

Managing Unified Storage with IBM System Storage N series Operation Manager

Learning about N series Provisioning
Manager

Using N series Performance
Advisor

Using N series Protection
Manager to protect your data



Alex Osuna
George Lane
Naren Rajasingam
Sudheer N Shivakumar

Redbooks



International Technical Support Organization

**Managing Unified Storage with IBM System Storage N
series Operation Manager**

August 2009

Archived

Note: Before using this information and the product it supports, read the information in “Notices” on page vii.

Archived

First Edition (August 2009)

This edition applies to DataFabric Manager Version 3.7.1 and Data ONTAP Version 7.3

© Copyright International Business Machines Corporation 2009. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Notices	vii
Trademarks	viii
Preface	ix
The team that wrote this book	ix
Become a published author	x
Comments welcome	xi
Part 1. Introduction and Operations Manager core installation	1
Chapter 1. Introduction	3
1.1 Operations Manager UI	9
1.1.1 Features provided by Operations Manager UI	9
1.1.2 Tabs provided by Operations Manager UI	10
1.1.3 IBM N series Management Console	11
1.1.4 DataFabric Manager Server	11
1.2 Host Agent	14
1.3 Protection Manager	16
1.3.1 Provisioning Manager overview	17
1.3.2 Provisioning Manager features	17
1.4 What is FSRM	18
1.5 Naming conventions	19
Chapter 2. Installing Operations Manager: Windows 2003 32-bit operating system ..	21
2.1 Host prerequisites	22
2.2 License requirements	22
2.3 N series prerequisites	24
2.3.1 Data ONTAP requirements	24
2.4 Installing Operations Manager	25
2.4.1 Deploying Operations Manager software	25
Chapter 3. Installing Operations Manager: Windows 2003 64-bit operating system ..	37
3.1 Host prerequisites	38
3.2 N series prerequisites	38
3.3 Plug-ins	41
3.4 Installation	45
Chapter 4. Installing Operations Manager: Linux	55
4.1 Overview	56
4.2 Host prerequisites	56
4.2.1 Upgrading from DataFabric Manager V3.5.1 or earlier	57
4.2.2 License requirements	58
4.3 N series prerequisites	59
4.4 Installation	59
4.4.1 Locating the installation files	60
4.4.2 Installation options	60
4.4.3 Installation steps	60
4.4.4 Switching off Autosupport notifications for Operations Manager	75
Part 2. Host Agent installation	77

Chapter 5. Host Agent installation for Windows 2003	79
5.1 Host prerequisites	80
5.2 What Operations Manager Host Agent can do	80
5.3 N series prerequisites	81
5.4 Overview of configuration steps	82
5.5 N series installation	83
Chapter 6. Host Agent installation for Linux	91
6.1 Host prerequisites	92
6.2 DFM Server prerequisites	92
6.3 N series prerequisites	93
6.4 Installation	93
6.5 Reviewing and configuring host agent settings	93
6.6 Starting and stopping the service	96
6.7 Limitations of Host Agent	96
Part 3. N series Management Console and applications	97
Chapter 7. N series Management Console installation on Linux	99
7.1 Host prerequisites	100
7.2 N series prerequisites	100
7.3 Installation	100
7.3.1 Linux command-line installation	102
Chapter 8. N series Management Console installation for Windows	107
8.1 Host prerequisites	108
8.2 N series prerequisites	108
8.2.1 Function of N series Management Console	108
8.2.2 License requirements	109
8.3 Installation	110
Part 4. Configuring Operations Manager	125
Chapter 9. Configuring Operations Manager	127
9.1 What is Operations Manager	128
9.2 Operations Manager licensing considerations	129
9.3 Discovery	132
9.4 Operations Manager discovery methods	134
9.5 Grouping	137
9.5.1 Homogeneous groups	138
9.6 Reports	141
9.7 Monitoring and alerting	153
9.8 Role Based Access Controls (RBAC)	158
9.9 Storage system management	167
Chapter 10. Performance Advisor operation and configuration	199
10.1 Performance Advisor overview	200
10.2 N series Management Console	204
10.3 Setting up Performance Advisor	215
10.4 Working with the N series Management Console interface	217
10.5 Events, alarms, and thresholds	238
Chapter 11. File Storage Resource Manager	269
11.1 Tracking file system usage and capacity information	271
11.2 About FSRM	271

11.3	How FSRM works	273
11.4	Prerequisites	274
11.5	Quick reference for FSRM tasks	274
11.5.1	Host Agent management tasks	275
11.5.2	Path management tasks	284
11.6	Identifying SRM host agents	290
11.7	Managing host agents	292
11.8	Configuring host agent administration access	296
11.9	Managing FSRM search paths	299
11.10	Scenario: identifying the oldest files in a storage network	303
Chapter 12.	Protection Manager setup	307
12.1	Host prerequisites	308
12.2	License requirements	309
12.3	Installing the license	310
12.4	Running N series Management Console	312
Chapter 13.	Protecting your data with Protection Manager	313
13.1	Introduction to Protection Manager concepts	314
13.1.1	Data sets	316
13.1.2	Protection policies	317
13.1.3	Resource pools	320
13.1.4	Protection Manager Dashboard	321
13.1.5	Protection of discovered data	323
13.2	Restoring data	325
13.3	Setting up the environment for demonstration	326
13.3.1	Overview of the protection policies templates	326
13.3.2	Creating a new protection policy	332
13.4	Demonstration of protecting data with a data set	332
13.4.1	Identifying unprotected data	333
13.4.2	Creating the data set	334
13.4.3	Reviewing the resources to be protected	340
13.4.4	Assigning a protection policy to the data set	344
13.4.5	Reviewing background jobs that ran in order to create the data set	356
13.4.6	Protecting the data set manually or on demand	359
13.5	SnapVault, Open Systems SnapVault, and SnapMirror management	360
13.5.1	Adding an OSSV host to a DFM Server	361
13.5.2	Adding a suitable storage host as an OSSV secondary	373
13.5.3	Creating a OSSV data set	381
13.5.4	Protecting the OSSV data set	386
13.6	Demonstration of restoring data	397
13.7	Creating resource pools, provisioning policies, and data sets	405
13.7.1	Creating a resource pool	405
13.7.2	Creating the provisioning policies	414
13.7.3	Provisioning the data sets using provisioning policies	418
Chapter 14.	Provisioning Manager setup	437
14.1	Host prerequisites	438
14.2	N series prerequisites	438
14.3	Setup	438
14.3.1	Adding the Provisioning Manager license key for Operations Manager	438
14.4	Using the N series Management Console to view Provisioning Manager	442
14.5	Dashboards provisioning information	443
14.5.1	Data sets information	445

14.5.2 Data resource pools information	473
14.5.3 Policies provisioning information.	484
14.5.4 Policies vFilers Templates information	504
14.5.5 Host storage system information.	514
14.5.6 Host vFile Units information.	536
14.5.7 Provisioning Manager configuration help	538
Chapter 15. Operations Manager with VMware ESX server	539
15.1 Introduction	540
15.2 Operations Manager for VMware ESX Server host using N series storage	542
15.3 Running a DataFabric Manager Server on a VMware guest.	543
15.4 Running N series Management Console on a VMware guest.	544
15.5 Host Agent	545
Related publications	547
IBM Redbooks	547
Other publications	547
Online resources	547
How to get Redbooks.	548
Help from IBM	548
Index	549

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:


This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

AIX®
DS4000®
DS8000®

IBM®
Redbooks®
Redbooks (logo) ®

System Storage™
Tivoli®

The following terms are trademarks of other companies:

AMD, AMD Opteron, the AMD Arrow logo, and combinations thereof, are trademarks of Advanced Micro Devices, Inc.

Vfiler, Snapshot, RAID-DP, WAFL, SnapVault, SnapMirror, SnapDrive, NetCache, NearStore, MultiStore, FlexVol, FilerView, DataFabric, Data ONTAP, NetApp, and the NetApp logo are trademarks or registered trademarks of NetApp, Inc. in the U.S. and other countries.

SUSE, the Novell logo, and the N logo are registered trademarks of Novell, Inc. in the United States and other countries.

Oracle, JD Edwards, PeopleSoft, Siebel, and TopLink are registered trademarks of Oracle Corporation and/or its affiliates.

Red Hat, and the Shadowman logo are trademarks or registered trademarks of Red Hat, Inc. in the U.S. and other countries.

SAP, and SAP logos are trademarks or registered trademarks of SAP AG in Germany and in several other countries.

VMotion, VMware, the VMware "boxes" logo and design are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions.

Java, JDBC, JRE, Solaris, Sun, Sun Java, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Active Directory, Microsoft, MS, Windows NT, Windows Server, Windows Vista, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Microsoft product screen captures(s) reprinted with permission from Microsoft Corporation.

Intel, Pentium 4, Pentium, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Preface

IBM® System Storage™ N series with Operations Manager software offers comprehensive monitoring and management for N series enterprise storage and content delivery environments. Operations Manager is designed to provide alerts, reports, and configuration tools from a central control point, helping you keep your storage and content delivery infrastructure in-line with business requirements for high availability and low total cost of ownership.

We focus especially on Protection Manager, which is designed as an intuitive backup and replication management software for IBM System Storage N series unified storage disk-based data protection environments. The application is designed to support data protection and help increase productivity with automated setup and policy-based management.

This IBM Redbooks® publication demonstrates how Operation Manager manages IBM System Storage N series storage from a single view and remotely from anywhere. Operations Manager can monitor and configure all distributed N series storage systems, N series gateways, and data management services to increase the availability and accessibility of their stored and cached data. Operations Manager can monitor the availability and capacity utilization of all its file systems regardless of where they are physically located. It can also analyze the performance utilization of its storage and content delivery network. It is available on Windows®, Linux®, and Solaris™.

The team that wrote this book

This book was produced by a team of specialists from around the world working at the International Technical Support Organization, Austin Center.

Alex Osuna is a project leader at the International Technical Support Organization, Tucson Center. He writes extensively and on all areas of storage. Before joining the ITSO 4 years ago, Alex worked for the Tivoli® Western Region as a Principal SE in storage. Alex has over 30 years in the IT industry, 28 of them with IBM, mainly focused on storage. He holds certifications from IBM, Microsoft®, Red Hat®, and the Open Group.

George Lane is an IT Specialist with the IBM Storage and Technology Group (STG). He has 30 years experience in the IT industry. George has been a part of STG for the last 8 years. His current responsibilities include Worldwide Pre-Sales Technical Support for all of the IBM storage portfolio. He has co-authored three Redbooks publications dealing with NAS, iSCSI, and N series.

Naren Rajasingam is a Consulting IT Specialist for IBM. He has 22 years of experience in the IT industry (9 years with IBM) in the areas of software development, systems development, network programming, network administration, operating systems development, technical instruction, consulting on technical issues, IT systems deployment and support, rapid response, and storage systems. Until recently, he was an IT Infrastructure Architect for IBM GTS. In his current role, Naren works for the Systems & Technology Group in IBM as a storage architect and Consulting Systems Engineer.

Sudheer N Shivakumar is an Technical Advocate System Storage Specialist with the STG/CSI/GSAI Team in India. He has 9 years of experience with IT hardware (servers and IBM System Storage products), and in the areas of education, services, support, and presales. He has certifications from IBM and HP, has Bachelor's degree in Engineering, and a diploma in Electronics and Communication.



Figure 1 Sudheer, George, Alex, and Naren

Thanks to Mark A. Taylor, IBM Americas ATS for his contributions.

Become a published author

Join us for a two- to six-week residency program! Help write a book dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You will have the opportunity to team with IBM technical professionals, Business Partners, and Clients.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you will develop a network of contacts in IBM development labs, and increase your productivity and marketability.

Find out more about the residency program, browse the residency index, and apply online at: ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks publications in one of the following ways:

- Use the online **Contact us** review Redbooks form found at:

ibm.com/redbooks

- Send your comments in an e-mail to:

redbooks@us.ibm.com

- Mail your comments to:

IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400

Archived



Part 1

Introduction and Operations Manager core installation

In this part, we give a architectural introduction and discuss prerequisites, requirements, and recommendations, as well as outline the installation and setup for Operations Manager.

Archived



Introduction

In this chapter, we introduce the components that make up IBM System Storage N series and Operations Manager.

Before Operations Manager (Figure 1-1), administrators of enterprise storage and content delivery networks had to administrate and configure each device individually. This meant relying on spreadsheets and scripts to track system information and device configuration. In this scenario, keeping the storage and content delivery infrastructure in line with business requirements is difficult at best. Common problems include:

- ▶ Configuring and collecting data for large, geographically-diverse networks
- ▶ No simple way of getting a centralized and global status of all devices
- ▶ Configuring and administering each device locally



Figure 1-1 Before Operations Manager

Operations Manager overcomes these issues by consolidating all management activities into a remote and centralized location (see Figure 1-2). This location, known as the DataFabric® Manager (DFM) server, provides a platform from which you can monitor and administer your N series storage network. The monitoring capacity of Operations Manager includes a event log that contains system events, as well as the ability to alert the administrator when specific types of events occur. On the administration side, Operations Manager provides the ability to organize N series devices into hierarchal groups for ease of management. Operations Manager also provides detailed device reporting and usage forecasting. Most importantly, Operations Manager performs auto discovery of your network and then adds any new devices to its database.




Figure 1-2 After Operations Manager

Operations Manager provides infrastructure services for various applications through the IBM N series Management Console (see Figure 1-3). Examples of IBM N series Management Console applications are Performance Advisor, Protection Manager, and Provisioning Manager (see Figure 1-5 on page 7). The DataFabric Manager Server is the Windows or Linux system on which Operations Manager is installed. Services running on the DataFabric Manager Server pass data to the client applications that run in IBM N series Management Console. Performance Advisor allows viewing of historical and real-time performance data collected from IBM N series systems.

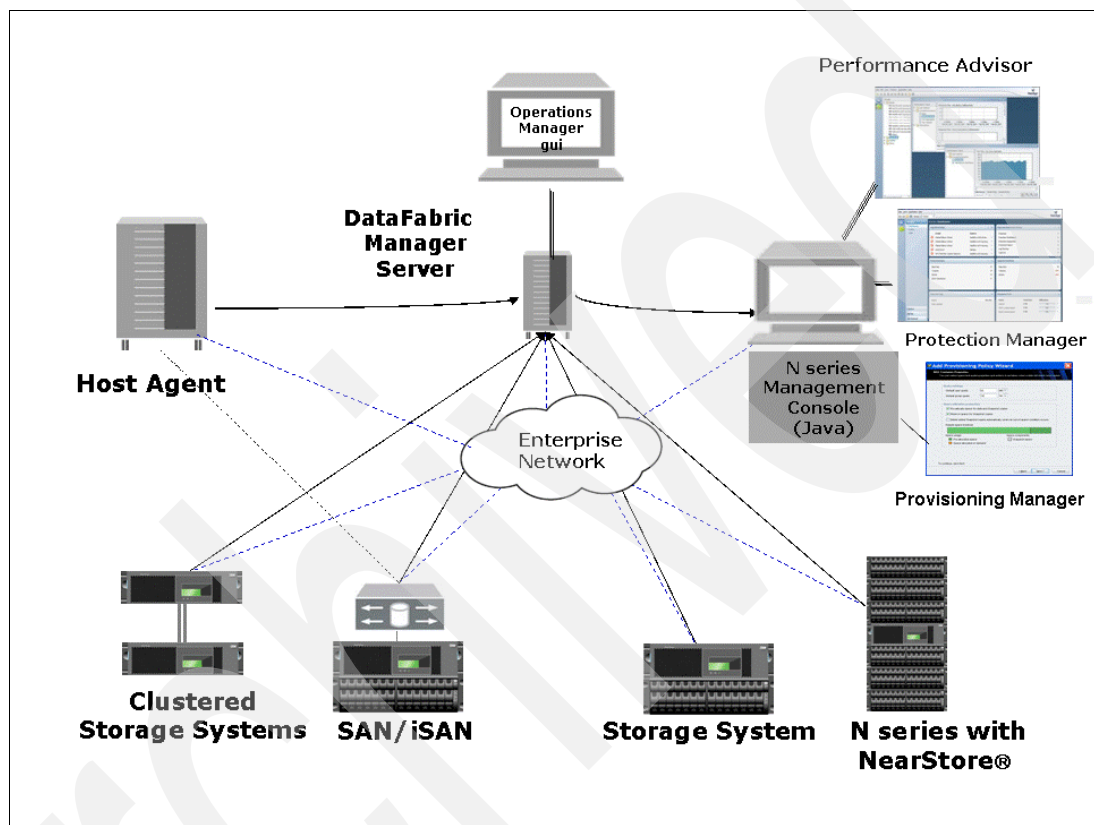


Figure 1-3 Operations Manager architecture

Protection Manager simplifies the managing and monitoring of SnapVault® and SnapMirror® data protection relationships (Figure 1-4 on page 7). Provisioning Manager simplifies and automates provisioning and managing storage for NAS and SAN access (Figure 1-5 on page 7). The Operations Manager GUI is the Web-based user interface of Operations Manager (Figure 1-6 on page 8), from which you can monitor and manage multiple storage systems and active/active configurations on storage systems. Operations Manager is used for day-to-day monitoring, alerting, and reporting about storage infrastructure.

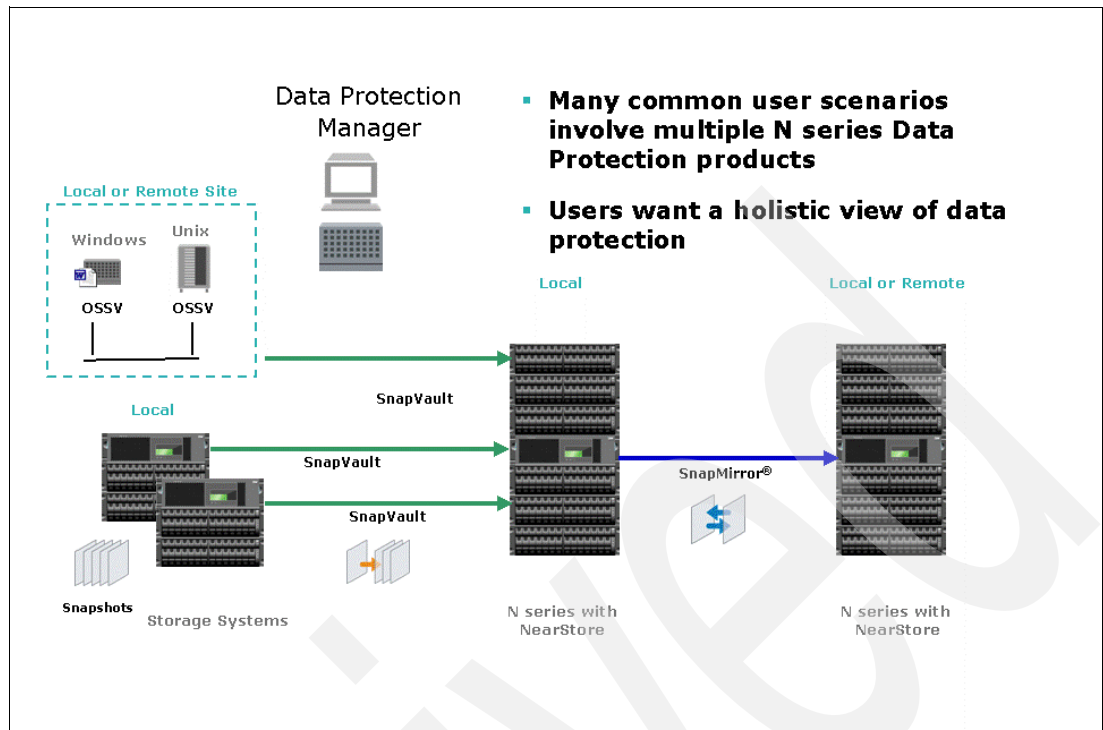


Figure 1-4 Protection Manager

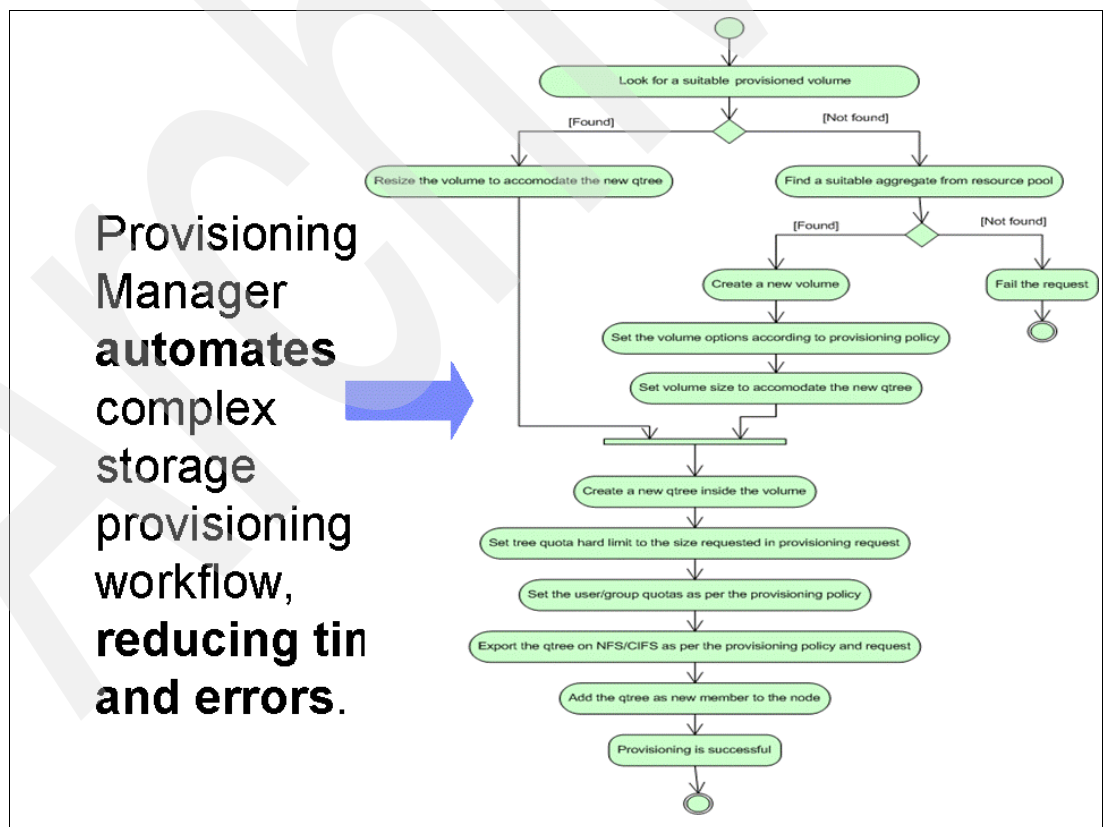


Figure 1-5 Provisioning Manager automation



Figure 1-6 Operations Manager

1.1 Operations Manager UI

Operations Manager has a Web-based UI. It is used for day-to-day monitoring, alerting, and reporting on storage and storage system infrastructure.

1.1.1 Features provided by Operations Manager UI

The Operations Manager UI offers the following features:

- ▶ Discovery
- ▶ Monitoring and reporting
- ▶ Alerting
- ▶ Management

Discovery

You can use the Operations Manager UI to configure Operations Manager for the discovery of storage systems, volumes, qtrees, logical unit numbers (LUNs), disks, and quotas, and then view the results filtered by RBAC and user-defined grouping.

Monitoring and reporting

You can monitor device or object health, capacity utilization, and performance. You can also view or export reports with the relevant information and create custom reports.

Alerting

You can configure alerts and thresholds for event management. Operations Manager issues alerts and Operations Manager generates event reports for monitored systems, volumes, and qtrees. Operations Manager sends the alerts through e-mail, pager, or by generating SNMP traps to be sent to other monitoring applications.

Management

Use Operations Manager UI to perform the following actions:

- ▶ Group devices, vFiler units, host agents, volumes, qtrees, and LUNs into meaningful groups for ease of management. The groups are stored within Operations Manager and are shared with other applications.
- ▶ Configure RBAC settings.
- ▶ Define group configuration management templates and apply those templates to one or more systems.
- ▶ Edit volume, qtree, or user quotas.
- ▶ Run Data ONTAP® CLI commands simultaneously on multiple systems.

- Manage host users, user groups, domain users, local users (see Figure 1-7)



Figure 1-7 Operations Manager managing users

1.1.2 Tabs provided by Operations Manager UI

Operations Manager UI is organized into the following different tabs, each supporting a different type of administrative task:

- Control Center
- Backup
- Disaster Recovery
- File SRM

Control Center

Allows users to configure and view results for discovery, monitoring, reporting, and alerting for storage systems. This tab is enabled when the Operations Manager license is installed.

Backup

Allows users to monitor and manage SnapVault and Open Systems SnapVault disk-to-disk backups. This tab is displayed only when the Business Continuity Option license is installed.

Disaster Recovery

Allows users to monitor and manage SnapMirror disk-to-disk mirroring. This tab is displayed only when the Business Continuity Option license is installed.

File SRM

Allows users to view file system reports filtered and sorted by size, age, owner, access date, and modified date. This subtab is displayed only when the File Storage Resource Manager (File SRM) license is installed.

1.1.3 IBM N series Management Console

IBM N series Management Console (see Figure 1-8) is a client software that contains a number of storage system management applications. IBM N series Management Console consists of the following items:

- ▶ Performance Advisor allows viewing of historical and real-time performance data collected from IBM N series storage systems.
- ▶ Protection Manager provides policy-based data protection by using IBM N series storage systems that have SnapVault, Open Systems SnapVault, or SnapMirror licenses.
- ▶ Provisioning Manager improves efficiency in storage utilization, and automates provisioning and managing storage for NAS and SAN access.

To use the preceding features, you must download and install IBM N series Management Console

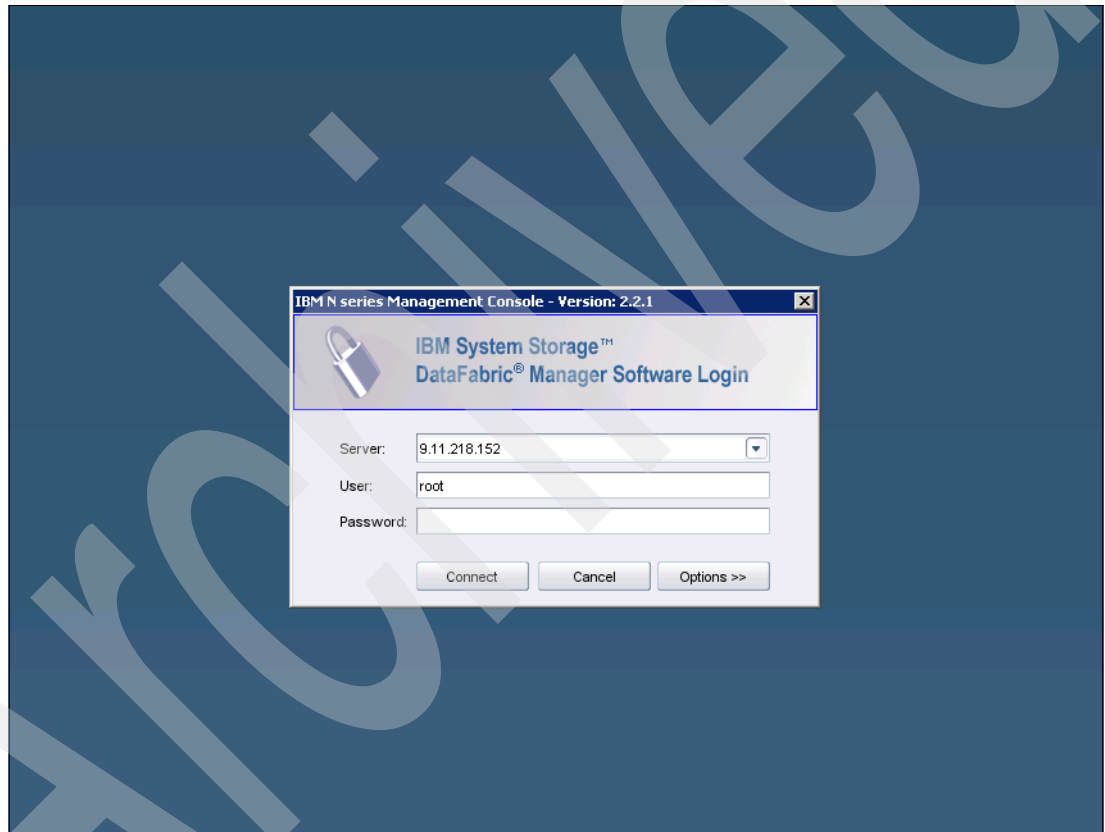


Figure 1-8 N series Management Console

1.1.4 DataFabric Manager Server

The DataFabric Manager Server provides infrastructure services, such as discovery, monitoring, role-based access control (RBAC), auditing, and logging for products in the Storage and Data suites. You can script commands using the command-line interface (CLI) of the DataFabric Manager Server software that runs on a separate server. The software does not run on the storage systems.

Figure 1-9 shows the Operation Manager main window.

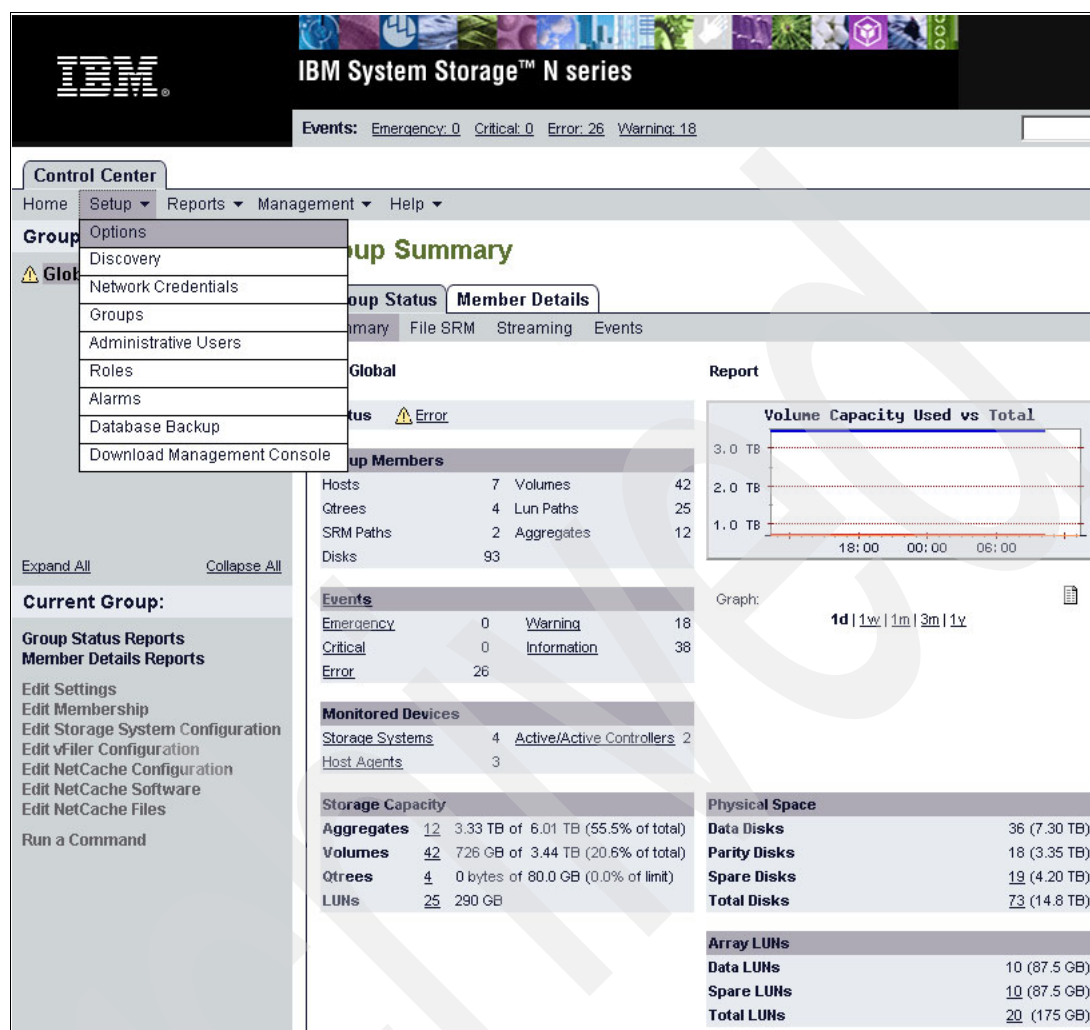


Figure 1-9 Control Center Setup options

In the figure, we show you how to access the Options menu. Select the **Control Center** tab, and then select **Setup** → **Options**.

We discuss the available options in the next section.

Options

The Options windows enable you to configure and customize global options for discovering N series storage systems, for storage monitoring and management, authentication, for user quota monitoring, and for SAN monitoring.

- ▶ The lists of the options
 - Audit Log options
 - Autosupport options
 - Backup Default Thresholds options
 - Backup Discovery options
 - Chargeback options
 - Credential Cache options

- Custom Comment Fields options
- Default Thresholds options
- Database Backup options
- Discovery options
- Display options
- Distribution ACL options
- Events and Alerts options
- Storage System Configuration options
 - Host Agent options
 - LDAP options
 - Licensed Features options
 - Monitoring options
- IBM N series Management Console options
 - Performance Advisor options
 - Reports options
 - Script plug-ins options
 - Security options
 - SNMP trap listener options
 - SRM options
 - Users options

Operation Manager licenses

The Operation Manager licenses you currently have installed on the Operation Manager server workstation are listed in the bottom left-pane area of Operations Manager. Licenses allow you to use certain features and therefore determine the parts of Operations Manager.

The following Operation Manager licenses and the features are as follows.

- ▶ Operations Manager license (formerly known as DataFabric Manager Core license)

Operations Manager license provides the following features:

- Appliance discovery, monitoring, and notifications that alert you of status and problems.
- Monitoring of volumes, qtrees, user quotas, and LUNs on storage systems.
- Monitoring and management of Fibre Channel switches.
- Cluster console for managing active/active configurations.
- Real-time monitoring of streaming protocols.

- ▶ Storage Resource Management option

In addition to the core features, the Storage Resource Management Option provides:

- Monitoring of host agents.
- Monitoring of paths on remote systems.
- File-level and directory-level statistics collection.

► Protection Manager option

The Protection Manager option provides:

- Backup and disaster recovery operations with a policy-based management tool.
- Integration of SnapVault, SnapMirror, and Open Systems SnapVault to help you manage large scale deployments easily.

► Provisioning Manager option

The Provisioning Manager option provides:

- Policy-based provisioning of storage in data sets.
- Conformance of storage to the provisioning policy attached to the data set.
- Space management.

If you have both Protection Manager and Provisioning Manager licensed, then the following features are enabled:

- Assigning provisioning policies to nonprimary nodes of a data set.
- Policy-based provisioning of nonprimary nodes.

1.2 Host Agent

Host Agent is software that resides on a Windows, Linux, or Solaris host. It collects information, such as OS name, version, HBA information, and file system meta data, and then sends that information back to the DataFabric Manager Server. Users can create reports of the collected information by using the Operations Manager UI or the DataFabric Manager Server CLI.

Figure 1-10 on page 15 shows the Host Agent interaction.

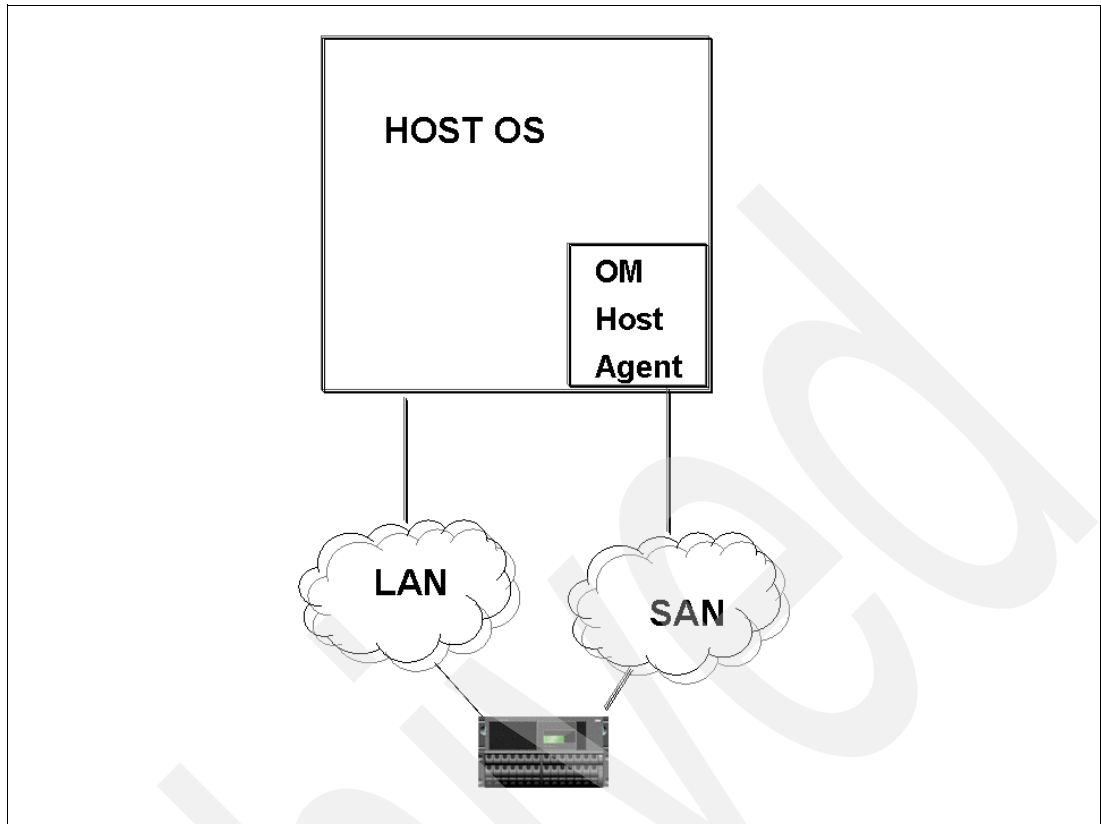


Figure 1-10 Host Agent interaction

To enable a target host to communicate with the DataFabric Manager Server, install and configure the Host Agent software on that host. After the DataFabric Manager Server discovers that instance of Host Agent, no further configuration is required.

Host Agent does not initiate any management actions on the Windows, Linux, or Solaris host. It is strictly a passive agent. It acts only on requests from external management applications, such as the DataFabric Manager Server.

Operations Manager is able to report on data usage by hosts with the installation of a host agent. After you install Host Agent on a non-IBM host, you can use Operations Manager to perform a variety of SAN and FSRM functions.

- ▶ **SAN capabilities:** Using Host Agent and Operations Manager, you can perform the following SAN tasks:

- Monitor basic system information for the SAN hosts.
- View detailed HBA and LUN information.

For more information about Host Agent and SANs, see Chapter 5, “Host Agent installation for Windows 2003” on page 79.

- ▶ **FSRM capabilities:** Using Host Agent and Operations Manager, you can perform the following FSRM tasks:

- Collect storage usage data at the file and directory level.
- Identify a variety of file-related information, for example, largest files, oldest files, or space consumed per file type.

If the Host Agent is located in the same subnet as the DataFabric Manager Server, the DataFabric Manager Server is able to locate it by means of a periodic discover broadcast. If the agent is on another subnet, either the agent of DataFabric Manager Server can be configured to get this agent to report to a designated DataFabric Manager Server.

The Linux host agent listens to communications on port 409.

1.3 Protection Manager

Protection Manager simplifies data protection through automated policy based management of data sets. Protection Manager is a component of the IBM System Storage N series Manageability Software Family (see Figure 1-11). It is a member of the Data Suite, whose products are targeted at backup, replication, and data migration administrators.

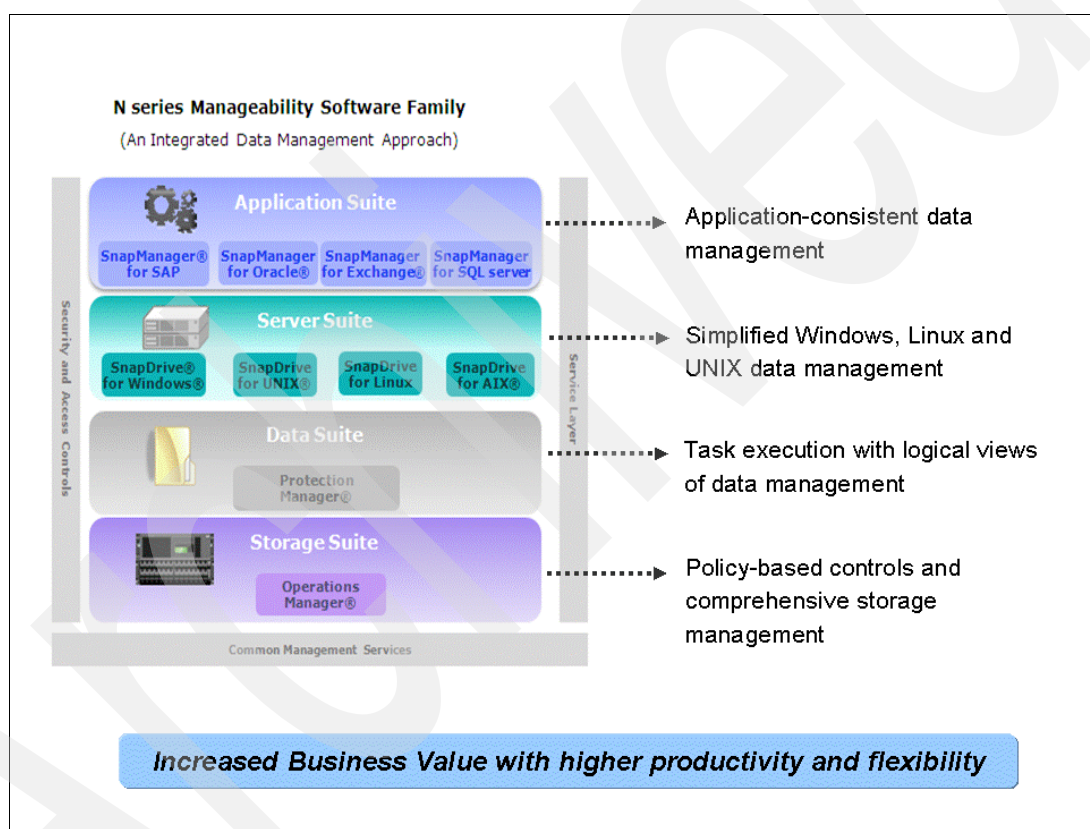


Figure 1-11 N series Manageability Software Family

Protection Manager provides users with an easy way to manage large amounts of data protection relationships by aggregating those volumes and qtrees into data sets (groups). Additionally, Protection Manager provides the ability to create backup/replication policy templates that can be applied to data sets and that automate the replication relationship creation process. Protection Manager is primarily designed to allow cascaded backup to replication policies, to provide end-to-end reporting on the entire data protection pipeline, and to allow users to restore data from any data copy, regardless of how that copy was created.

Protection Manager also provides better management of OSSV client software by automating distribution of OSSV updates and configuration settings.

1.3.1 Provisioning Manager overview

Provisioning Manager, a policy based provisioning tool that allows you to categorize, provision, and monitor storage, and export it using NAS and SAN technologies. This tool allows for the organization and categorization of storage systems and aggregates into resource pools. The created pools can then be associated with data set(s) along with a policy that specifies the rules for provisioning and monitoring the storage (see Figure 1-12). Once the data set is configured, storage can be provisioned and exported easily and monitored throughout the lifetime of the storage

Provisioning Manager runs periodic monitoring and conformance, which leads to appropriate alerts and events, thereby enabling necessary corrective measures.

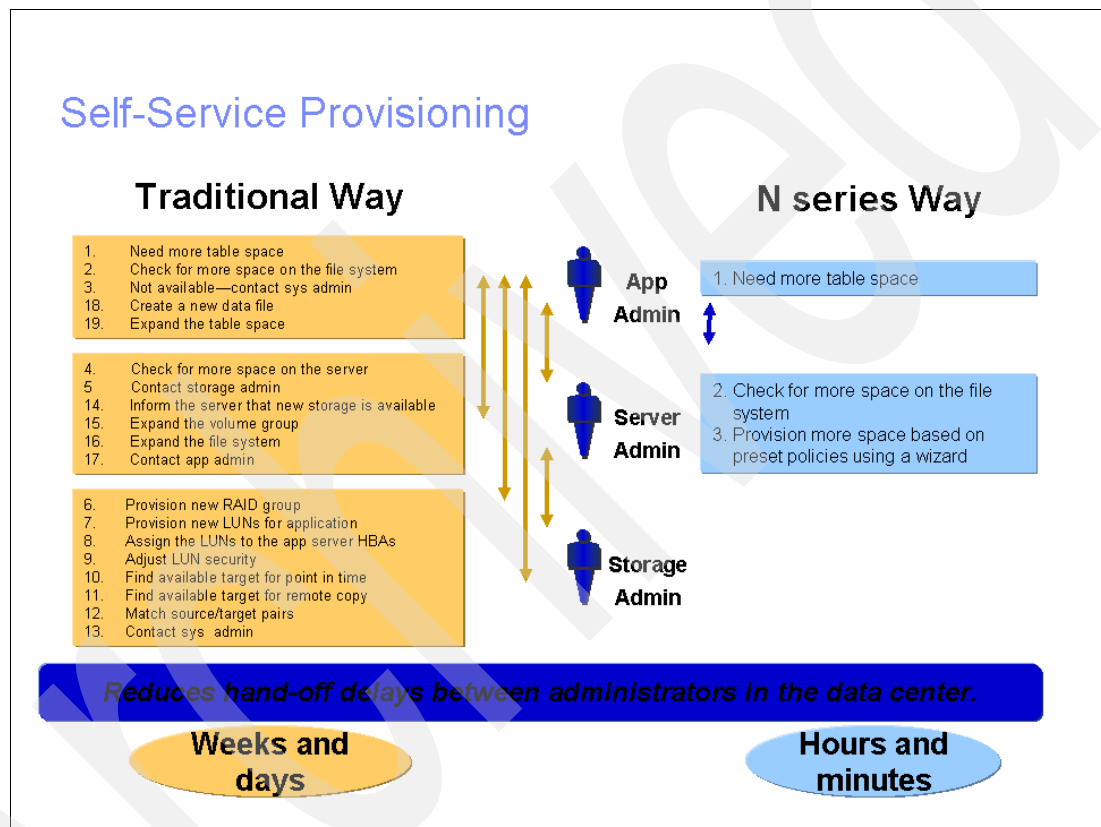


Figure 1-12 Self-Service Provisioning window

1.3.2 Provisioning Manager features

The goal of the Provisioning Manager product is to provide policy-based storage provisioning automation for N series customers with large SAN and NAS environments. It does that by offering the following features:

Provisioning Manager is a new product in Operation Manager Server V3.7. It runs in the N series Management Console. Provisioning Manager helps you simplify and automate the tasks of provisioning and managing storage for data.

