:

Number of network interfaces per controller: 2

Number of FCP targets per controller: 1

Number of processors per controller: 2

With the above setup, the amount of space required to monitor one Data ONTAP 7.2.3 controller is 529MB per year.

For monitoring N hosts the space requirements would be $(529 \times N)$ MB of disk space.

Other assumptions:

- There are no custom views.
- Qtree basic counter group is disabled.
- DataFabric Manager version is 3.7.
- Percentage of volumes with priority queues: 10%.

8. RECOMMENDED OPTIONS

NetApp recommends the following options when using Protection Manager:

- `dpmaxActiveDataTransfers=x` for all storage systems (use `dfm host set`)

The value of “x” is the value that depends on Data ONTAP version, replication protocol, and platform. The correct value should always be less than or equal to the maximum supported value for the factors above. The 7.3 releases have a table on the NOW™ site at http://now.netapp.com/NOW/knowledge/docs/ontap/rel731/html/ontap/onlinebk/protecting/reference/r_oc_prot_max-simultaneous-replication-ops.html

There is an NDMP stream count limit of 40 streams for Data ONTAP 7.2.4 and earlier. The stream count is 128 for Data ONTAP 7.3 and increases for newer Data ONTAP releases. Depending on the Data ONTAP version, you can increase the number.

- `dpScheduledJobExpiration=12h`

Suppose that 150 jobs are started in Protection Manager, and that 100 of those jobs finish in 12 hours. The remaining 50 jobs are dropped, and an error is logged in the `dfmscheduler.log` file. There is no other record of those dropped jobs. No retry will be done after 12 hours. So depending on number of jobs you might want to set this value high for retries.

To enable scalable architecture, it is recommended to run Operations Manager and Performance Advisor on one server and install Protection Manager and Provisioning Manager on a second server. On the second server, disable Performance Advisor and do not perform any activities using OM. Optionally, on the second server, the following options can be disabled in Operations Manager through the command-line interface.

- `agentMonInterval`
- `ccMonInterval`
- `envMonInterval`
- `fcMonInterval`
- `hostRBACMonInterval`
- `opsMonInterval`
- `SANHostMonInterval`
- `srmMonInterval`
- `userQuotaMonInterval`

Effectively reduce considerable load on the second server by only discovering storage systems and MultiStore instances that are involved only with respect to data protection.

RECOMMENDED CONFIGURATION FOR A LARGE INFRASTRUCTURE (100 + STORAGE SYSTEMS):

HARDWARE CONFIGURATION

- 2 Intel quad Core Xeon 5440 at 2.83 GHz CPU
- DataFabric Manager server Database 10GB of Free Space
 - 20GB free space
 - DataFabric Manager server database to reside on NetApp FC LUN
- Minimum 8GB RAM

SOFTWARE CONFIGURATION

Windows 2000/2003/2008, RedHat Linux ES, SUSE

HARDWARE CONFIGURATION FOR NMC

- Xeon 3.20GHz
- 1GB RAM

OPERATING SYSTEM FOR NMC

Windows XP SP 2

For software requirements, see the compatibility matrix on NOW (NetApp on the Web) at http://now.netapp.com/NOW/knowledge/docs/olio/guides/dfm_compatibility.

9. CONCLUSIONS

Results might be different with earlier or later releases of DataFabric Manager. When new releases substantially alter these results, this document will be updated accordingly. The tests described here show a slight constant increase in the access times (negligible) in the Web UI (file systems, volumes, and disk report

pages) of Operations Manager using DataFabric Manager Server 3.8, when compared with DataFabric Manager Server 3.7. However, the response time of NMC in DataFabric Manager Server 3.8 and 4.0 has improved significantly.

In addition to this, with DataFabric Manager 4.0, following was observed:

- It is recommended to disable PA when Operations Manager, Protection Manager, and Provisioning Manager are enabled on a single server for a scaled environment as described in section 6.2.1.
- Linux platform provided better performance when all licenses were enabled.
- For a scaled setup as in section 6.2.1, number of data sets needs to be restricted to 120 per DataFabric Manager Server instance.

NetApp provides no representations or warranties regarding the accuracy, reliability or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.



© 2010 NetApp, Inc. All rights reserved. Specifications are subject to change without notice. NetApp, the NetApp logo, Go further, faster, DataFabric, Data ONTAP, MultiStore, NearStore, NOW, SnapMirror, Snapshot, SnapVault, and vFiler are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. Intel and Pentium are registered trademarks and Xeon is a trademark of Intel Corporation. Linux is a registered trademark of Linus Torvalds. VMware is a registered trademark of VMware, Inc. Microsoft and Windows are registered trademarks of Microsoft Corporation. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. TR-3440