Provisioning Manager provides the following capabilities:

- Provisioning policies that manage provisioning and exporting of storage.
- Automatic provisioning of a data set when you assign a provisioning policy to it.

- ▶ Periodic checking of provisioned storage for conformance to the policy.
- ▶ Manual controls for adding volumes or qtrees to a data set on existing storage and newly provisioned storage.
- ▶ Manual controls for resizing volumes and for deleting old Snapshot™ Copies on existing and newly provisioned storage.

1.4 What is FSRM

The File Storage Resource Manager (FSRM) feature of Operations Manager provides monitoring and management of storage resources, including applications, files, file systems, and networks. The DataFabric Manager Server interacts with the Host Agent software residing on remote Windows, Solaris, or Linux hosts to recursively examine the directory paths you have specified in the Operations Manager FSRM configuration options. The results of these directory examinations are used to generate a variety of useful file-level and directory-level reports, as shown in Figure 1-13 on page 19.

Host Agent enables you to perform the following FSRM functions through Operations Manager:

- ▶ Set up path walk schedules for collecting file-level storage usage statistics.
- ▶ Identify file-level statistics, such as the following:
 - Largest files
 - Oldest files
 - Stalest files
 - Newest files
 - Largest directories
 - Files by owner
 - Files by type

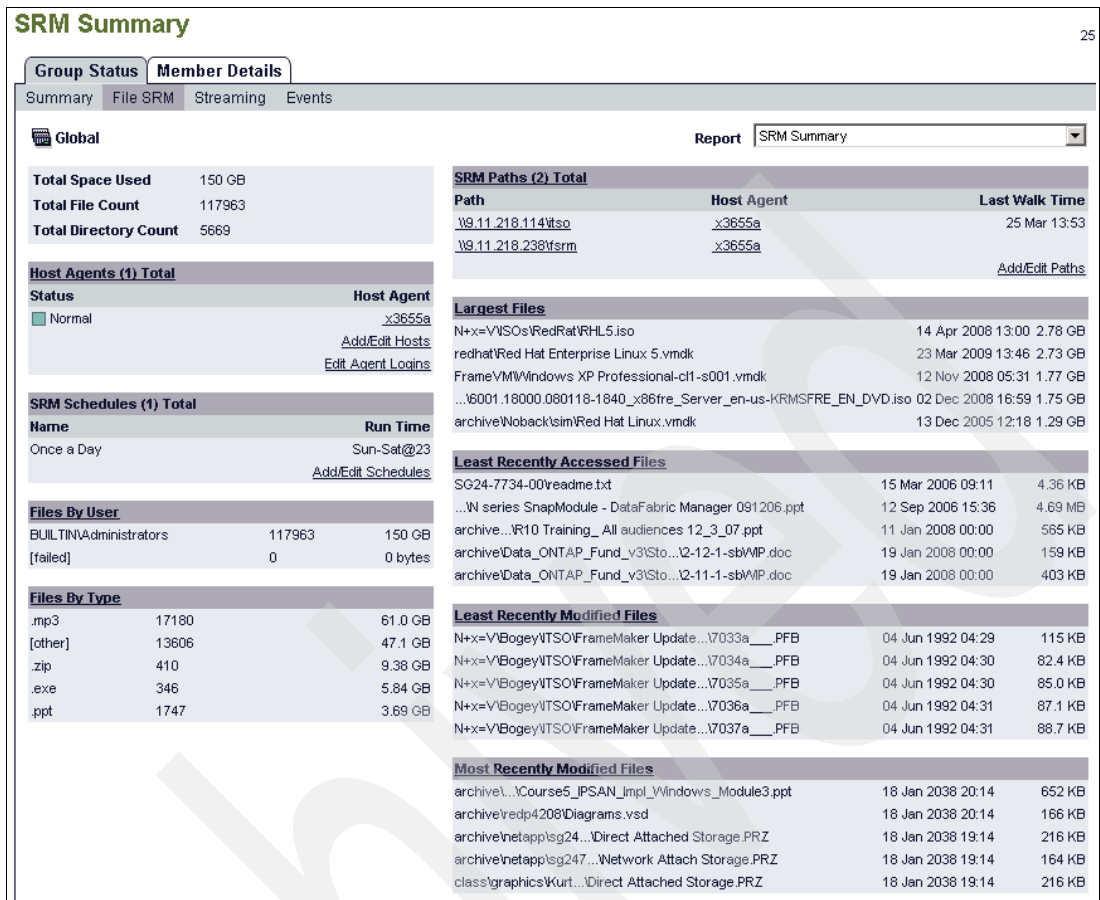


Figure 1-13 FSRM example

Through Operations Manager, you can monitor only directory paths that are already visible to the Host Agent software. Therefore, if you want to enable FSRM monitoring of an IBM System Storage N series storage system, the host agent must mount that storage system share using NFS or CIFS, or the host agent must use a LUN on that storage system.

Note: The DataFabric Manager Server cannot obtain file system data for files located in a storage system's volumes that are not exported by CIFS or NFS. Host agents can also gather FSRM data about file system paths that are not on a storage system, for example, local disk and third-party storage.


1.5 Naming conventions

Previously, Operations Manager Core was referred to as DataFabric Manager Core. There is also a Operations Manager UI that was previously referred to as Operations Manager. In Version 3.7.1, the setup screens and executables will still refer to DataFabric Manager, and many of the commands still begin with a DataFabric Manager acronyms like *dfm*.

Table 1-1 explains this conversion.

Table 1-1 Naming conventions

Component	DataFabric Manager (DFM)	Operations Manager
Core	DFM	Operations Manager
Host Agent	DFM Host Agent	Host Agent or Operations Manager Host Agent
User Interface (UI)	Operations Manager	Operations Manager UI
DataFabric Manager Server	DataFabric Manager Server	DataFabric Manager Server
N series Management Console	N series Management Console	N series Management Console
DataFabric Manager Server Command Line	DataFabric Manager command line and DataFabric Manager CLI	DataFabric Manager Server command line and DataFabric Manager CLI



Installing Operations Manager: Windows 2003 32-bit operating system

This chapter will discuss installing the Operations Manager on the Microsoft Windows 2003 32-bit operating system.

2.1 Host prerequisites

Your DataFabric Manager Server must meet the requirements outlined in Table 2-1 and Table 2-2 before you install Operations Manager. For more information, refer to the *IBM System Storage N series Operations Manager Sizing and Installation Guide*, REDP-4270. This document provides information that can help you determine the correct configuration for a system to host the DataFabric Manager Server.

Table 2-1 gives the recommended requirements for environments with 1 to 25 nodes.

Table 2-1 Requirements

Hardware requirements	Software requirements
<ul style="list-style-type: none">▶ Intel®-based PC with single 2 GHz (Xeon or Pentium® 4)▶ 4 GB of free disk space minimum, 8 GB recommended▶ 1 GB of memory minimum	Windows 2003 Server 32-bit on x86

Table 2-2 gives the recommended requirements for environments with 25 to 100 nodes.

Table 2-2 Requirements for 25 to 100 nodes

Hardware requirements	Software requirements
<ul style="list-style-type: none">▶ PC based on Intel with single 2 GHz CPU (Xeon or Pentium 4™)▶ 6 GB of free disk space, 12 GB recommended▶ iSCSI, FC LUN, or internal RAID disk▶ 2 GB of memory minimum	Windows 2003 Server 32-bit on x86

Note: Operations Manager V3.7.1 is not supported on Windows NT® 4.0, Windows 2000, and Windows XP.

2.2 License requirements

You must have a valid DataFabric Manager Server license key to complete the Operations Manager installation (see Table 2-3 on page 23). You can access your license key at the following address:

<http://www.ibm.com/storage/nas/>

After you have accessed the site, follow the options provided to access license keys. After you complete the installation, you can enter additional license keys on the Options window in Operations Manager.

You can install (or upgrade to) Operations Manager V3.7.1 using the core license key. However, if you do not have the core license key, you need the following licenses to monitor and manage your storage systems:

- ▶ DataFabric Manager Server
- ▶ Additive

DataFabric Manager Server license

The DataFabric Manager Server license is the server license with a unique serial number that tracks the number of Operations Manager installations. You must have this license to enable features. The node count is one.

Additive license

The additive license is an additional license with a unique serial number that is used to increase the node count and enable the features.

Table 2-3 Required licenses

To user	To install this license or product	That enables these features
Operations Manager	<ul style="list-style-type: none">▶ DataFabric Manager Server license▶ Operations Manager license <p>Note: Required for all licensed Operation manager installations. Sets the maximum number of storage systems that the DataFabric Manager Server can monitor in this installation.</p>	<ul style="list-style-type: none">▶ Automated policy-based data protection for NAS and SAN storage systems SnapVault, Open Systems SnapVault, and SnapMirror management▶ Policy conformance checking and alerting Monitoring Reports Storage usage and availability, such as qtrees, volumes, aggregates, LUNs, and disks▶ Storage systems▶ vFiler units▶ Real-time streaming events▶ Managing▶ Storage system configuration▶ Scripts▶ Monitoring and managing storage system▶ Clusters using Cluster Console▶ Displaying historical and real-time performance data using Performance Advisor in IBM Nseries Management Console

2.3 N series prerequisites

At the time of publication, the N series models shown in Figure 2-1 were supported on Operations Manager.

N3000 System Hardware	N3300 2859-(A10, A20), N3600 2862-(A10, A20)
N3700 System Hardware	N3700 2863-(A10, A20)
N5000 System Hardware	N5200 2864-(A10, A20, G10, G20), N5500 2865-(A10, A20, G10, G20) N5300 2869-(A10, A20, G10, G20), N5600 2868-(A10, A20, G10, G20)
N6000 System Hardware	N6040 2858-(A10, A20) or with Gateway Feature Code 9551 N6070 2858-(A11, A21) or with Gateway Feature Code 9551
N7000 System Hardware	N7600 2866-(A10, A20, G10, G20), N7800 2867-(A10, A20, G10, G20) N7700 2866-(A11, A21, G11, G21), N7900 2867-(A11, A21, G11, G21)
N3000, N3700, N5000, N6000, N7000 Licensed Functions	2870-(591, 592, 621, 622, 631, 632, 641, 642, 645, 646, 651, 652, 655, 656, 661, 662, 665, 666, 663, 664, 667, 668, 671, 672, 675, 676, 673, 674, 677, 678, 681, 682, 685, 686, 691, 692, 695, 696)

Figure 2-1 N series products compatible with Operations Manager V3.7.1

2.3.1 Data ONTAP requirements

You must be running Data ONTAP Version 7.1 or later.

Note: You must have a DataFabric Manager plug-in for each version of Data ONTAP that you are running across your system. Operations Manager automatically includes the plug-ins for Data ONTAP. To list the versions of the plug-ins for Data ONTAP, use the **dfm plugin list** command at the DFM Server command line. You do not need to download a plug-in unless you are using a different version of Data ONTAP.

DataFabric Manager Server Data ONTAP plug-in installation instructions

To download and install the DataFabric Manager Server Data ONTAP plug-in software on Windows that you want to use as the DataFabric Manager Server, complete the following steps:

1. Download `filerconfig_Windows.zip`, which should contain the version of the DataFabric Manager Server Data ONTAP plug-in that you need. You can download this file at the following address:
<http://www-947.ibm.com/systems/support/supportsite.wss/supportresources?brandind=5000029&familyind=5329833&taskind=2>
2. From the DataFabric Manager Server command line, run:

```
dfm plugin add [-f] {path | URL}
```

You can omit the `-f` switch, unless you want Operations Manager to reinstall a plug-in that has already been installed.

Any path or URL that can be interpreted or resolved by your system and network is valid. For example, you could run the following command:

```
dfm plugin add http://web.company.com/~jsmith/filerconfig_Windows.zip
```
3. To confirm that the plug-in is installed from the DataFabric Manager Server command line, run:

```
dfm plugin list
```

2.4 Installing Operations Manager

Prior to installing the Operations Manager software, you must ensure that you have met requirements in the following areas:

- ▶ Hardware and software requirements
- ▶ License requirements
- ▶ Data ONTAP requirements

To install Operations Manager on your Windows server, complete the following steps.

1. Insert the CD into the CD-ROM drive of your management station.
2. Launch dfmsetup-3-2-win32.exe in the win32 directory on your CD.
3. Follow the Operations Manager Setup prompts to complete the installation, or get the Operations Manager V3.7.1 installer from:

<http://www.ibm.com/storage/support/nas/>

2.4.1 Deploying Operations Manager software

You must deploy Operations Manager on a system that is running no other applications.

To start the deployment, launch dfmsetup-3-7-1-win32.exe.

Figure 2-2 shows the Windows 2003 R2 32-bit operating system desktop, where you can find the dfmsetup file, which is the Operations Manager installation setup file.




Figure 2-2 Windows 2003 and Operations Manager setup file window

Note: In a standard Windows 2003 installation, the default security setting for the browser is “high.” This setting can cause the browser to block certain actions that can interfere with your Operations Manager upgrade. To ensure the best browsing experience when using Operations Manager, you might need to adjust your browser security setting to “medium.”

Follow these steps to install the Operations Manager:

1. Execute the Operations Manager's setup file. Figure 2-3 shows the first window of the Operations Manager.

Click **Cancel** if you want to cancel the installation of Operations Manager

Click **Next** if you want to continue the installation of Operations Manager.



Figure 2-3 Operations Manager installation window

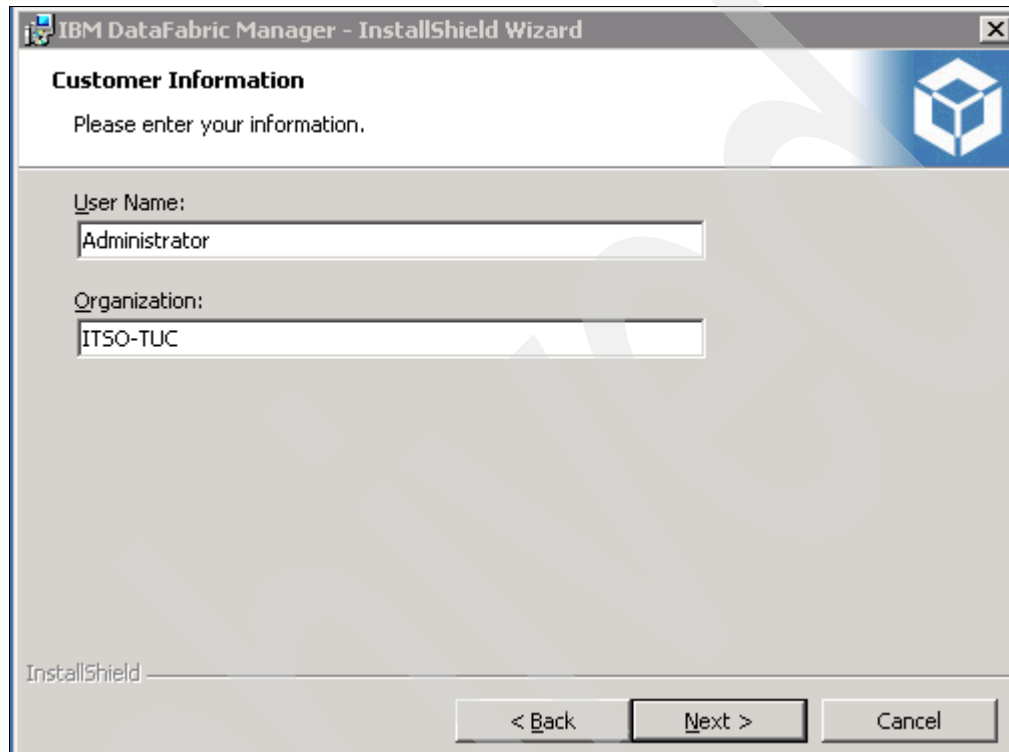
2. Figure 2-4 shows you the IBM Autosupport authorization Information for the Operations Manager. Select **I accept**, which allows you to proceed with the Installation of the Operations Manager.



Figure 2-4 Operations Manager Autosupport Notice window

3. Figure 2-5 prompts you to provide the User Name and the Organization. By default, it captures the computer name as the User Name. You can change the User Name and the Organization as needed; here we use Administrator as the User Name and ITSO-TUC as the Organization.

Once you have provided the User Name and Organization, click **Next** to proceed with the installation.



The screenshot shows a Windows-style dialog box titled "IBM DataFabric Manager - InstallShield Wizard". The window has a standard title bar with a close button (X) in the top right corner. Below the title bar, the text "Customer Information" is displayed in bold. Underneath, a prompt reads "Please enter your information." To the right of this text is the IBM logo. The main area of the window contains two text input fields. The first field is labeled "User Name:" and contains the text "Administrator". The second field is labeled "Organization:" and contains the text "ITSO-TUC". At the bottom of the window, there is a status bar that says "InstallShield". To the right of the status bar are three buttons: "< Back", "Next >", and "Cancel". The "Next >" button is highlighted with a black border.

Figure 2-5 Operations Manager Customer Information window

4. Figure 2-6 prompts you for the 14 character License Key for Operations Manager. Without the authorization key, you cannot proceed with the installation.

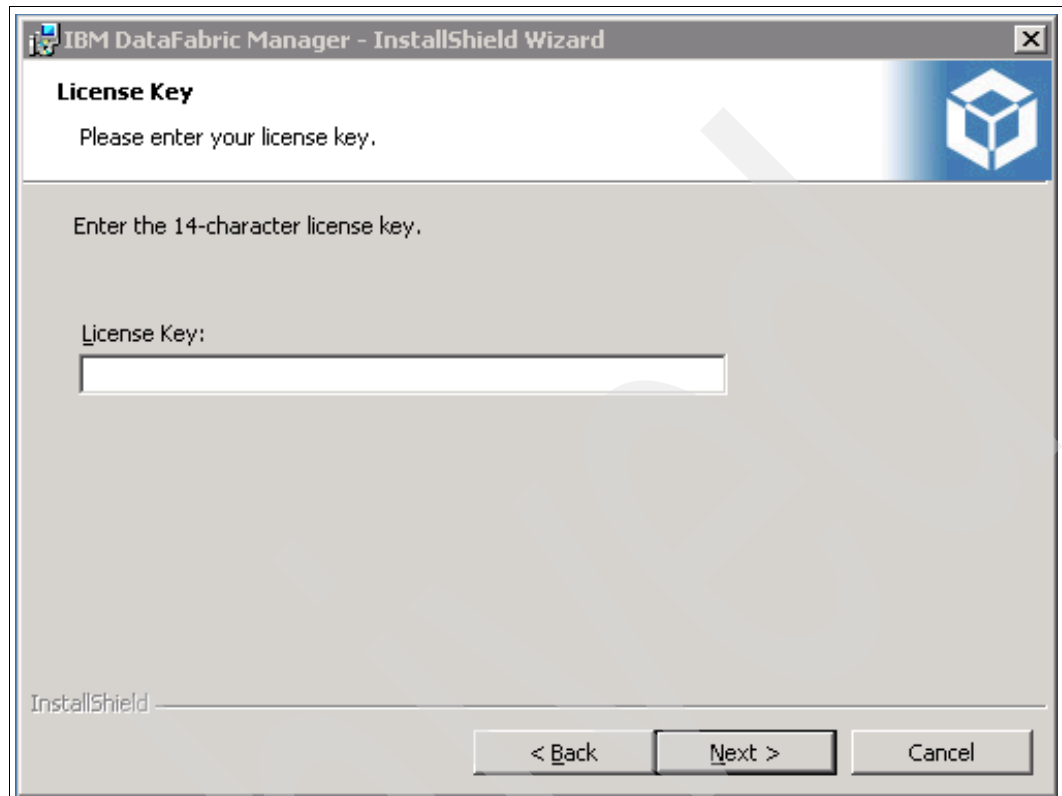


Figure 2-6 Operations Manager authorized license key window

5. Figure 2-7 will prompt you to select the path where you want to install the Operations Manager. By default, the path is C:\Program Files\IBM\DataFabric. You can change the path by clicking **Change** and browsing to the path you wish to use. Here we used the default path.




Figure 2-7 Operations Manager Destination Folder window

6. Figure 2-8 shows that our system is ready for the installation of the Operations Manager. Click **Install**.

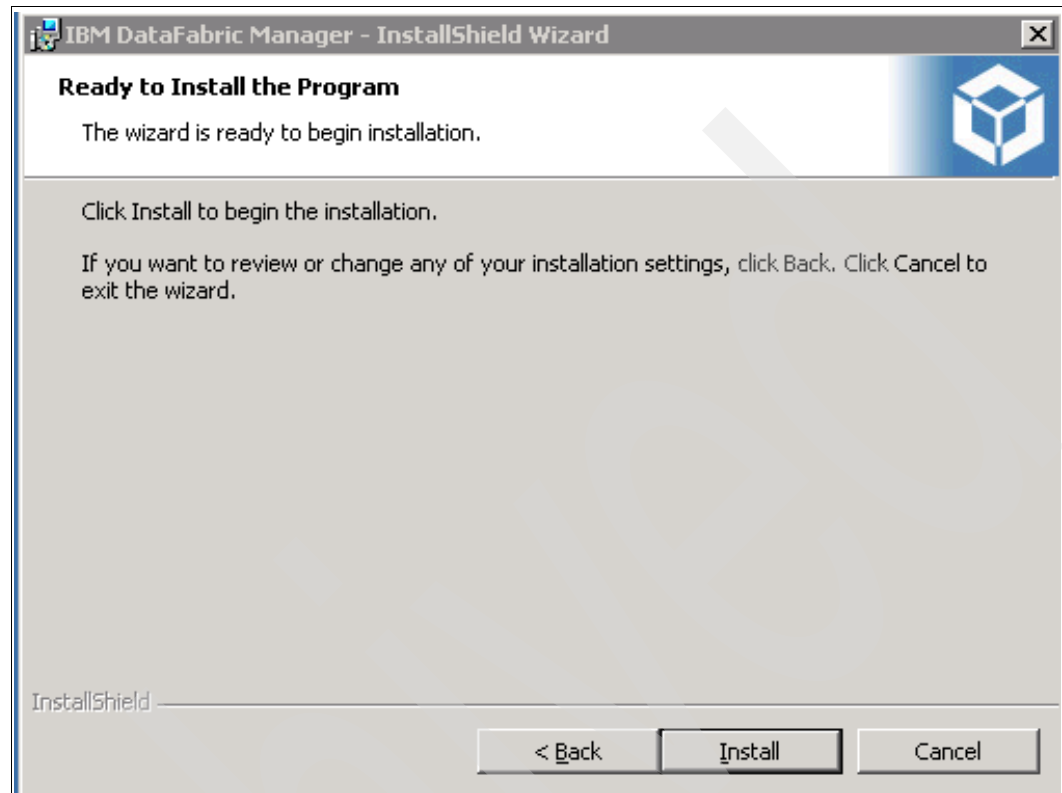


Figure 2-8 Operations Manager Ready to Install window

7. Figure 2-9 shows that Operations Manager is copying all the installation files required for the Operations Manager to run. This process takes approximately 3 minutes.

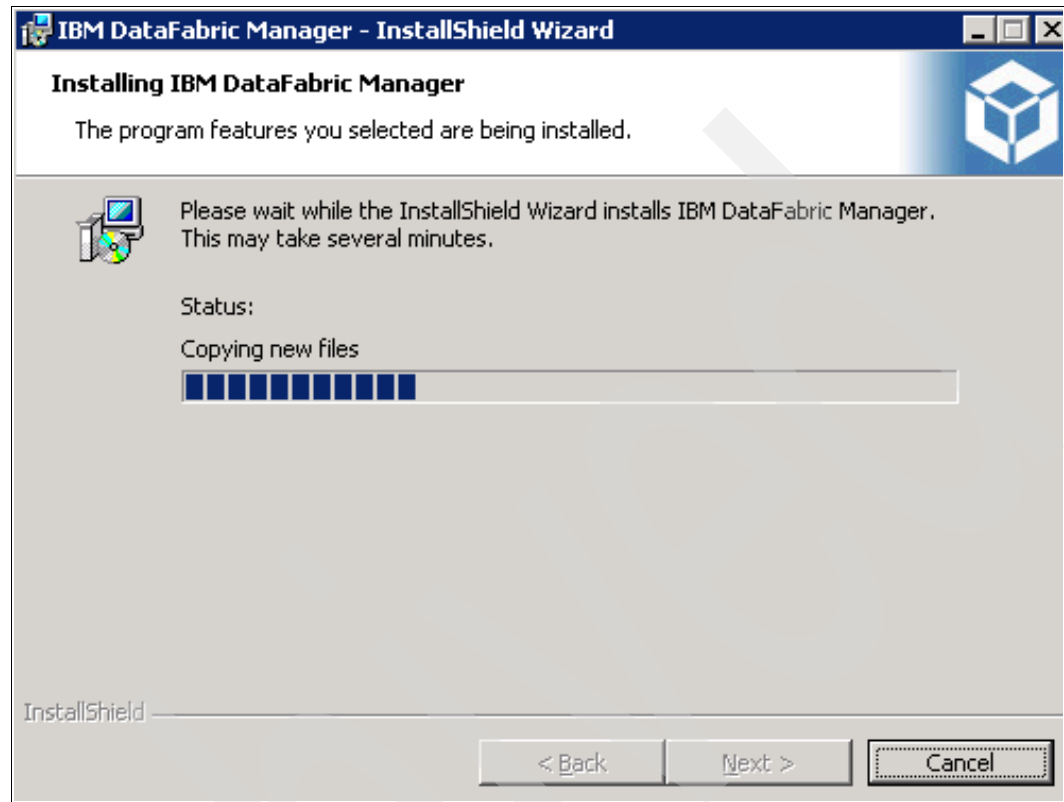


Figure 2-9 Operations Manager installation window

8. Figure 2-10 show that once the copying of the installation files of the Operations Manager completes, the services start, which takes about 5 minutes to complete.




Figure 2-10 Operations Manager installation window

9. Figure 2-11 shows that the installation has finished. Click **Finish**.
The total installation of the Operations Manager takes approximately 15 minutes.




Figure 2-11 Operations Manager installation completion window

10. Figure 2-12 prompts you to restart the system to complete the installation. Click **Yes** if you want to restart now or click **No** to restart the system later.




Figure 2-12 Operations Manager system restart prompt window

11. To open the N series Management Console after the system is restarted (or without restart), click **Start** → **All Programs** → **IBM** → **DataFabric Manager** → **Show Appliance Summary Page**. You should then see the window shown in Figure 2-13 on page 35.

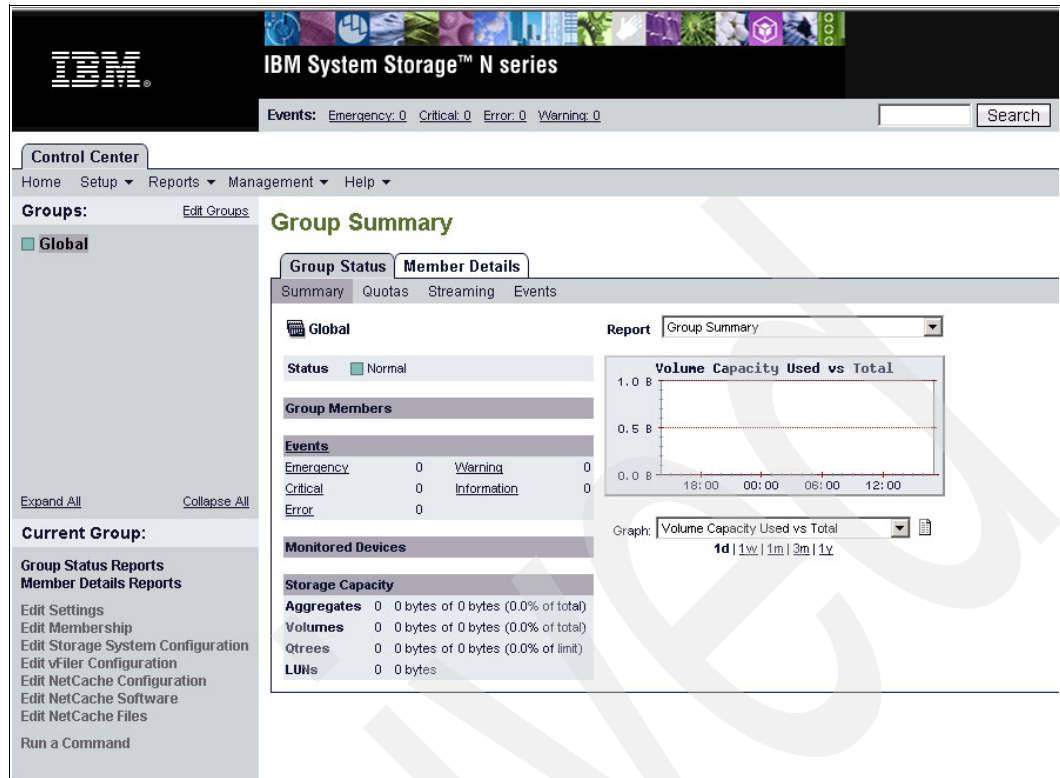


Figure 2-13 Operations Manager window

Archived

Installing Operations Manager: Windows 2003 64-bit operating system

This chapter covers the basic prerequisites and installation of Operations Manager on a Windows 2003 64-bit server. You need to install the following items:

- ▶ Operations Manager for the appropriate OS
- ▶ License keys for the options needed
- ▶ Plug-in for the appropriate DataFabric Manager Server and Data ONTAP version
- ▶ Agents for any hosts to be monitored
- ▶ IBM N series Management Console (This will be covered in Chapter 8, “N series Management Console installation for Windows” on page 107.)

3.1 Host prerequisites

The following prerequisites are for 1 to 25 nodes. These requirements are for a Operations Manager installation with only the basic system monitoring enabled. If you enable additional features and monitor additional objects, such as additional storage systems, qtrees, user quotas, and if use the Storage Resource Manager, Performance Advisor, Business Continuance Option, Provisioning Manager, or Protection Manager features, a more powerful platform may be required.

A server should be dedicated to Operations Manager and its services. No other applications or services should be running on the server.

Table 3-1 Requirements matrix

Windows 2003 server	
Hardware requirements	Software requirements
Intel-based PC with a single 2 GHz CPU (Xeon or Pentium 4)	Windows 2003 server, 64-bit on x64 (in WOW64 mode)
4 GB of free disk space minimum, 8 GB recommended	
1 GB of memory minimum	

In our lab, the server we used was an IBM System x3655 with a Dual Core AMD™ Opteron™ 2.60 GHz processor with 4 GB of RAM. For sizing purposes, refer to the *IBM System Storage N series Operations Manager Sizing and Installation Guide*, REDP-4270, which can be found on the IBM Redbooks Web site at the following address:

<http://www.redbooks.ibm.com/abstracts/redp4270.html?Open>

This document provides information that can help you determine the correct configuration for a system to host the DataFabric Manager Server.

3.2 N series prerequisites

There are several licensing and Data ONTAP requirements depending on the options you plan to use. The following information is a guide, but be sure to check the IBM System Storage N series Operations Manager Interoperability Matrix on the IBM Web site and the “About Operations Manager and the DataFabric Manager Server” section. These items can be found at the following address:

<http://www.ibm.com/totalstorage/nas>

You may also find this information in the *Installation and Upgrade Guide for use with DataFabric Manager Server 3.7*, GC26-7892.

Note: All of the features of Operations Manager V3.7 are supported with Data ONTAP V7.0 and above, except for SNMP V3 discovery and monitoring, which requires Data ONTAP V7.3.

Figure 3-1 on page 39 show a list of hardware that supports Operations Manager.

N3000 System Hardware	N3300 2859-(A10, A20), N3600 2862-(A10, A20)
N3700 System Hardware	N3700 2863-(A10, A20)
N5000 System Hardware	N5200 2864-(A10, A20, G10, G20), N5500 2865-(A10, A20, G10, G20) N5300 2869-(A10, A20, G10, G20), N5600 2868-(A10, A20, G10, G20)
N6000 System Hardware	N6040 2858-(A10, A20) or with Gateway Feature Code 9551 N6070 2858-(A11, A21) or with Gateway Feature Code 9551
N7000 System Hardware	N7600 2866-(A10, A20, G10, G20), N7800 2867-(A10, A20, G10, G20) N7700 2866-(A11, A21, G11, G21), N7900 2867-(A11, A21, G11, G21)
N3000, N3700, N5000, N6000, N7000 Licensed Functions	2870-(591, 592, 621, 622, 631, 632, 641, 642, 645, 646, 651, 652, 655, 656, 661, 662, 665, 666, 663, 664, 667, 668, 671, 672, 675, 676, 673, 674, 677, 678, 681, 682, 685, 686, 691, 692, 695, 696)

Figure 3-1 Operations Manager supported hardware

Table 3-2 will give you an idea of what license is needed to activate a particular function within Operations Manager. You must install the core Operations Manager license first before installing any other Operations Manager components.

Table 3-2 Features table

To use	Install this license or product	That enables these features
Operations Manager	DataFabric Manager Server license	Automated policy-based data protection for N series NAS and SAN storage systems
	Operations Manager license	SnapVault, Open Systems SnapVault, and SnapMirror management
		Policy conformance checking and alerting
		Monitoring
		Reports
		Storage usage and availability, such as qtrees, volumes, aggregates, LUNs, and disks
		Storage systems
		Vfiler™ Units
		Real-time streaming events
		Managing <ul style="list-style-type: none"> ► Storage system configuration ► Scripts
		Monitoring and managing storage system clusters using Cluster Console
	Required for all licensed Operations Manager installations. Sets the maximum number of appliances that the DataFabric Manager Server can monitor.	Displaying historical and real-time performance data using Performance Advisor in N series Management Console

To use	Install this license or product	That enables these features
Protection Manager	DataFabric Manager Server license	Automated policy-based provisioning for N series SAN and NAS storage systems
	Operations Manager license	Space management policies and capacity reporting
	Protection Manager license	Policy conformance checking and alerting
	N series Management Console	Enables management of Provisioning Manager and Protection Manager
Provisioning Manager	DataFabric Manager Server license	Automated policy-based provisioning for N series SAN and NAS storage systems
	Operations Manager license	Space management policies and capacity reporting
	Provisioning Manager license	Policy conformance checking and alerting
	N series Management Console	If you have both Protection Manager and Provisioning Manager licensed, then the following features are enabled: <ul style="list-style-type: none"> ▶ Assigning provisioning policies to nonprimary nodes ▶ Policy-based provisioning of primary storage ▶ Assigning protection policies to provisioned data sets
Protection Manager with Disaster Recovery	DataFabric Manager Server license	Failover and manual failback for N series NAS and SAN storage systems
	Operations Manager license	
	Protection Manager license	
	Protection Manager Disaster Recovery license	
	N series Management Console	
File Storage Resource Manager (File SRM)	DataFabric Manager Server license	Tracking file system usage and capacity information
	Operations Manager license	
	File SRM Option	
Business Continuance	DataFabric Manager Server license	Backup Manager
	Operations Manager license	This space left intentionally blank.
	Business Continuance Option (BCO)	Configuring and scheduling disk-to-disk backups of all systems enabled with SnapVault, including Open Systems SnapVault
		Disaster Recovery Manager <ul style="list-style-type: none"> ▶ Monitoring SnapMirror relationships ▶ Configuring and scheduling disk-to-disk mirrors of all systems enabled with SnapMirror

To use	Install this license or product	That enables these features
Storage Area Network	DataFabric Manager Server license	Monitoring and managing SAN hosts and FC switches
	Operations Manager license	Managing LUNs on Windows SAN hosts
	Storage Area Network Option	

3.3 Plug-ins

An important item to keep in mind is that for each version of Operations Manager you install, you need to check the level of plug-in for the version of Data ONTAP you are running. For example, we are running Data ONTAP V7.3.1, so we need to install the plug-in for that version. For example, we download the DFM Server plug-in for Data ONTAP V7.3.1 and performed the following steps:

1. Select **DataFabric Manager** → **Continue**.
2. Select **DataFabric Manager Plug-ins for Data ONTAP 7.3.1**.
3. Select **DataFabric Manager Data ONTAP 7.3.1 Plug-in for Windows**.
4. Click **I Agree**.
5. Click **I Confirm**.

The first time you perform the download, the system will download a small piece of code that you can configure for all downloads of the plug-ins that follow. You can set, for example, the directory to save to and unpacking options.

You do not need to automatically unpack the plug-ins. The **dfm** command to install the plug-in expects the code to be in a zipped format.

To install the plug-ins, do these steps:

1. Select **Setup** in the Download Directory window, as shown in Figure 3-2.

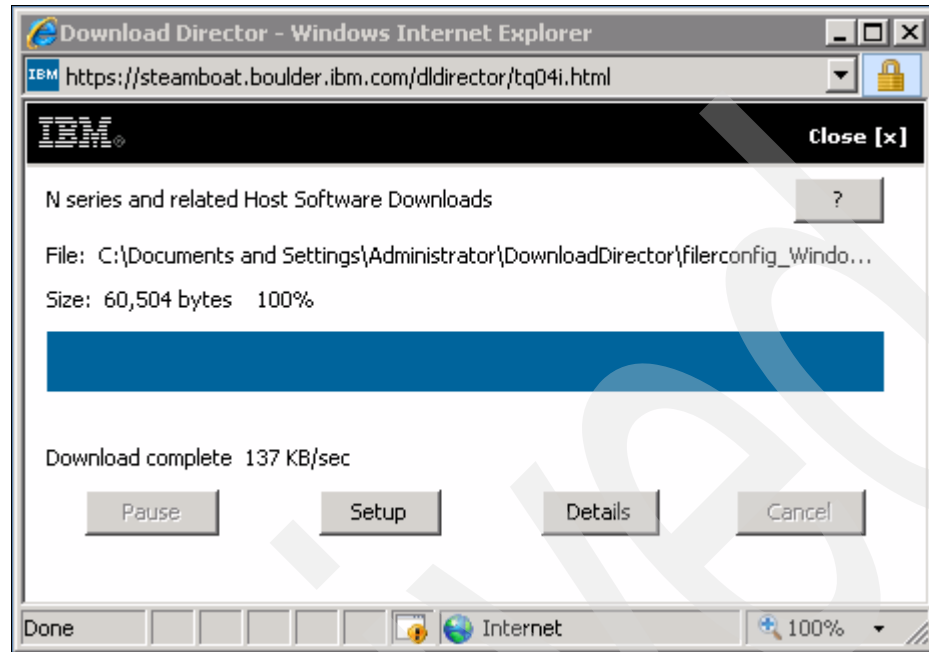


Figure 3-2 Select setup window

2. The next window will give you the option to set a default download directory or request a download directory each time you do a download. It will also give you a chance to set the unpack options.

We suggest that you do not have the file unpacked, as the Operations Manager expects the file to be in the packed format.

Select the **Set unpack options** bar, as shown in Figure 3-3.

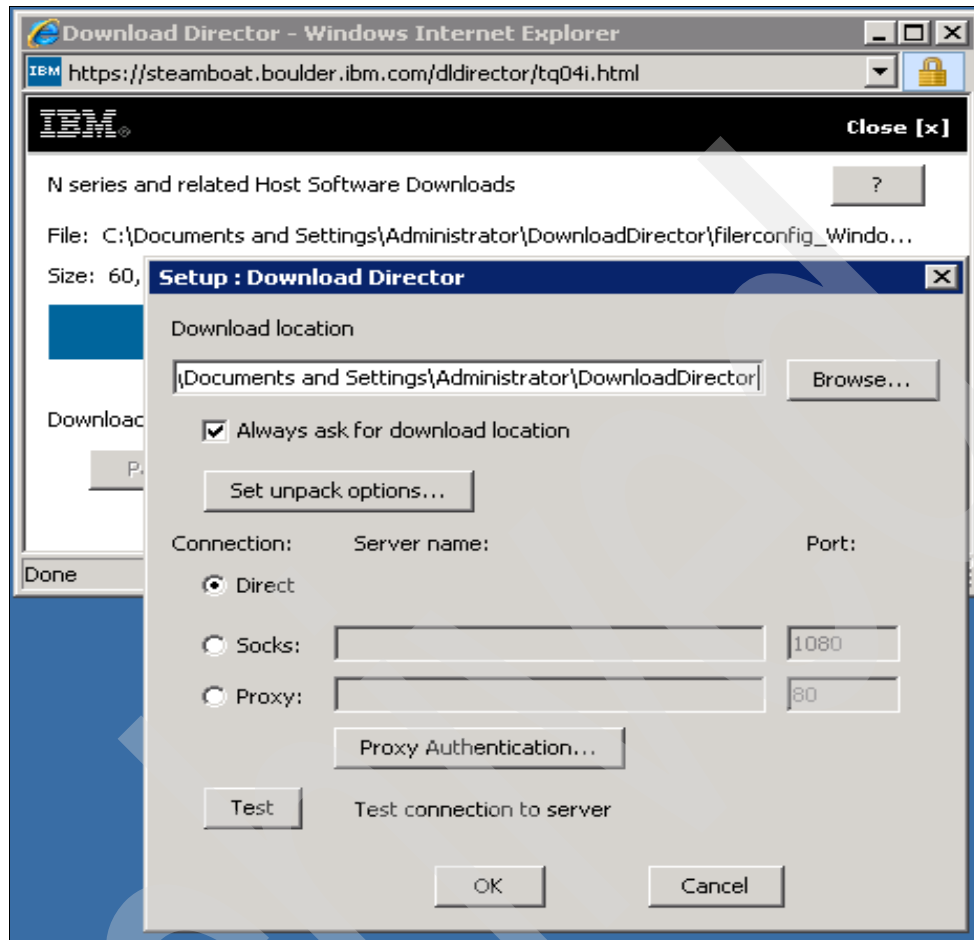


Figure 3-3 Select Set unpack options window

3. Figure 3-4 shows you the different unpacking options.

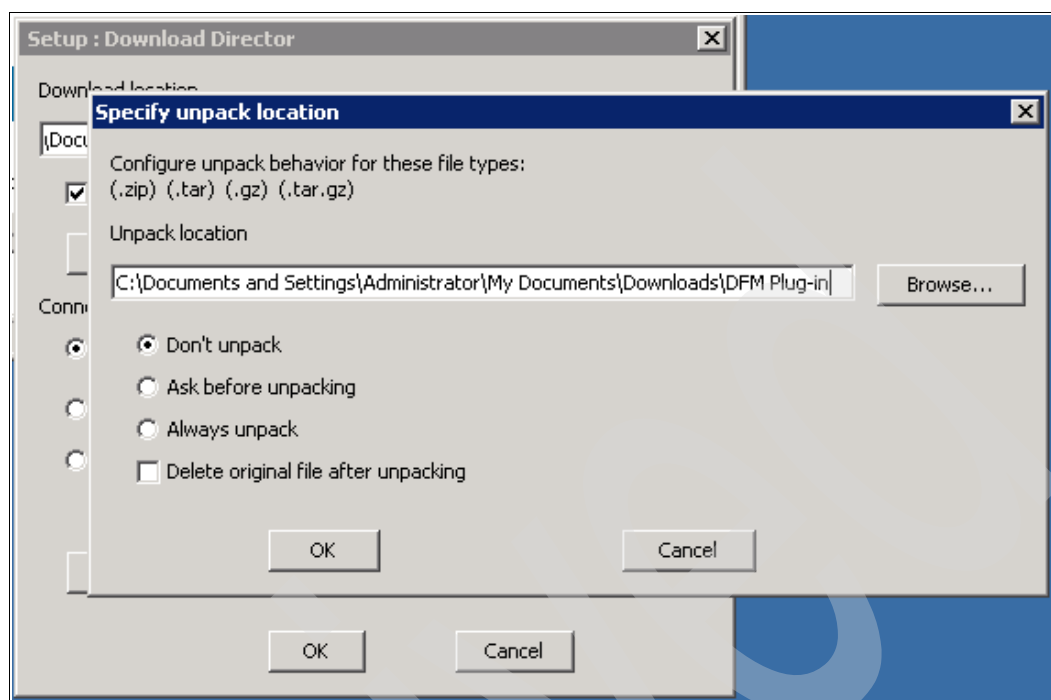


Figure 3-4 Unpacking options window

Once the file has been downloaded, you need to install the plug-in using the **dfm** command. The process to install the file is as follows:

a. From the DataFabric Manager Server command line, enter:

```
dfm plug-in add [-f] {path | URL }
```

You can omit the -f switch, unless you want DataFabric Manager Server to reinstall a plug-in that has already been installed.

b. Any path or URL that can be interpreted or resolved by your system and network is valid. For example, you can set paths as follows:

- **dfm plugin add http://web.company.com/~jsmith/filerconfig_Windows.zip**
- **dfm plugin add /server/jsmith/filerconfig_Windows.zip**
- **C:\Documents and Settings\Administrator>dfm plugin add C:\Documents and Settings\Administrator>DownloadDirector\filerconfig_Windows.zip**

c. To confirm that the plug-in is installed, from the DataFabric Manager Server command line, enter

```
dfm plugin list
```

For example:

```
C:\>dfm plugin list
```

Plugin Type	Version	Release	Description
Storage System Config 7.2.4		7.2.4	storage systems and vFilers
Storage System Config 7.3.1		7.3.1	storage system and vFilers

It is important to remember that anytime the DataFabric Manager Server is updated, the equivalent plug-in needs to be installed.

3.4 Installation

The installation of Operations Manager is quite simple even though the installation was done on a Windows 2003 64-bit operating system, Though the requirements indicate that Windows is supposed to be in WoW64 mode, we do not change the settings in the OS to make it install in that mode. In our environment, we use the dfmsetup-3-7-1-win32.exe file to run the installation instead of a CD. Download the file from <http://www-947.ibm.com/systems/support/supportsite.wss/supportresources?brandind=5000029&familyind=5329833&taskind=2> and place it in the directory of your choice.

The following steps show the installation procedure:

1. Double-click the dfmsetup-3-7-1-win32.exe file.
2. The initial installation window appears, as shown in Figure 3-5. Press **Next**.




Figure 3-5 Initial installation window

3. Press **I accept** on the Autosupport Notice window, as shown in Figure 3-6, and press **Next**.

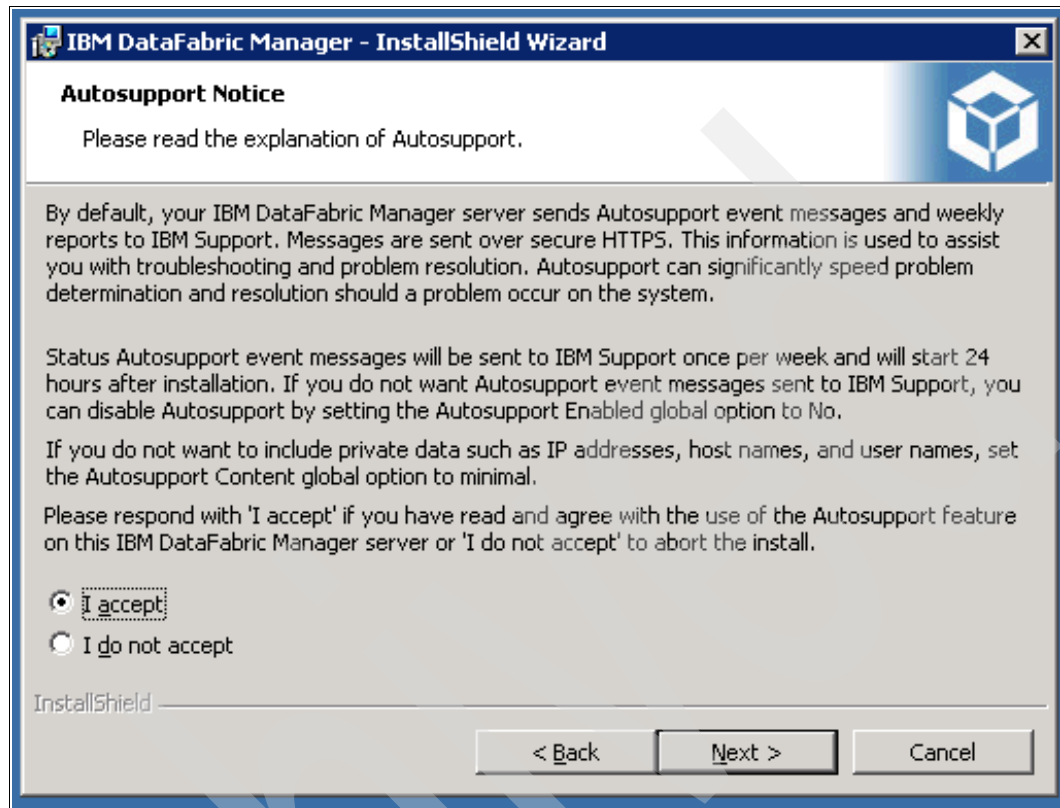
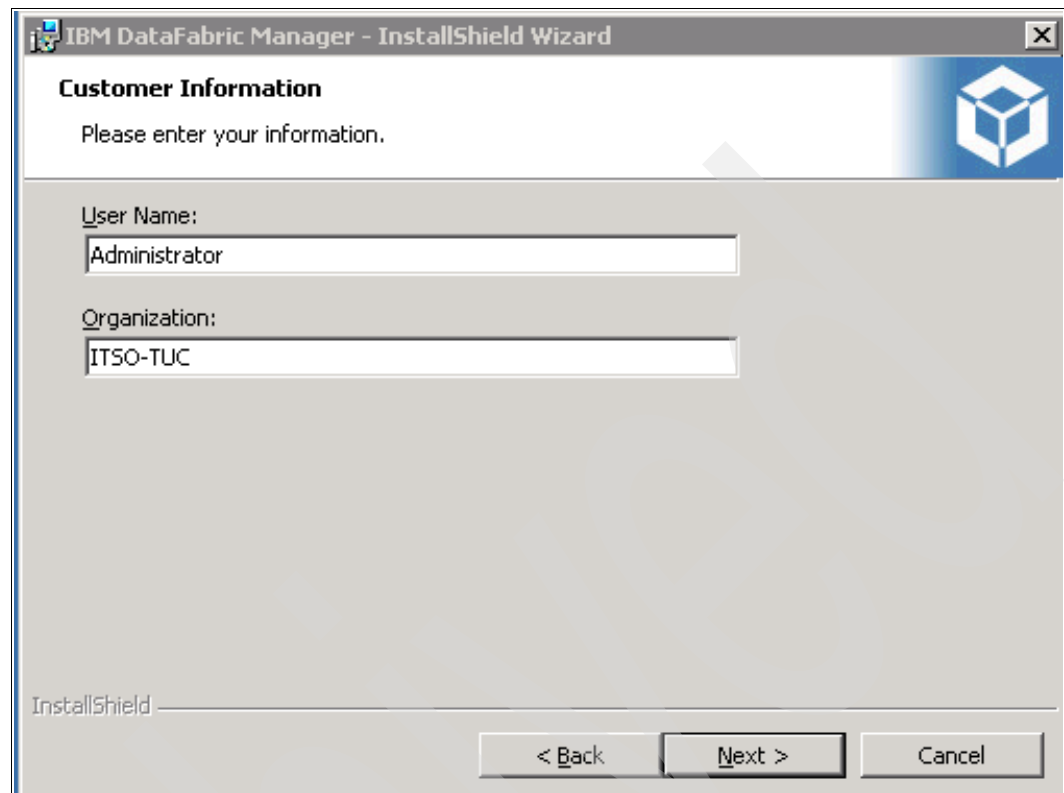


Figure 3-6 Autosupport notice

4. Enter the User Name and Organization and press **Next**, as shown in Figure 3-7.



The screenshot shows a Windows-style dialog box titled "IBM DataFabric Manager - InstallShield Wizard". The window has a standard title bar with a close button (X) in the top right corner. Below the title bar, the text "Customer Information" is displayed in a bold font, followed by the instruction "Please enter your information." in a smaller font. To the right of this text is a blue square icon containing a white cube. The main area of the window contains two text input fields. The first field is labeled "User Name:" and contains the text "Administrator". The second field is labeled "Organization:" and contains the text "ITSO-TUC". At the bottom of the window, there is a horizontal bar with the text "InstallShield" on the left. To the right of this bar are three buttons: "< Back", "Next >", and "Cancel". The "Next >" button is highlighted with a black border.

Figure 3-7 Customer Information

5. Enter your license key for Operations Manager and press Next, as shown in Figure 3-8.

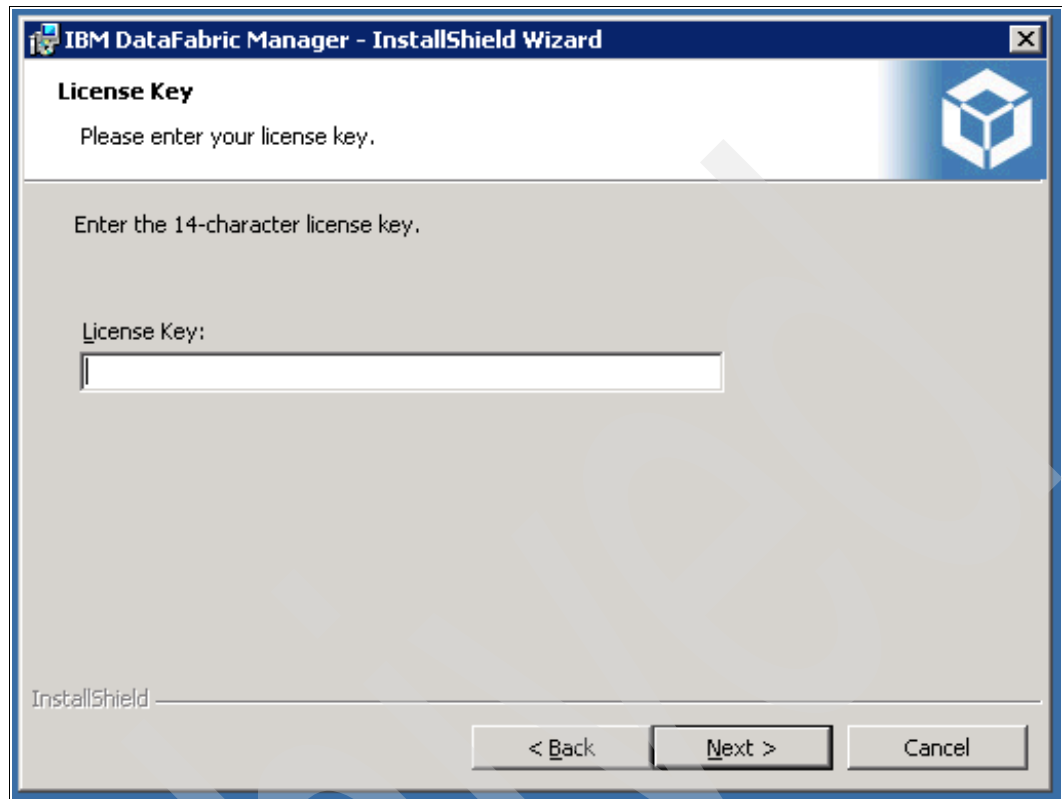


Figure 3-8 License Key

6. You have the option to change the destination directory. We use the default, as shown in Figure 3-9 on page 49.




Figure 3-9 Destination Folder

Press **Next** to continue.

7. Press **Install** to continue, as shown in Figure 3-10.




Figure 3-10 Begin the Installation

8. The program will now copy the Operations Manager files to the chosen directory, as shown in Figure 3-11.




Figure 3-11 Copying files

9. After copying the files, the installation will continue, as shown in Figure 3-12.




Figure 3-12 Installing help files

10. The program is now installed, as shown in Figure 3-13 on page 53. Press **Finish**.




Figure 3-13 Completion

11. Once the installation is finished, the computer needs to be restarted, as shown in Figure 3-14. However, this window may be covered by another window, as Operations Manager will begin automatically after you press **Finish**. Be aware that the system should be restarted first before using Operations Manager.




Figure 3-14 Restart window

12. Figure 3-15 shows the starting window for Operations Manager once you press **Finish**.

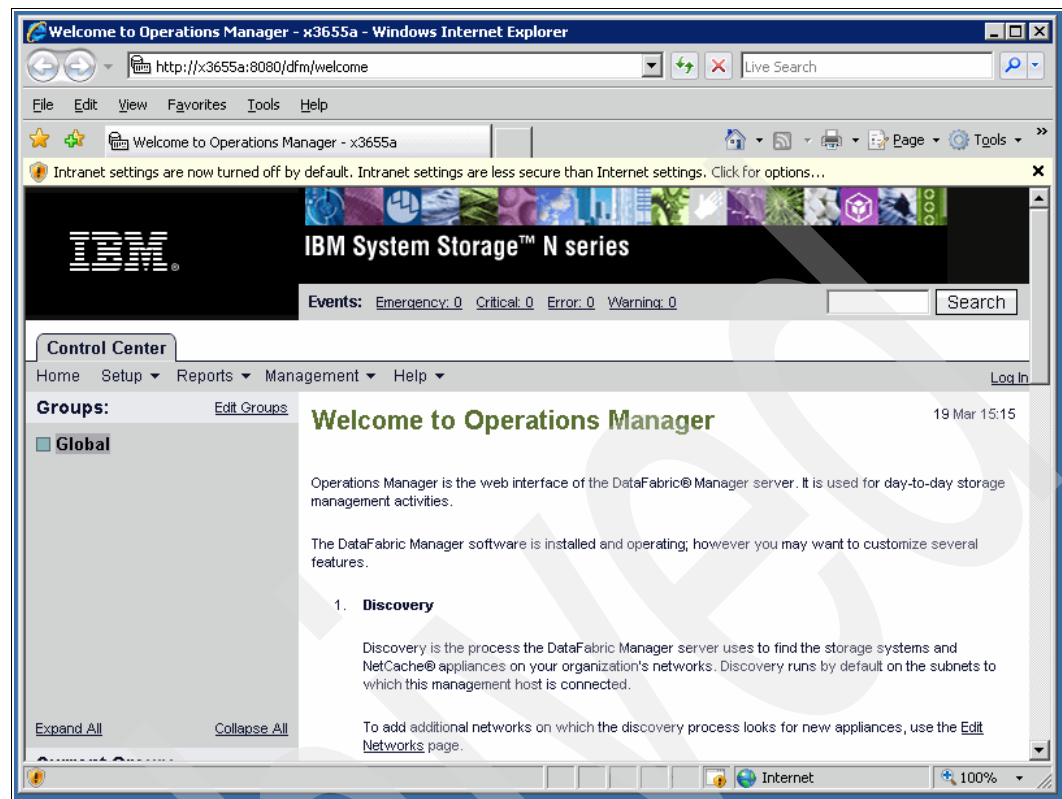


Figure 3-15 Operations Manager starting window

Note: When Windows 2003 is installed, the default security setting for the browser is “high”. This setting can cause the browser to block certain actions that can interfere with your Operations Manager operation. To ensure the best browsing experience when using Operations Manager, you might need to adjust your browser security setting to “medium”.

This concludes the steps necessary to install Operations Manager.



Installing Operations Manager: Linux

This chapter gives you a overview and detailed instructions about installing Operations Manager on Linux.

4.1 Overview

Operations Manager is a multipart, client-server platform. The core components of Operations Manager are shown in Figure 4-1. One Linux install file is used to install all these components. The requirements for the host running DFM Server are detailed in 4.2, “Host prerequisites” on page 56. We highly recommend that this host is set aside exclusively for DFM Server operations. Another computer, running either Microsoft Windows or Linux, should be used to connect to the DFM Server through a Web browser.

Operations Manager also has a dashboard facility called the *N series Management Console*. This console should also be installed on a computer other than the host designated for the DFM Server.

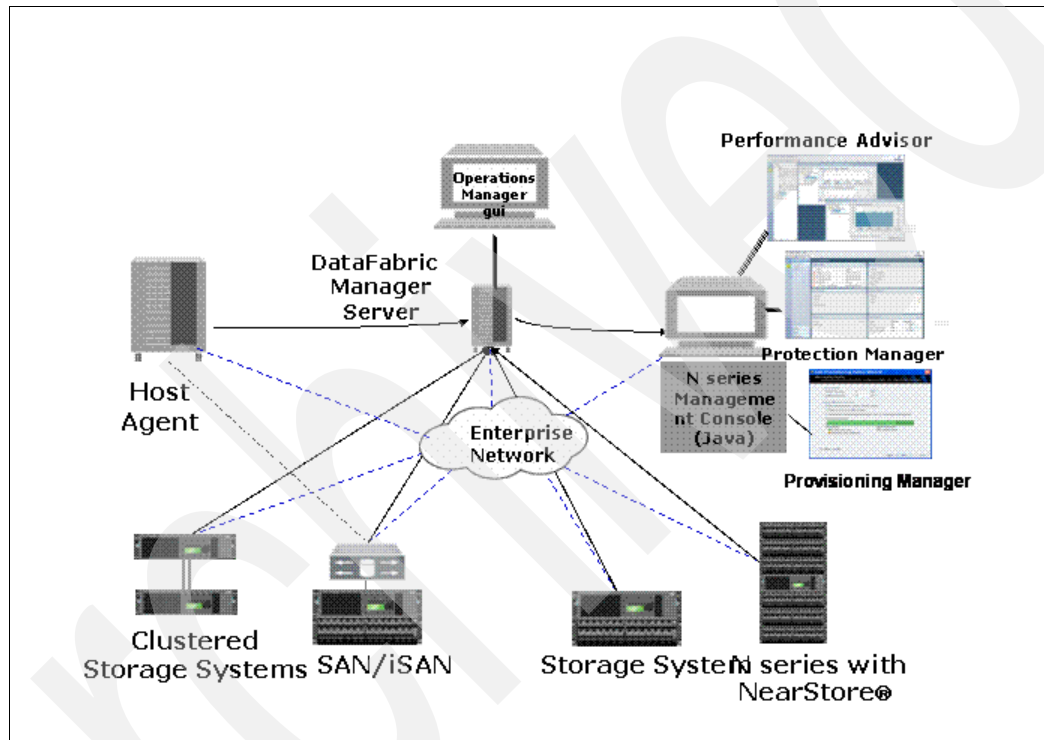


Figure 4-1 Overview of Operations Manager components

4.2 Host prerequisites

Table 4-1 on page 57 details the hardware and software prerequisites for installing IBM System Storage N series Operations Manager on a Linux System on physical hardware or in a VMware® ESX guest image.

Table 4-1 Installation requirements for Operations Manager on Linux hosts

Linux workstation or server	
Hardware requirements	Software requirements
Intel-based PC with single 2 GHz CPU (Xeon or Pentium 4) 4 GB of free disk space minimum, 8 GB recommended 1 GB of memory minimum	Red Hat Enterprise Linux AS 4 (Update 3 or later) for x86, 32-bit and 64-bit Red Hat Enterprise Linux Advanced Platform 5 for x86, 32-bit and 64-bit SUSE® Linux Enterprise Server 9 (Service Pack 2 or later) for x86, 32-bit and 64-bit SUSE Linux Enterprise Server 10 for x86, 32-bit and 64-bit
Linux servers on VMware ESX server 3.0.1 or later	
Hardware requirements	Software requirements
Intel-based PC with single 2 GHz CPU (Xeon or Pentium 4) 4 GB of free disk space minimum, 8 GB recommended 1 GB of memory minimum Single dedicated network interface	Red Hat Enterprise Linux AS 4 (Update 3 or later) for x86, 32-bit and 64-bit

Note: DataFabric Manager (DFM) Server V3.7 is not supported on Windows NT 4.0, Windows 2000, Windows XP, Solaris 8, or distributions of Linux not listed in the preceding table.

Operations Manager V3.7 does not support VMware VMotion® and VMware High Availability features.

These requirements are for a Operations Manager installation with only basic system monitoring enabled. If you enable additional features and monitor additional objects, a more powerful platform is probably required. Examples of objects and features that might require a more powerful platform include additional storage systems, qtrees, user quotas, and use of the Storage Resource Management, Performance Advisor, Business Continuity Option, Provisioning Manager, or Protection Manager features.

It is important that you run Operations Manager on a system that is running no other applications. Running other applications that will take away CPU, I/O, and memory bandwidth from Operations Manager may cause Operations Manager to be unable to carry out its tasks properly or reliably.

4.2.1 Upgrading from DataFabric Manager V3.5.1 or earlier

If you are upgrading from DataFabric Manager V3.5.1 or earlier to DataFabric Manager V3.6.1 or later, it takes a long time to upgrade the performance data files (there are 20 GB or more of data to install). The length of time depends on the platform used. The space used by the performance data files increases by about 65% during the upgrade.

4.2.2 License requirements

You must have a valid DataFabric Manager Server license key to complete the Operations Manager installation. You can access your license key at <http://www.ibm.com/storage/nas/>. After you have accessed the Web site, follow the options provided to access the license keys.

After you complete the installation, you can enter additional license keys on the Options window in Operations Manager. You can install (or upgrade to) Operations Manager V3.7 using the core license key.

However, if you do not have the core license key, you need the following licenses to monitor and manage your storage systems:

- ▶ DataFabric Manager Server license
- ▶ Additive license

DataFabric Manager Server license

The DataFabric Manager Server license is a server license with a unique serial number that tracks the number of Operations Manager installations. You must have this license to enable features. The node count is one.

Additive license

The additive license is an additional license with a unique serial number that is used to increase the node count and enable the features.

Table 4-2 shows the license requirements for each feature you might require

Table 4-2 License requirements

To Use	Install This product	That enables these features
Operations Manager	DataFabric Manager Server license Operations Manager license Note: Required for all licensed DFM installations. Sets the maximum number of storage systems that the DataFabric Manager Server can monitor in this installation.	<ul style="list-style-type: none">▶ Managing<ul style="list-style-type: none">– Storage system configuration– Scripts– NetCache® software and configuration▶ Monitoring and managing storage system clusters using Cluster Console▶ Displaying historical and real-time performance data using Performance Advisor in IBM N series Management Console
Protection Manager	<ul style="list-style-type: none">▶ DataFabric Manager Server license▶ Operations Manager license▶ Protection Manager license▶ Management Console	<ul style="list-style-type: none">▶ Automated policy-based data protection for NAS and SAN storage systems▶ SnapVault, Open Systems SnapVault, and SnapMirror management▶ Policy conformance checking and alerting monitoring<ul style="list-style-type: none">– Reports– Storage usage and availability, such as qtrees, volumes, aggregates, LUNs, and disks– Storage systems– vFiler units– NetCache appliances– Real-time streaming events

To Use	Install This product	That enables these features
Provisioning Manager	<ul style="list-style-type: none"> ▶ DataFabric Manager Server license ▶ Operations Manager license ▶ Provisioning Manager license ▶ Management Console 	<ul style="list-style-type: none"> ▶ Automated policy-based provisioning for SAN and NAS storage systems ▶ Space management policies and capacity reporting ▶ Policy conformance checking and alerting <p>If you have both Protection Manager and Provisioning Manager licensed, then the following features are enabled:</p> <ul style="list-style-type: none"> ▶ Assigning provisioning policies to nonprimary nodes ▶ Policy-based provisioning of primary storage ▶ Assigning protection policies to provisioned data sets
Protection Manager with Disaster Recovery	<ul style="list-style-type: none"> ▶ DataFabric Manager Server license ▶ Operations Manager license ▶ Protection Manager license ▶ Protection Manager Disaster Recovery license ▶ Management Console 	Failover and manual failback for NAS and SAN storage systems
File Storage Resource Manager (File SRM)	<ul style="list-style-type: none"> ▶ DataFabric Manager Server license ▶ Operations Manager license ▶ File SRM Option 	Tracking file system usage and capacity information through managed hosts running DFM Host agents

4.3 N series prerequisites

You must be running Data ONTAP Version 7.1 with DataFabric Manager V3.3.1 or later to manage IBM N series storage systems. IBM will not support earlier versions of Data ONTAP.

Note: You must have a DataFabric Manager plug-in for each version of Data ONTAP that you are running across your system. DataFabric Manager Server automatically includes the plug-ins for Data ONTAP. To list the versions of the plug-ins for Data ONTAP, use the **dfm plugin list** command at the DFM Server command line. You do not need to download a plug-in unless you are using a different version of Data ONTAP.

4.4 Installation

When you download Operations Manager, you need to make sure you have the Operations Manager install code, Host Agent code, and the relevant DFM plug-ins for the environment you plan to manage. In this chapter, we focus on setting up the DFM Server and the plug-ins.

The DFM Server utilizes plug-ins to communicate with and manage N series storage systems. These plug-ins are Data ONTAP version specific and must correctly match the Data ONTAP version running on the N series storage systems or gateways you plan to manage.

4.4.1 Locating the installation files

The latest version of the Operations Manager installation files for Linux can be downloaded from the IBM Systems Storage Support Web site at the following address:

<http://www-947.ibm.com/systems/support/supportsite.wss/supportresources?brandind=5000029&familyind=5329833&taskind=2>

You will need to log in with your IBM ID; access to the files are granted if you have the appropriate permissions. Contact your IBM Support representative or Business Partner to get access to this site if you experience any difficulty.

Note: It is very important that you download and install all the plug-ins necessary for the N series environment you manage. Operations Manager V3.7.1 comes packaged with plug-ins to support Data ONTAP versions up to and including V7.2.4.1. For our environment, we had to download plug-ins for Versions 7.2.6.1 and 7.3.1.

Note: If at any time you upgrade Data ONTAP, you must locate and download the correct plug-in that works fully with that version of Data ONTAP. While you may be able to run DFM Server and Operations Manager to query and manage the upgraded N series storage system without the upgraded plug-in, you run the risk of making invalid API calls to the kernel and possibly panic the Data ONTAP kernel. Before you call IBM support, you will need to have installed the correct plug-in version and proven that the problem still exists.

4.4.2 Installation options

You can install Operations Manager using the provided defaults or you can install using the `-d <directory>` switch to specify a different location. If you do specify a separate installation directory, no IBMdfm directory will be created.

4.4.3 Installation steps

As mentioned before, Operations Manager comprises the DataFabric Manager Server and the Operations Manager Web Application Interface, collectively called Operations Manager. The following steps are for installing Operations Manager on a Linux Red Hat server. Figure 4-2 on page 61 describes the files that will be used. In this chapter, we focus on installing Operations Manager only. Host Agent and N series console will be addressed in subsequent chapters.

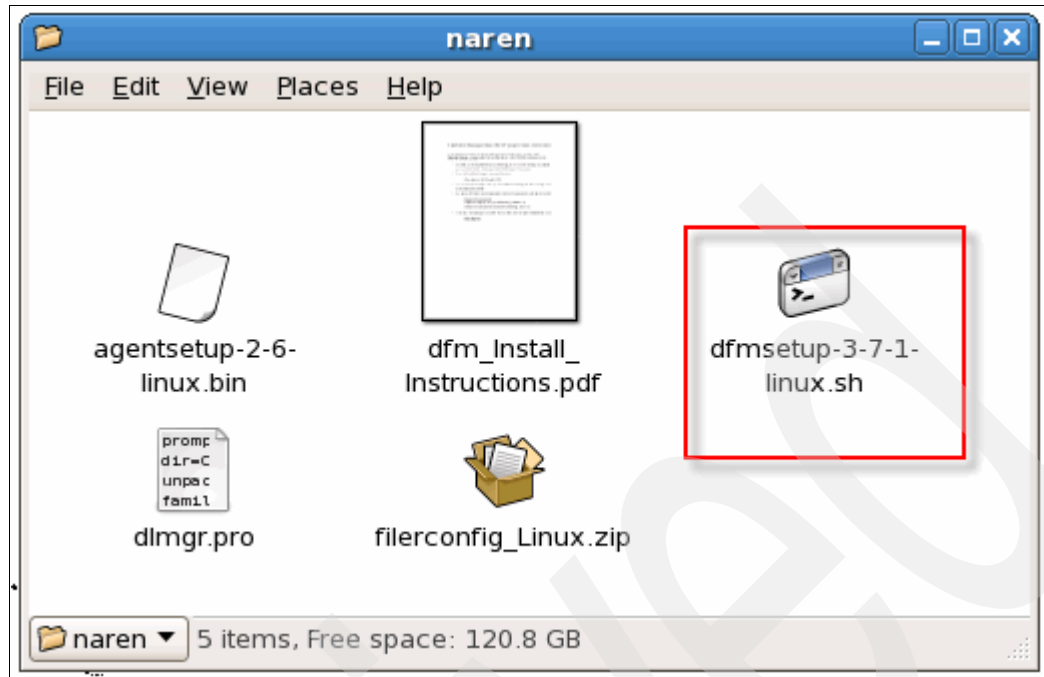


Figure 4-2 List of files for Linux installation

Installing DFM Server

From a terminal window, navigate to the folder where the installation files are located. In our case, we downloaded the files to a folder called /tmp/dfmInstaller.

A directory listing should show the files shown in Example 4-1.

Example 4-1 Directory listing of Operations Manager installation files

```
[root@vegemite dfmInstaller]#
[root@vegemite dfmInstaller]# ll
total 268660
-rwxr--r-- 1 root root 1650400 Mar 17 16:40 agentsetup-2-6-linux.bin
-rwxr--r-- 1 root root 14890 Mar 17 16:41 dfm_Install_Instructions.pdf
-rwxr--r-- 1 root root 259502985 Mar 17 16:39 dfmsetup-3-7-1-linux.sh
-rwxr--r-- 1 root root 52852 Mar 17 16:41 filerconfig_Linux.zip
[root@vegemite dfmInstaller]#
```

Review the dfm_Install_Instructions.pdf file prior to installing the software.

Running dfmsetup-3-7-1-linux.sh

As shown in Example 4-2, run the **dfmsetup-3-7-1-linux.sh** command from the command prompt. You will be prompted to confirm that you understand that DFM will want to send periodic (weekly) Autosupport and performance information to IBM support.

Note: You will also be asked to provide your Operations Manager license key.

Example 4-2 Execution window for dfmsetup

```
[root@vegemite dfmInstaller]# ./dfmsetup-3-7-1-linux.sh
```

```
Unpacking files needed for the installation ...
```

By default, the DataFabric Manager Server sends Autosupport event messages and weekly reports to IBM Support. Messages are sent over secure HTTPS. This information is used to assist you with troubleshooting and problem resolution. Autosupport will significantly speed problem determination and resolution should a problem occur on the system. Status Autosupport event messages will be sent to IBM once per week and will start 24 hours after installation. If you do not want Autosupport event messages sent to IBM, you can disable Autosupport by setting the Autosupport Enabled global option to "No". This option will initially be set to "Unknown" unless manually set, and will automatically become "Yes" after 24 hours if not set manually to "No". If you do not want to include private data such as IP addresses, host names, and user names, set the Autosupport Content global option to "minimal". Respond with "Yes" if you have read and agree with the use of the Autosupport feature on this DataFabric Manager Server or "No" to abort the installation. [y,n,?]: y

```
Enter your IBM DataFabric Manager license key [?,q]: rayehkrucoulic
```

```
Beginning the installation ...
```

```
Preparing... ##### [100%]
```

```
1:IBMdfm ##### [100%]
```

```
Installing wrappers in /usr/bin directory.
```

```
Installing scripts in /etc/init.d directory.
```

```
Configuring DataFabric Manager Server services.
```

```
Setting up sql ...
```

```
Database "/opt/IBMdfm/data/monitordb.db" created successfully
```

```
Starting SQL ...
```

```
Setting up DBA user ...
```

```
Setting up transaction log management...
```

```
Defining SQL schema ...
```

```
Defining SQL Views ...
```

```
Setting up DFM user ...
```

```
Creating or updating sample backup schedules and throttles.
```

```
Stopping SQL ...
```

```
Adaptive Server Anywhere Stop Engine Utility Version 9.0.2.3397
```

```
Enabled Operations Manager license.
```

```
Installing the online help.
```

```
Starting DataFabric Manager Server services.
```

```
Service: sql started.
```

```
Service: http started.
```

```
Service: eventd started.
```

```
Service: monitor started.
```

```
Service: scheduler started.
```

```
Service: server started.
```

```
Service: watchdog started.
```


Installing Filer plug-ins.

New plug-in files installed to /opt/IBMdfm/plugins/filer/6.5.1.
New plug-in files installed to /opt/IBMdfm/plugins/filer/6.5.2.
New plug-in files installed to /opt/IBMdfm/plugins/filer/6.5.3.
New plug-in files installed to /opt/IBMdfm/plugins/filer/6.5.4.
New plug-in files installed to /opt/IBMdfm/plugins/filer/6.5.5.
New plug-in files installed to /opt/IBMdfm/plugins/filer/6.5.6.
New plug-in files installed to /opt/IBMdfm/plugins/filer/6.5.7.
New plug-in files installed to /opt/IBMdfm/plugins/filer/7.0.
New plug-in files installed to /opt/IBMdfm/plugins/filer/7.0.1.
New plug-in files installed to /opt/IBMdfm/plugins/filer/7.0.2.
New plug-in files installed to /opt/IBMdfm/plugins/filer/7.0.3.
New plug-in files installed to /opt/IBMdfm/plugins/filer/7.0.4.
New plug-in files installed to /opt/IBMdfm/plugins/filer/7.0.5.
New plug-in files installed to /opt/IBMdfm/plugins/filer/7.0.6.
New plug-in files installed to /opt/IBMdfm/plugins/filer/7.1.
New plug-in files installed to /opt/IBMdfm/plugins/filer/7.1.1.
New plug-in files installed to /opt/IBMdfm/plugins/filer/7.1.2.
New plug-in files installed to /opt/IBMdfm/plugins/filer/7.2.
New plug-in files installed to /opt/IBMdfm/plugins/filer/7.2.1.
New plug-in files installed to /opt/IBMdfm/plugins/filer/7.2.2.
New plug-in files installed to /opt/IBMdfm/plugins/filer/7.2.3.
New plug-in files installed to /opt/IBMdfm/plugins/filer/7.2.4.

You may now point your browser to access Operations Manager at

<http://vegemite.tucson.ibm.com:8080/dfm/welcome>

for an overview of product features, and for a list of steps to take next.

For help on the dfm command, enter

dfm help

```
[root@vegemite dfmInstaller]#  
[root@vegemite dfmInstaller]#
```

After the installation of Operations Manager finishes, you can use your web browser to go to the provided link. If you are logged on locally to the installation server, you may use the link provided; otherwise, you may need to substitute the name `localhost` for the proper DNS or IP address of the DFM Server.

You will be asked to log in. Here you can provide the appropriate credentials, as shown in Figure 4-3, to set up the initial installation and create additional users/roles as necessary.

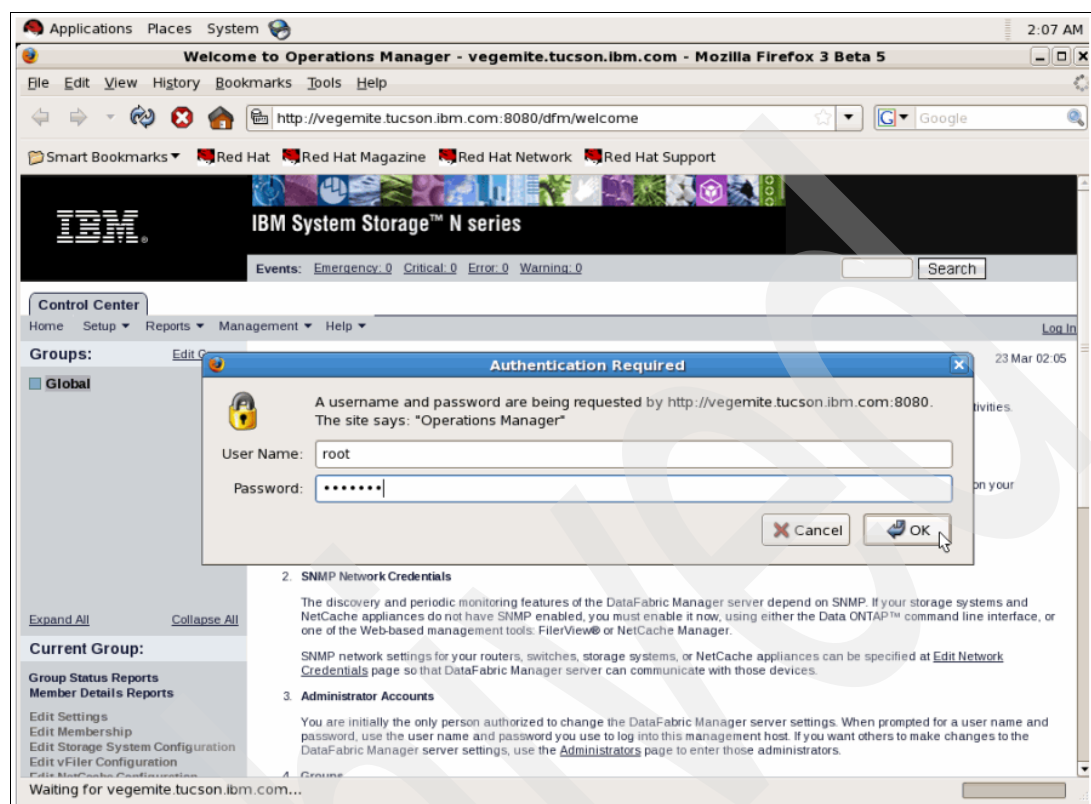


Figure 4-3 Login to Operations Manager

At this point, you have successfully installed DFM Server, but still need to install the plug-ins relevant to your environment. Notice that in the previous example the DFM Server package only had plug-ins up to Data ONTAP V7.2.4.

Installing the DFM plug-ins

On the IBM System Storage support Web site, we located plug-ins for Version 7.2.6.1 and Version 7.1.3. We found that there was a problem with the download manager mechanism when attempting to download plug-ins and had to revert to http download for each file.

Note: The file names of the plug-ins are the same for each version. While you are able to select multiple versions of the plug-ins on the download window using the download manager, it is not possible to successfully download the files, as you would end up with each version file in the download batch overwriting the previous file in the download directory on your computer.

For downloading of plug-ins, use the *http://*based download mechanism instead.

Example 4-3 on page 65 shows the plug-ins we downloaded, highlighted in bold. The output of the screen has been modified for brevity.

Example 4-3 Plug-ins listing after downloading

```
[root@vegemite dfmInstaller]#  
[root@vegemite dfmInstaller]# ll  
total 268848  
-rwxr--r-- 1 root root 1650400 Mar 17 16:40 agentsetup-2-6-linux.bin  
-rwxr--r-- 1 root root 14890 Mar 17 16:41 dfm_Install_Instructions.pdf  
-rwxr--r-- 1 root root 259502985 Mar 17 16:39 dfmsetup-3-7-1-linux.sh  
-rwxr--r-- 1 root root 52852 Mar 26 19:00 filerconfig_DFMDDataONTAP7.1.3plug-in.zip  
-rwxr--r-- 1 root root 55483 Mar 26 18:59 filerconfig_DFMDDataONTAP7.2.6.1plug-in.zip  
-rwxr--r-- 1 root root 53605 Mar 26 18:58 filerconfig_DFMDDataONTAP7.3.1plug-in.zip  
-rwxr--r-- 1 root root 1770 Mar 17 16:41 JDateChooserWrapper.java  
-rwxr--r-- 1 root root 6853589 Mar 17 16:41 jfreechart-1.0.4.zip  
-rwxr--r-- 1 root root 1528283 Mar 17 16:41 jpf-bin-0.10.zip  
-rwxr--r-- 1 root root 552960 Mar 17 16:41 pthreads-000813.tar  
-rwxr--r-- 1 root root 3610 Mar 17 16:41 SlidingPanelDemo.java  
-rwxr--r-- 1 root root 4589449 Mar 17 16:41 swinglabs-0.8.0-bin.zip  
[root@vegemite dfmInstaller]#
```

Once you have downloaded the relevant plug-ins to your folder on the Linux DFM Server, you can run the **dfm** command to install the plug-in. The **dfm** command is function rich, and in our case we are interested in the **dfm plugins** part of the command, as shown in Example 4-4.

Example 4-4 dfm plugins command

```
[root@vegemite dfmInstaller]#  
[root@vegemite dfmInstaller]# dfm plugins help
```

NAME

plugin -- install and manage plugins for configuration management

SYNOPSIS

```
dfm plugin list [ -t type ] [ path ]  
dfm plugin add [ -f ] { path | url }  
dfm plugin delete [ -t type ] { all | <versions> ... }
```

DESCRIPTION

The plugin command manages the storage system and NetCache plug-ins.

```
[root@vegemite dfmInstaller]#
```

Example 4-5 shows the install sequence we went through. Commands are highlighted in bold and important responses from the application immediately follow those commands in bold. Note also what happens when we attempt to install a plug-in that has already been installed.

Example 4-5 Adding plug-ins to DFM Server

```
[root@vegemite dfmInstaller]#  
[root@vegemite dfmInstaller]# dfm plugin list
```

Plug-in	Type	Version	Release	Description
Storage System Config	6.5.1	6.5.1	storage systems and vFilers	
Storage System Config	6.5.2	6.5.2	storage systems and vFilers	
Storage System Config	6.5.3	6.5.3	storage systems and vFilers	
Storage System Config	6.5.4	6.5.4	storage systems and vFilers	
Storage System Config	6.5.5	6.5.5	storage systems and vFilers	
Storage System Config	6.5.6	6.5.6	storage systems and vFilers	
Storage System Config	6.5.7	6.5.7	storage systems and vFilers	

```

Storage System Config 7.0      7.0.0.1      storage systems and vFilers
Storage System Config 7.0.1    7.0.1.1      storage systems and vFilers
Storage System Config 7.0.2    7.0.2        storage systems and vFilers
Storage System Config 7.0.3    7.0.3        storage systems and vFilers
Storage System Config 7.0.4    7.0.4        storage systems and vFilers
Storage System Config 7.0.5    7.0.5        storage systems and vFilers
Storage System Config 7.0.6    7.0.6        storage systems and vFilers
Storage System Config 7.1      7.1.0.1      storage systems and vFilers
Storage System Config 7.1.1    7.1.1.1      storage systems and vFilers
Storage System Config 7.1.2    7.1.2.1      storage systems and vFilers
Storage System Config 7.1.3    7.1.3        storage systems and vFilers
Storage System Config 7.2      7.2          storage systems and vFilers
Storage System Config 7.2.1    7.2.1.1      storage systems and vFilers
Storage System Config 7.2.2    7.2.2        storage systems and vFilers
Storage System Config 7.2.3    7.2.3        storage systems and vFilers
Storage System Config 7.2.4    7.2.4        storage systems and vFilers
[root@vegemite dfmInstaller]#
[root@vegemite dfmInstaller]#
[root@vegemite dfmInstaller]#
[root@vegemite dfmInstaller]# dfm plugin add filerconfig_DFMDDataONTAP7.2.6.1plug-in.zip
New plugin files installed to /opt/IBMdfm/plugins/filer/7.2.6.
[root@vegemite dfmInstaller]#
[root@vegemite dfmInstaller]# dfm plugin add filerconfig_DFMDDataONTAP7.3.1plug-in.zip
New plugin files installed to /opt/IBMdfm/plugins/filer/7.3.1.
[root@vegemite dfmInstaller]#
[root@vegemite dfmInstaller]#
[root@vegemite dfmInstaller]# dfm plugin add filerconfig_DFMDDataONTAP7.1.3plug-in.zip
Error: The existing installation of plugin version 7.1.3 is identical
to the one you are trying to install.
Use the '-f' flag to force the overwrite.
[root@vegemite dfmInstaller]#

```

The -f flag can be used to reinstall the plug-in. This is appropriate if your plug-in is corrupted or has a bug in the code. The alternative to replacing a plug-in is to first delete the plug-in and then add it in again.

After installing the plug-ins, we have the plug-ins lists shown in Example 4-6.

Example 4-6 Plug-ins list

```

[root@vegemite dfmInstaller]# dfm plugins list

```

Plugin Type	Version	Release	Description
Storage System Config	6.5.1	6.5.1	storage systems and vFilers
Storage System Config	6.5.2	6.5.2	storage systems and vFilers
Storage System Config	6.5.3	6.5.3	storage systems and vFilers
Storage System Config	6.5.4	6.5.4	storage systems and vFilers
Storage System Config	6.5.5	6.5.5	storage systems and vFilers
Storage System Config	6.5.6	6.5.6	storage systems and vFilers
Storage System Config	6.5.7	6.5.7	storage systems and vFilers
Storage System Config	7.0	7.0.0.1	storage systems and vFilers
Storage System Config	7.0.1	7.0.1.1	storage systems and vFilers
Storage System Config	7.0.2	7.0.2	storage systems and vFilers
Storage System Config	7.0.3	7.0.3	storage systems and vFilers
Storage System Config	7.0.4	7.0.4	storage systems and vFilers
Storage System Config	7.0.5	7.0.5	storage systems and vFilers

Storage System Config 7.0.6	7.0.6	storage systems and vFilers
Storage System Config 7.1	7.1.0.1	storage systems and vFilers
Storage System Config 7.1.1	7.1.1.1	storage systems and vFilers
Storage System Config 7.1.2	7.1.2.1	storage systems and vFilers
Storage System Config 7.1.3	7.1.3	storage systems and vFilers
Storage System Config 7.2	7.2	storage systems and vFilers
Storage System Config 7.2.1	7.2.1.1	storage systems and vFilers
Storage System Config 7.2.2	7.2.2	storage systems and vFilers
Storage System Config 7.2.3	7.2.3	storage systems and vFilers
Storage System Config 7.2.4	7.2.4	storage systems and vFilers
Storage System Config 7.2.6	7.2.6.1	storage systems and vFilers
Storage System Config 7.3.1	7.3.1	storage systems and vFilers

```
[root@vegemite dfmInstaller]#
[root@vegemite dfmInstaller]#
```

Reviewing Operations Manager through a Web browser

Once you have installed and verified the plug-ins, you can use your Web browser, either from your Linux host, another Linux host, or a Windows desktop, to connect to the DFM Server. The DFM Server runs the Apache Tomcat Web server on port 8080 and you can get to it using the following address:

```
http://<dfm sever>:8080/dfm/welcome
```

In our case, the address was:

```
http://vegemite.tucson.ibm.com:8080/dfm/welcome
```

as shown in Figure 4-3 on page 64. You will be prompted to log in with root privileges.

The main window shown in Figure 4-4 is an introduction window and provides you with an overview of what you need to do to configure the DFM Server. This will be covered in more detail in Chapter 9, “Configuring Operations Manager” on page 127.

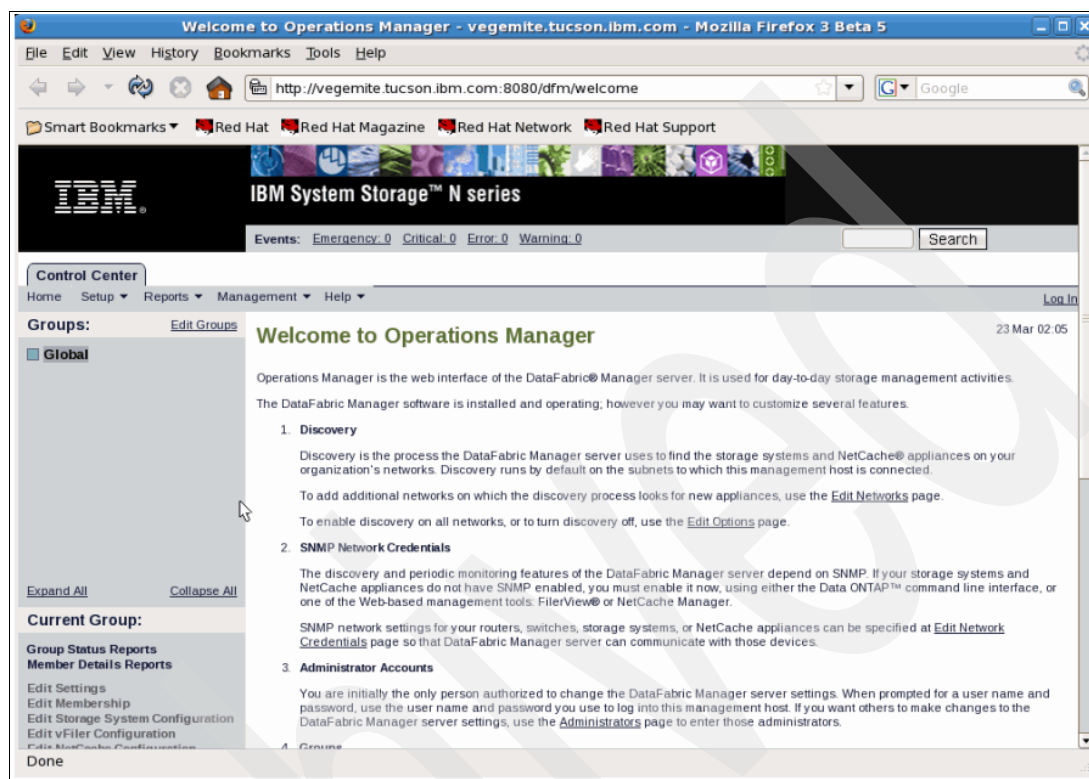


Figure 4-4 Operations Manager overview

Operations Manager provides a quick graphical view into the current health of your N series storage system. An example is shown in Figure 4-5 on page 69, where we see the Events tab with the heading “Information or Worse” and a list of warnings and errors that require immediate attention. These warnings probably would not be noticed without diligent manual monitoring if Operations Manager is not used.

As you might have observed, we have not yet manually entered the IP addresses of any N series storage arrays, yet Operations Manager is able to provide some information. This is because Operations Manager issues a broadcast or multicast to locate N series storage arrays in the network.

The information provided here is what can be obtained through a read only query of the N series arrays Operations Manager located. For more detailed management of each N series storage array, you need to provide Operations Manager with the root (or equivalent) credentials of each storage system you want it to manage. This will be described further in Chapter 9, “Configuring Operations Manager” on page 127.

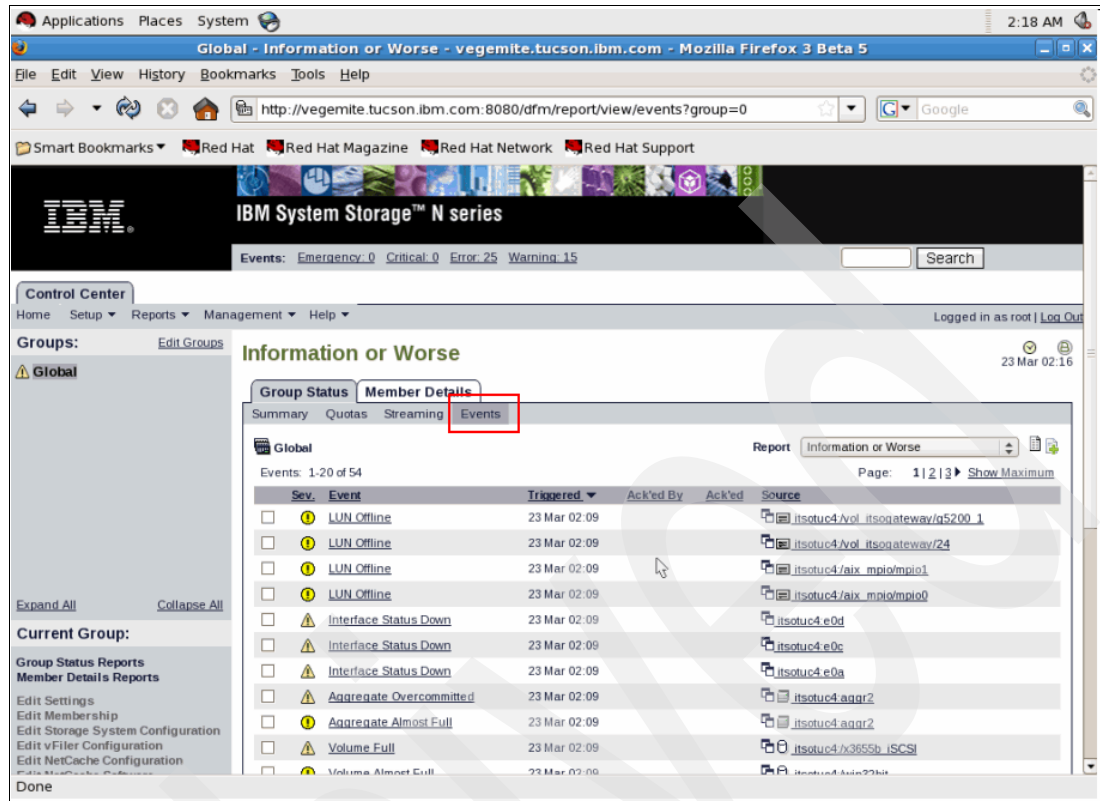


Figure 4-5 Viewing events within Operations Manager

Installing the Sun Java Runtime Environment plug-in for Firefox

If you have just commissioned a new Linux server to serve as the DFM Server, it is possible that not all plug-ins have yet been installed for your Web browser. You will require a current version of Sun™ Java™ Runtime Environment (Sun JRE™) at a minimum.

If you are unable to get your browser to automatically install the plug-in, you may have to manually install Sun Java Runtime Environment and then update the Web browser's plugins folder.

Here is an example of the steps we took to install Sun JRE on Mozilla Firefox on our Red Hat DFM Server:

1. Figure 4-6 shows that there is a missing plug-in for Firefox.

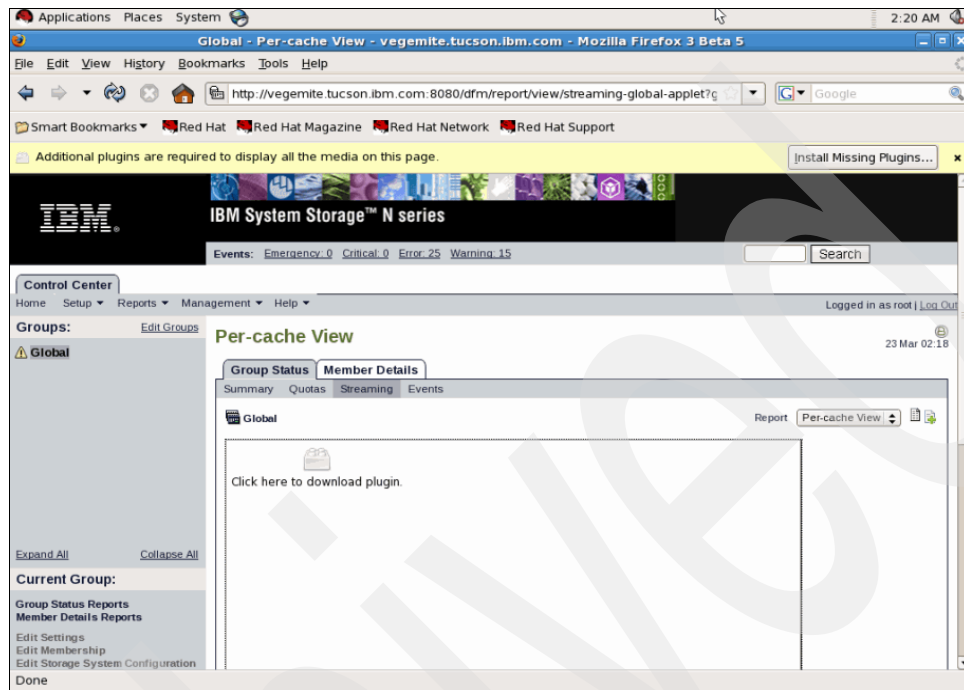


Figure 4-6 Missing plug-in for Firefox

2. Figure 4-7 on page 72 shows that Firefox is searching for the missing plug-in.




Figure 4-7 Firefox is searching for the Sun JRE plug-in

3. Figure 4-8 shows that Firefox has found the missing plug-in, but the installation fails, so we have to switch to manual install.




Figure 4-8 Sun JRE plug-in installation fails

4. Figure 4-9 on page 73 shows the download of the Sun JRE plug-in from the Sun Web site.

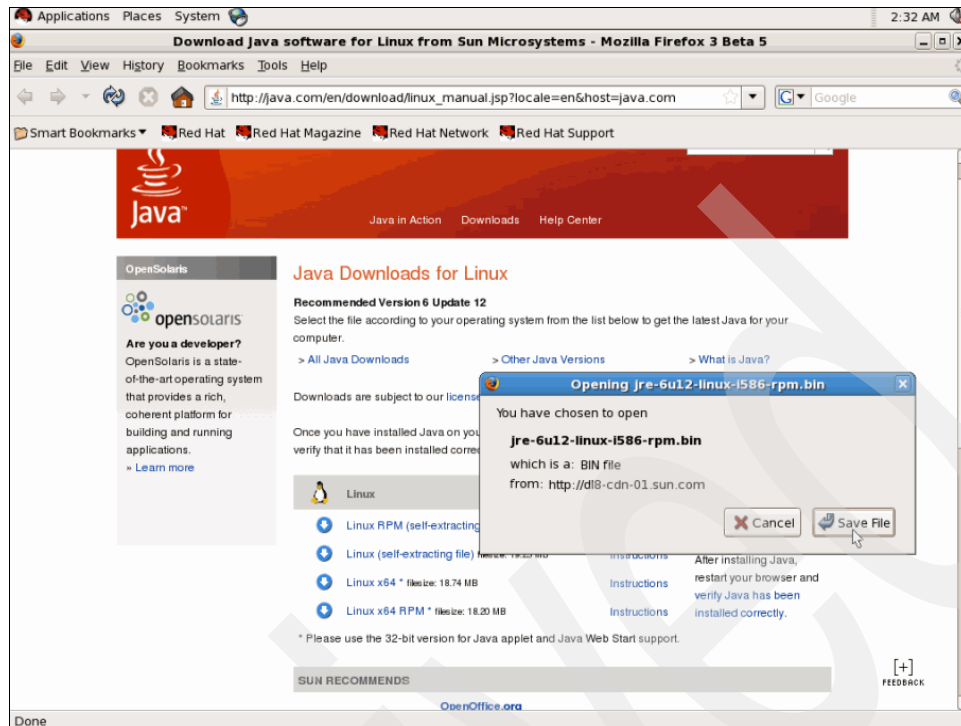


Figure 4-9 Manually download Sun JRE from the SUN Web site

5. For the manual install, we save the file to /usr/java, then, as root, change the bin file's mode to executable through the **chmod** command, as shown in Example 4-7.

Example 4-7 Changing the file to executable

```
[root@localhost java]# ll
total 19204
-rw-r--r-- 1 root root 19635520 Mar 25 14:55 jre-6u12-linux-i586-rpm.bin
[root@localhost java]# chmod a+x jre-6u12-linux-i586-rpm.bin
[root@localhost java]# ll
total 19204
-rwxr-xr-x 1 root root 19635520 Mar 25 14:55 jre-6u12-linux-i586-rpm.bin
[root@localhost java]#
```

6. We then execute the file to install the Sun JRE, as shown in Example 4-8.

Example 4-8 Executing and installing Sun JRE from the Linux command line

```
[root@localhost java]# ll
total 19204
-rwxr-xr-x 1 root root 19635520 Mar 25 14:55 jre-6u13-linux-i586-rpm.bin
[root@localhost java]# ./jre-6u13-linux-i586-rpm.bin
Sun Microsystems, Inc. Binary Code License Agreement
```

for the JAVA SE RUNTIME ENVIRONMENT (JRE) VERSION 6 and
JAVAFX RUNTIME VERSION 1

SUN MICROSYSTEMS, INC. ("SUN") IS WILLING TO LICENSE THE
SOFTWARE IDENTIFIED BELOW TO YOU ONLY UPON THE CONDITION
THAT YOU ACCEPT ALL OF THE TERMS CONTAINED IN THIS BINARY
CODE LICENSE AGREEMENT AND SUPPLEMENTAL LICENSE TERMS

(COLLECTIVELY "AGREEMENT"). READ THE AGREEMENT CAREFULLY. BY DOWNLOADING OR INSTALLING THIS SOFTWARE, YOU ACCEPT THE TERMS OF THE AGREEMENT. INDICATE ACCEPTANCE BY SELECTING THE "ACCEPT" BUTTON AT THE BOTTOM OF THE AGREEMENT. IF YOU ARE NOT WILLING TO BE BOUND BY ALL THE TERMS, SELECT THE "DECLINE" BUTTON AT THE BOTTOM OF THE AGREEMENT AND THE DOWNLOAD OR INSTALL PROCESS WILL NOT CONTINUE.

1. DEFINITIONS. "Software" means the identified above in binary form, any other machine readable materials (including, but not limited to, libraries, source files, header files, and data files), any updates or error corrections provided by Sun,

...

<< This part of the license agreement truncated for brevity >>

....

For inquiries contact: Sun Microsystems, Inc., 4150
Network Circle, Santa Clara, California 95054, U.S.A.

Do you agree to the above license terms? [yes or no]

yes

Unpacking...

Checksumming...

Extracting...

UnZipSFX 5.50 of 17 February 2002, by Info-ZIP (Zip-Bugs@lists.wku.edu).

inflating: jre-6u13-linux-i586.rpm

UnZipSFX 5.50 of 17 February 2002, by Info-ZIP (Zip-Bugs@lists.wku.edu).

inflating: jre-6u13-linux-i586.rpm

Preparing...

[100%]

1:jre

[100%]

Unpacking JAR files...

rt.jar...

jsse.jar...

charsets.jar...

localedata.jar...

plugin.jar...

javaws.jar...

deploy.jar...

Done.

[root@localhost java]#

-
7. The final step is to go to the Firefox plugins folder and create a link to the Sun JRE plug-in, as shown in Example 4-9 on page 75.

Example 4-9 Linking the Sun JRE plug-in

```
[root@localhost java]# cd /usr/lib/mozilla/plugins
[root@localhost plugins]# ll /usr/java
[root@localhost plugins]# ln -s /usr/java/latest/plugin/i386/ns7/libjavaplugin_oji.so
[root@localhost plugins]# ll
total 4
lrwxrwxrwx 1 root root 53 Mar 25 15:23 libjavaplugin_oji.so ->
/usr/java/latest/plugin/i386/ns7/libjavaplugin_oji.so
[root@localhost plugins]#
```

We note that there are also other ways to install the plug-in, depending on the version of Mozilla or Firefox being used.

4.4.4 Switching off Autosupport notifications for Operations Manager

The Operations Manager installation and upgrade process automatically installs the Autosupport feature with Autosupport enabled and displays a message about how to disable the feature. To disable Autosupport, use the following command:

```
dfm option set autosupportEnabled=no
```

Note: While switching Autosupport off is appropriate in our lab, we strongly advise that it is left on for your production environment, as it provides IBM Support with up-to-date information of the health and state of your Operations Manager environment and N series servers being managed.

At this point, you have installed the DFM Server and its plug-ins successfully and have connected to the management interface through a Web browser. In the next chapter, we look at installing the N series Management Console, which is a dashboard application and is typically installed on a desktop or management console, as well as the steps for installing Operations Manager agents on Linux hosts. The process for installing Operations Manager agents on Windows hosts is covered in Chapter 5, “Host Agent installation for Windows 2003” on page 79.

Archived



Part 2

Host Agent installation

In this part, we discuss Host Agent installation for Windows and Linux.

Archived

Host Agent installation for Windows 2003

You can use the IBM System Storage N series Operations Manager Host Agent software to monitor hosts through Operations Manager and the Operations Manager interface. This chapter describes how to install this software on both a Windows 2003 64-bit and 32-bit host.

This chapter emphasizes the preparations and product installation that you should carry out on Windows 2003 64-bit hosts to enable their monitoring by the DataFabric Manager Server.

5.1 Host prerequisites

Operations Manager Host Agent is application software that resides on a Windows or Linux host. In this chapter, we discuss only the Windows 64- or 32-bit host. The Host Agent collects information, such as OS name, version, HBA information, and file system meta data, and then sends that information back to the DataFabric Manager Server. Users can create reports of the collected information by using Operations Manager or the DataFabric Manager Server CLI.

To enable a target host to communicate with the DataFabric Manager Server, install and configure the Operations Manager Host Agent software on that host. After the DataFabric Manager Server discovers that instance of Operations Manager Host Agent, no further configuration is required.

Operations Manager Host Agent does not initiate any management actions on the Windows or Linux host. It is strictly a passive agent. It acts only on requests from external management applications, such as the DataFabric Manager Server.

You need Operations Manager Host Agent *only* if you want to monitor SAN hosts or FSRM-generated file system data through Operations Manager.

Operations Manager Host Agent uses a Web-based interface for configuration. You can access it either from the computer on which Operations Manager Host Agent is installed or from any other computer in the network. To access the Operations Manager Host Agent remotely, be sure to precede the IP address with `http://`.

5.2 What Operations Manager Host Agent can do

After you install Operations Manager Host Agent on a non-IBM host, you can use Operations Manager to perform a variety of SAN and FSRM functions.

- ▶ **SAN capabilities:** Using Operations Manager Host Agent and Operations Manager, you can perform the following SAN Tasks:
 - Monitor basic system information for the SAN hosts.
 - View detailed HBA and LUN information.
- ▶ **FSRM capabilities:** Using Operations Manager Host Agent and Operation manager, you can perform the following FSRM tasks:
 - Collect storage usage data at the file and directory level (see Figure 5-1 on page 81).
 - Identify a variety of file-related information, for example, largest files, oldest files, or space consumed per file type.

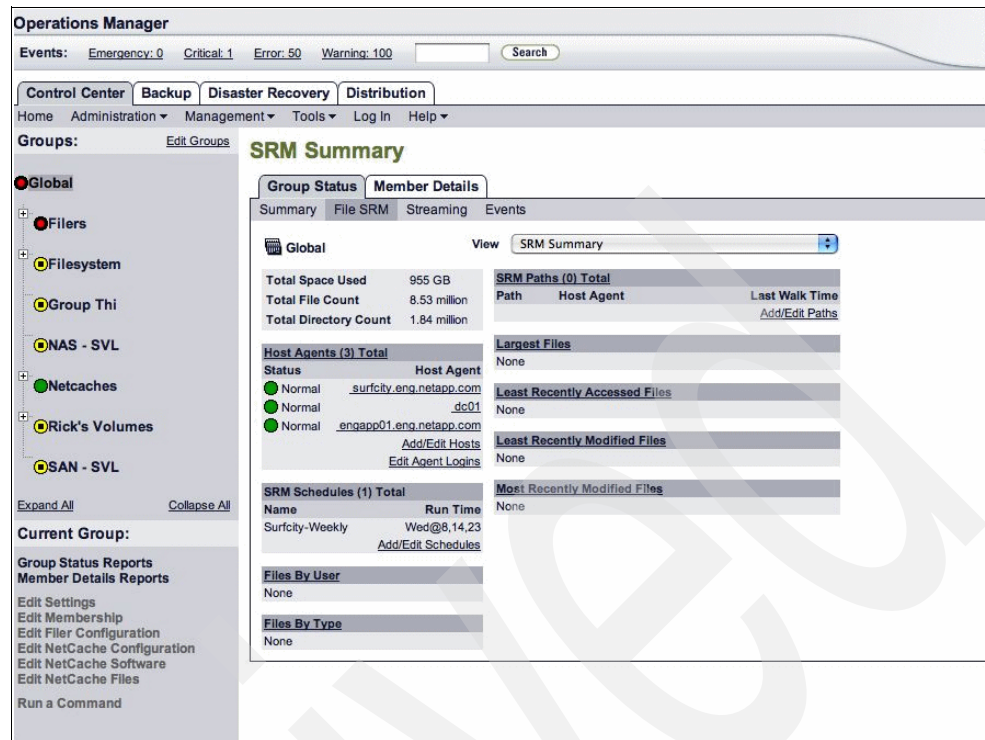


Figure 5-1 Operations Manager Host Agent

5.3 N series prerequisites

In order to access the SAN capabilities, you need to load the Host Agent on the server to be monitored. No other configuration needs to be done.

On the other hand, to be able to use the features gained from FSRM, you must do a great deal of configuration.

You must meet the following prerequisites to use the FSRM feature:

- ▶ You must have a valid FSRM license installed on your DataFabric Manager Server. Select **Setup** → **Discovery** → **Licensed Features** → **New License Key** to see if you have the license installed.

Note: The Quotas subtab is visible in the Operations Manager user interface (under **Control Center** → **Home** → **Group Status**) until you install the File SRM license. After you install the license, the Quotas subtab is renamed “File SRM,” and all of the FSRM features become visible when you click it.

- ▶ All hosts to be managed through Operations Manager must be connected to a TCP/IP network either known to or discoverable by the DataFabric Manager Server. The hosts must be connected to the network through an Ethernet port and must have a valid IP address.
- ▶ All directory paths to be monitored must be visible to the host agent. For example, to enable FSRM host monitoring by the DataFabric Manager Server, the host agent must mount a storage system share using NFS or CIFS, or the host agent must use a LUN on the storage system.

- Before setting up FSRM paths and schedules, you must enable administrative access to your host agents.

You must enable administrative access to your host agents before you can perform FSRM tasks with them. To enable administrative access, you must ensure that the password specified on the DataFabric Manager Server matches those set in the Operations Manager Host Agent software. Table 5-1 describes the options that must be set to enable administrative access.

Table 5-1 Enable management access

Access type	Operations Manager options	Operations Manager host agent software
Monitoring only	Host Agent Login=guest Host Agent Monitoring Password	Monitoring API Password
Management	Management Host Agent Login=admin Host Agent Management Password	Management API Password

5.4 Overview of configuration steps

To begin gathering file level information, you need to perform the following tasks:

1. Identify FSRM host agents.

If you have installed an FSRM license, Operations Manager will automatically discover all host agents.

2. Add new host agents manually in Operations Manager, if they have not been discovered.
3. Set up Host Agent administrative access on the hosts to be monitored.

You must enable administrative access to your host agents before you can perform FSRM tasks with them.

4. Verify Operations Manager Host Agent software administrative access.
5. Add paths in Operations Manager.

After host agents have been discovered, you must define paths to them.

When you use the Operations Manager GUI, you see that the host name is already provided for you. Therefore, you only need to provide the path name in a shortened UNC format as follows:

`\\sharename\path\filename`

Another example might be:

`\\ipaddress\path\filename`

In our example (Figure 5-2 on page 83), we did not specify a file name.

Figure 5-2 SRM Path example

You will notice that we did not need to include the host name in the path. (It was already included in the SRM Host window.)

6. Set up path-walk schedules in Operations Manager.

5.5 N series installation

Installation of the Operations Manager Host Agent is very easy. Our installation was on a Windows 2003 64-bit server. Currently, the Host Agent is a 32-bit application. There is no 64-bit agent as of the writing of this publication. We downloaded the agent and started the installation by double-clicking the Host Agent file, as shown in Figure 5-3.

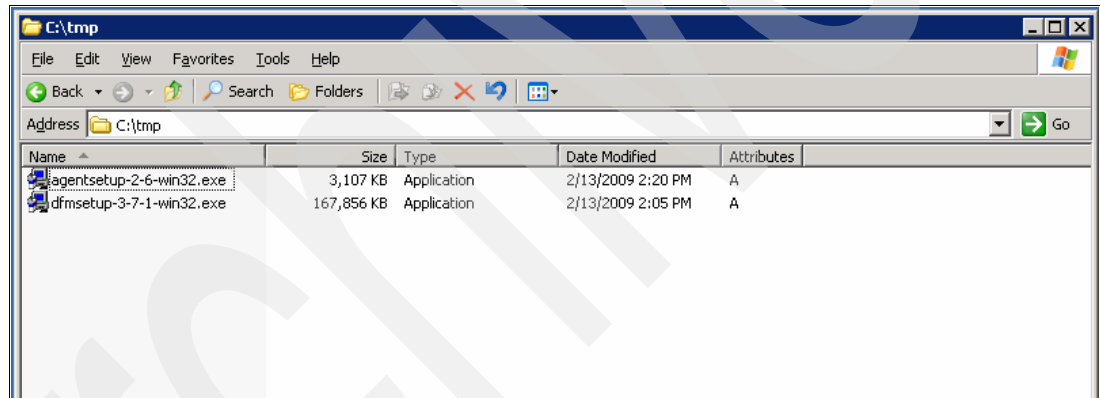


Figure 5-3 Double-click the agent.exe file to begin installation

Here are the steps to install the Operations Manager Host Agent:

1. Once you have double-clicked the install file, you will see the window shown in Figure 5-4. Click **Next** to continue.



Figure 5-4 Operations Manager Host Agent installation welcome window

2. You can change the installation location here. We used the default location in our installation, as shown in Figure 5-5. Click **Next** to continue.




Figure 5-5 Destination Folder window

3. You can return to any previous window and make changes before you commit to the installation, as shown in Figure 5-6. Click **Install** to continue.




Figure 5-6 Ready to Install the Program window

4. Once the installation is complete, you will see the window shown in Figure 5-7 on page 87. Click **Finish** to complete the installation.




Figure 5-7 Installation is complete

5. Once the installation is complete, you will see the window shown in Figure 5-8. Select **Edit Settings**.

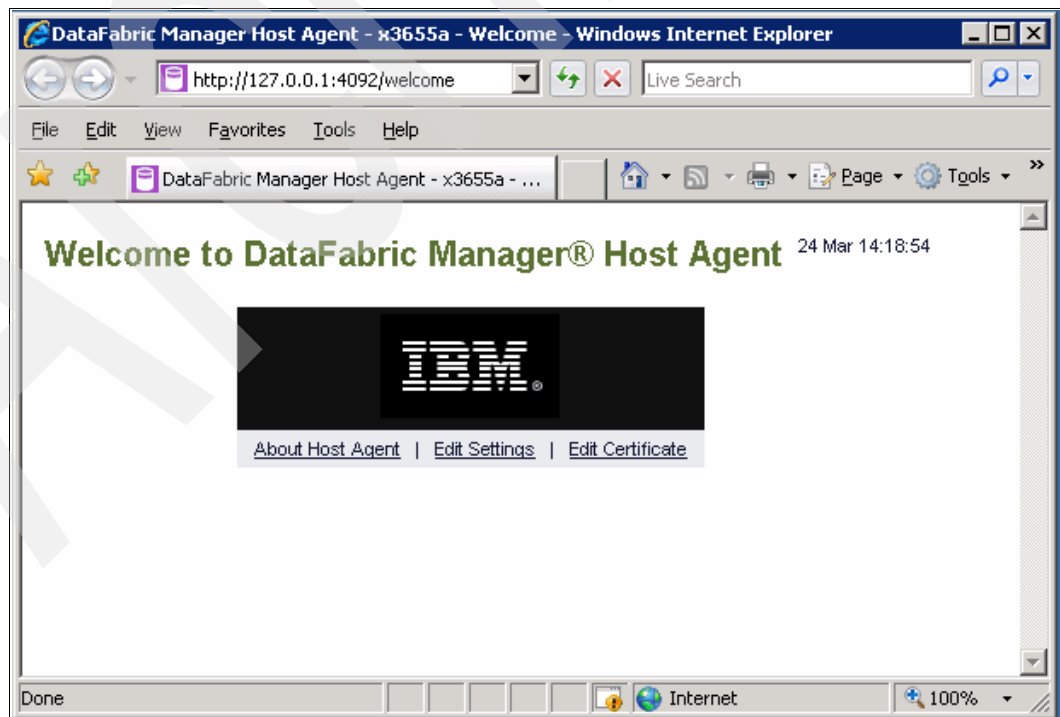


Figure 5-8 DataFabric Manager Host Agent welcome window

6. You will now see a window (Figure 5-9) where you can enter passwords for the functions you want the host agent to perform.

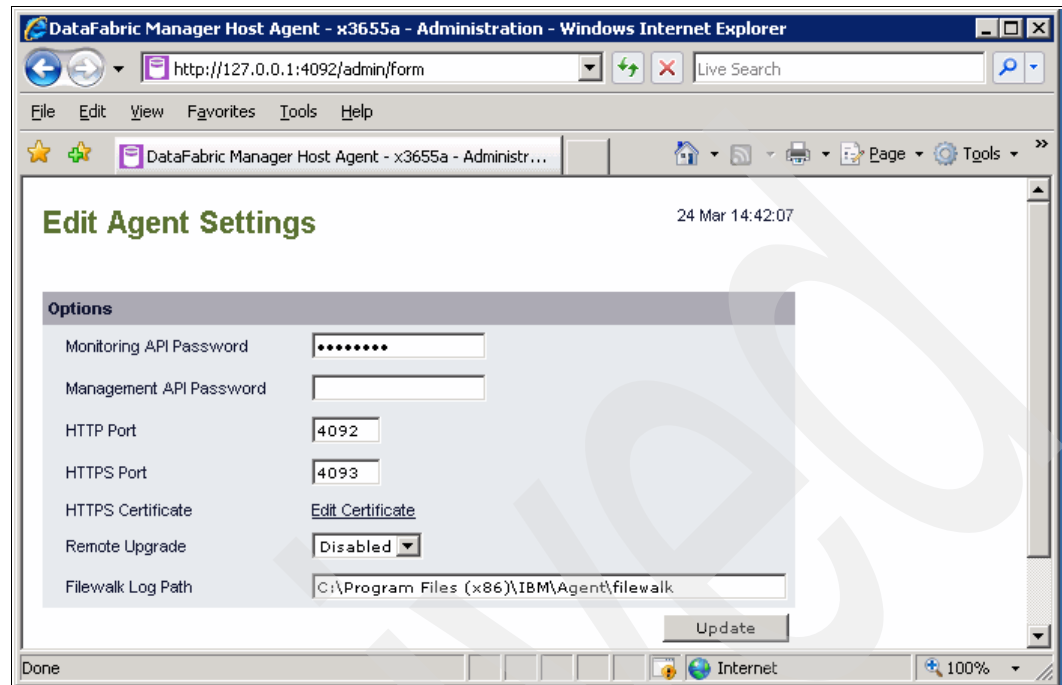


Figure 5-9 Edit Agent Settings - enter passwords

7. You can enter passwords for both functions. After doing so, click **Update**, as shown in Figure 5-10.

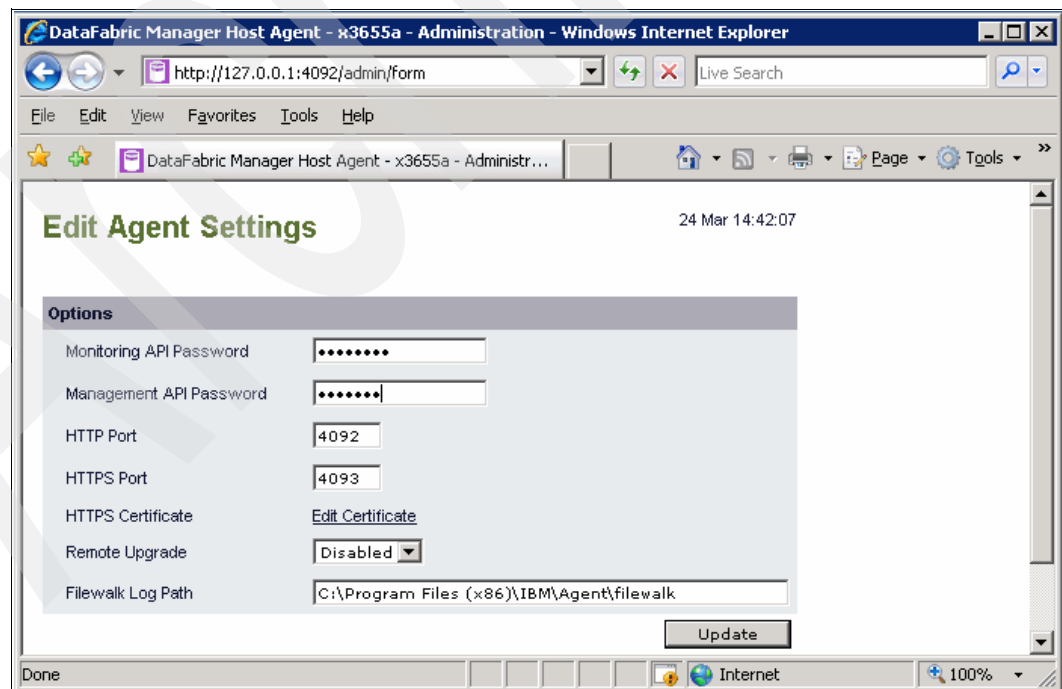


Figure 5-10 Edit Agent Settings - update

Earlier in this chapter, we discussed how you can access the Operations Manager Host Agent to update passwords, HTTP ports, or firewall path logs. If you are on the host system that contains the host agent, you will need to point your browser to `http://127.0.0.1` or `http://localhost`. If you are working somewhere else in the network and want to access the host agent on another system, be sure to point your browser to that IP address or DNS name. Also, when putting the IP address or DNS name in the address field of your browser, be sure to use “http/” before the IP or DNS address.

This concludes the steps for installing the host agent.

Archived



Host Agent installation for Linux

This chapter will cover the preparation, installation, and setup of the Operations Manager Host Agent for Linux.

6.1 Host prerequisites

The Operations Manager Host Agent can be installed on Windows and Linux. In this chapter, we look at the Linux prerequisites. Table 6-1 shows the prerequisites for Operations Manager Host Agent.

Table 6-1 Prerequisites

Platform	Supported operating systems
Linux workstation or server	<ul style="list-style-type: none">▶ Red Hat Enterprise Linux, Version 4, update 3 or later (32-bit and 64-bit x86)▶ Red Hat Enterprise Linux, Version 5 (32-bit and 64-bit x86)▶ SUSE Linux Enterprise Server 9, SP 2 or later (32-bit and 64-bit x86)▶ SUSE Linux Enterprise Server 10 (32-bit and 64-bit x86)
VMware ESX Server, Standard or Enterprise Edition, Version 3	Windows 2003 Server (32-bit) Red Hat Enterprise Linux AS, Version 4

6.2 DFM Server prerequisites

Table 6-2 shows the software levels support for this host agent (V2.6).

Table 6-2 Supported software

Software	Supported levels
DFM Server	DataFabric Manager Server V3.6 or later

You must meet the following prerequisites to use the Operations Manager FSRM feature:

- ▶ You must have a valid FSRM license installed on your DataFabric Manager Server. Contact your sales representative to obtain a File SRM license.

Note: The Quotas subtab is visible in the Operations Manager user interface (select **Control Center** → **Home** → **Group Status** tabs to see it) until you install the File SRM license. After you install the license, the Quotas subtab is renamed to “File SRM”, and all of the FSRM features become visible when you click it.

- ▶ All hosts to be managed through Operations Manager must be connected to a TCP/IP network either known to or discoverable by the DataFabric Manager Server. The hosts must be connected to the network through an Ethernet port and must have a valid IP address.
- ▶ All directory paths to be monitored must be visible to the host agent. For example, to enable FSRM host monitoring by the DataFabric Manager Server, the host agent must mount a storage system share using NFS or CIFS, or the host agent must use a LUN on the storage system.
- ▶ Before setting up FSRM paths and schedules, you must enable administrative access to your host agents. For more information, see Chapter 11, “File Storage Resource Manager” on page 269.

6.3 N series prerequisites

The N Series storage system must be at Data ONTAP (DoT) V7.1 or higher.

6.4 Installation

The host agent code can be downloaded from the IBM system support Web site at the following address:

<http://www.ibm.com/storage/support/nas>

To install the Host Agent software, proceed to the folder where the agent is installed and execute the agent `setup-2-6-linux.bin` file or the latest version available after this version. This will automatically unpack the files and install Host Agent to the default directory, as well as start it up, as shown in Example 6-1.

Example 6-1 Installation of the host agent for Linux

```
[root@localhost dfmAgent]# ./agentsetup-2-6-linux.bin
Unpacking files needed for the installation ...
Stopping agent.
Beginning the installation ...
Starting agent.
```

You may now point your browser to

`http://127.0.0.1:4092/welcome`

to configure the agent software.

```
[root@localhost dfmAgent]#
```

6.5 Reviewing and configuring host agent settings

In order to manage the host agent and set passwords, simply point a Web browser (from any machine) to the host agent residing on port 4092 of the server you just installed it on.

Figure 6-1 shows the link if you chose to run the Web browser from the same host. If you are going to use another host, such as your desktop, replace the address 127.0.0.1 with the IP address of the fully qualified DNS name of the target host in your Web browser.

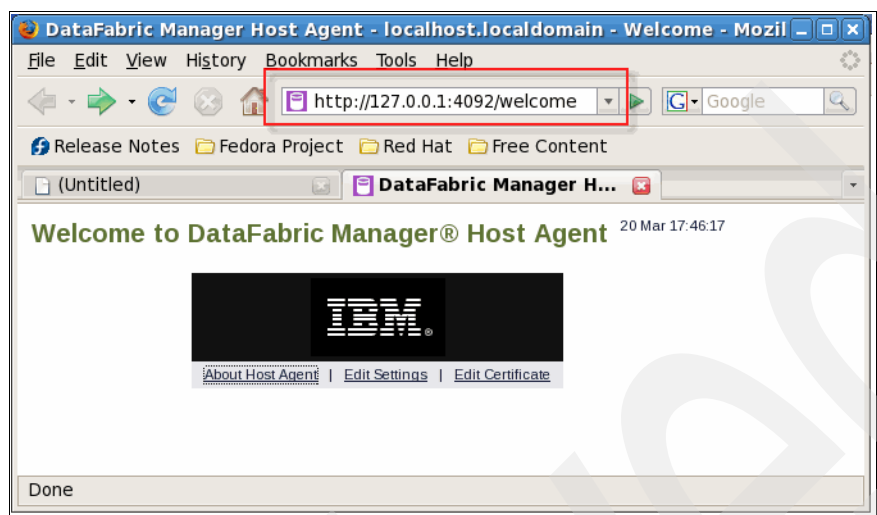


Figure 6-1 View of the Host Agent from a Web browser

The format of the link is:

`http://your host:4092/welcome`

Clicking the **About Host Agent** link will give you a window detailing the Host Agent version number, host name, and other details, as shown in Figure 6-2.

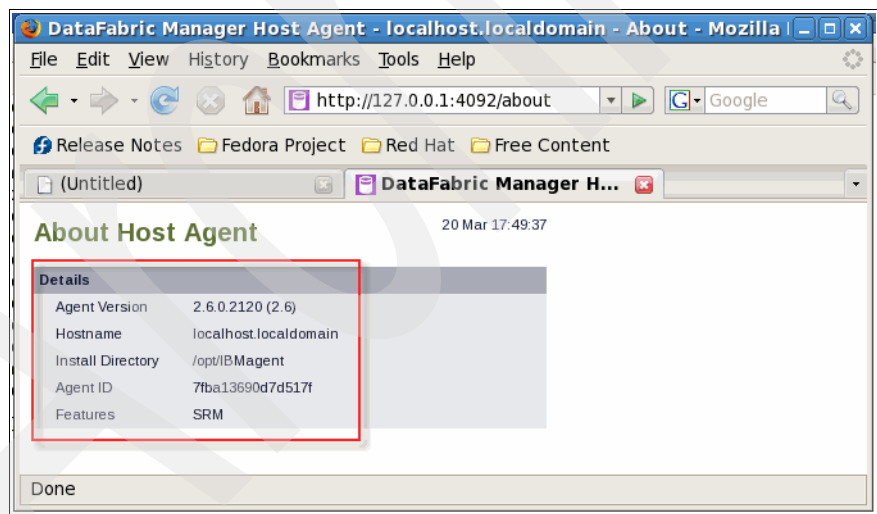


Figure 6-2 Reviewing Host Agent information

Clicking the **Edit Settings** link on the main Web page will give you a page where you can edit the access passwords and if, necessary, change the host ports the agent will listen on, as shown in Figure 6-3 on page 95.

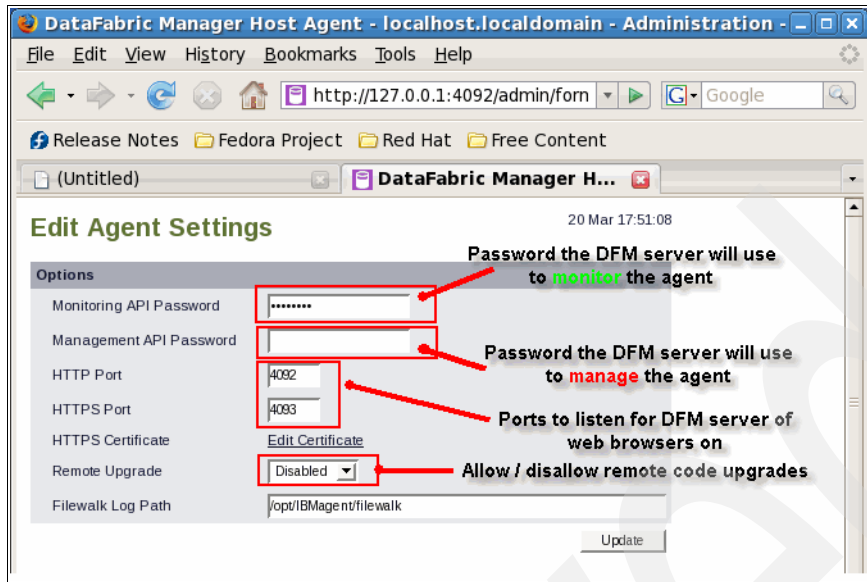


Figure 6-3 Editing Host Agent settings

Note: If you change any of the settings here, you must update your DFM Server to reflect this information or the DFM Server will not be able to connect to this agent. This is particularly important for the HTTP port and HTTPS port.

The default passwords are shown in Table 6-3:

Table 6-3 Default user names and passwords

Description	Default value
Guest user name	guest
Guest password	public
Admin user name	admin
Admin password	Not specified. This is disabled until the user specifies a password.

In Operations Manager Host Agent V2.5 and later, the `-o` option lets you configure passwords without having root privileges. Example 6-2 shows an example syntax of that function.

Example 6-2 `sudo dfm_agent -o command`

```
[root@localhost dfmAgent]# sudo dfm_agent -o Admin-Password=secret
[root@localhost dfmAgent]#
```

The default port addresses are listed in Table 6-4.

Table 6-4 Default port addresses

Description	Default value
HTTP port	4092
HTTPS port	4093

If your host is properly configured in your network, you can install a HTTPS certificate, generated by your internal Certificate Authority or an external one here. Refer to your organization's certificate and key management infrastructure and policies to see if you need to install a certificate on your host agent.

At this point, the host agent is discoverable by the DFM Server and should be able to managed by the DFM Server once the relevant credentials are entered in the DFM Server.

6.6 Starting and stopping the service

On Linux hosts, Operations Manager Host Agent is a daemon started by the init service.

You can also start and stop the daemon manually, as follows:

- ▶ To start the daemon, at the Linux command line, run **dfm_agent start**.
- ▶ To stop the daemon, at the Linux command line, run **dfm_agent stop**.

Note: These commands will only work after you have successfully installed the Host Agent on the your host.

6.7 Limitations of Host Agent

The Host Agent will create a unique system ID to represent the host. This information is maintained in the System-ID variable in the following file:

`/opt/IBMAgent/dfm_agent.cfg`

Operations Manager uses this information to uniquely track each host. If you clone the host, the DFM Server will be unable to distinguish the clones and continue to treat them as one and the same host.

You can force the cloned host to regenerate a new unique System ID key. To do this, simply do the following steps:

- ▶ Stop the dfm agent by running the following command:
`dfm_agent stop`
- ▶ Remove *only* the System-ID variable from the file `/opt/IBMAgent/dfm_agent.cfg`.
- ▶ Start the dfm agent by running the following command:
`dfm_agent start`



Part 3

N series Management Console and applications

This part discusses installation of the N series Management Console as well as its administration.

Archived



N series Management Console installation on Linux

Here we discuss the steps necessary to install the IBM System Storage N series Management Console on Linux.

7.1 Host prerequisites

The N series Management Console for Linux is a client dashboard application that is typically installed on a dedicated management computer accessible by the IT operations staff. This tool requires Java, and we do not recommend running it on the DFM Server for performance reasons.

7.2 N series prerequisites

The N series Management Console prerequisites for a Linux host are the same as the requirements for a Linux host running the Operations Manager interface. These requirements are described in Chapter 4, “Installing Operations Manager: Linux” on page 55.

7.3 Installation

The N series Management Console can be downloaded from the DFM Server through the Operations Manager Web interface.

Do the following steps to install the N series Management Console:

1. Open a Web browser. We use the Mozilla Firefox Web browser to map to the DFM Server.
2. Log on to the DFM Server.
3. On the Control Center tab, we select the **Setup** menu and then click **Download Management Console**, as shown in Figure 7-1.

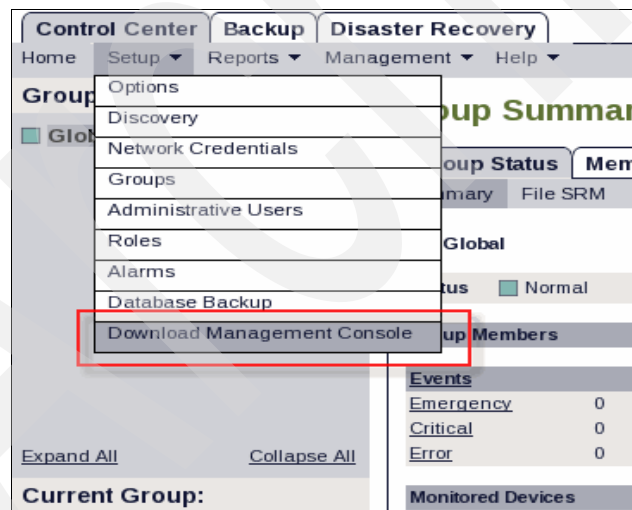


Figure 7-1 Download Management Console menu option

4. We then click **Download Linux Installation (version 2.2.1)**, as shown in Figure 7-2 on page 101.




Figure 7-2 Download Linux Installation (version 2.2.1) window

5. We download the file to a folder (in our example, /tmp/nconsole), as shown in Figure 7-3 and Figure 7-4.

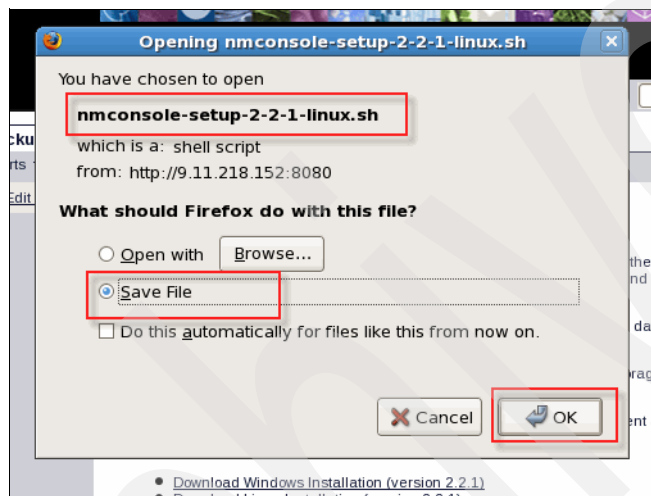


Figure 7-3 Save the file to a folder

6. The file can be viewed by navigating to the /tmp/nconsole folder. We decide to install this file through the Linux command-line interface (CLI) in order to specify a target installation directory.

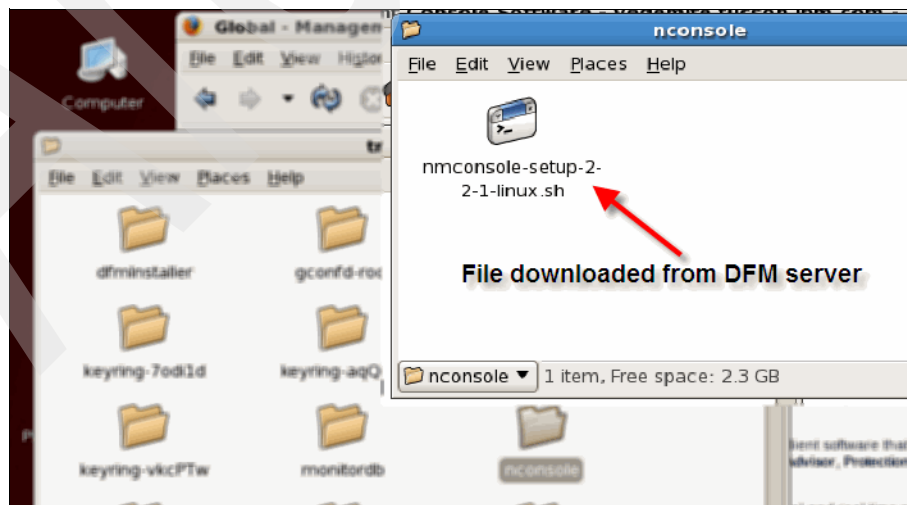


Figure 7-4 Location of downloaded installation file

7.3.1 Linux command-line installation

Do the following steps to install the N series Management Console, make a link file, and paste the link file onto the desktop.

1. Launch a terminal window and navigate to the folder where we saved the downloaded installation file.
2. Run **ll** to list the directory contents. You can see that the file is not yet marked as executable, as shown in Example 7-1.

Example 7-1 Permissions checking

```
[root@vegemite nconsole]# ll
total 63108
-rw-r--r-- 1 root root 64547812 Apr  2 14:45 nmconsole-setup-2-2-1-linux.sh
[root@vegemite nconsole]#
```

3. Execute **chmod** to make the installation file executable, as shown in Figure 7-2 on page 101.

Example 7-2 Making the install file executable

```
[root@vegemite nconsole]#
[root@vegemite nconsole]# chmod 744 nmconsole-setup-2-2-1-linux.sh
[root@vegemite nconsole]# ll
total 63108
-rwxr--r-- 1 root root 64547812 Apr  2 14:45 nmconsole-setup-2-2-1-linux.sh
[root@vegemite nconsole]#
```

4. Obtain the usage parameters, as shown in Example 7-3.

Example 7-3 Usage parameters

```
[root@vegemite nconsole]# ./nmconsole-setup-2-2-1-linux.sh -?
Unpacking files needed for the installation ...
nmconsole-setup.sh:
Usage: nmconsole-setup.sh [ -d <directory> ]
[root@vegemite nconsole]#
[root@vegemite nconsole]#
```

5. Use the **-d** switch to specify the install directory, as shown in Example 7-4.

Example 7-4 Using the -d switch

```
[root@vegemite nconsole]# mkdir /opt/nconsole
[root@vegemite nconsole]# ./nmconsole-setup-2-2-1-linux.sh -d /opt/nconsole
Unpacking files needed for the installation ...
Beginning the installation ...
./
./nmconsole
./boot.properties
..
<< Output truncated for brevity >>
..
./help/Working_with_data_in_views.htm
./help/zoom_icn.gif
[root@vegemite nconsole]#
```

6. Navigate to the installation folder to make the nconsole file executable, as shown in Example 7-5.

Example 7-5 Navigating to the installation folder

```
[root@vegemite nconsole]# cd /opt/nconsole/
[root@vegemite nconsole]# ll
total 80
-rw-r--r-- 1 20041 gopher 1051 Oct 27 12:28 boot.properties
drwxr-xr-x 2 20041 gopher 4096 Oct 27 12:29 css
drwxr-xr-x 7 20041 gopher 4096 Oct 27 12:29 help
drwxr-xr-x 4 20041 gopher 4096 Oct 27 12:29 images
drwxr-xr-x 8 20041 gopher 4096 Oct 27 12:28 jre-1.5.0
drwxr-xr-x 2 20041 gopher 4096 Oct 27 12:29 lib
-rw-r--r-- 1 20041 gopher 1428 Oct 27 12:28 log4j.properties
-rw-r--r-- 1 root root 162 Apr 2 15:21 nmconsole
drwxr-xr-x 30 20041 gopher 4096 Oct 27 12:29 plugins
-rw-r--r-- 1 20041 gopher 10 Oct 27 12:28 vendor.properties
[root@vegemite nconsole]#
[root@vegemite nconsole]# chmod 744 nmconsole
[root@vegemite nconsole]#
[root@vegemite nconsole]# ll
total 80
-rw-r--r-- 1 20041 gopher 1051 Oct 27 12:28 boot.properties
drwxr-xr-x 2 20041 gopher 4096 Oct 27 12:29 css
drwxr-xr-x 7 20041 gopher 4096 Oct 27 12:29 help
drwxr-xr-x 4 20041 gopher 4096 Oct 27 12:29 images
drwxr-xr-x 8 20041 gopher 4096 Oct 27 12:28 jre-1.5.0
drwxr-xr-x 2 20041 gopher 4096 Oct 27 12:29 lib
-rw-r--r-- 1 20041 gopher 1428 Oct 27 12:28 log4j.properties
-rwxr--r-- 1 root root 162 Apr 2 15:21 nmconsole
drwxr-xr-x 30 20041 gopher 4096 Oct 27 12:29 plugins
-rw-r--r-- 1 20041 gopher 10 Oct 27 12:28 vendor.properties
[root@vegemite nconsole]#
```

7. At this point, the N series Management Console is installed and can be launched by opening the nconsole file in the directory /opt/nconsole.

8. We set up a link (or shortcut) to the N series Management Console from the graphical desktop, as shown in Figure 7-5.

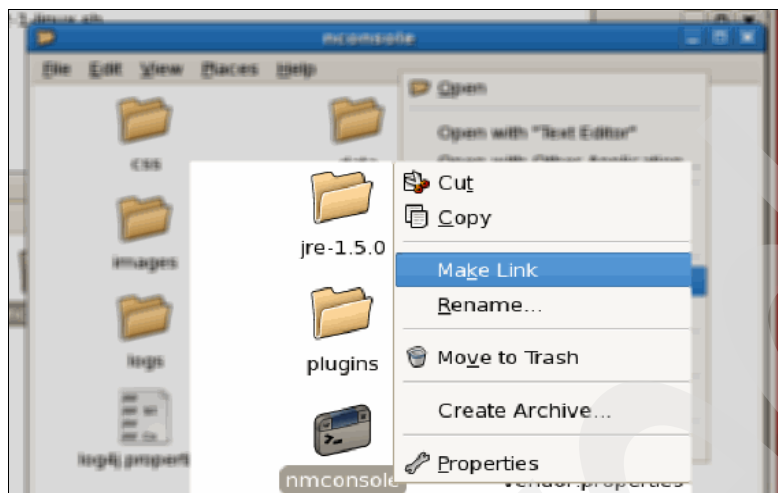


Figure 7-5 Create a link file of the file `nconsole` in `/opt/nconsole`

9. The file called link to `nconsole` should be moved to the desktop using cut and paste, as shown in Figure 7-6.




Figure 7-6 Move "link to `nconsole`" to the desktop

10. Finally, right-click the `nconsole` link icon to get to the Properties menu for the link file and make it an executable, as shown in Figure 7-7 and in Figure 7-8 on page 105.




Figure 7-7 Properties of the link file

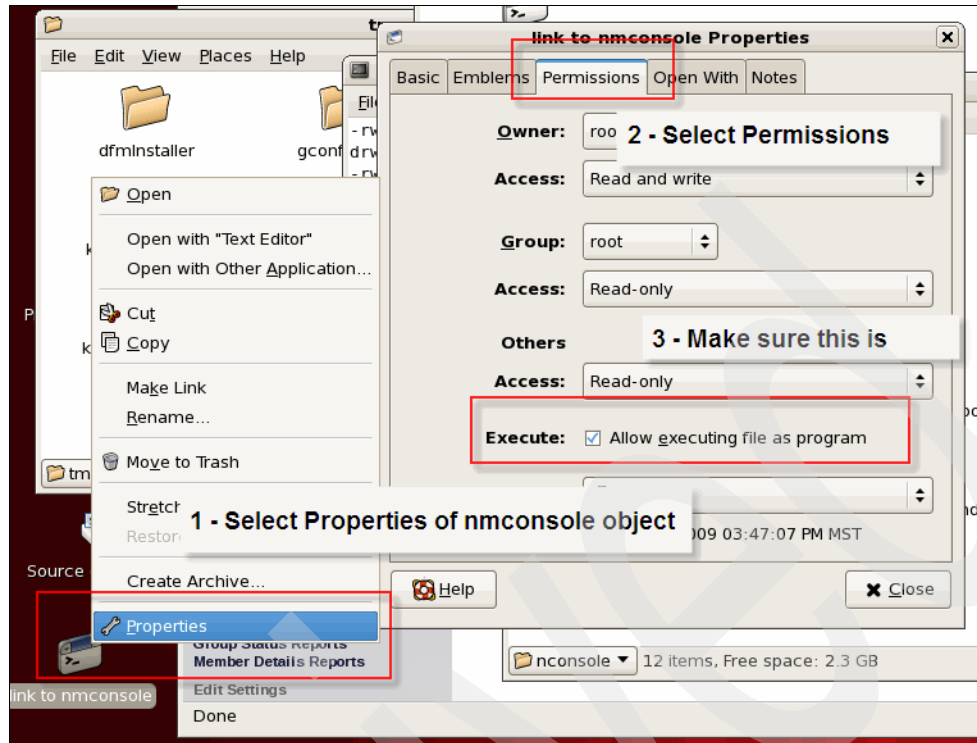


Figure 7-8 Select the Permissions tab and ensure that the Execute option is selected

The N series Management Console is now ready to be launched from the desktop.

Launch the application and you will be presented with the connect and login window, as shown in Figure 7-9. Here you will need to enter the name or IP address of your DFM Server as well as your login credentials. In our lab, we use root, but in a normal production environment, you really should use a user account that is dedicated to this task.

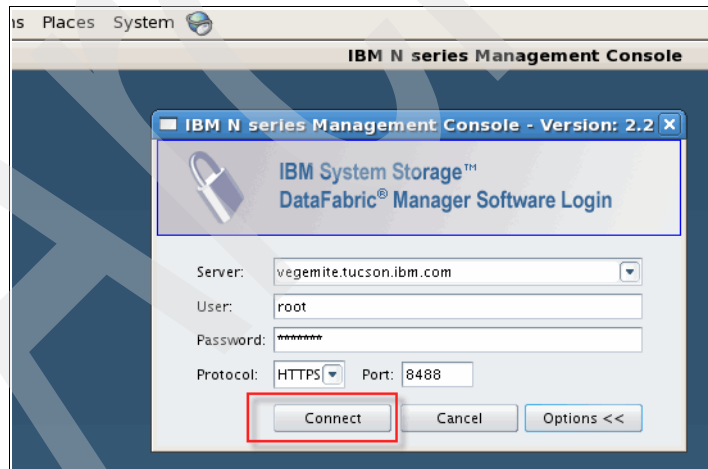


Figure 7-9 Connect and login window for N series Management Console

DFM Server supports Role Based Access Control (RBAC) and you can set up roles to represent different administrative functions that match the method of administration your organization supports. This capability in DFM Server is very sophisticated and facilitates the delegation of tasks to different entities within your organization. Once you have set up the roles, you can then assign user IDs to these roles. User IDs can be either locally created IDs or IDs listed in LDAP Directory servers, such as Microsoft Active Directory® or IBM Directory Server.

After you have logged in, you should get a window similar to the one shown in Figure 7-10.

Details about how to use the N series Management Console are covered in Chapter 9, “Configuring Operations Manager” on page 127.

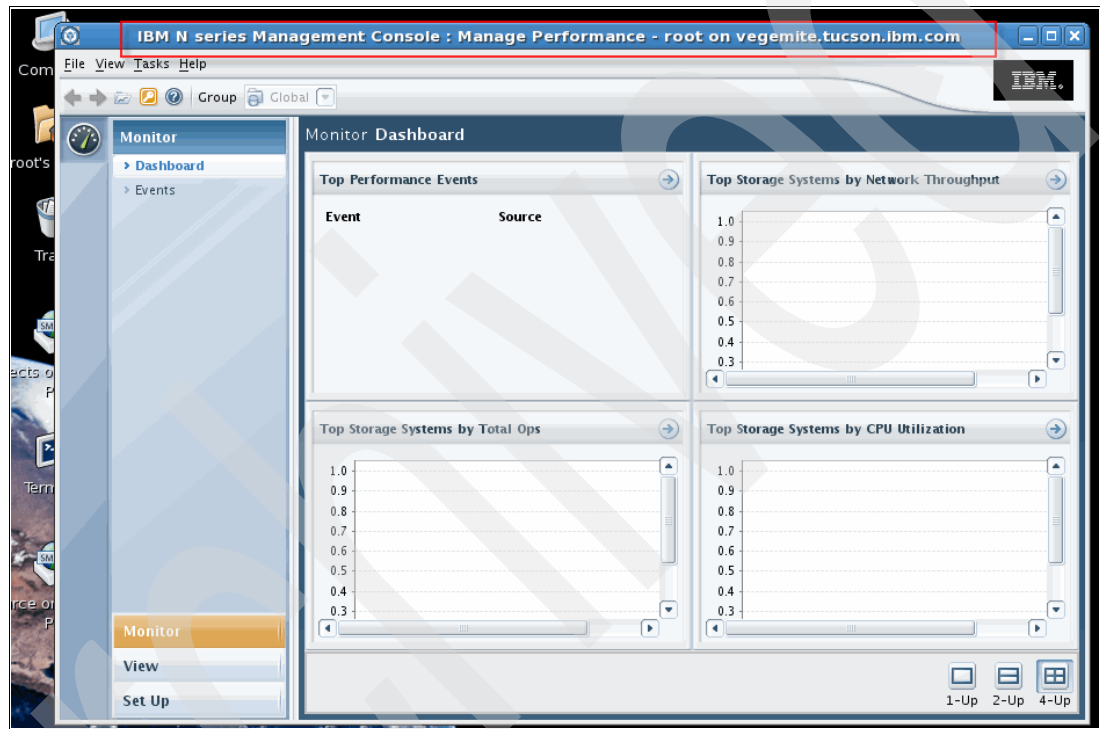


Figure 7-10 N series Management Console as viewed from a Linux desktop



N series Management Console installation for Windows

Here we discuss the steps necessary to install the IBM System Storage N series Management Console on Windows.

8.1 Host prerequisites

The N series Management Console is a client platform for Java-based applications running on Windows 2003, Windows XP, Windows Vista®, Red Hat Enterprise Linux, and Suse Linux platforms.

When you install N series Management Console on Windows 2003 R2 32-bit and on a Windows XP host, they operate the same. The Windows 2003 64-bit installation is the same and will not be shown in this chapter.

Note: N series Management Console V2.2 does not support Solaris.

You should not install N series Management Console on the DataFabric Manager Server. Installing the console on the server can have a negative impact on server performance.

Note: N series Management Console V2.2 is supported on DataFabric Manager V3.6.1 or later.

8.2 N series prerequisites

The N series Management Console is the console for the Operations Manager, so the N series requirements are the same as shown in Chapter 2, “Installing Operations Manager: Windows 2003 32-bit operating system” on page 21.

8.2.1 Function of N series Management Console

The N series Management Console is a client software that contains a number of storage system management applications.

N series Management Console incorporates the following applications, as shown in Figure 8-1 on page 109:

- ▶ The Performance Advisor application allows viewing of historical and real-time performance data collected from N series storage systems.
- ▶ The Protection Manager application provides policy-based data protection by using N series storage systems that have SnapVault, Open Systems SnapVault, or SnapMirror licenses.
- ▶ The Provisioning Manager application improves efficiency in storage utilization, and automates provisioning and managing storage for NAS and SAN access.

To use these features, you must download and install N series Management Console.

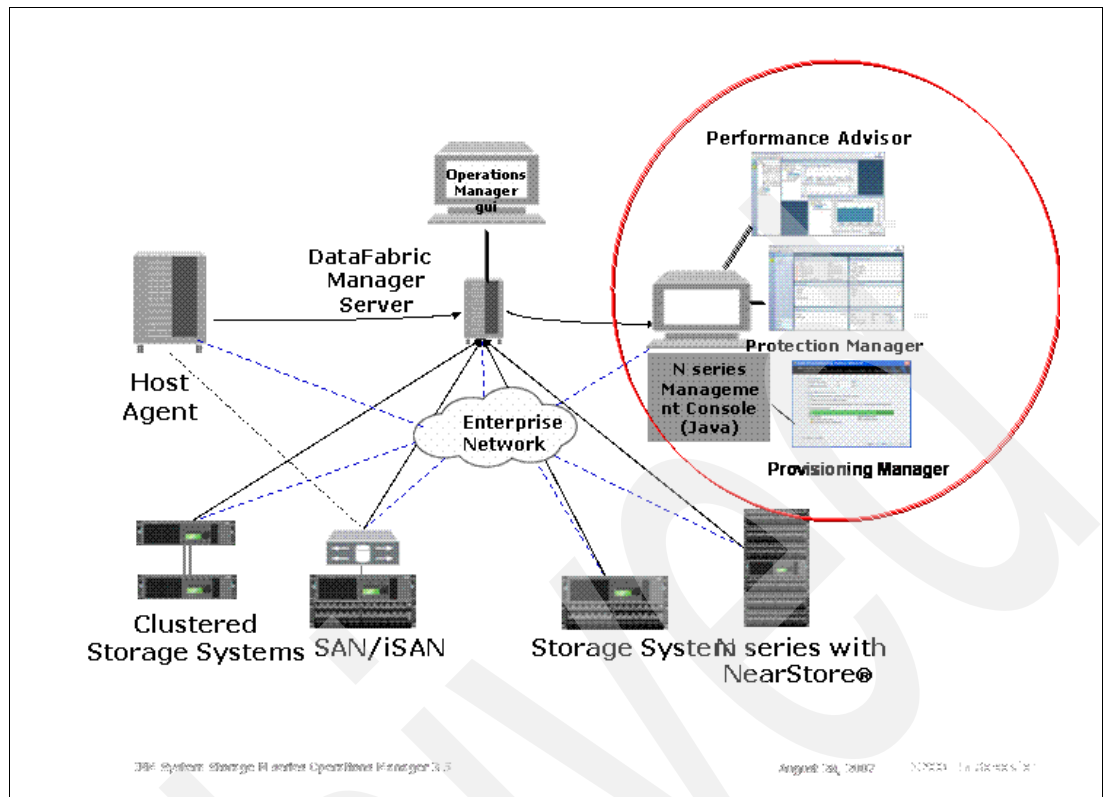


Figure 8-1 N series Management Console

8.2.2 License requirements

The license requirements for N series Management Console are:

- ▶ Performance Advisor application license.
- ▶ Protection Manager application provides policy-based data protection by using N series storage systems that have SnapVault, Open Systems SnapVault, or SnapMirror licenses.
- ▶ Provisioning Manager application license.

8.3 Installation

Follow these steps to download and install the N series Management Console:

1. In the Control Center tab of the Operations Manager main window, select the **Setup** tab and then select **Download Management Console**, as shown in Figure 8-2.

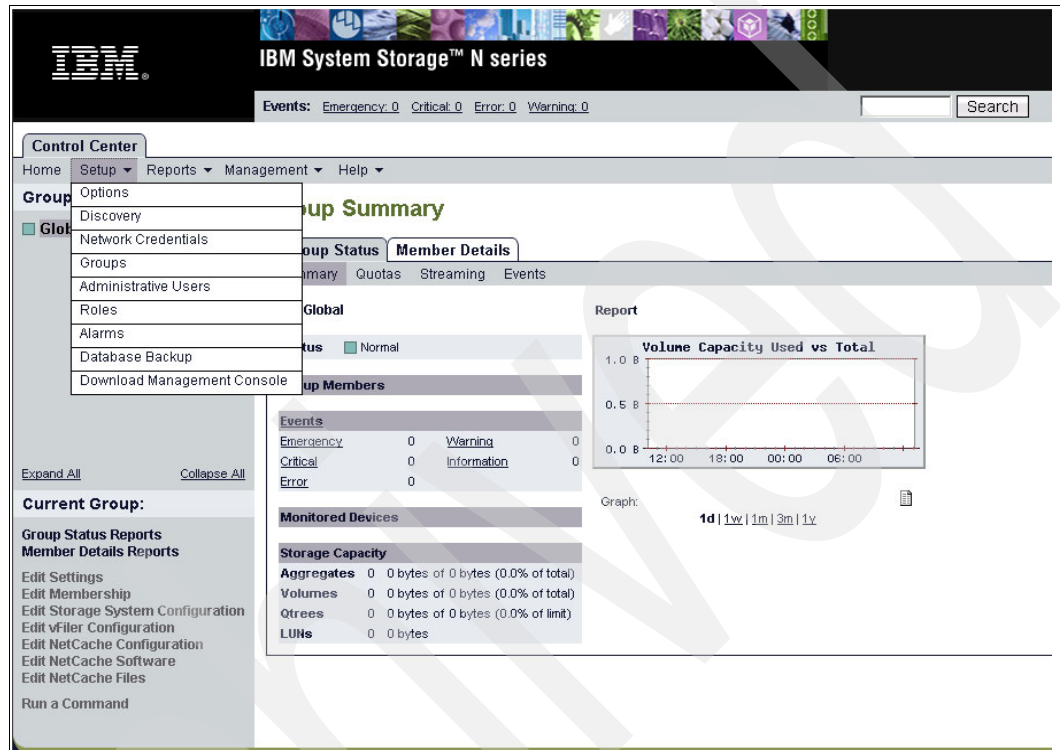


Figure 8-2 N series Management Console setup window

2. Figure 8-3 on page 111 shows the links to the Windows and Linux downloads. Click **Download Windows Installation (version 2.2.1)**.

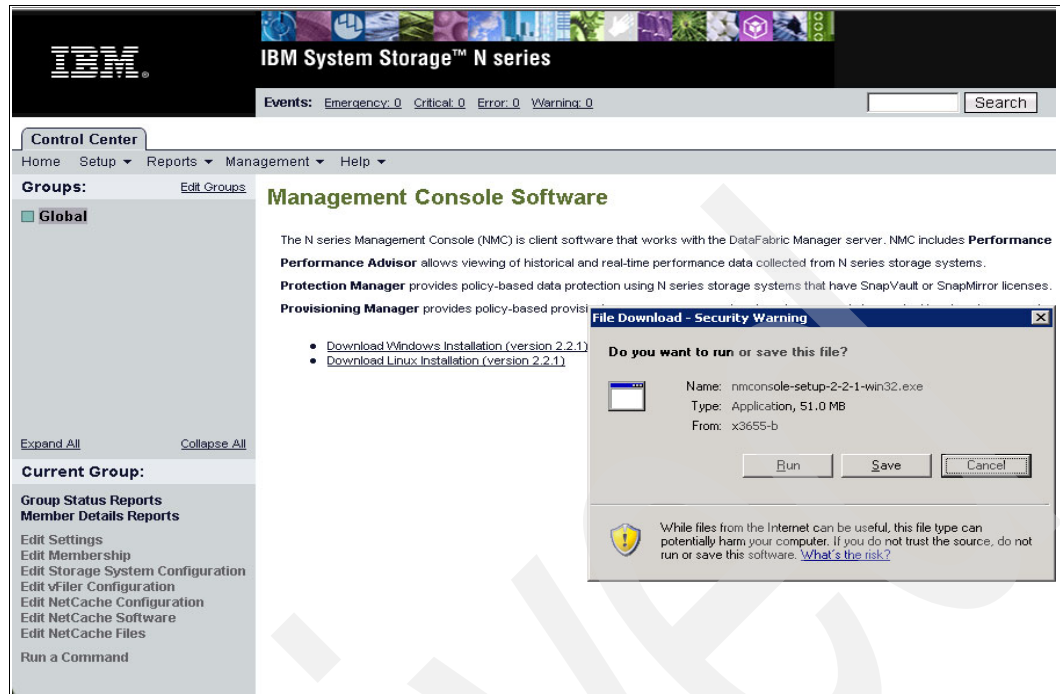


Figure 8-3 N series Management Console Windows 2003 download window

3. Click **Save** in the File Download - Security Warning window. You can then save the installation files to a specific location, as shown in Figure 8-4. If you want to install the N series Management Console on a separate server (recommended), you can save it on an USB drive, CD, or on a network share.

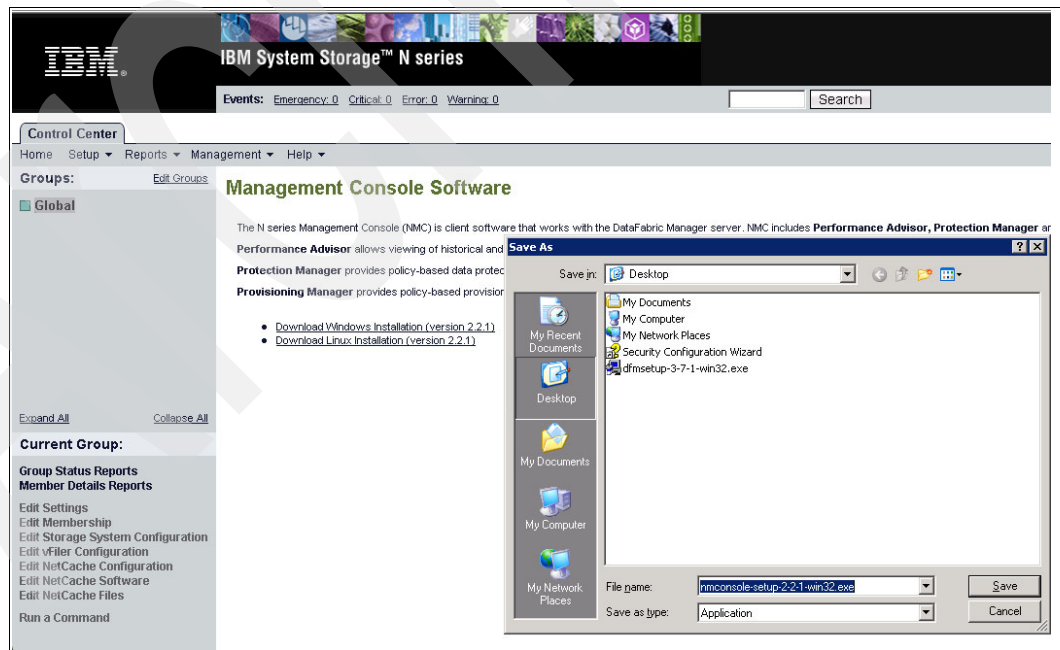


Figure 8-4 N series Management Console download and save window

We save the files to the Windows desktop, as shown in Figure 8-5.

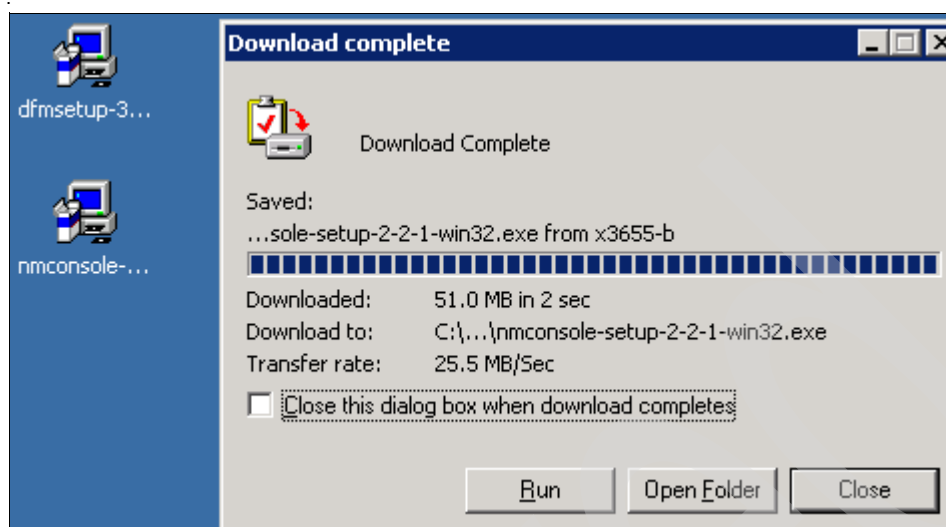


Figure 8-5 N series Management Console download on a Windows 2003 R2 desktop window

4. Execute the N series Management Console setup file. The Preparing to Install window (Figure 8-6) shows that the N series Management Console is checking the operating system compatibility of the system on which we are installing the N series Management Console.

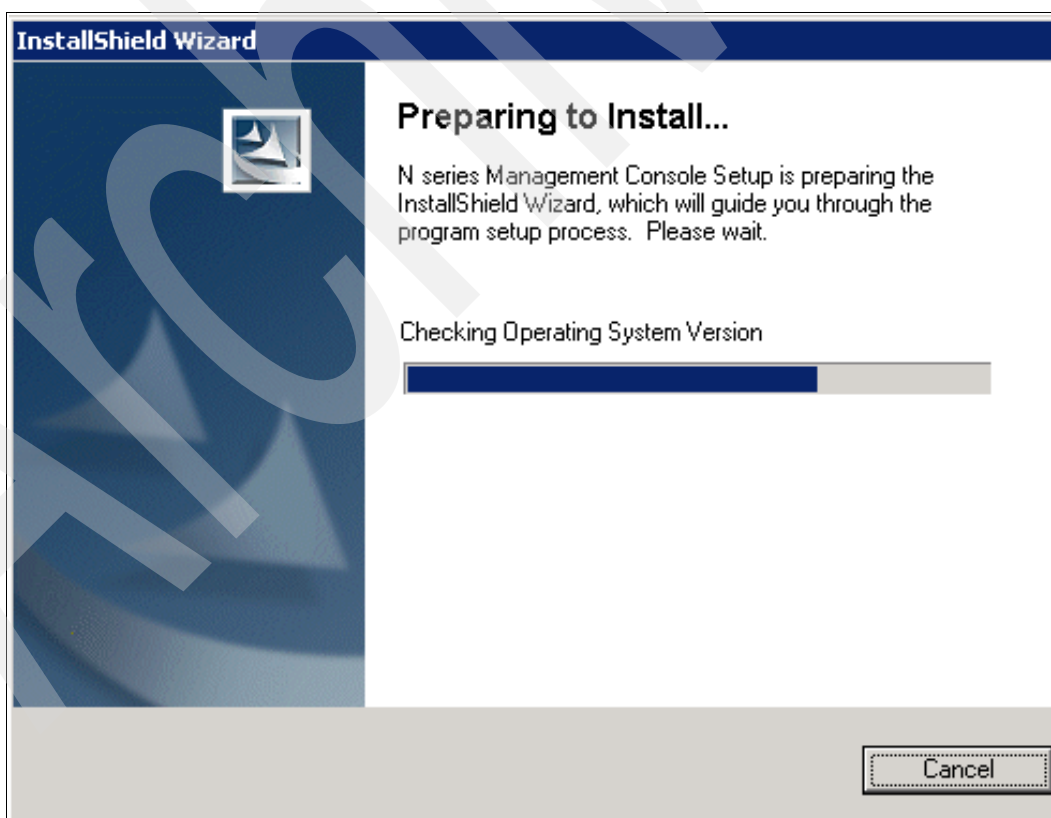


Figure 8-6 Checking the operating system version for installation of N series Management Console

Figure 8-7 shows that the Windows installer is preparing to install the N Series Management Console.




Figure 8-7 Configuring Windows installer for installation of N series Management Console

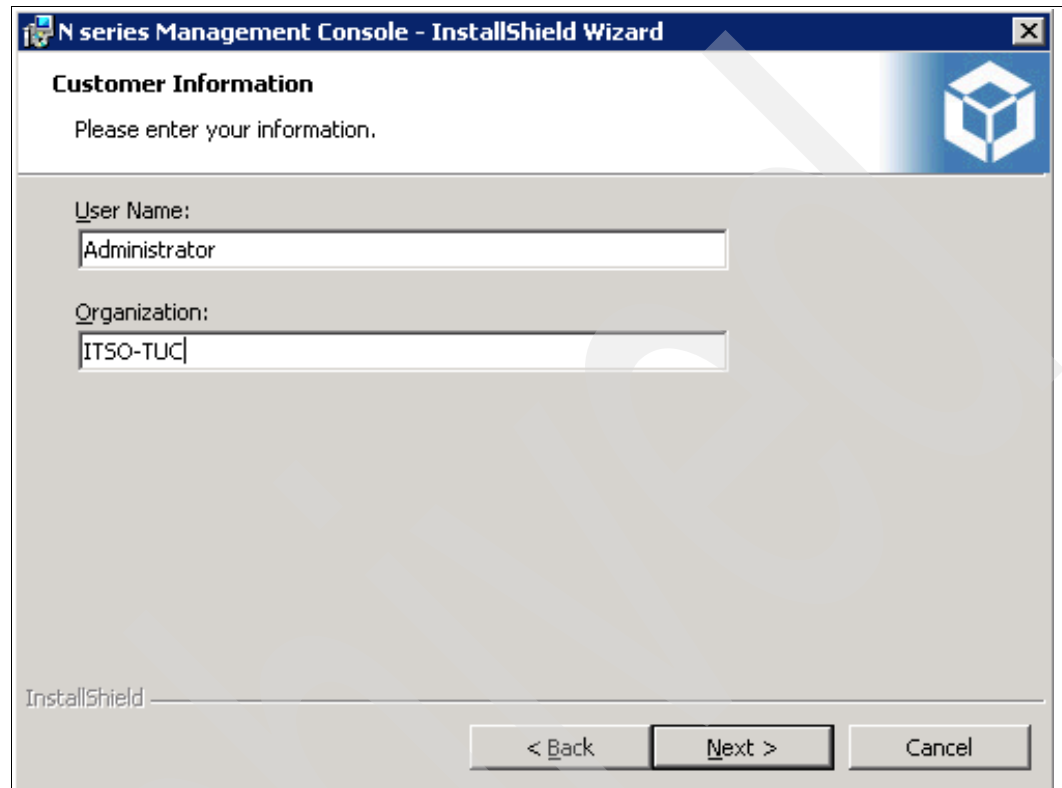
5. Figure 8-8 shows the welcome window for the N series Management Console. Select **Next** to continue the installation of the N series Management Console.



Figure 8-8 The welcome window for the installation of the N series Management Console

6. Figure 8-9 shows the Customer Information window for the N series Management Console. By default, the installer uses the Windows system name as the User Name and Organization. You have the option to modify these fields if needed.

Select **Next** to continue the installation of the N series Management Console.



The screenshot shows a Windows-style dialog box titled "N series Management Console - InstallShield Wizard". The window has a blue header bar with the title and a close button. Below the header, the text "Customer Information" is displayed in bold, followed by the instruction "Please enter your information." and a blue cube icon. The main area contains two text input fields: "User Name:" with the text "Administrator" and "Organization:" with the text "ITSO-TUC". At the bottom, there is a status bar with the text "InstallShield" and three buttons: "< Back", "Next >", and "Cancel".

Figure 8-9 Customer Information window

7. Figure 8-10 shows the Destination Folder for the N series Management Console files. It is set to the path shown on the window by default, but you have option to change the path by selecting **Change...** and browsing to the path you wish to use.

Select **Next** to continue the installation of the N series Management Console.




Figure 8-10 Destination Folder window

8. Figure 8-11 shows that the N series Management Console is ready to install.
Select **Install** to continue the installation of the N series Management Console.

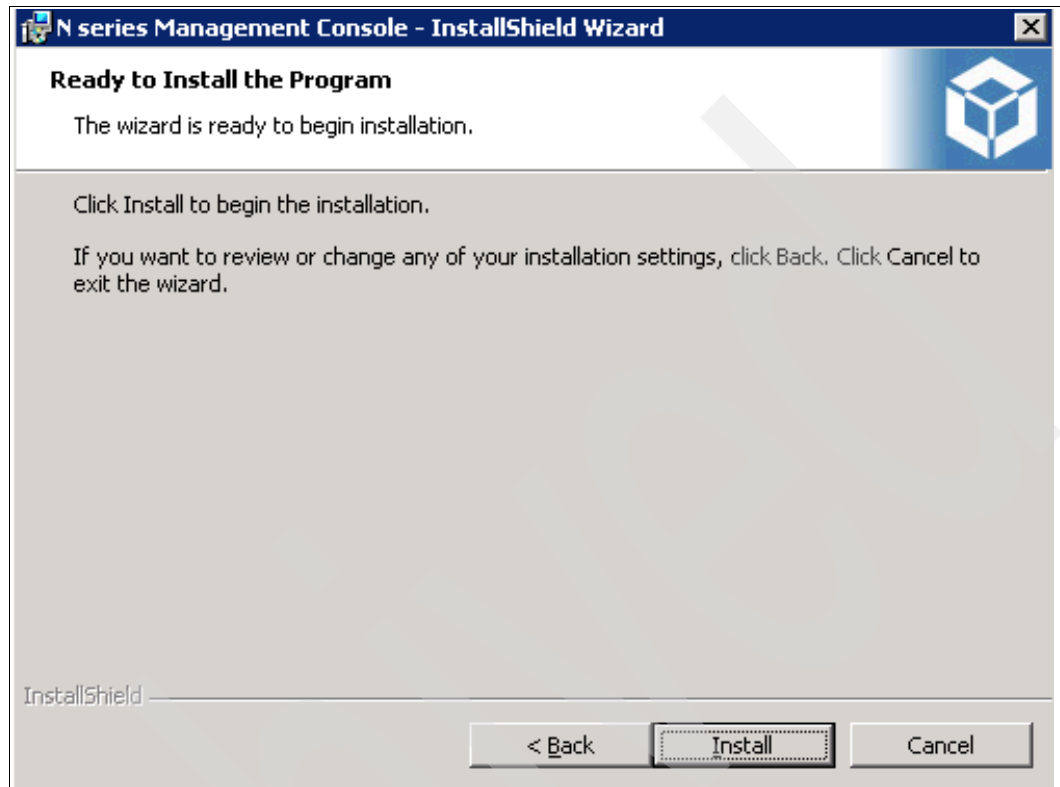


Figure 8-11 Ready to install the N series Management Console window

9. Figure 8-12 shows that the N series Management Console is checking the files before installing them to make sure all the requirements have been met.
Select **Cancel** if you want to cancel the installation of the N series Management Console.




Figure 8-12 N series Management Console installation status window

Figure 8-13 shows that the N series Management Console is copying the new files to the specified path.




Figure 8-13 N series Management Console copying new files window

10. Figure 8-14 shows that the N series Management Console has been successfully installed. Click **Finish** to exit the wizard.

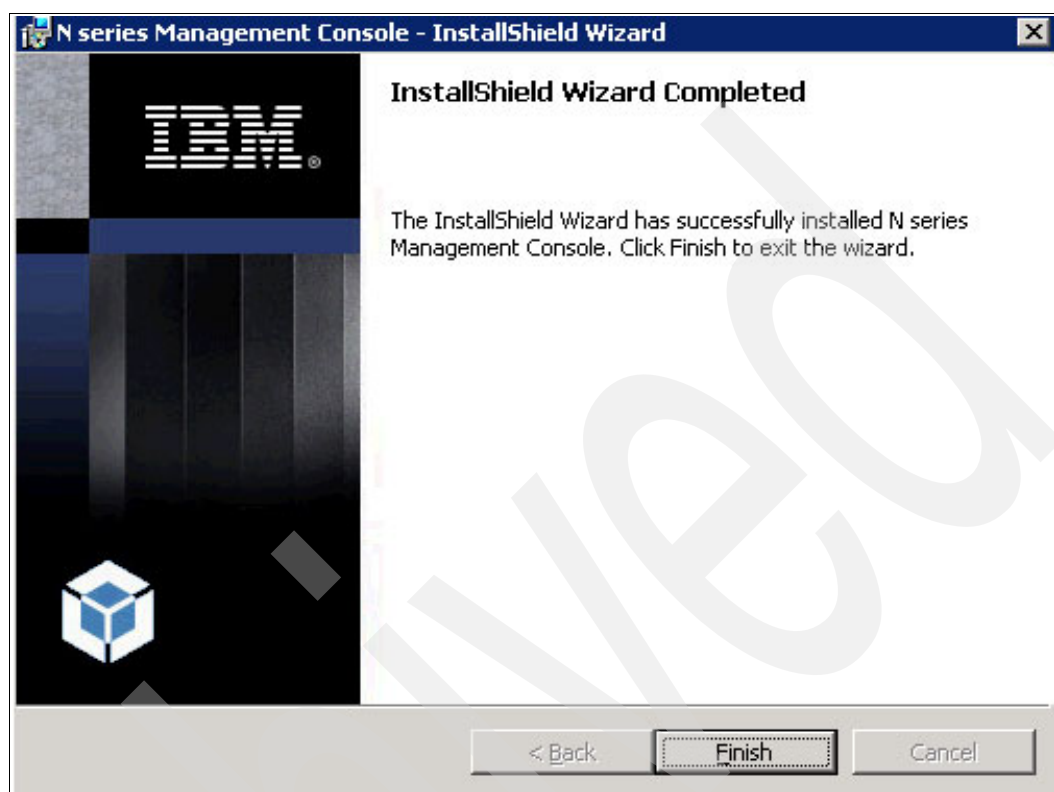


Figure 8-14 N series Management Console installation completed window

11. Figure 8-15 shows that the N series Management Console icon is now available on the Windows desktop.



Figure 8-15 N series Management Console

12. Double-click the N series Management Console icon and you will get the window shown in Figure 8-16.

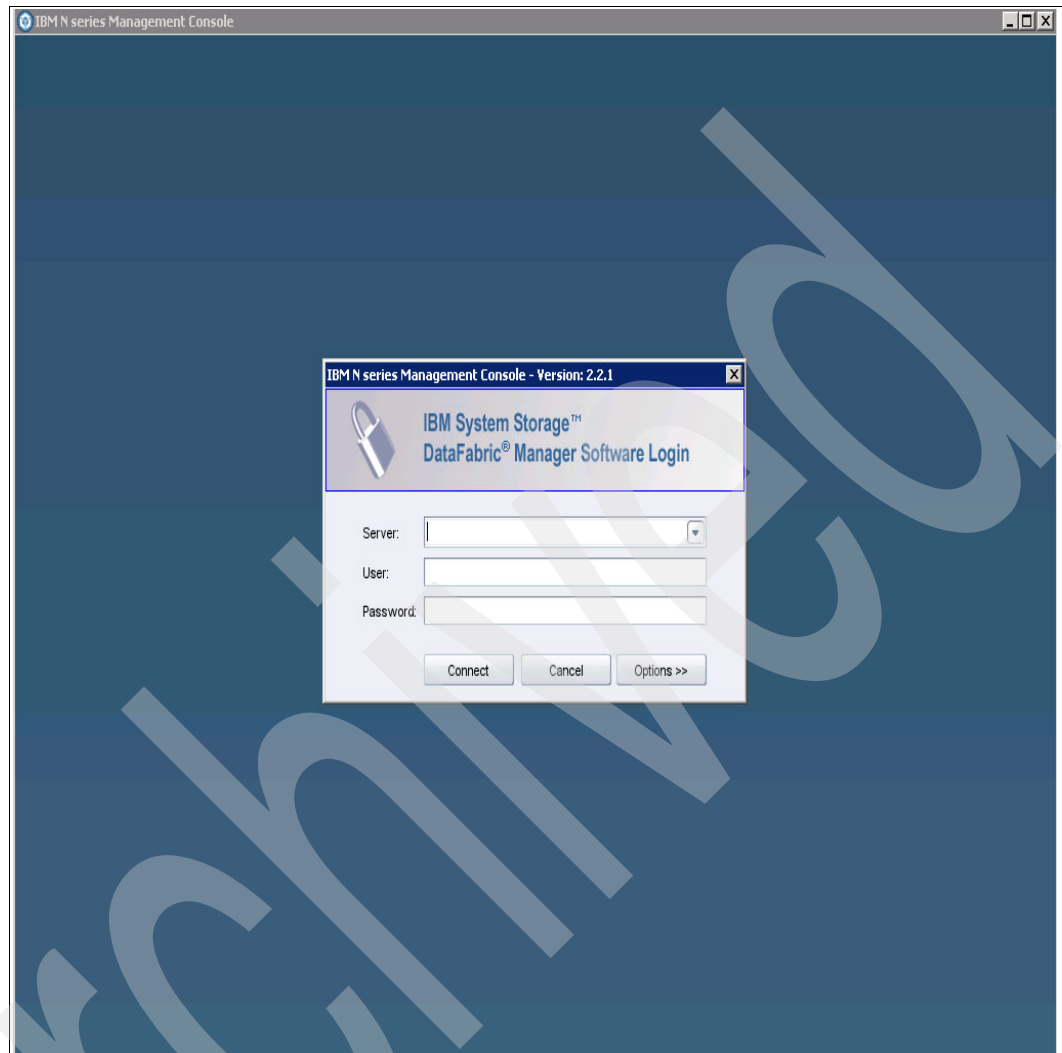


Figure 8-16 N series Management Console main window

In this window, you have three options:

- Server

You have to enter the IP address or DNS name of the Operations Manager if the Operations Manager is installed on a separate server. We recommend this approach because installing the N series Management Console and the Operations Manager on the same server impacts the performance of Operations Manager. We installed Operations Manager and the N series Management Console on the same server because this is a test environment, so we use localhost.

- User

Specify the user name of the Operations Manager. We use the administrator ID.

- Password

Specify the Operations Manager password.

Once you provide with the above information, you will get the window shown in Figure 8-17, which shows the Dashboard window before the discovery of the N series storage systems by the Operations Manager.

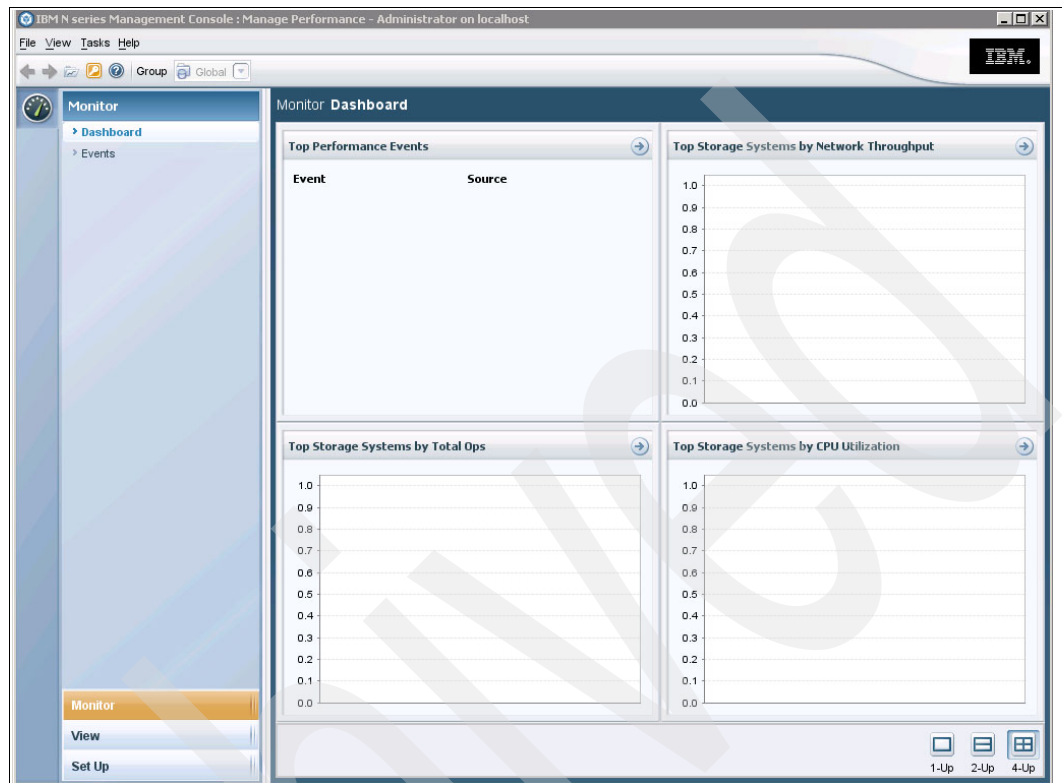


Figure 8-17 N series Management Console Dashboard window

Figure 8-18 shows the Events window before the discovery of the N series storage systems by the Operations Manager.

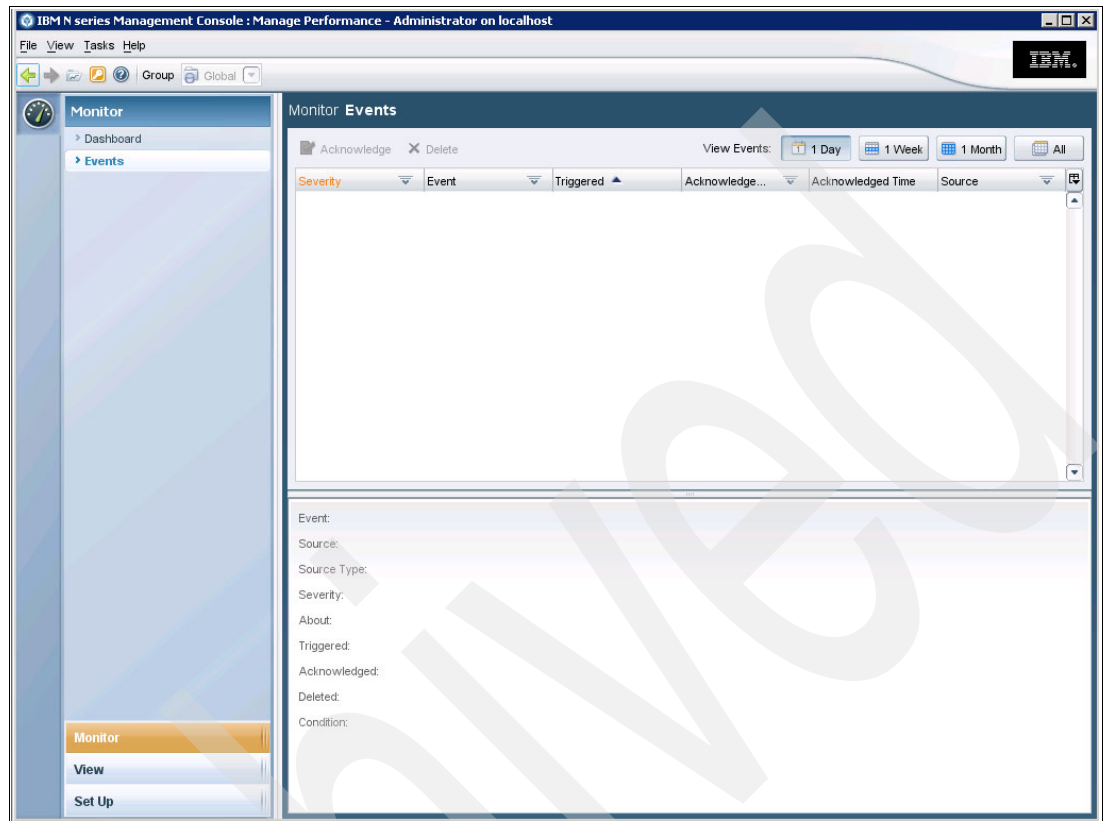


Figure 8-18 N series Management Console Events window

This concludes our discussion about the installation of the N series Management Console installation for Windows.



Part 4

Configuring Operations Manager

This part will cover the setup of Protection Manager, Provisioning Manager, and File Storage Resource Manager.

Archived



Configuring Operations Manager

The Operations Manager installation is very simple and is covered in several chapters of this book, depending on the operating system you chose. However, the application itself has many options and supports many aspects of managing and monitoring a NetApp® and IBM System Storage N series environment. In this chapter, we will cover many of the features and functions of Operations Manager at a configuration level, but we recommend that you refer to the *Operations Manager Administration Guide for Use With DataFabric Manager Server 3.7*, GC26-7889, for more detailed information.

9.1 What is Operations Manager

Operations Manager is a management platform that will give you a better look into your IBM N series and NetApp infrastructure. You will be able to access such services as discovery and grouping, monitoring and alerting, reporting, and many other activities needed to effectively manage a complex storage environment. You can manage this environment from either a graphical user interface or from the DataFabric Manager Server CLI. DataFabric Manager Server is part of the Operations Manager infrastructure. It is activated by using the core license for Operations Manager. Operations Manager does not run on storage systems themselves, but is run from a dedicated server properly sized for the environment you have. Monitoring of your storage infrastructure is accomplished by Operations Manager detecting N series storage devices as well as loading a Host Agent on clients that reside in your N series environment. For complete sizing information, refer to the *IBM System Storage N series Operations Manager Sizing and Installation Guide*, REDP-4270.

Figure 9-1 shows an overview of the Operations Manager infrastructure.

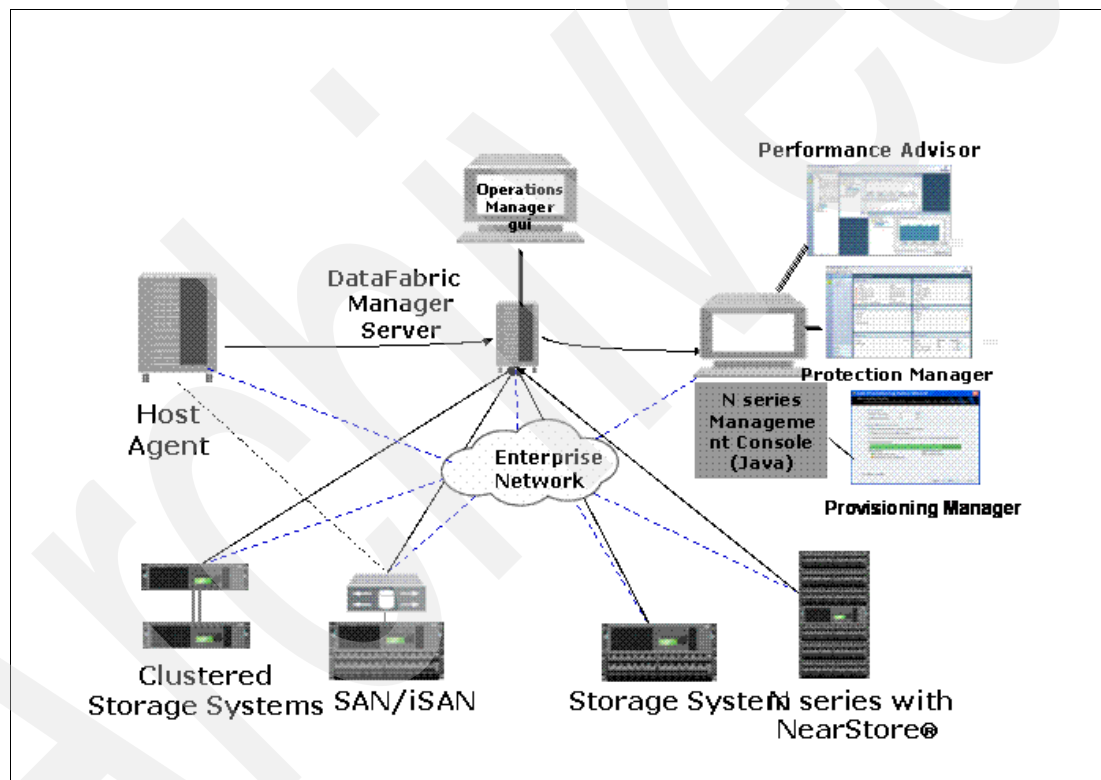


Figure 9-1 Operations Manager infrastructure

It is not within the scope of this IBM Redbooks publication to cover the process of upgrading to Operations Manager V3.7.1 from a previous version, but the following suggestions apply before proceeding with an upgrade.

- ▶ Use the **dfm backup create** command to back up the Operations Manager database before upgrading.
- ▶ Dedicate a machine to Operations Manager and all of its components.
- ▶ Set the browser to support Java Applets if you are planning to use Disaster Recovery Manager.

9.2 Operations Manager licensing considerations

Before using the Operations Manager application, you need to consider what licenses you will need. There is a core license for Operations Manager that will give you the ability to view and monitor N series and NetApp storage devices, but that is all. You will not be able to access any of the file and volume data beyond monitoring. You must add several licenses that will give you additional abilities beyond monitoring alone.

Table 9-1 gives you a list of features activated by each license.

Table 9-1 Licenses for features

To use	Install this license or product	That enables these features
Operations Manager	DataFabric Manager Server license	Automated policy-based data protection for N series NAS and SAN storage systems
	Operations Manager license	SnapVault, Open Systems SnapVault, and SnapMirror management
		Policy conformance checking and alerting
		Monitoring
		Reports
		Storage usage and availability, such as qtrees, volumes, aggregates, LUNs, and disks
		Storage systems
		Vfiler units
		Real-time streaming events
		Managing <ul style="list-style-type: none"> ► Storage system configuration ► Scripts
		Monitoring and managing storage system clusters using Cluster Console
	Required for all licensed Operations Manager installations. Sets the maximum number of storage systems that the DataFabric Manager Server can monitor	Displaying historical and real-time performance data using Performance Advisor in N series Management Console
Protection Manager	DataFabric Manager Server license	Automated policy-based provisioning for N series SAN and NAS storage systems
	Operations Manager license	Space management policies and capacity reporting
	Protection Manager license	Policy conformance checking and alerting
	N series Management Console	Enables management of Provisioning Manager and Protection Manager

To use	Install this license or product	That enables these features
Provisioning Manager	DataFabric Manager Server license	Automated policy-based provisioning for N series SAN and NAS storage systems
	Operations Manager license	Space management policies and capacity reporting
	Provisioning Manager license	Policy conformance checking and alerting
	N series Management Console	If you have both Protection Manager and Provisioning Manager licensed, then the following features are enabled: <ul style="list-style-type: none"> ▶ Assigning provisioning policies to nonprimary nodes ▶ Policy-based provisioning of primary storage ▶ Assigning protection policies to provisioned data sets
Protection Manager with Disaster Recovery	DataFabric Manager Server license	Failover and manual failback for N series NAS and SAN storage systems
	Operations Manager license	
	Protection Manager license	
	Protection Manager Disaster Recovery license	
	N series Management Console	
File Storage Resource Manager (File SRM)	DataFabric Manager Server license	Tracking file system usage and capacity information
	Operations Manager license	
	File SRM Option	
Business Continuance	DataFabric Manager Server license	Backup Manager
	Operations Manager license	This space left intentionally blank.
	Business Continuance Option (BCO)	Configuring and scheduling disk-to-disk backups of all systems enabled with SnapVault, including Open Systems SnapVault
		Disaster Recovery Manager <ul style="list-style-type: none"> ▶ Monitoring SnapMirror relationships ▶ Configuring and scheduling disk-to-disk mirrors of all systems enabled with SnapMirror
Storage Area Network	DataFabric Manager Server license	Monitoring and managing SAN hosts and FC switches
	Operations Manager license	Managing LUNs on Windows SAN hosts
	Storage Area Network Option	

Figure 9-2 on page 131 shows what the license window looks like when all the licenses have been installed.

Options

Licensed Features	
Operations Manager	demo expires 20 Jul 2009
Business Continuance Option	demo expires 20 Jul 2009
File SRM Option	demo expires 20 Jul 2009
Protection Manager	demo expires 18 May 2009
Provisioning Manager	demo expires 18 May 2009
Licensed Node Limit	250
New License Key	<input type="text"/>

Figure 9-2 Fully licensed Operations Manager

The node count for the server license is one. You must use the server license to monitor storage systems. The server license sums up the node count of each additive license to arrive at the total node count.

An additive license enable features and increases the node count without replacing the installed licenses. The serial numbers of the core license and the additive licenses do not have to match. You can install multiple additive licenses to increase the node count for a feature

Refer to Figure 9-3 for an example of determining node count.

Licensing

- When you purchase DataFabric Manager, IBM provides DataFabric Manager server license and Operations Manager additive license that allow you to install the software and access a specific set of features. You need these licenses to monitor and manage storage systems.
- Example to monitor a N3600A20, a N6060A12 and a N3300A10 you will need the following licenses:
 - 2 FC 8262 Ops Mgr Core tier 2 license
 - 1 FC 8265 Ops Mgr Core tier 5 license
 - 1 FC 8261 Ops Mgr Core tier 1 license

N3300	2859	A10	8261	Ops Mgr Core tier 1 license
		A20	8261	Ops Mgr Core tier 1 license
N3600	2862	A10	8262	Ops Mgr Core tier 2 license
		A20	8262	Ops Mgr Core tier 2 license
N6040	2858	A10	8264	Ops Mgr Core tier 4 license
		A20	8264	Ops Mgr Core tier 4 license
N6060	2858	A12	8265	Ops Mgr Core tier 5 license
		A22	8265	Ops Mgr Core tier 5 license
N6070	2858	A11	8265	Ops Mgr Core tier 5 license
		A21	8265	Ops Mgr Core tier 5 license
N7700	2866	A21	8266	Ops Mgr Core tier 6 license
N7900	2867	A21	8267	Ops Mgr Core tier 7 license

Figure 9-3 Node count example

9.3 Discovery

Operations Manager uses two methods for locating and identifying storage systems and networks: autodiscovery and manual addition. Autodiscovery is typically the primary process Operations Manager uses to locate and identify storage systems and networks.

In the discovery process, Operations Manager and systems (storage systems and vFiler units) communicate automatically. Discovery reduces the work required to locate and identify systems by replacing manual input with an automated process.

The first thing that happens when you start Operations Manager is an automatic discovery of all of the N series and NetApp storage systems within view of your server. Discovery is the process of locating device, including N series and NetApp storage systems, host agents, and Open Systems SnapVault (OSSV). Discovery is enabled at the volume, qtree, and LUN level.

Manual addition is secondary to the discovery process. You typically only need it for storage systems and networks that are added after the existing infrastructure has been discovered.

Host discovery

Operations Manager automatically locates and identifies storage systems and networks through its built-in discovery function. Refer to Figure 9-6 on page 134 for the Discovery Options window.

Discovery method and topology

The most efficient method of discovery depends on the topology of the networks.

Host discovery is typically used for all storage systems that are in the same geographical area as the server on which Operations Manager is installed.

When you install Operations Manager, the Host Discovery option is enabled by default. Host discovery allows Operations Manager to automatically discover storage systems and vFiler units. To include storage systems on other networks in Operations Manager, you can manually add networks and systems.

Figure 9-4 on page 133 shows how to add additional networks to be monitored by Operations Manager.

Discovery and Grouping

- Discovers All N series Devices through Automatic and Manual Discovery
- By default, all devices are grouped under “Global” group
- Under Automatic discovery, desired subnets can be added

The screenshot shows the 'Discovery' configuration page. On the left, there are several sections with expandable options:

- Host (Appliance) Discovery Options:** Host Discovery (Enabled), HTTP Discovery (Enabled), HTTP Discovery Port (3132), Appliance-Initiated Discovery (Enabled), SNMP On HTTP-Discovered Appliances (Disabled).
- SAN Discovery Options:** SAN Device Discovery (Enabled).
- Agent Discovery Options:** Host Agent Discovery (Enabled), vFilter Discovery Options (vFilter Discovery (Enabled)).
- Network Discovery Options:** Network Discovery (Enabled), Network Discovery Limit (in hops) (15), Networks To Discover (null), Settings for All Discovery Options (except for vfilters) (Discovery Interval: 15 minutes, Discovery Timeout: 5 seconds, Network Credentials: null).

On the right, there is a 'Networks To Discover' section with a table:

Network Address	Network Mask	Hop Count	Last Searched	Edit	Delete
10.73.88.0	255.255.254.0	0	08 Jul 14 33	edit	delete
10.73.191.52	255.255.255.252	1	08 Jul 14 34	edit	delete
10.73.191.56	255.255.255.252	1	08 Jul 14 31	edit	delete
192.168.190.0	255.255.254.0	1	08 Jul 14 29	edit	delete

Buttons at the bottom include 'Back' and 'Update'.

Figure 9-4 Add additional subnets to discover storage systems and hosts

To add networks, select **Setup** → **Discovery** on the Control Center tab, as shown in Figure 9-5.

The screenshot shows the 'Control Center' interface. The top navigation bar includes 'Home', 'Setup', 'Reports', 'Management', and 'Help'. The 'Setup' menu is expanded, showing a list of options: 'Options', 'Discovery', 'Network Credentials', 'Groups', 'Administrative Users', 'Roles', 'Alarms', 'Database Backup', 'Download Management Console', 'Distribution ACL', and 'Events and Alerts'. The 'Discovery' option is highlighted. The right side of the screen shows the 'Discovery' configuration page, similar to the one in Figure 9-4.

Figure 9-5 Selecting the Discovery menu

In the window shown in Figure 9-6, edit the Networks to Discover field.

The screenshot shows the 'Options' window with the 'Discovery' tab selected. The 'Networks To Discover' field is highlighted with an arrow pointing to an 'edit' link. The 'edit' link is located next to the 'Networks To Discover' field, which currently contains the value '15'. The 'edit' link is a small text link that says 'edit'.

Figure 9-6 Edit Networks to Discover field

Add networks and subnets to discover additional storage systems and Hosts Agents, as shown in Figure 9-7.

The screenshot shows the 'Networks To Discover' window. It has a form to add a new network with fields for 'Network Address' and 'Network Mask'. Below the form is a table listing discovered networks. The table has columns for 'Network Address', 'Network Mask', 'Hop Count', 'Last Searched', 'Edit', and 'Delete'. The first row shows a network address of '9.11.218.0' and a network mask of '255.255.255.0'.

Network Address	Network Mask	Hop Count	Last Searched	Edit	Delete
9.11.218.0	255.255.255.0	0	07 Apr 10:32	edit	<input type="checkbox"/>

Figure 9-7 Add networks and subnets for additional systems and hosts

9.4 Operations Manager discovery methods

The Discovery Options window gives you several methods of discovery. The storage systems in your network determines the discovery method used. Figure 9-8 on page 135 shows the different “discoveries” that can take place in your storage system network.

Options

Discovery

Host (Appliance) Discovery Options

- Host Discovery: Enabled
- HTTP Discovery: Enabled
- HTTP Discovery Port: 3132
- Appliance-Initiated Discovery: Enabled
- SNMP On HTTP-Discovered Appliances: Disabled

Agent Discovery Options

- Host Agent Discovery: Enabled

vFiler Discovery Options

- vFiler Discovery: Enabled

Network Discovery Options

- Network Discovery: Disabled
- Network Discovery Limit (in hops): 15
- Networks To Discover: [edit](#)

Settings for All Discovery Options (except for vFilers)

- Discovery Interval: 15 minutes
- Discovery Timeout: 5 seconds
- Network Credentials: [edit](#)

Figure 9-8 Different discoveries

Operations Manager uses SNMP queries for host discovery. Host discovery supplies the information required for storage systems to be monitored and managed through Operations Manager. You must have *SNMP enabled* on your storage systems and routers for Operations Manager to monitor and manage systems. By default, SNMP is enabled on storage systems.

Ping methods might include ICMP echo, HTTP, SNMP, NDMP, or ICMP echo and SNMP.

As soon as Operations Manager is installed, it attempts to discover storage systems on the local subnet. This means that you need to ensure that SNMP is enabled on your storage systems before Operations Manager is installed, if you want Operations Manager to discover them immediately.

You can also wait until after installing Operations Manager to enable SNMP on storage systems, but this causes a delay before Operations Manager can discover those storage systems.

For more detailed information about discovery methods and processes, including step by step information about how discovery is performed by Operations Manager, refer to Chapter 3 in *Operations Manager Administration Guide for Use With DataFabric Manager Server 3.7*, GC26-7889.

SNMPv1 or SNMPv3

Now that you understand that Operations Manager uses SNMP for Storage System discovery, you need to decide whether to use SNMPv1 or SNMPv3.

When Operations Manager is installed for the first time or updated, the global and network setting uses SNMPv1 as the preferred version by default. However, you can configure the global and network setting to use SNMPv3 as the default version.

You can set SNMPv3 as the preferred version for the storage system discovery on a specific network by completing the following procedure:

1. From the Control Center window, select **Setup** → **Network Credentials**, as shown in Figure 9-9. Alternatively, you can select **Setup** → **Discovery** to access the same options window, as shown in Figure 9-10.

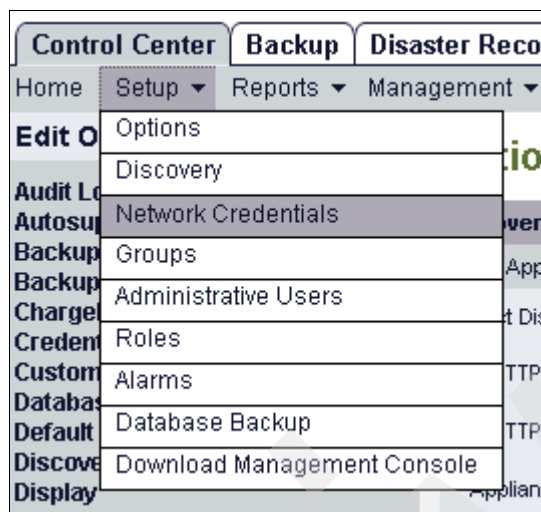


Figure 9-9 Select Setup and Network Credentials

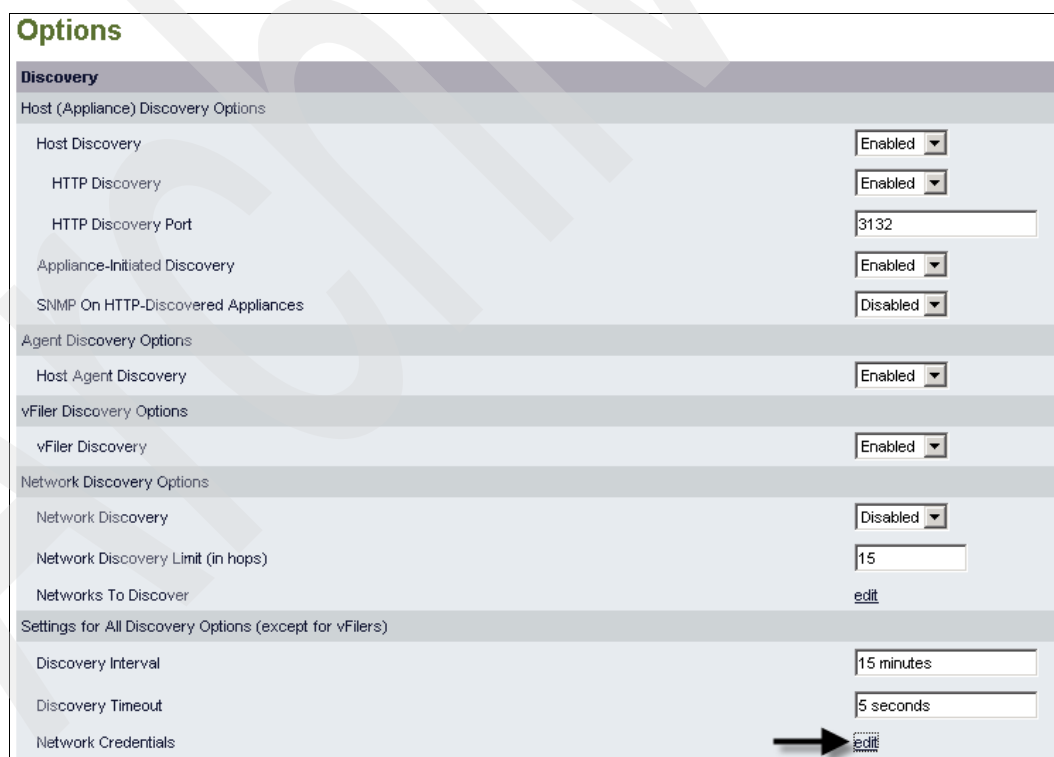


Figure 9-10 Select Setup and Discovery

2. From the Network Credentials window, you can choose which SNMP version to use, as shown in Figure 9-11 on page 137.

Network Credentials

Add Network Credentials

IP Address

Network Mask

Preferred SNMP Version

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)

IP Address	Network Mask	Preferred SNMP Version	SNMP Community
default		SNMPv1	public

Figure 9-11 Select which SNMP protocol you want to use

All storage systems in the selected network will be monitored using SNMPv3. By default, all networks are monitored using SNMPv1.

Because of the way discovery takes place, when all or most of the storage systems in a network are running only a particular SNMP version, we recommend specifying only that version as the preferred SNMP version for the network. This speeds up the discovery of storage systems running only a particular SNMP version.

We suggest using SNMPv3 over SNMPv1, because SNMPv3 is more secure. SNMPv1 is the earliest implementation of SNMP and therefore does not have the safeguards for security implemented in SNMPv3. SNMPv3 offers three things not found in version 1: authentication, privacy, and access control.

More information about SNMP implementation with Operations Manager can be found in *Operations Manager Administration Guide for Use With DataFabric Manager Server 3.7*, GC26-7889.

9.5 Grouping

Once discovery is complete, you can group storage objects in several ways. This can be helpful in managing a complex storage environment. You can group according to applications, locations of the storage devices, type of devices, departments they serve, or any way that makes sense to your environment.

Management can be done by group and any action you need to be allied to a group can be accomplished quickly. You can pull reports by group and run simultaneous Operations Manager commands by group.

Figure 9-12 shows an example of different groups that can be created.

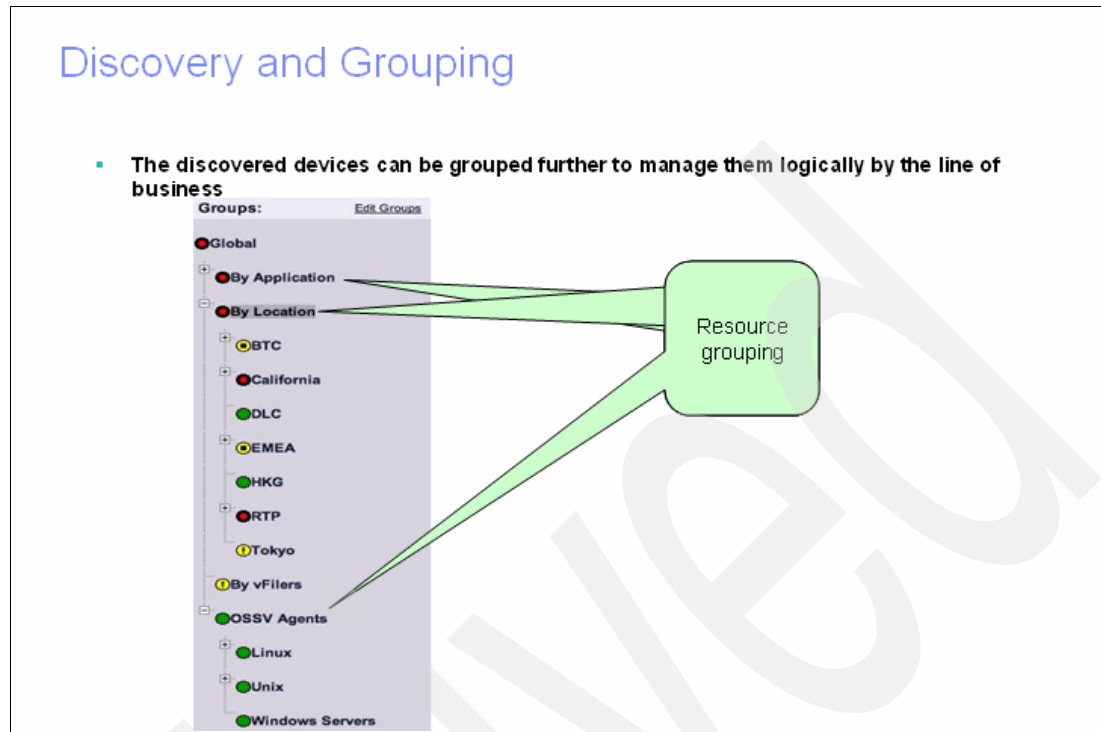


Figure 9-12 Example of groups

Operations Manager automatically determines the type of a group based on the objects it contains. If you place your cursor over an icon to the left of a group name on the left side of Operations Manager main window, you can quickly discover the type of objects the group contains.

9.5.1 Homogeneous groups

You can group objects into sets of objects with common characteristics. They might, for example, have the same operating system or belong to a specific project or group in your organization.

You can create the following types of groups:

- ▶ Appliance Resource group: Contains storage systems, vFiler units, or host agents
- ▶ Aggregate Resource group: Contains aggregates only
- ▶ File System Resource group: Contains volumes, qtrees, or both
- ▶ LUN Resource group: Contains LUNs only
- ▶ Configuration Resource group: Contains storage systems associated with one or more configuration files
- ▶ SRM path group: Contains SRM paths only
- ▶ Data set: The data stored in a collection of primary storage containers, including all the copies of the data in those containers
- ▶ Resource pool: A collection of storage objects from which other storage containers are allocated

When you bring up the Edit Groups Membership window, you can see, in the drop-down menu, the different categories you can use to group the devices. The items mentioned above are listed here, as shown in Figure 9-13.

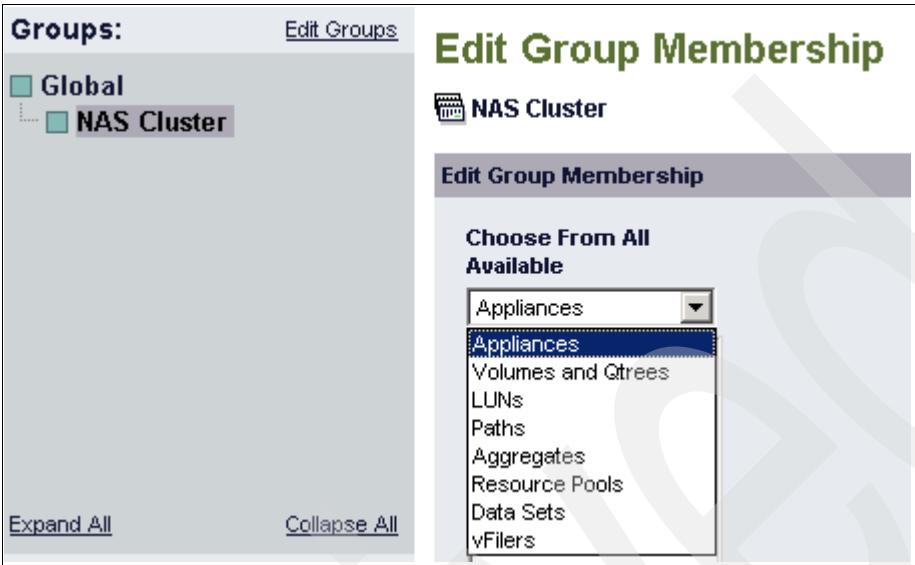


Figure 9-13 Group types

Though there may be many reasons to create groups, you decide how you would like to group your storage environment. Remember, all items are a member of at least one group (Global), but can just as easily be a member of two or more groups. For example, you could group IBM System Storage N series 5200 gateways in the same group, as shown in Figure 9-14. Notice that they are in the Appliance group.

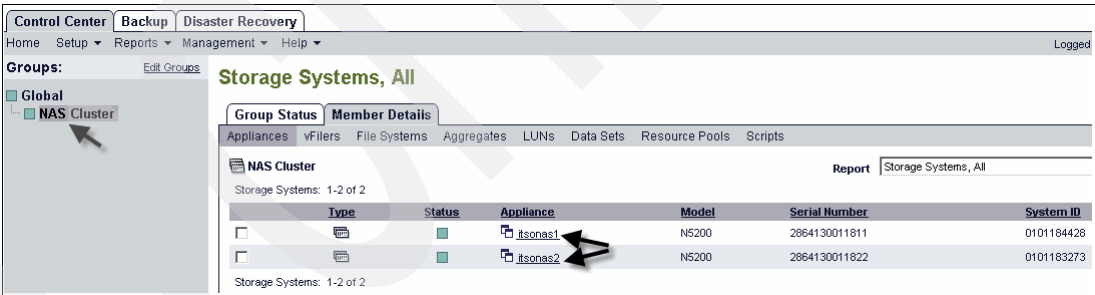


Figure 9-14 Group example

When discovery is first done, all storage devices and hosts fall under the Global group. From there, you can create your individual groups. You can either add to a current group or create a new group. Figure 9-15 shows an example.

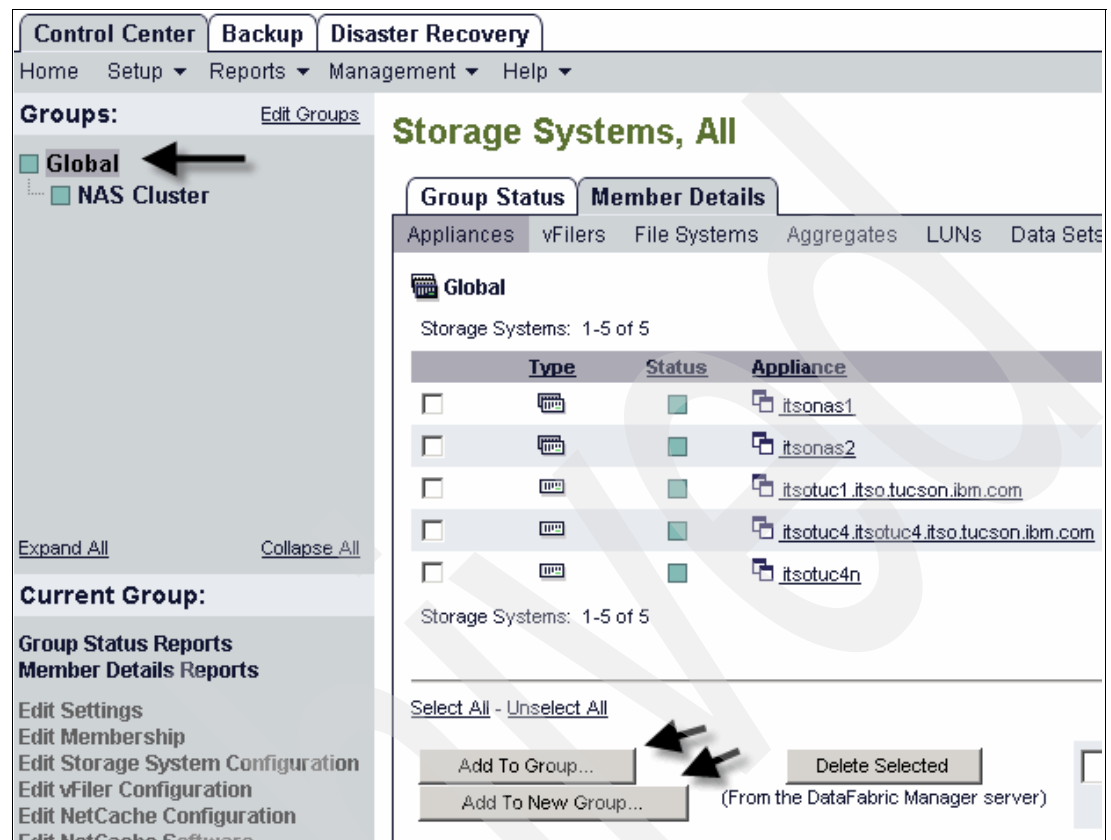


Figure 9-15 Select Add to Group or Add to New Group

Use the following guidelines when you create groups:

- ▶ You can group similar or mix-types objects in a group.
- ▶ An object can be a member of any number of groups.
- ▶ You can group a subset of group members to create a new group.
- ▶ You cannot create a group of groups.
- ▶ You can create any number of groups.
- ▶ You can copy a group or move a group in a group hierarchy.

.For more detailed information about groups, refer to Chapter 5 in *Operations Manager Administration Guide for Use With DataFabric Manager Server 3.7*, GC26-7889.

9.6 Reports

Operations Manager provides standard reports that you can view from the DataFabric Manager Server CLI or Operations Manager interface. You can run reports and create custom reports from the DFM Server CLI. However, Operations Manager is designed to provide reports in easy-to-use Operations Manager interface, in which you can do the following:

- ▶ View a report.
- ▶ Save a report in CSV format.
- ▶ Print a report.
- ▶ Create a report.
- ▶ Delete a custom report. You cannot delete a standard report.
- ▶ Use a custom report as a template to create a report.

The Reports drop-down menu will give you quick access to several reporting features, as shown in Figure 13-15.

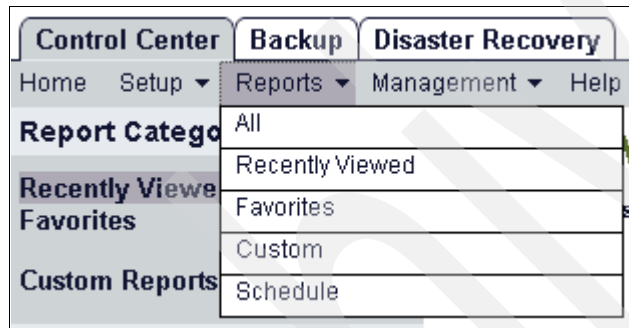


Figure 9-16 Quick access to your favorite reports

From this drop-down menu, you can select All reports, Recently Viewed, Favorites, Custom and a time to schedule when reports are run. Using DataFabric Manager V3.6.1 or later, you can search all the reports from Reports menu All. All the reports are divided into the following categories:

- ▶ Recently Viewed
- ▶ Favorites
- ▶ Custom Reports
- ▶ Logical Objects
- ▶ Physical Objects
- ▶ Monitoring
- ▶ Performance
- ▶ Backup
- ▶ Disaster Recovery
- ▶ Miscellaneous

You can run reports against a single entity or an entire group of storage systems. Figure 9-17 gives you some idea of the many that are available.

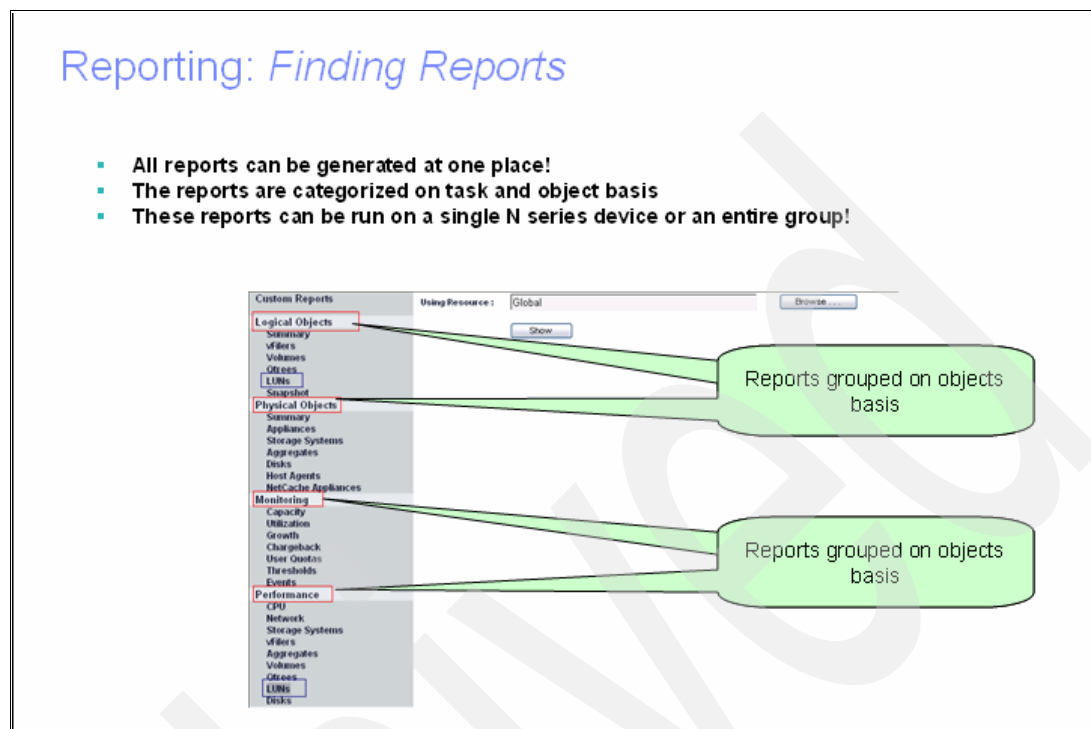


Figure 9-17 Sample list of reports

Let us take a look for a moment at some reports that are available. Say we want to get a summary of our logical objects. To do this task, select **Summary** under Logical Objects, as shown in Figure 9-18.



Figure 9-18 Summary report for Logical Objects

When you select the report you want to view, you receive a list of optional reports available under that tab, as shown in Figure 9-19 on page 143.

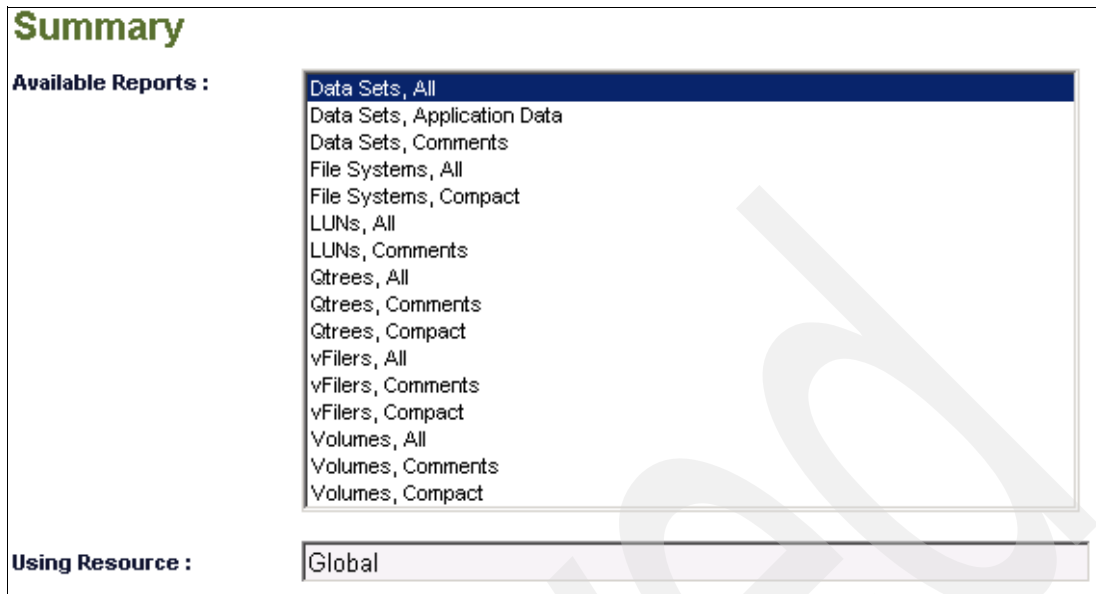


Figure 9-19 Summary reports available

Here we select the File Systems (Compact) report, as shown in Figure 9-20.

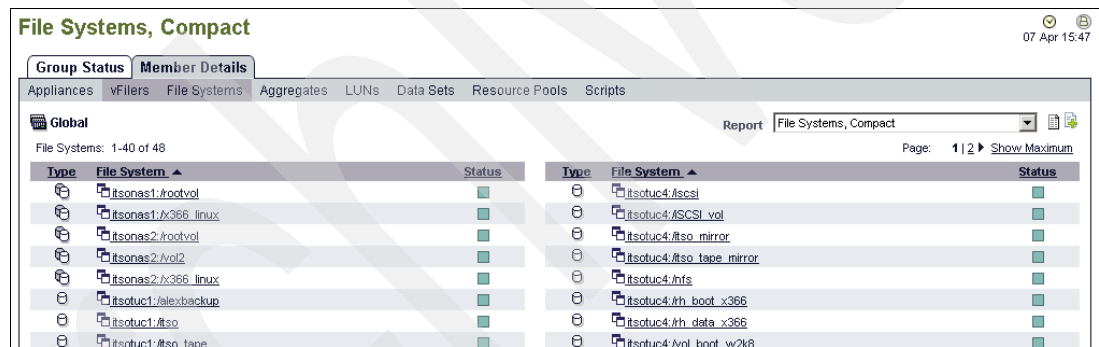


Figure 9-20 File Systems (Compact) report

As you can see, there is a broad range of reports already built into Operations Manager. let us take a look at a few more standard reports.

An example might be a business unit that has grouped its storage systems based on applications. For example, we want to create an asset report for an Oracle® application that we refer to as “App1”. All we need to do is select the group “App1” and select **Appliances** under Member Details. Now we have the asset report for the Apps1 group, as shown in Figure 9-21.



Figure 9-21 Asset Report for Apps1

Operations Manager is helpful for space planning by producing reports useful for trending consumption. With the use of flexible volumes, it is very important to keep track of flexvols that might put the aggregate in danger of being overcommitted. A requirement of thin provisioning, for example, is that the administrator be aware of the growth of any flexvol(s) that might put the aggregate in danger of running out of space.

Effective storage management means having as much detail about storage utilization as possible. It also means being proactive in addressing potential problems. Operations Manager lets you see at a glance which volumes may need attention and how much space is taken by Snapshots, space reservation, and other allocations.

Operations Manager presents trending reports showing, based on calculated daily growth rate, when a volume will become full. This report also exists for aggregates and qtrees, as shown in Figure 9-22 on page 145.

Note the “Days to Full” column. Based on the daily growth rate, this will give you some idea when the aggregate will be full. This gives you time to take action before a problem arises, as shown in Figure 9-23 on page 145.

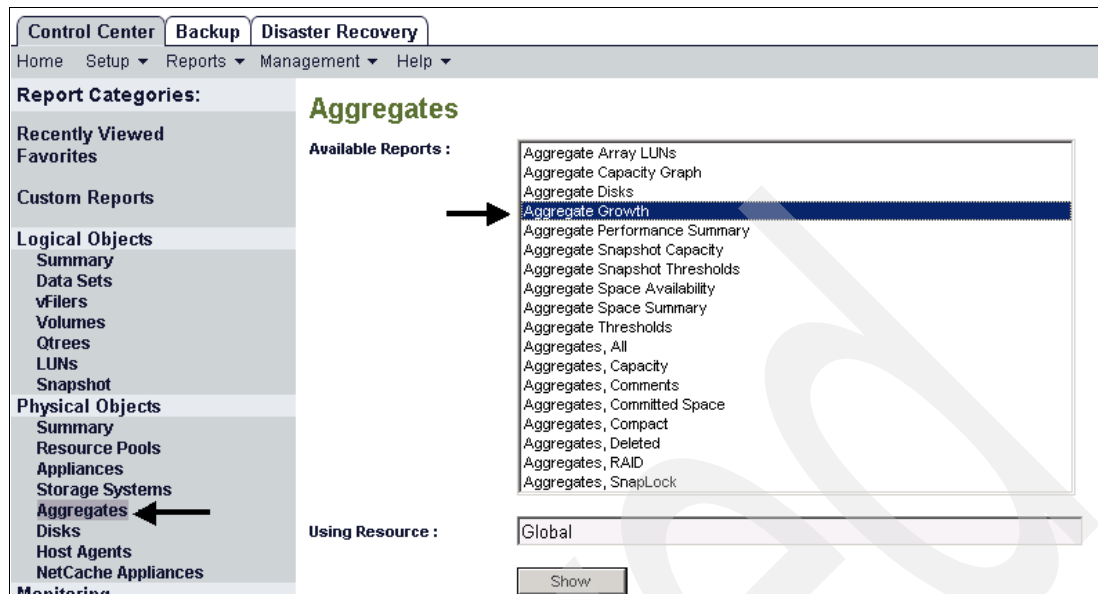


Figure 9-22 Where to find the Aggregate Growth report



Figure 9-23 Growth and Trending example

Efficient storage management is more than tracking how much storage is used. It is important to know how much storage is taken up by Snapshots and whether or not that consumption is growing.

When space is at a premium, it is valuable to have the option of deleting Snapshots in order to make space for data or other Snapshots.

Two things are necessary in order to decide which Snapshots can be deleted. First, there can be no dependencies, for example, SnapVault relationships, that would prohibit a Snapshot from being deleted. Operations Manager will inform you of any dependencies and give you the necessary steps to remove them if desired.

It is also convenient to know how old the Snapshots are and how much space would be recovered upon their deletion. Figure 9-24 shows an example of how you can discover this information from the reports already provided in Operations Manager.

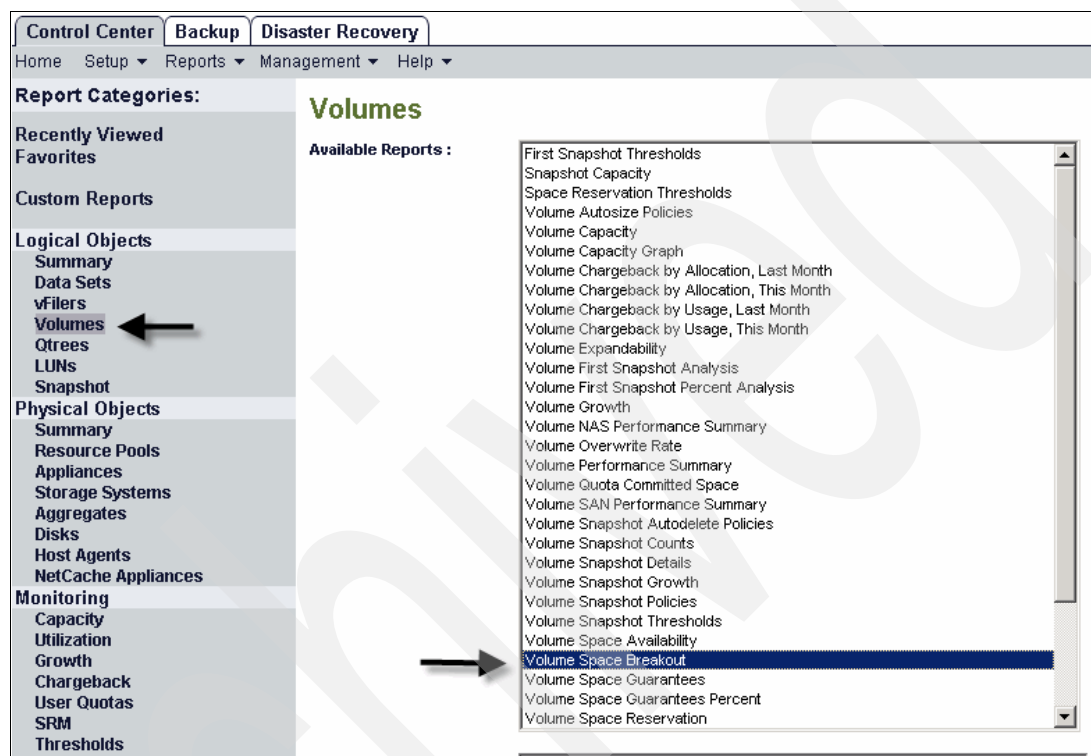


Figure 9-24 Volume Space Breakout location

Space breakout

The bar graph shown in Figure 9-25 on page 147 indicates the percentage of space used for Snapshots and volume data. If the volume has a Snapshot reserve, the vertical bar on the right indicates the size of the Snapshot reserve as a proportion of total volume size. If the Snapshot reserve size is zero (the recommended setting for SAN volumes), the vertical bar is not visible.

Reporting: *Space Availability Reports*

- Operations Manager provides various capacity reports which are helpful to understand space allocation and reclamation.

Volume Space Breakout

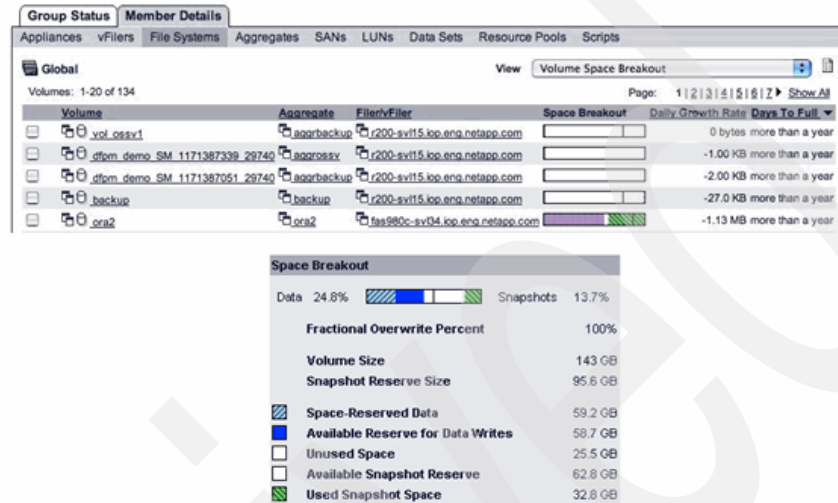


Figure 9-25 Volume Space Breakout window

The data displayed in this area is different for NAS and SAN volumes:

- ▶ A NAS volume does not contain space-reserved files.
- ▶ A SAN volume contains one or more space-reserved files.

The above example for space breakout (Figure 9-25) shows a SAN volume.

The Space Breakout window also shows the following information:

- ▶ Space-Reserved Data

The amount of data used to store LUNs and other space-reserved files.

- ▶ Available Reserve for Data Writes

The amount of reserved space available for write operations. This field appears only if there is at least one Snapshot copy, because space is not allocated to the reserve until the first Snapshot copy is made.

- ▶ Unused Space

The amount of free space remaining in the volume. This field appears only if its value is not zero.

- ▶ Available Snapshot Reserve

The amount of free space remaining in the Snapshot reserve. This field appears only when there is some space remaining, which is not recommended for SAN volumes.

- ▶ Used Snapshot Space

The amount of volume Snapshot reserve space currently storing Snapshot copies.

Another standard report available from Operations Manager is the Chargeback report, which is shown in Figure 9-26. You can find this report by selecting the **Chargeback** item in the Report Categories pane, as shown in Figure 9-27.

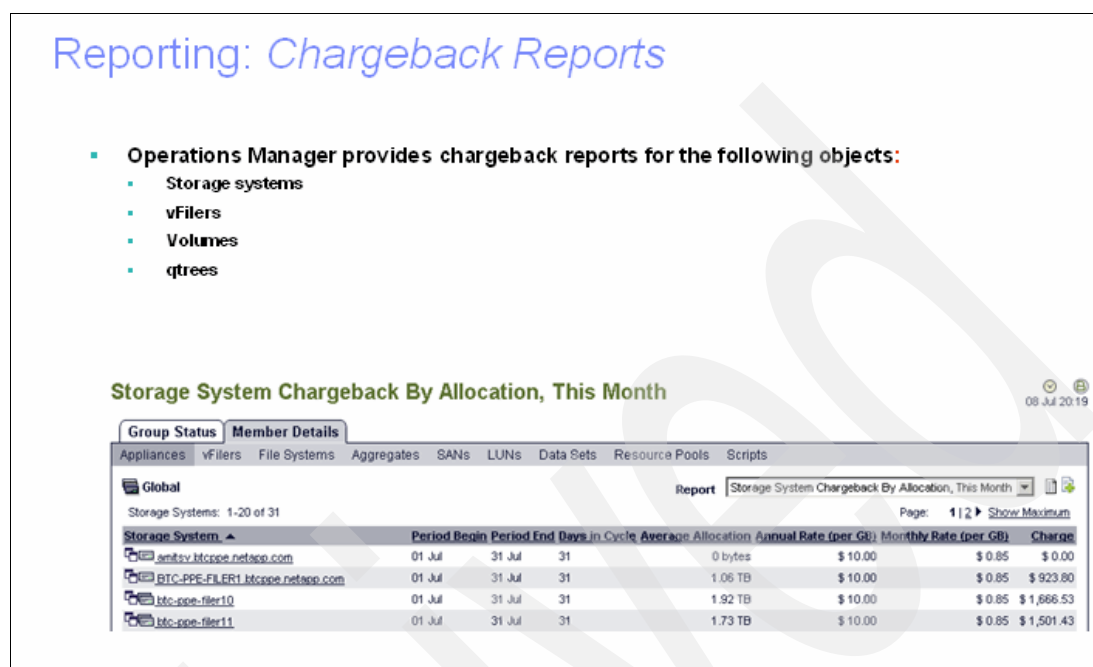


Figure 9-26 Sample Chargeback report

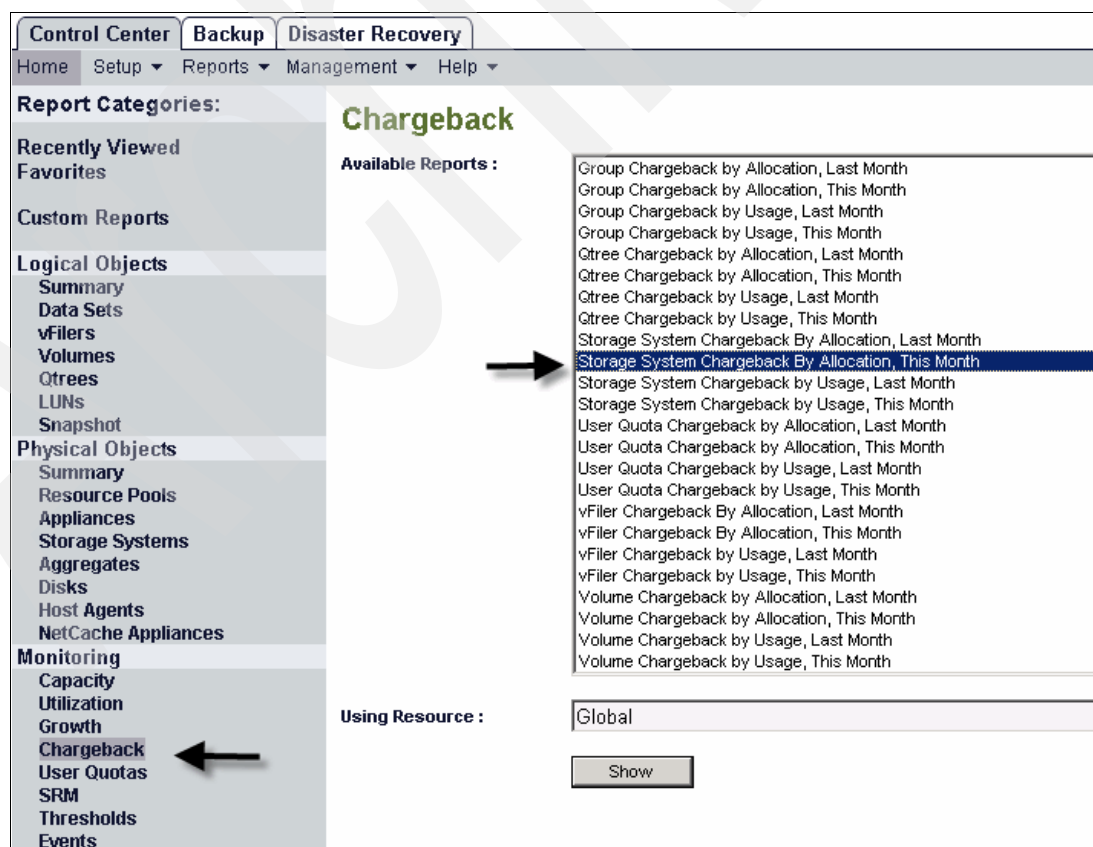


Figure 9-27 Where to find the Chargeback report

There are many standard reports within Operations Manager. However, it is still possible that a report you need is not available. Or you may want to use a report as a template because it either does not have all the information you want on it or it has too much. Operations Manager gives you the ability to create custom reports to meet your needs, as shown in Figure 9-28.

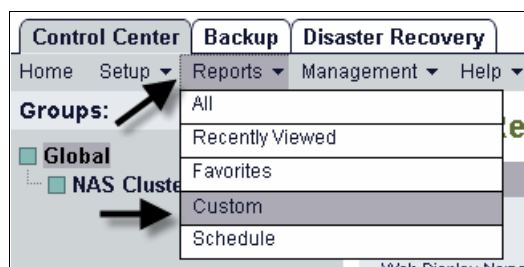


Figure 9-28 Creating a custom report

Operations Manager provides report catalogs that you use to customize reports. You can set basic report properties from the DFM Server CLI or Operations Manager interface, such as a short report name (for DFM Server CLI output), long report name (for Operations Manager output), field description, the fields to display, and the report catalog it was created from.

Every report generated by Operations Manager, including those you customize, is based on the catalogs that are listed in Figure 9-29.

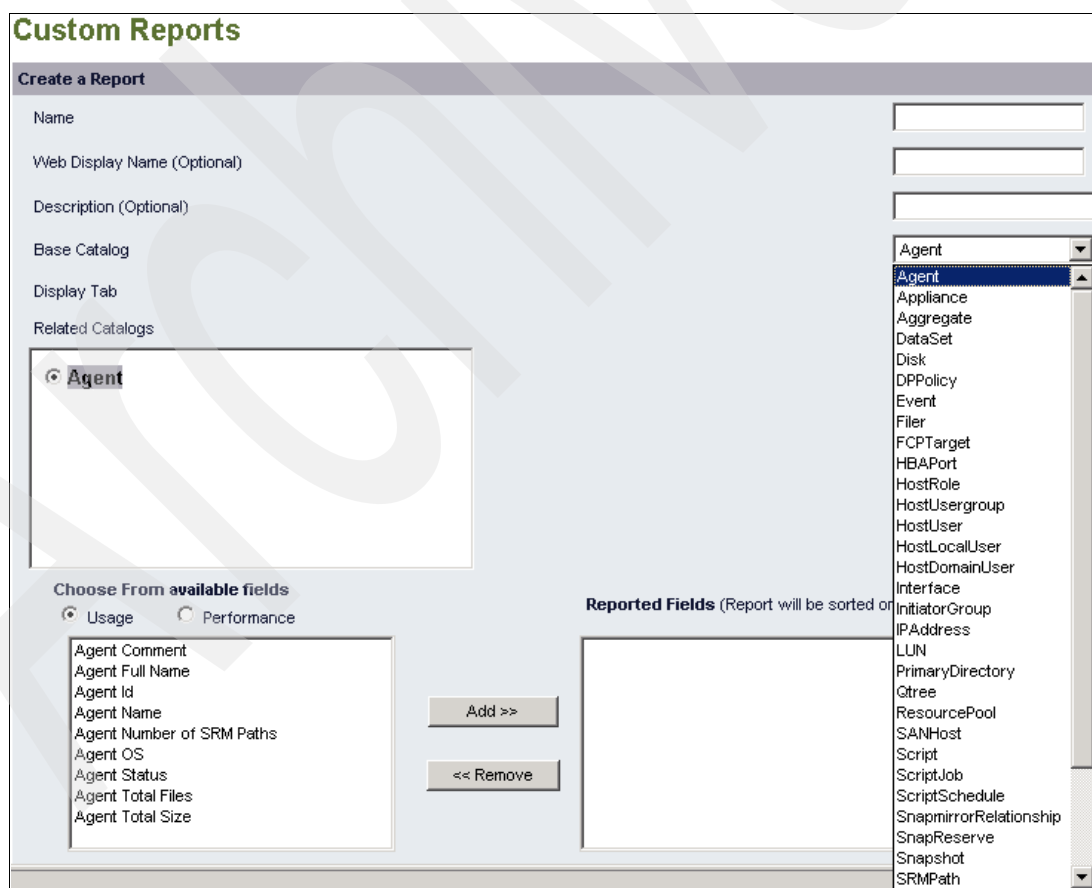


Figure 9-29 Report Catalogs

There are many items you can view based on the catalogs available. Figure 9-30 shows a minor example of some of the things you can do with custom reports.

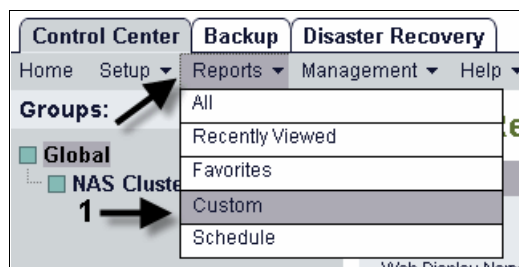


Figure 9-30 Select custom report

The following steps describe how to use Operations Manager to configure custom reports:

1. Select **Custom** from the Reports menu. The Custom Reports window is displayed.
2. Enter a (short) name for the report as you want it to display in the DFM Server CLI.
3. Optionally, enter a (long) name for the report as you want it to display in Operations Manager.
4. Optionally, add comments to the report description.
5. Select the catalog from which the available report fields are based.
6. Select where you want Operations Manager to display this report.
7. Select the related catalog from which you want to choose fields. You might need to expand the list in order to display the catalog you want to select.
8. You can view two types of information fields in the “Choose From available fields” section:
 - To view fields related to the usage and configuration metrics of the object, click **Usage**.
 - To view fields related to performance metrics of the object, click **Performance**.
9. Select a field from the drop-down list.
10. Optionally, enter the name for the field as you want it displayed on the report. Make your field name as short and clear as possible. You must be able to look at a field name in the reports and determine which field the information relates to.
11. Optionally, specify the format of the field. If you choose not to format the field, the default format displayed is used.
12. Click **Add** to move the field to the Reported Fields list.
13. Repeat Steps 8 to 12 for each field you want to include in the report.
14. Optionally, click **Move Up** or **Move Down** to reorder the fields.
15. If you clicked **Performance**, select the required data consolidation method from the list.
16. Click **Create**.
17. To view this report, locate this report in the list at the lower part of the window and click the display tab name. Find the report from the Report drop-down list.

Figure 9-31 on page 151 gives an overview of these steps.

Custom Reports

Create a Report

Name 3 →

Web Display Name (Optional)

Description (Optional)

Base Catalog 4 →

Display Tab 5 →

Related Catalogs 6 →

☒ **Appliance**

Choose From available fields 8

☒ Usage ☐ Performance

9 12

Appliance CPU %
Appliance CPU Threshold
Appliance CPU Threshold Interval
Appliance Comment
Appliance Config Group
Appliance Config Matches
Appliance Console Address
Appliance Contact
Appliance Deleted By
Appliance Deleted When

Enter Field Name Displayed on Report 10

Appliance CPU %

Choose Formatting to Apply on Field 11

Reported Fields (Report will be sorted on the first field)

14 16

Figure 9-31 Creating custom reports

Figure 9-31 shows you the drop-down menu with many of the available catalogs. A more exhaustive list can be found in *Operations Manager Administration Guide for Use With DataFabric Manager Server 3.7*, GC26-7889.

Once created, these reports are available just like any of the standard reports.

Exporting reports

One of the features of Operations Manager is its ability to export report data. The *Operations Manager Administration Guide for Use With DataFabric Manager Server 3.7*, GC26-7889 gives detailed information about performing this operation. The Operations Manager database can be connected to third-party reporting tools through JDBC™ or ODBC interfaces.

Figure 9-32 shows an example of data to be exported and some things to check to successfully export data to third-party programs.




Figure 9-32 Example of exported data

Exporting data is desirable to creating the information from scratch. However, to accomplish this task, certain Operations Manager commands and options need to be set.

Using the DFM Server CLI to enable the export of the Operations Manager database

You want to export data of storage systems and resource groups from Operations Manager and Performance advisor in order to load them to user specific databases. How can this be accomplished?

Operations Manager has the capability to export data from the DataFabric Manager Server and Performance Advisor to text files. For example, you want to export the data of an N series storage system. Here are the DFM Server CLI commands used to accomplish the task:

- First, make sure that dfmDataExportEnabled and perfDataExportEnabled are enabled:

First, change the Operations Manager data export enabled to Yes:

```
dfm option set dfmDataExportEnabled=yes
```

Then change the Performance Advisor data export enabled to Yes:

```
dfm option set perfDataExportEnabled=yes
```

- Now execute the following command:

```
dfm data export run -t "15 secs" -f min -d tab -h "15 secs"
btc-ppe-filer7.btcppe.N series.com
```

For more detailed information and instructions about transferring data to third-party applications, refer to Chapter 6 of *Operations Manager Administration Guide for Use With DataFabric Manager Server 3.7*, GC26-7889.

9.7 Monitoring and alerting

Events are generated based on either the crossing of a threshold or the receipt of a trap from an N series storage system.

Monitoring overview

Monitoring involves several processes. First, Operations Manager discovers the storage systems supported on your network. It periodically monitors data that it collects from the discovered storage systems, such as CPU usage, interface statistics, free disk space, qtree usage, and chassis environment. Operations Manager generates *events* when it discovers a storage system, when the status is abnormal, or when a predefined threshold is breached. If configured to do so, Operations Manager sends a notification to a recipient when an event triggers an *alarm*.

Figure 9-33 illustrates the Operations Manager monitoring process.

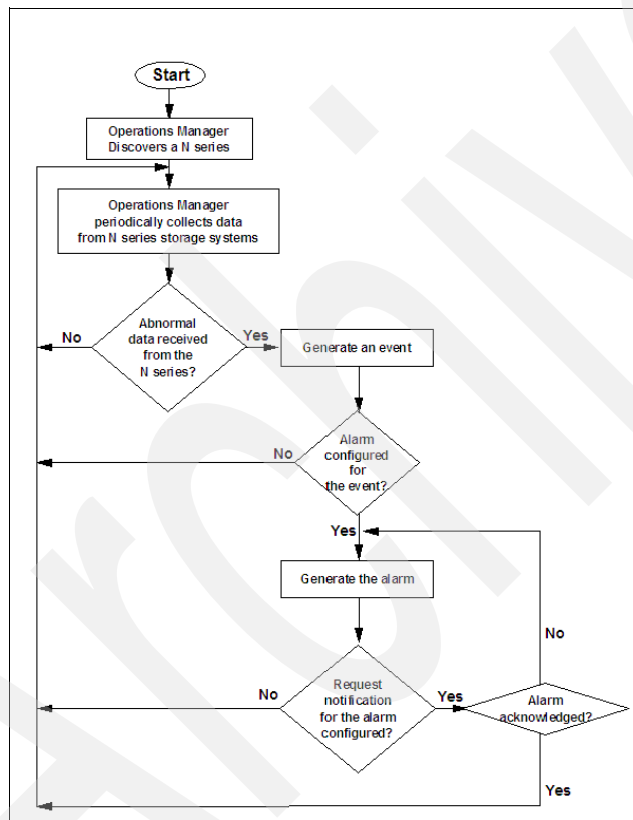


Figure 9-33 Event monitoring process

SNMP queries

Operations Manager uses periodic SNMP queries to collect data from the storage systems it discovers. The data is reported by Operations Manager in the form of tabular and graphical reports and event generation.

The time interval at which an SNMP query is sent depends on the data being collected. For example, although Operations Manager pings each storage system every minute to ensure that the storage system is reachable, the amount of free space on the disks of a storage system is collected every 30 minutes.

Global monitoring options

The SNMP query time intervals are specified by the global monitoring options located in the Monitoring Options section of the Options window. Although you should generally keep the default values, you might need to change some of the options to suit your environment.

All of the monitoring option values apply to all storage systems in all groups.

Considerations before changing monitoring intervals

There are advantages and disadvantages to changing the monitoring intervals. If you decrease the monitoring intervals, you receive more real-time data; however, Operations Manager queries the storage systems more frequently, thereby increasing the network traffic and the load on the server on which Operations Manager is installed and the storage systems responding to the queries. Similarly, if you increase the monitoring interval, the network traffic and the storage system load are reduced; however, data reported might not reflect the current status or condition of a storage system, as shown in Figure 9-34.



Figure 9-34 Monitoring Options settings

SNMP trap listener

In addition to periodically sending out SNMP queries, Operations Manager includes an SNMP trap listener as part of the server service. Event generation and alerting is faster than with SNMP queries because the proper monitoring mechanism is started immediately after the SNMP trap is received. In addition, monitoring is performed asynchronously, instead of waiting for the monitoring interval.

The SNMP trap listener listens for SNMP traps from monitored storage systems, if they have been manually configured to send traps to the DataFabric Manager Server (over UDP port 162).

Note: Currently, the SNMP trap listener can receive SNMP traps only from storage systems that are supported on Operations Manager. Traps from other sources are dropped.

When the SNMP trap listener receives an SNMP trap, Operations Manager issues an information event, but does not change the status of the host. Instead, the corresponding monitor associated with the trap generates the proper event and continues to monitor the host to report status changes.

The Management Information Block (MIB) information for all of these traps if found in the *Data ONTAP Network Management Guide*, GC52-1280. Not only can traps be created within Operations Manager (refer to Figure 9-37 on page 157) but also from FilerView® on the storage appliance through the DFM Server CLI. (Figure 9-35) The same MIBs used for the N series appliance is also used for Operations Manager.

Figure 9-35 Trap window in FilerView

The name associated with the SNMP trap information event indicates the severity of the trap, for example, Warning or Worse. (See Figure 9-36 on page 156 for some examples.) The trap severities are deduced from the last digit of the trap ID, as specified in the custom MIB.

The following list describes the SNMP trap information event types:

- ▶ Emergency Trap Received
- ▶ Alert Trap Received
- ▶ Critical Trap Received
- ▶ Error Trap Received
- ▶ Warning Trap Received
- ▶ Notification Trap Received
- ▶ Information Trap Received

If the severity of a trap is unknown, Operations Manager drops the trap.

SNMP trap listener configuration requirements

The following configuration requirements must be met to enable reception of SNMP traps from managed storage systems.

On Operations Manager

No configuration is needed to start the SNMP trap listener on Operations Manager (the trap listener is automatically started after installation). The SNMP trap global options are also configured with the default settings, although you might want to modify these settings.

On managed storage systems

You must manually add the DataFabric Manager Server as a trap destination on all supported systems to be monitored. The traps must be sent to the DataFabric Manager Server over UDP port 162.

Events

Events are automatically generated if a predefined condition, such as a disabled storage system, is met or an object crosses a predefined threshold, such as Aggregate Full. For some threshold events, you can configure a threshold interval, which prevents the event from being generated unless the condition persists for the configured amount of time. All events are automatically logged and reported by Operations Manager.

All events are associated with a severity level. Operations Manager defines the following severity levels (in decreasing order of severity):

- ▶ **Emergency:** Indicates that an object has stopped performing unexpectedly and has experienced unrecoverable data loss. You must take corrective action immediately to avoid extended downtime.
- ▶ **Critical:** Indicates that a problem has occurred that might lead to service disruption, if corrective action is not taken immediately.
- ▶ **Error:** Indicates that an object is still performing, but corrective action is required to avoid service disruption.
- ▶ **Warning:** Indicates that the object has experienced an occurrence that you should be aware of. Such events do not cause service disruption, and corrective action might not be required.
- ▶ **Information:** Indicates a normal occurrence, just like the Normal severity level event, and does not require you to take any action.
- ▶ **Normal:** Indicates that an object is in normal status and is operating within the desired thresholds.

Figure 9-36 is an example of some of the events you will see generated by Operations Manager.

- ▶ Operations Manager generates events based on “threshold breach” and “SNMP traps”
- ▶ Events are categorized on severity levels and the reports window filters them based on severity level.
- ▶ The user can acknowledge these events and they will be recorded by Operations Manager.

Warning or Worse

Group Status

Member Details

Summary

File SRM

Streaming

Events

Global

Report

Warning or Worse

Events: 1-20 of 47

Page: 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 | 24 | 25 | 26 | 27 | 28 | 29 | 30 | 31 | 32 | 33 | 34 | 35 | 36 | 37 | 38 | 39 | 40 | 41 | 42 | 43 | 44 | 45 | 46 | 47 | 48 | 49 | 50 | 51 | 52 | 53 | 54 | 55 | 56 | 57 | 58 | 59 | 60 | 61 | 62 | 63 | 64 | 65 | 66 | 67 | 68 | 69 | 70 | 71 | 72 | 73 | 74 | 75 | 76 | 77 | 78 | 79 | 80 | 81 | 82 | 83 | 84 | 85 | 86 | 87 | 88 | 89 | 90 | 91 | 92 | 93 | 94 | 95 | 96 | 97 | 98 | 99 | 100 | 101 | 102 | 103 | 104 | 105 | 106 | 107 | 108 | 109 | 110 | 111 | 112 | 113 | 114 | 115 | 116 | 117 | 118 | 119 | 120 | 121 | 122 | 123 | 124 | 125 | 126 | 127 | 128 | 129 | 130 | 131 | 132 | 133 | 134 | 135 | 136 | 137 | 138 | 139 | 140 | 141 | 142 | 143 | 144 | 145 | 146 | 147 | 148 | 149 | 150 | 151 | 152 | 153 | 154 | 155 | 156 | 157 | 158 | 159 | 160 | 161 | 162 | 163 | 164 | 165 | 166 | 167 | 168 | 169 | 170 | 171 | 172 | 173 | 174 | 175 | 176 | 177 | 178 | 179 | 180 | 181 | 182 | 183 | 184 | 185 | 186 | 187 | 188 | 189 | 190 | 191 | 192 | 193 | 194 | 195 | 196 | 197 | 198 | 199 | 200 | 201 | 202 | 203 | 204 | 205 | 206 | 207 | 208 | 209 | 210 | 211 | 212 | 213 | 214 | 215 | 216 | 217 | 218 | 219 | 220 | 221 | 222 | 223 | 224 | 225 | 226 | 227 | 228 | 229 | 230 | 231 | 232 | 233 | 234 | 235 | 236 | 237 | 238 | 239 | 240 | 241 | 242 | 243 | 244 | 245 | 246 | 247 | 248 | 249 | 250 | 251 | 252 | 253 | 254 | 255 | 256 | 257 | 258 | 259 | 260 | 261 | 262 | 263 | 264 | 265 | 266 | 267 | 268 | 269 | 270 | 271 | 272 | 273 | 274 | 275 | 276 | 277 | 278 | 279 | 280 | 281 | 282 | 283 | 284 | 285 | 286 | 287 | 288 | 289 | 290 | 291 | 292 | 293 | 294 | 295 | 296 | 297 | 298 | 299 | 300 | 301 | 302 | 303 | 304 | 305 | 306 | 307 | 308 | 309 | 310 | 311 | 312 | 313 | 314 | 315 | 316 | 317 | 318 | 319 | 320 | 321 | 322 | 323 | 324 | 325 | 326 | 327 | 328 | 329 | 330 | 331 | 332 | 333 | 334 | 335 | 336 | 337 | 338 | 339 | 340 | 341 | 342 | 343 | 344 | 345 | 346 | 347 | 348 | 349 | 350 | 351 | 352 | 353 | 354 | 355 | 356 | 357 | 358 | 359 | 360 | 361 | 362 | 363 | 364 | 365 | 366 | 367 | 368 | 369 | 370 | 371 | 372 | 373 | 374 | 375 | 376 | 377 | 378 | 379 | 380 | 381 | 382 | 383 | 384 | 385 | 386 | 387 | 388 | 389 | 390 | 391 | 392 | 393 | 394 | 395 | 396 | 397 | 398 | 399 | 400 | 401 | 402 | 403 | 404 | 405 | 406 | 407 | 408 | 409 | 410 | 411 | 412 | 413 | 414 | 415 | 416 | 417 | 418 | 419 | 420 | 421 | 422 | 423 | 424 | 425 | 426 | 427 | 428 | 429 | 430 | 431 | 432 | 433 | 434 | 435 | 436 | 437 | 438 | 439 | 440 | 441 | 442 | 443 | 444 | 445 | 446 | 447 | 448 | 449 | 450 | 451 | 452 | 453 | 454 | 455 | 456 | 457 | 458 | 459 | 460 | 461 | 462 | 463 | 464 | 465 | 466 | 467 | 468 | 469 | 470 | 471 | 472 | 473 | 474 | 475 | 476 | 477 | 478 | 479 | 480 | 481 | 482 | 483 | 484 | 485 | 486 | 487 | 488 | 489 | 490 | 491 | 492 | 493 | 494 | 495 | 496 | 497 | 498 | 499 | 500 | 501 | 502 | 503 | 504 | 505 | 506 | 507 | 508 | 509 | 510 | 511 | 512 | 513 | 514 | 515 | 516 | 517 | 518 | 519 | 520 | 521 | 522 | 523 | 524 | 525 | 526 | 527 | 528 | 529 | 530 | 531 | 532 | 533 | 534 | 535 | 536 | 537 | 538 | 539 | 540 | 541 | 542 | 543 | 544 | 545 | 546 | 547 | 548 | 549 | 550 | 551 | 552 | 553 | 554 | 555 | 556 | 557 | 558 | 559 | 560 | 561 | 562 | 563 | 564 | 565 | 566 | 567 | 568 | 569 | 570 | 571 | 572 | 573 | 574 | 575 | 576 | 577 | 578 | 579 | 580 | 581 | 582 | 583 | 584 | 585 | 586 | 587 | 588 | 589 | 590 | 591 | 592 | 593 | 594 | 595 | 596 | 597 | 598 | 599 | 600 | 601 | 602 | 603 | 604 | 605 | 606 | 607 | 608 | 609 | 610 | 611 | 612 | 613 | 614 | 615 | 616 | 617 | 618 | 619 | 620 | 621 | 622 | 623 | 624 | 625 | 626 | 627 | 628 | 629 | 630 | 631 | 632 | 633 | 634 | 635 | 636 | 637 | 638 | 639 | 640 | 641 | 642 | 643 | 644 | 645 | 646 | 647 | 648 | 649 | 650 | 651 | 652 | 653 | 654 | 655 | 656 | 657 | 658 | 659 | 660 | 661 | 662 | 663 | 664 | 665 | 666 | 667 | 668 | 669 | 670 | 671 | 672 | 673 | 674 | 675 | 676 | 677 | 678 | 679 | 680 | 681 | 682 | 683 | 684 | 685 | 686 | 687 | 688 | 689 | 690 | 691 | 692 | 693 | 694 | 695 | 696 | 697 | 698 | 699 | 700 | 701 | 702 | 703 | 704 | 705 | 706 | 707 | 708 | 709 | 710 | 711 | 712 | 713 | 714 | 715 | 716 | 717 | 718 | 719 | 720 | 721 | 722 | 723 | 724 | 725 | 726 | 727 | 728 | 729 | 730 | 731 | 732 | 733 | 734 | 735 | 736 | 737 | 738 | 739 | 740 | 741 | 742 | 743 | 744 | 745 | 746 | 747 | 748 | 749 | 750 | 751 | 752 | 753 | 754 | 755 | 756 | 757 | 758 | 759 | 760 | 761 | 762 | 763 | 764 | 765 | 766 | 767 | 768 | 769 | 770 | 771 | 772 | 773 | 774 | 775 | 776 | 777 | 778 | 779 | 780 | 781 | 782 | 783 | 784 | 785 | 786 | 787 | 788 | 789 | 790 | 791 | 792 | 793 | 794 | 795 | 796 | 797 | 798 | 799 | 800 | 801 | 802 | 803 | 804 | 805 | 806 | 807 | 808 | 80

Figure 9-36 Example traps of Warning or worse

It is possible to create alarms of items you are interested in as well. If the default alarms and events do not meet your needs, you can create a custom alarm. Follow these steps to create a custom alarm:

1. From the Alarms window, select the group that you want Operations Manager to monitor. You might need to expand the list to display the one you want to select.
2. Specify what triggers the alarm: an event or the severity of event.
3. Specify the recipient of the alarm notification. If you want to specify more than one recipient or configure a repeat notification, continue to Step 5.
4. Click **Add** to set the alarm. If you want to configure additional options, continue with step 5.
5. Click **Advanced Version**. The Alarms window displays the additional options.
6. Optionally, if you want to specify a class of events that should trigger this alarm, specify the event class. You can use regular expressions.
7. Optionally, specify the recipients of the alarm notification. Formats include administrator names, e-mail addresses, pager addresses, or an IP address of the system to receive SNMP traps (or port number to send the SNMP trap to).
8. Optionally, specify the time period that Operations Manager sends alarm notifications.
9. Optionally, select **Yes** to resend the alarm notification until the event is acknowledged or **No** to notify the recipients only once.
10. Optionally, set the interval (in minutes) that Operations Manager waits before it tries to resend a notification.
11. Activate the alarm by selecting **No** in the Disable field.
12. Click **Add**.

Figure 9-37, Figure 9-38 on page 158, and Figure 9-39 on page 158 give an overview of the steps to create a custom alarm.

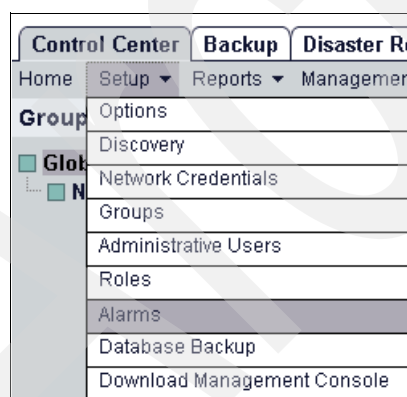


Figure 9-37 Select Setup and Alarms from the Control Center

Figure 9-38 General Alarms setup window

Figure 9-39 Advanced options for alarms

9.8 Role Based Access Controls (RBAC)

Operations Manager uses RBAC for administrative user login. If you have not changed Operations Manager's default settings for administrative user access, you *do not* need to log in to view information using Operations Manager. However, the information you see is very limited, as shown in Figure 9-40 on page 159. When you initiate an operation that requires specific privileges, Operations Manager prompts you to log in, as shown in Figure 9-41 on page 159.

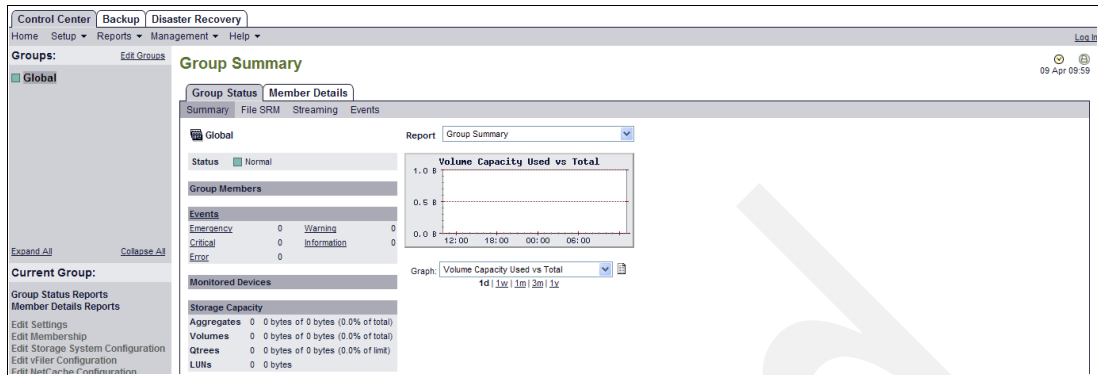


Figure 9-40 Default window before administrator login

Any operation that requires Administrative access will pop up a login window. Here we select **Options** under the Setup menu and get a login window.

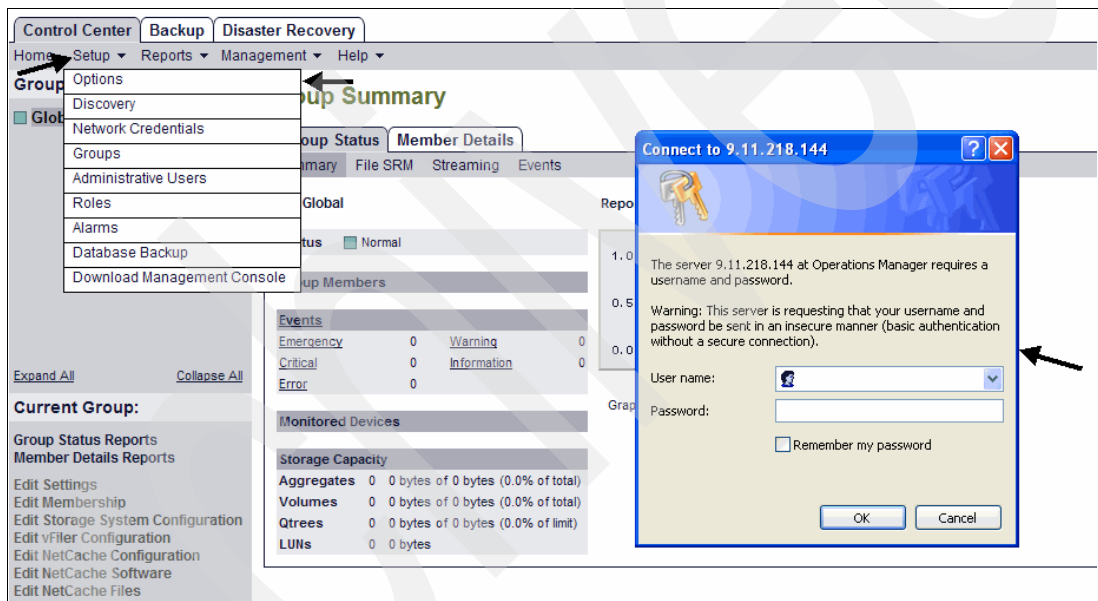


Figure 9-41 Actions require Administrative Login

When you open the Operations Manager window for the first time (remember, you *must* use "http://" before the "ipaddress:8080" to open the browser window), you can log in with your credentials, as shown in Figure 9-42.



Figure 9-42 Log in link

When you log in using the Log In link, the window shown in Figure 9-43 appears.

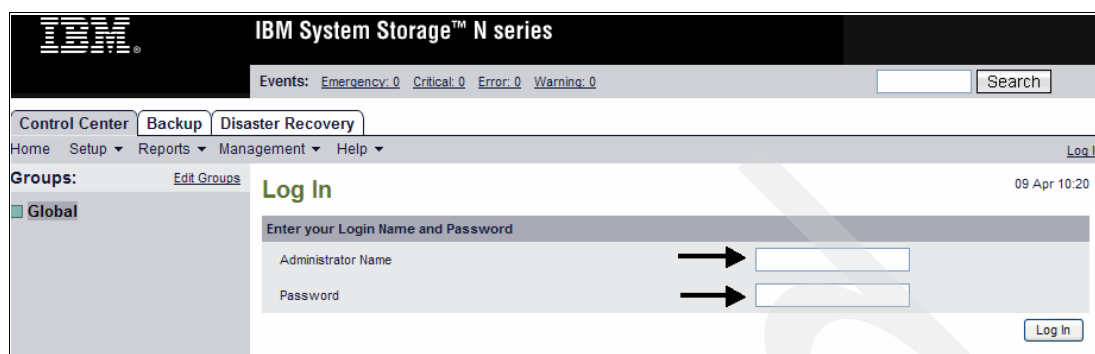
The screenshot shows the IBM System Storage N series Operations Manager interface. At the top, there's a black header with the IBM logo and the text 'IBM System Storage™ N series'. Below this, a status bar shows 'Events: Emergency: 0 Critical: 0 Error: 0 Warning: 0' and a search box. The main navigation bar includes tabs for 'Control Center', 'Backup', and 'Disaster Recovery', along with a menu for 'Home', 'Setup', 'Reports', 'Management', and 'Help'. A 'Log In' link is visible in the top right. On the left, under 'Groups:', there's a list with 'Global' selected. The central area is titled 'Log In' and contains a form with the heading 'Enter your Login Name and Password'. It has two input fields: 'Administrator Name' and 'Password', each with a black arrow pointing to it. A 'Log In' button is at the bottom right of the form. The date and time '09 Apr 10:20' are displayed in the top right corner.

Figure 9-43 Key in log in credentials for administrator

To create administrator accounts, you need to log in with administrator account access.

Operations Manager uses administrator accounts to manage access control and maintain security. When you install Operations Manager software, the default administrator accounts Administrator and Everyone are created. Both accounts are global and have predefined roles assigned to them.

Administrator account

The administrator has super-user privileges and can perform any operation in the Operations Manager database and add other administrators. The administrator account is given the same name as the name of the administrator who installed the software. Therefore, if you install Operations Manager on a UNIX® workstation, the administrator account is called root.

Everyone account

After installing Operations Manager, you must log in as the Administrator and set up the Everyone account to grant view permission on this account.

A Operations Manager Administrator can assign roles to various Operations Manager users based on their nature of work. These roles can be provided with various capabilities on the DataFabric Manager Server. The process of assigning these roles, which have various capabilities on the DataFabric Manager Server, is termed RBAC for Operations Manager users. This process is shown in Figure 9-44 on page 161.

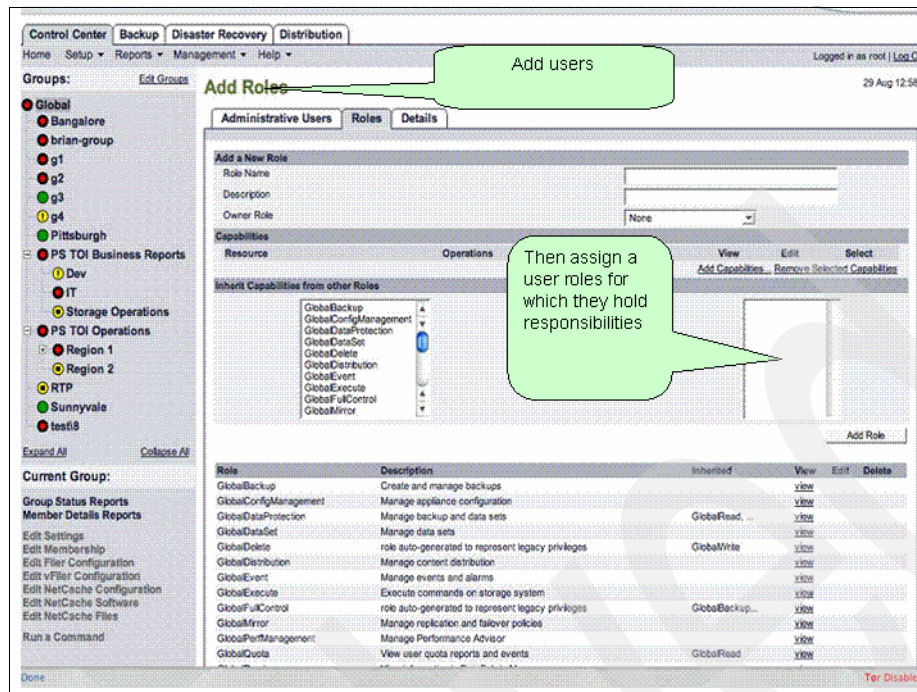


Figure 9-44 Administrator can assign roles based on user access requirements

The following procedure shows how to create and edit administrator accounts from Operations Manager. To set up administrator accounts using the DFM Server CLI, use the **dfm user** command.

First, select **Administrative Users** under the Setup menu in the Control Center tab, as shown in Figure 9-45.

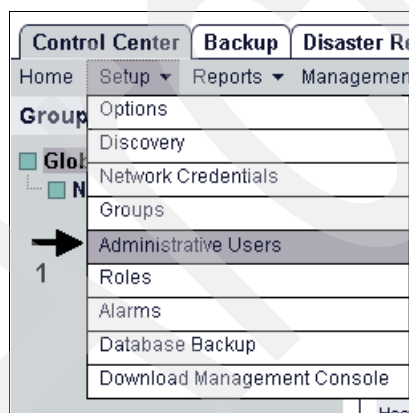


Figure 9-45 Select Administrative Users under the Setup menu

Then, follow these steps to create and edit the administrators:

1. Log in to the administrator account, as shown in Figure 9-42 on page 159.
2. Type the name for the administrator or the domain name for the group of administrators.
3. Optionally, enter the e-mail address for the administrator or administrator group.
4. Optionally, enter the window address, as an e-mail address or window number, for the administrator or administrator group.
5. Click **Add**.

Figure 9-46 shows the steps that were given above.

Administrators 09 Apr 11:0

Administrative Users Roles Details

Add a New Administrator

Administrator or Windows Group Name

Roles

Email Address

Pager Address

Add

Administrator Name	Roles	Email	Pager	View	Edit	Delete
Everyone				view	edit	
X3655A\Administrator	GlobalFullCo...			view	edit	<input type="checkbox"/>

Delete Selected

Figure 9-46 How to set up an administrator

Role management allows the administrator who logs in with super-user access to restrict the use of certain Operations Manager functions by other administrators. The super-user can assign roles to administrators on an individual basis, by group, or globally (and for all objects in Operations Manager).

In Figure 9-47 on page 163, we create additional administrative users, and restrict one user to Global Backup and Restore only and give the other one Global Full Control. We can also edit the level of control and delete administrative users in this window.

Administrators

09 Apr 11:

Administrative Users Roles Details

Add a New Administrator

Administrator or Windows Group Name

Roles

Email Address

Pager Address

GlobalBackup

GlobalConfigManagement

GlobalDataProtection

GlobalDataSet

GlobalDelete

GlobalDistribution

GlobalEvent

GlobalExecute

GlobalFailover

GlobalFullControl

>>

<<

Add

Administrator Name	Roles	Email	Pager	View	Edit	Delete
Chubaka	GlobalFullCo...			view	edit	<input type="checkbox"/>
Everyone				view	edit	<input type="checkbox"/>
Luke Skywalker	GlobalBackup...			view	edit	<input type="checkbox"/>
X3655AAdministrator	GlobalFullCo...			view	edit	<input type="checkbox"/>

Delete Selected

Figure 9-47 Add, delete, and edit administrative users

An operation must be specified for every role. You can assign multiple operations levels if you want the administrator to have more control than a specific role provides. For example, if you want an administrator to perform both the backup and restore operations, you must assign Back Up and Restore roles to the administrator.

Two specific roles require special mention:

- Global roles

Administrators assigned global roles can view information or configure settings for all groups in the Operations Manager database, including the Global group.

- Group roles

Administrators assigned group roles can view or configure settings for the group to which they belong. When you view roles for an administrator, the settings are those explicitly set for the administrator at the group level. For example, if administrators have the GlobalRead role, they implicitly have the Read role on all groups. Similarly, if administrators have the Read role on a parent group, they implicitly have the Read role on all of the subgroups of that parent group.

Several other factors also affect the group role granted to an administrator, such as the capabilities granted to the administrator "Everyone" or the administrator's membership to Active Directory (AD) user groups that have been added to the DataFabric Manager server database.

For more information about the roles Operations Manager provides, refer to the *Operations Manager Administration Guide for Use With DataFabric Manager Server 3.7*, GC26-7889.

Local users

Another level of access to the Operations Manager resources is the local user. Local users are the users created on storage systems and vFile units. A Operations Manager Administrator can assign roles to various Data ONTAP users based on their nature of work. These roles can be provided with various capabilities on the N series storage system. The process of assigning these roles that have various capabilities on the N series storage system is termed RBAC for Data ONTAP users.

This process is shown in Figure 9-48.

The administrator can also perform user management (add and delete users) and password management. The Operations Manager administrator can also push these users and their password changes to all N series storage systems listed under the Operations Manager.

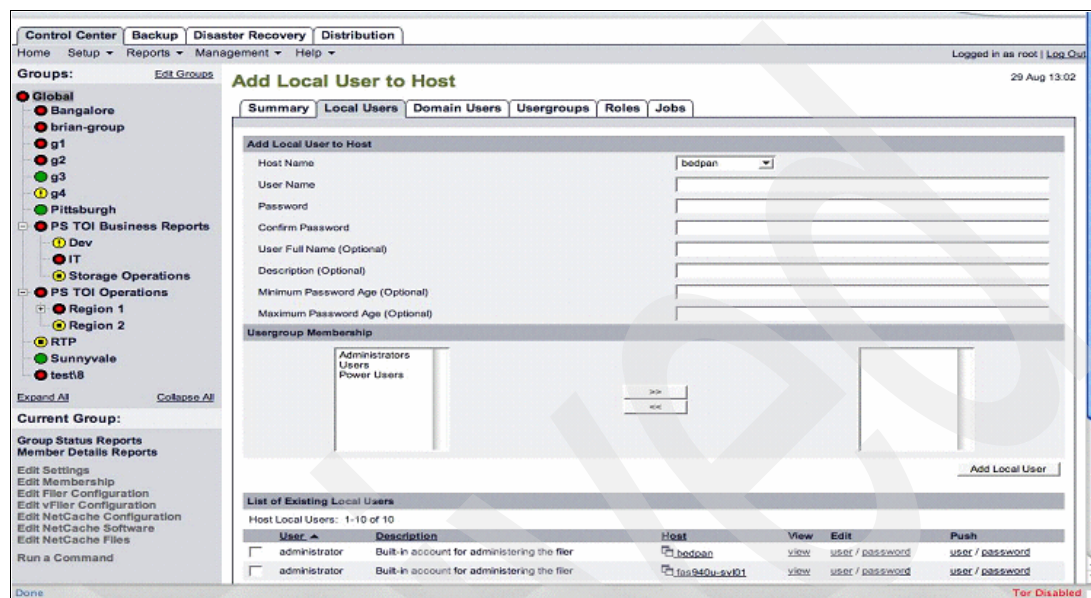


Figure 9-48 Adding local users

To add a local user to a storage system or vFiler unit, complete the following steps:

1. From any window, select **Management** → **Host Users** → **Local Users**. The Add Local User to Host window is displayed, as shown in Figure 9-49 on page 165.
2. Specify the parameters shown in Table 9-2.

Table 9-2 Local Users parameters

Parameter	Description
Host Name	Name of the storage system or vFiler unit from the drop-down list.
User Name	Name of the local user.
Password	Password of the local user.
Confirm Password	Confirm the password of the local user.
User Full-name (optional)	Full name of the local user.
Description (optional)	Description of the local user.
Minimum Password Age (optional)	Minimum number of days that a password must be used.
Maximum Password Age (optional)	Maximum number of days that a password can be used.
Usergroup Membership	User groups you want to be a member of.

3. Select one or more user groups from the list.

4. Click **Add Local User**.

Figure 9-49 and Figure 9-50 show the steps for adding local users.

1. First, in the Control Center tab, select **Management** → **Host Users**.




Figure 9-49 Select Management and Host Users

2. Then, using the steps shown above as your guide, perform the steps shown in Figure 9-50.

A screenshot of the 'Add Local User to Host' form. The form has tabs: Summary, Local Users (selected), Domain Users, Usergroups, Roles, and Jobs. Step 1 points to the 'Local Users' tab. Step 2 points to the 'Host Name' dropdown menu, which is set to 'itsnas1'. Step 3 points to the 'Usergroup Membership' section, which shows a list of usergroups: Administrators, Backup Operators, Compliance Administrators, Users, and Power Users. Step 4 points to the 'Add Loc' button at the bottom right. Below the form is a 'List of Existing Local Users' table.

User	Description	Host	View	Edit	Push
<input type="checkbox"/> administrator	Built-in account for administering the filer	itsnas1	view	user / password	user / password
<input type="checkbox"/> administrator	Built-in account for administering the filer	itsnas2	view	user / password	user / password
<input type="checkbox"/> administrator	Built-in account for administering the filer	itsotuc1	view	user / password	user / password
<input type="checkbox"/> mikeroll		itsotuc1	view	user / password	user / password

Figure 9-50 Steps to add users to host

Just as with administrative users, it is also possible to set local user roles. Figure 9-51 shows the window for establishing roles for a local user of various storage systems.

Once you have selected a host, you can add what capabilities you want for this role, as shown in Figure 9-51.

Add Role to Host 09 Apr 13:48

Summary Local Users Domain Users Usergroups **Roles** Jobs

Add Role to Host

Host Name: itsnas1

Role Name:

Description (Optional):

Selected Capabilities: None

[Add Capabilities](#) [Add Role](#)

List of Existing Roles

Host Roles: 1-15 of 15

Role	Description	Host	View	Edit	Push
<input type="checkbox"/> admin	Default role for administrator privileges.	itsnas1	view	edit	push
<input type="checkbox"/> admin	Default role for administrator privileges.	itsnas2	view	edit	push
<input type="checkbox"/> admin	Default role for administrator privileges.	itsotuc1	view	edit	push
<input type="checkbox"/> audit	Default role for audit privileges.	itsnas1	view	edit	push
<input type="checkbox"/> audit	Default role for audit privileges.	itsnas2	view	edit	push

Figure 9-51 Select host name, give role name, and add capabilities

3. Click **Add Capabilities** to see the list of options you can add, as shown in Figure 9-52.

Capabilities - Windows Internet Explorer

Capabilities:

☒ *

☐ api.*

☐ cli.*

☒ login.*

☒ login-console

☐ login-http-admin

☐ login-rsh

☒ login-ssh

☒ login-telnet

☐ login-snmp

☐ security.*

☐ security-passwd-change-others

☐ security-priv-advanced

☐ security-load-lclgroups

OK Cancel

Add Role to Host

Summary Local Users Domain Users Usergroups **Roles** Jobs

Add Role to Host

Host Name: itsnas1

Role Name:

Description (Optional):

Selected Capabilities:

[Add Capabilities](#)

List of Existing Roles

Host Roles: 1-15 of 15

Role	Description	Host	View	Edit	Push
<input type="checkbox"/> admin	Default role for administrator privileges.	itsnas1	view	edit	push
<input type="checkbox"/> admin	Default role for administrator privileges.	itsnas2	view	edit	push

Figure 9-52 Capabilities you can add for a role

Once you have confirmed what you want, click **Add Role** to complete the task.

9.9 Storage system management

As soon as Operations Manager is installed, it begins the process of discovering, monitoring, and generating data about your supported storage systems. However, before you can use the data to simplify your network administration tasks, you need to understand the different ways you can use Operations Manager to manage your storage system.

Using Operations Manager, you can:

- ▶ View the status of and obtain reports and information for a group of systems
- ▶ View information for individual systems
- ▶ Access the console of a storage system
- ▶ View the active/active configuration status and perform takeover and giveback operations if the storage system is an active/active controller
- ▶ View and respond to Operations Manager events
- ▶ Configure alarms that send you notification if Operations Manager logs a specific type of event or severity of event
- ▶ Edit the configuration settings of a storage system
- ▶ Link to FilerView for a selected storage system or vFiler unit
- ▶ Insert values into custom comment fields
- ▶ View user, qtree, and group quotas
- ▶ Edit user quotas

Though we have already covered several of these functions, we cover a few more here. However, because our focus is on configuration, some of these items may not be covered in this book; for those items, we highly recommend that you refer to *Operations Manager Administration Guide for Use With DataFabric Manager Server 3.7*, GC26-7889.

Storage system groups

Operations Manager is designed around the concept of *groups*. When a group is selected in the left pane of the Operations Manager main window, the windows change to display information relating to that group.

To display information about all of your storage systems, select the **Global** group in the Groups pane on the left side of Operations Manager. The Global group is the default group containing the super-set of all monitored objects and storage systems.

To display information about a specific group of systems, select the desired group name in the Groups pane on the left side of Operations Manager, as shown in Figure 9-53.

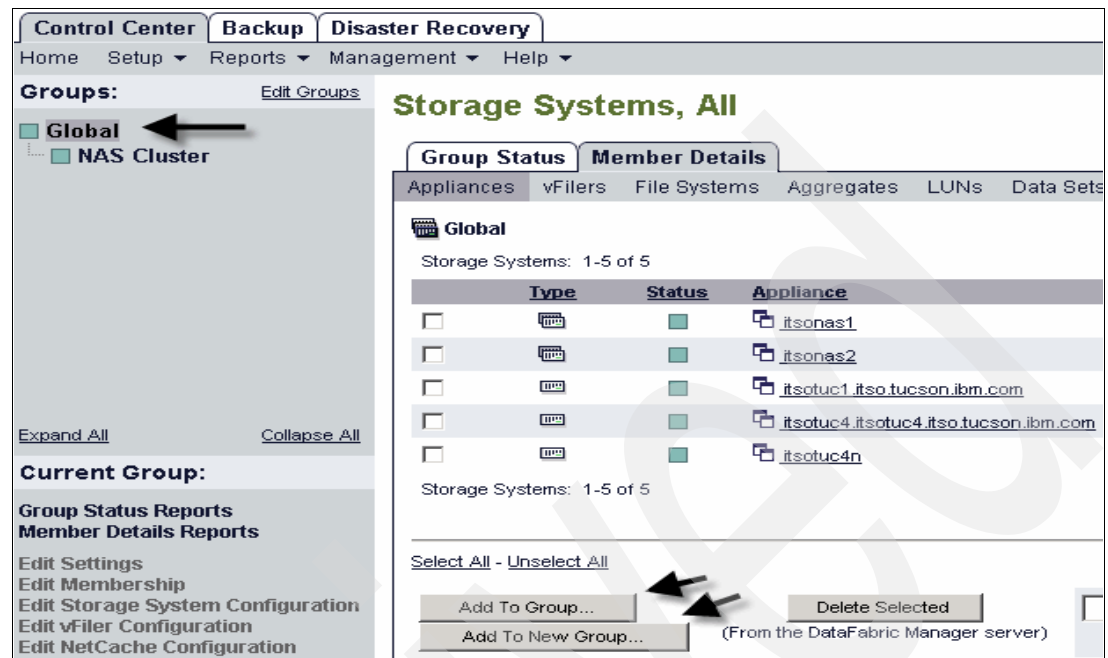


Figure 9-53 Creating groups

To manage your storage systems effectively, you should organize them into smaller groups so that you can view information only about objects in which you are interested. You can group your storage systems to meet your individual needs, for example, by geographic location, operating system version, and storage system platform, as shown in Figure 9-54 on page 169.

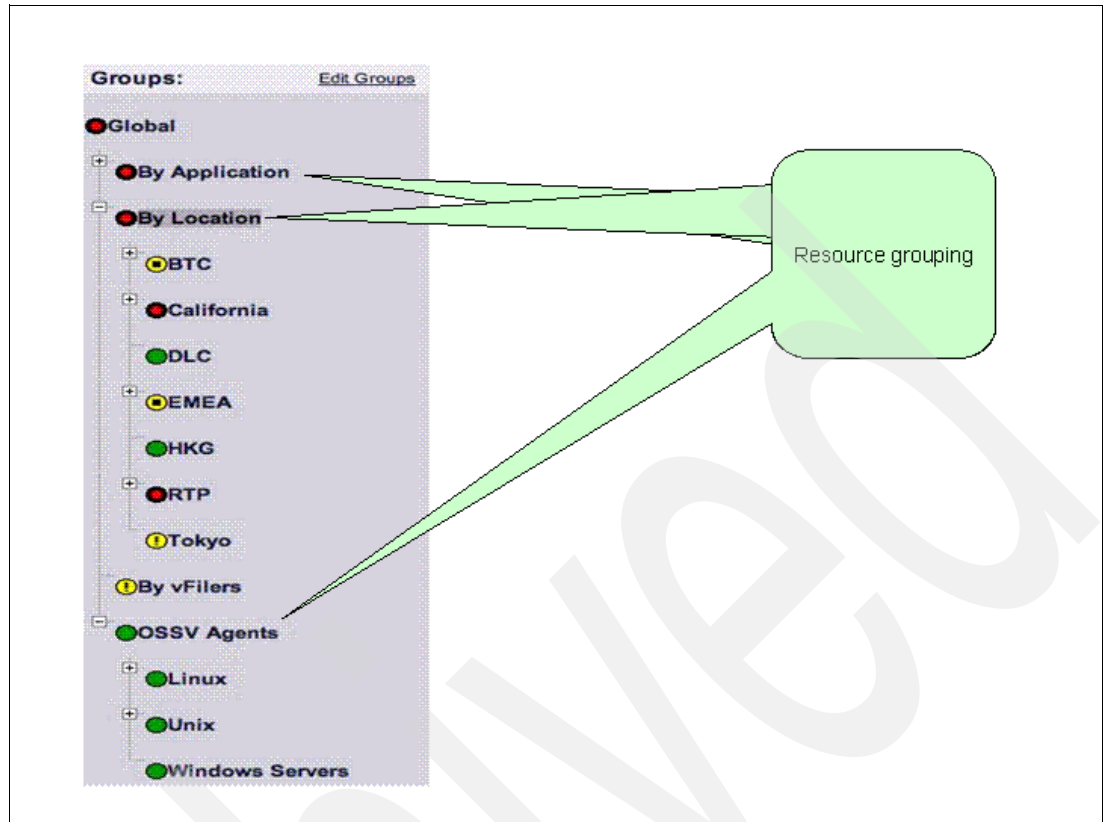


Figure 9-54 Groups by resource

Like the other Operations Manager Control Center tab windows, the Appliances and vFile reports enable you to view a wide variety of details in one place and to perform the following tasks:

1. View system data for all or a group of monitored systems, as shown in Figure 9-55.
2. Obtain more information about a specific storage system, as shown in Figure 9-56 on page 170.
3. Link to FilerView (storage systems only).

Appliances, All

Group Status Member Details 1

Appliances vFiles File Systems Aggregates LUNs Data Sets Resource Pools Scripts

Global Report Appliances, All

Appliances: 1-5 of 5

Type	Status	Appliance	Model	Serial Number	System ID
<input type="checkbox"/>	■	2 #sonas1	N5200	2864130011811	0101184428
<input type="checkbox"/>	■	#sonas2	N5200	2864130011822	0101183273
<input type="checkbox"/>	■	#sotuc1.#so.tucson.ibm.com	N5300	2869130005600	0118052508
<input type="checkbox"/>	■	#sotuc4.#sotuc4.#so.tucson.ibm.com	N5500	2865130008600	0101181370
<input type="checkbox"/>	■	#sotuc4n	N5500	2865130008600	0101181370

Appliances: 1-5 of 5

Select All - Unselect All

Add To Group... Delete Selected (From the DataFabric Manager server) Add (New appliance)

Figure 9-55 Select Member Details and the appliance of interest

To view more detail, click the storage system you are interested in viewing, as shown in Figure 9-56.

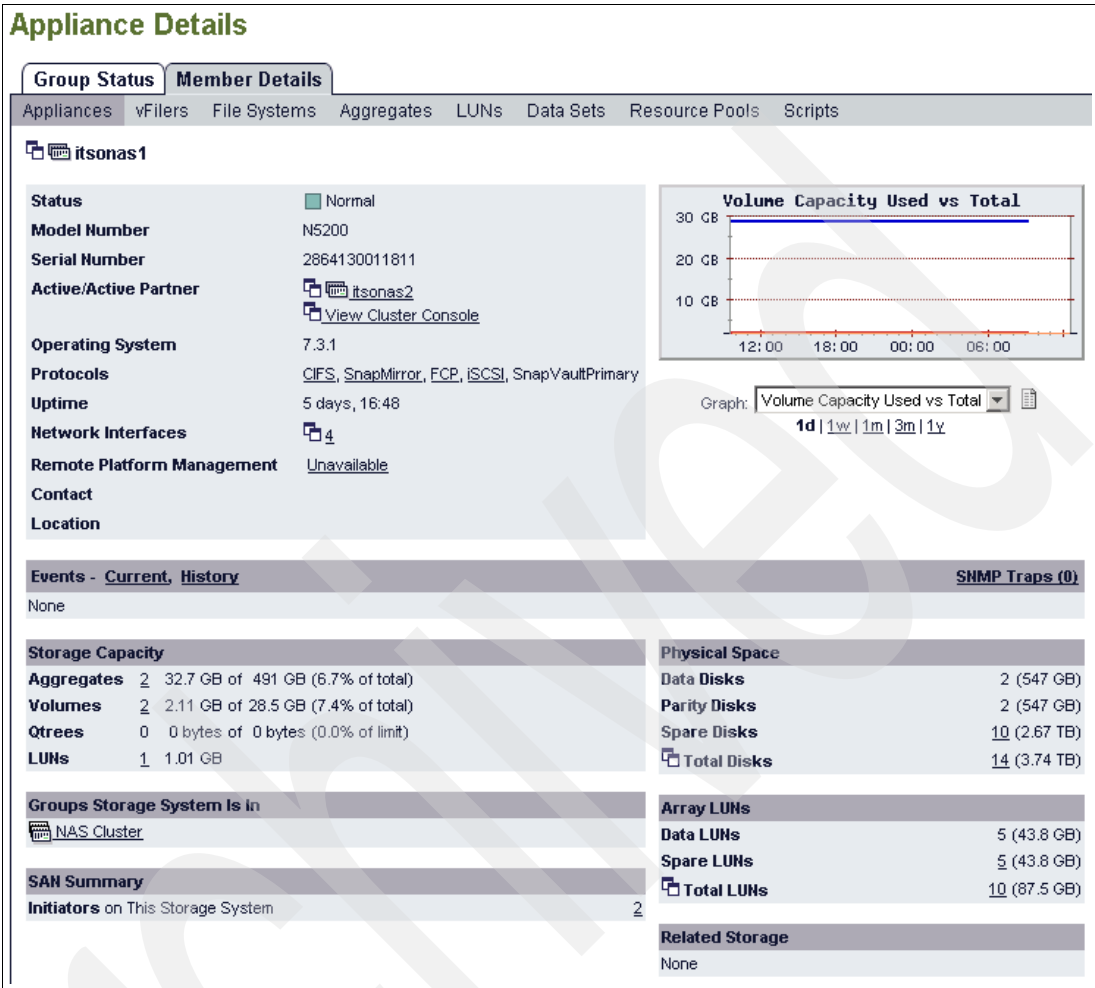


Figure 9-56 Appliance Details window

From this same view in the bottom left of the Operations Manager window, you can gain access to the Storage System GUI by selecting **Connect to Device Console** under Appliance Tools, as shown in Figure 9-57).

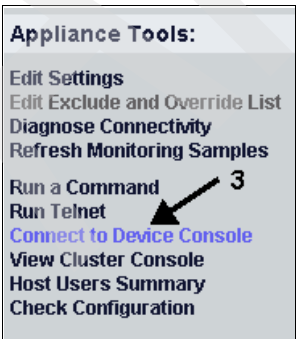


Figure 9-57 Connect to FilerView

Disk reports

Operations Manager associates disks with aggregates and provides reports that show a breakdown of the physical space versus usable space. All of the disk reports contain a column called “Aggregate”, as shown in Figure 9-58. The Aggregate column specifies the aggregate to which each disk is associated. Spare disks are not associated with any aggregate.

Aggregates, All

10 /

Group StatusMember Details

AppliancesvFilesFile SystemsAggregatesLUNsData SetsResource PoolsScripts

itsonas1

ReportAggregates, All

Aggregates: 1-2 of 2

	Aggregate ▲	Storage System	Type	RAID	State	Status
<input type="checkbox"/>	aggr1	itsonas1	Aggregate	raid_dp	online	<input checked="" type="checkbox"/>
<input type="checkbox"/>	aggr2	itsonas1	Aggregate	raid0	online	<input checked="" type="checkbox"/>

Aggregates: 1-2 of 2

Select All - Unselect All

Add To Group...Delete Selected

Add To New Group... (From the DataFabric Manager server)

Figure 9-58 Select an aggregate for detailed disk report

Select an aggregate to view a detailed disk report, as shown in Figure 9-59.

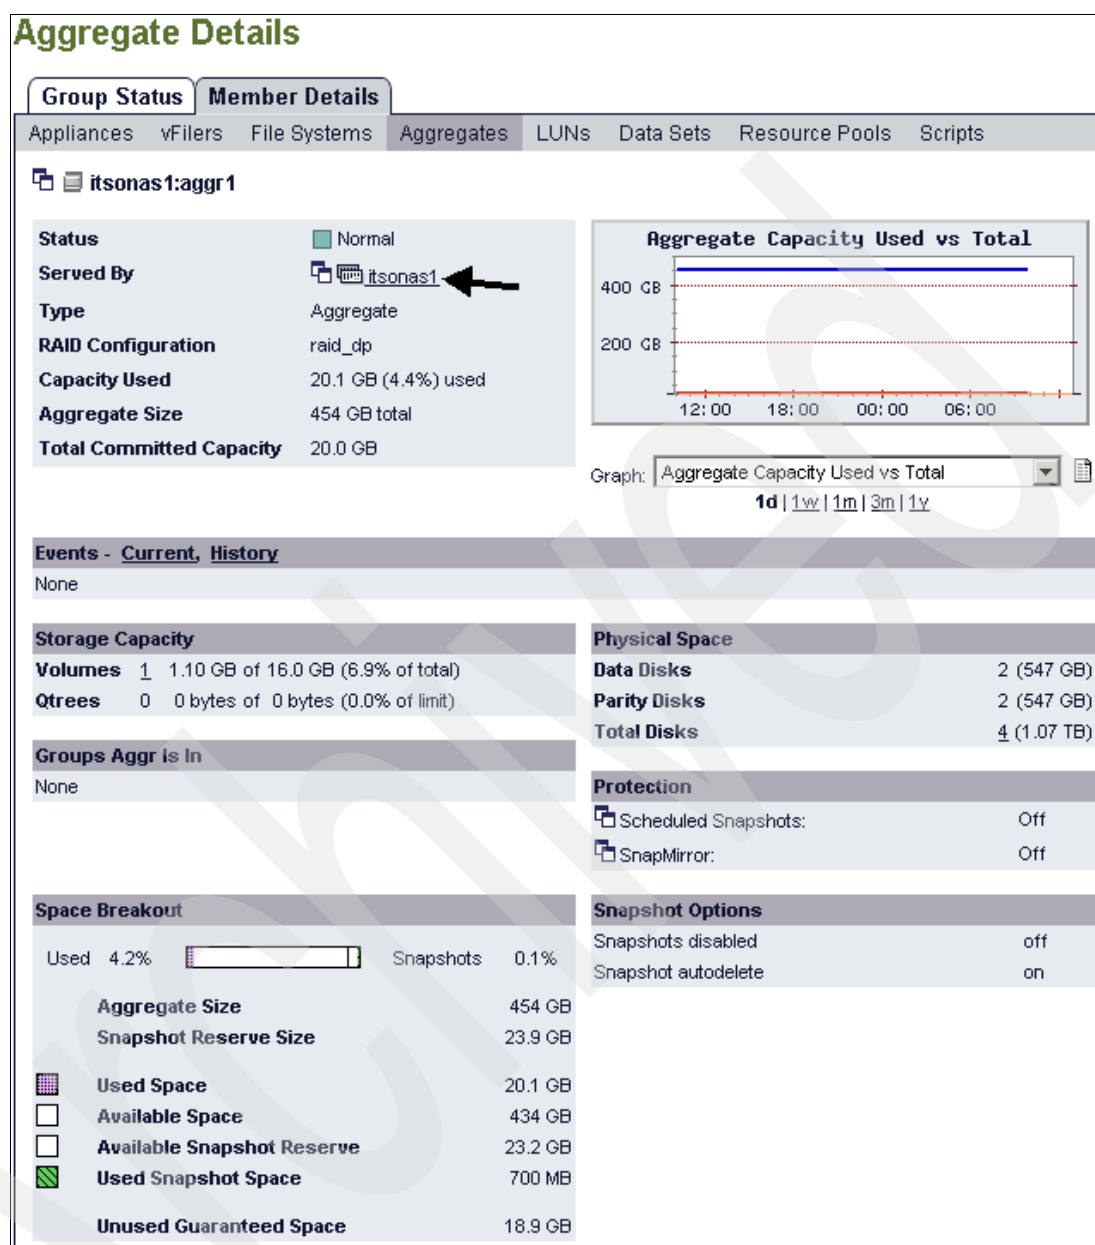


Figure 9-59 Disk report details

If you want to view additional details for a specific storage system or vFile unit, click the storage system or vFile unit name in any of the storage systems or vFile report windows to display the Details window for that system.

Operations Manager regularly refreshes monitoring data for the entire group within which a storage system or vFile unit resides, or you can click **Refresh Group Monitors** to manually refresh the data.

Clicking the appliance name gives you the Appliance Details window, shown in Figure 9-60 on page 173.

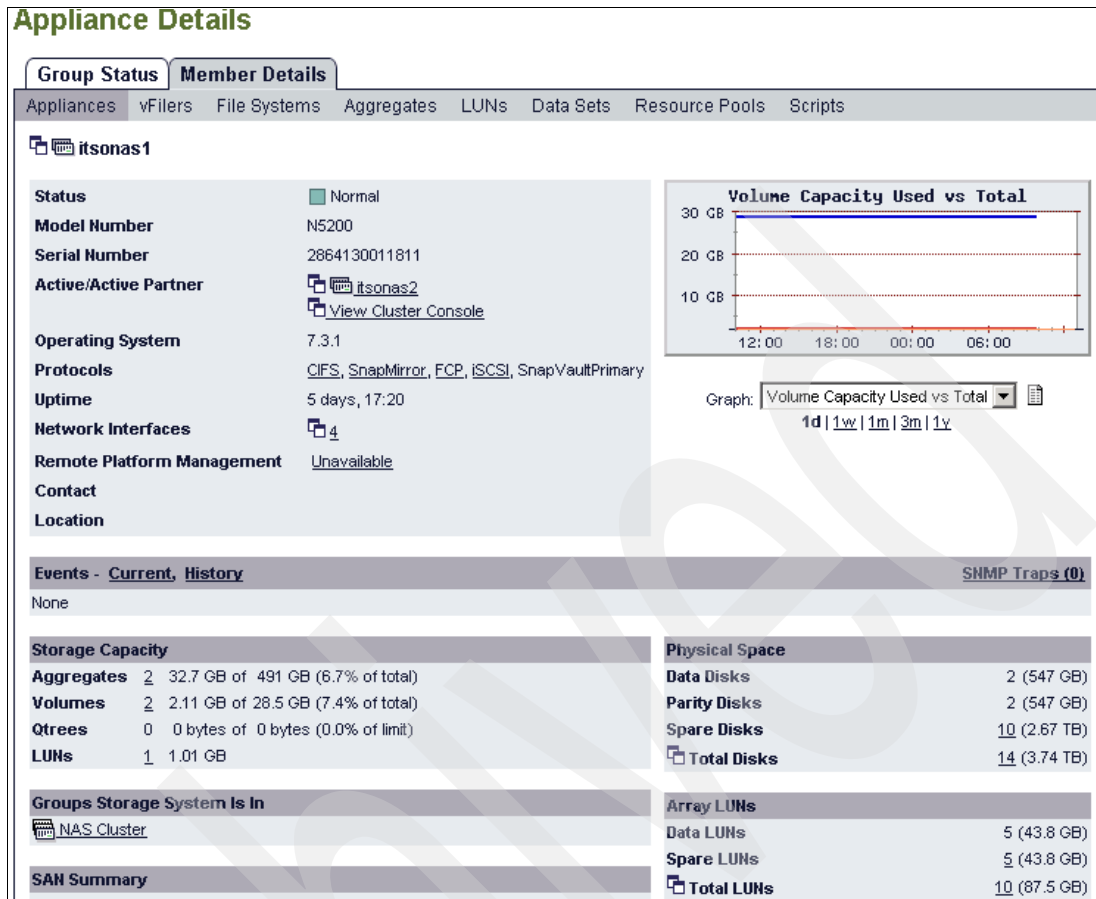


Figure 9-60 Appliance Details window

Tasks you can perform from the details window

You can perform the following storage system management tasks from the Appliance or vFiler Details window:

- ▶ View specific storage system or vFiler unit details.
- ▶ Edit the storage system or vFiler unit configuration using FilerView or the Appliance Manager.
- ▶ View the active/active configuration status and perform takeover and giveback operations using the cluster console (on active/active controllers only).
- ▶ Access the vFiler units that are hosted on a storage system.
- ▶ Check active/active controller configurations.
- ▶ Edit the storage system configuration using FilerView or Appliance Manager.
- ▶ Edit Remote LAN Module (RLM) port settings for the storage system.
- ▶ View events related to the storage system or vFiler unit.
- ▶ Access group summary windows.
- ▶ View graphing information specific to each type of storage system.

Many of these actions are available from the Appliance Tools and Current Group selections in the bottom left of the Operations Manager window while viewing the Appliance Details window, as shown in Figure 9-61.



Figure 9-61 Options available from Appliance Details window

About Appliance Tools

Appliance Tools enables you to set up the parameters needed to communicate with a specific storage system. You can access Appliance Tools from the Details window for the storage system or hosting storage system (of a vFiler unit). The Appliance Tools menu is located at the lower left display of Operations Manager, as shown in Figure 9-61.

You can use the Edit Appliance Settings window for your system to specify or change storage system or vFiler settings. Note, however, that you can set global values for many settings using the Options window, and that you do not need to modify storage system or vFiler-level settings unless they differ from your global values.

You can access the Edit Appliance Settings window by clicking Edit Settings, as shown in Figure 9-62 on page 175. You can then to modify the following information using the window shown in Figure 9-63 on page 175:

- ▶ IP address of the storage system or the hosting storage system that Operations Manager monitors
You might want to change the storage system IP address if you want to use a different interface for administrative traffic.
- ▶ Login and password
You configure a login and password if you want to use Operations Manager to run a command on a system. Operations Manager uses this information to authenticate itself to the storage system on which the command is run.
- ▶ Threshold values
The threshold values indicate the level of activity that must be reached on the storage system before an event is triggered. Using these options, you can set specific storage

system or group thresholds. For example, the Appliance CPU Too Busy threshold indicates the highest level of activity the CPU can reach before a CPU Too Busy event is triggered. Threshold values specified on this window supersede any global values specified on the Options window.

► **Threshold intervals**

The threshold interval is the period of time during which a specific threshold condition must persist before an event is triggered. For example, if the monitoring cycle time is 60 seconds and the threshold interval is 90 seconds, the event is generated only if the condition persists for two monitoring cycles. You can configure threshold intervals only for specific thresholds, as listed on the Options window.

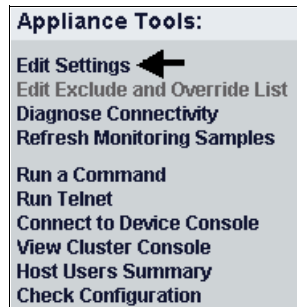


Figure 9-62 Select Edit Settings

Figure 9-63 Various Items can be changed from this window

Diagnose Connectivity

Use the Diagnose Connectivity tool to perform connectivity tests and review test outcome. You can access this tool by clicking the **Diagnose Connectivity** link under Appliance Tools, as shown in Figure 9-64.

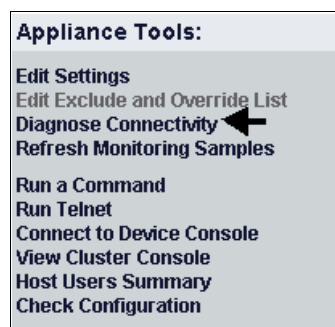


Figure 9-64 Diagnose Connectivity

There is a great deal of information in the window that appears that can be helpful in diagnosing connection issues, as shown in Figure 9-65.

Diagnostics - itsonas1

Network Connectivity

IP Address	9.11.218.163
Network	9.11.218.0/24 (last searched 10 Apr 10:37)
DNS Aliases	itsonas1
DNS Addresses	9.11.218.163
SNMP Version in Use	SNMPv1
SNMPv1	Passed (125 ms)
SNMP Community	public
SNMP sysName	itsonas1
SNMP sysObjectID	.1.3.6.1.4.1.789.2.3 (Clustered Filer)
SNMP productId	0101184428
SNMPv3	Failed: No SNMPv3 username specified.
SNMPv3 Username	
ICMP Echo	Passed (0 ms)
HTTP	Passed (0 ms)
NDMP (login not set)	Skipped
RSH	9.11.218.163: Permission denied. rsh: can't establish connection
SSH	FATAL ERROR: Network error: Connection refused
RLM	Skipped (hostRLMAddress is empty)
XML (http port 80)	HTTP POST - Authorization failed

Appliance Details

According to:	DataFabric Manager server	Host
Host Name	itsonas1	✓ itsonas1
System ID	0101184428	✓ 0101184428
Model	N5200	✓ N5200
Type	Clustered Storage System	✓ Clustered Storage System
OS Version	7.3.1	✓ 7.3.1
Revisions	310,7,3,2,0	✓ 310,7,3,2,0

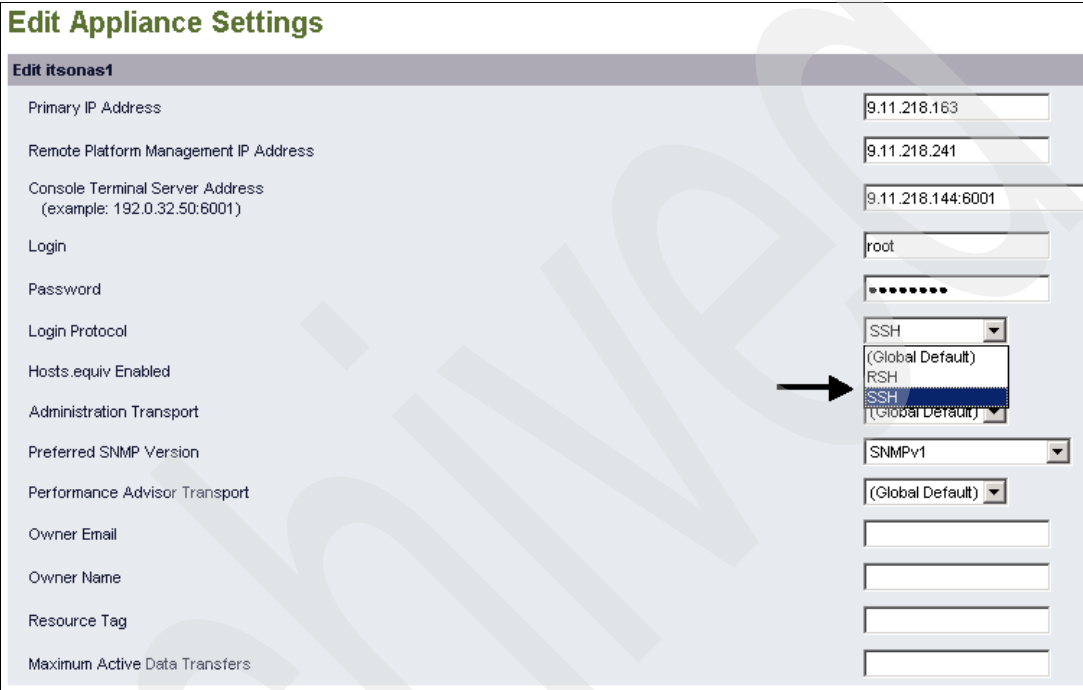
Monitoring Timestamps

Timestamp Name	Interval	Default	Last Updated	Error if older than ...
cacheTimestamp	5 minutes 5 minutes			10 Apr 10:35
ccTimestamp	4 hours 4 hours			10 Apr 06:40
cfTimestamp	5 minutes 5 minutes	10 Apr 10:40		10 Apr 10:35
cpuTimestamp	5 minutes 5 minutes	10 Apr 10:36		10 Apr 10:35
dfTimestamp	30 minutes 30 minutes	10 Apr 10:40		10 Apr 10:10
diskTimestamp	4 hours 4 hours	10 Apr 08:41		10 Apr 06:40
envTimestamp	5 minutes 5 minutes	10 Apr 10:37		10 Apr 10:35
fcTimestamp	5 minutes 5 minutes	10 Apr 10:38		10 Apr 10:35
fsTimestamp	15 minutes 15 minutes	10 Apr 10:40		10 Apr 10:25
hostPingTimestamp	1 minute 1 minute	10 Apr 10:39		10 Apr 10:39
ifTimestamp	15 minutes 15 minutes	10 Apr 10:40		10 Apr 10:25

Figure 9-65 Diagnostics window

Console connection through SSH

As we discovered in our testing, console connection is a very important aspect of the overall ability of Operations Manager to access and display pertinent information about the health and vitality of your storage systems and network connections. It is vitally important that certain functions be enabled on the storage systems before Operations Manager can access some functions. When you are setting options either globally or individually on your options screens for the storage systems you are interested in, you will find that you only have two choices when it comes to login protocol. The default is set to RSH. You have the option of setting the protocol to RSH or SSH, as shown in Figure 9-66.



The screenshot displays the 'Edit Appliance Settings' page for a device named 'itsonas1'. The page contains various configuration fields. The 'Login Protocol' dropdown menu is open, showing three options: 'SSH', 'RSH', and 'SSH (Global Default)'. An arrow points to the 'SSH' option, which is highlighted. Other visible settings include IP addresses, console terminal server address, login name, password, administration transport, preferred SNMP version, performance advisor transport, owner email, owner name, resource tag, and maximum active data transfers.

Setting	Value
Primary IP Address	9.11.218.163
Remote Platform Management IP Address	9.11.218.241
Console Terminal Server Address (example: 192.0.32.50:6001)	9.11.218.144:6001
Login	root
Password
Login Protocol	SSH
Hosts.equiv Enabled	(Global Default)
Administration Transport	RSH
Preferred SNMP Version	SNMPv1
Performance Advisor Transport	(Global Default)
Owner Email	
Owner Name	
Resource Tag	
Maximum Active Data Transfers	

Figure 9-66 Select RSH or SSH

It may seem like a small thing, but it is vitally important that this protocol choice be reflected on the storage device you are monitoring. If it is not, you will not be able to accomplish such tasks as viewing configurations, accessing the RLM, and others.

A good check to verify that you have these protocols enabled on the desired storage systems is to run the Diagnose Connectivity option associated with the storage device, as shown in Figure 9-67.



Figure 9-67 Run Diagnose Connectivity to check proper protocol access

This test takes just a little while to run, but gives you a wealth of information concerning the health of your network and the SSH and RSH status. In Figure 9-65 on page 176, you can see that we did not have SSH enabled on the storage device (itsonas1) when this picture was taken. After some trial and error we discovered the importance of having SSH protocol enabled on the storage system. Refer to Figure 9-68 to see the difference in the report after enabling the SSH protocol on the itsonas1 storage system.

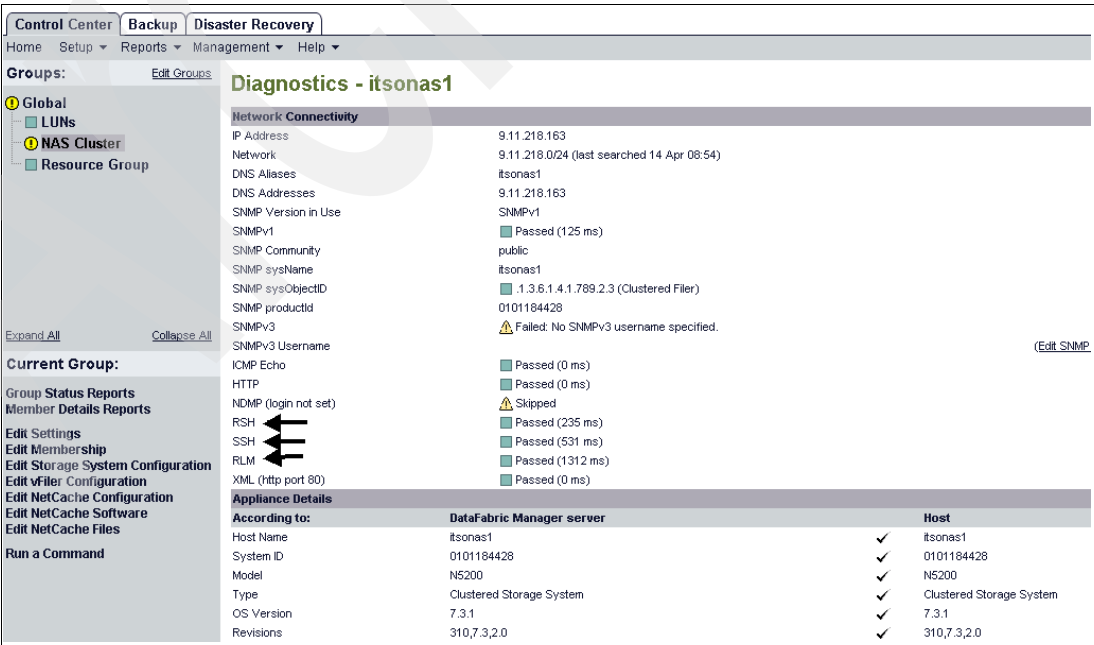


Figure 9-68 Network status after SSH enabled on storage system

In order to enable the SSH protocol on the storage system, you must log into the N series storage system and perform the following tasks.

1. Check the status of the SSH protocol by running the following command:

```
itsonas1> secureadmin status
```

2. If SSH is not enabled on your system, then you must run this command:

```
itsonas1>secureadmin setup ssh
```

The **secureadmin setup ssh** command will guide you through the setup of the SSH protocol for the device. (We used the defaults.) This command brings up the steps shown in Example 9-1.

Example 9-1 secureadmin setup

```
itsonas1> secureadmin setup ssh
SSH Setup
-----
Determining if SSH Setup has already been done before...no
SSH server supports both ssh1.x and ssh2.0 protocols.
SSH server needs two RSA keys to support ssh1.x protocol. The host key is
generated and saved to file /etc/sshd/ssh_host_key during setup. The server key is
re-generated every hour when SSH server is running.
SSH server needs a RSA host key and a DSA host key to support ssh2.0 protocol. The
host keys are generated and saved to /etc/sshd/ssh_host_rsa_key and
/etc/sshd/ssh_host_dsa_key files respectively during setup.
SSH Setup will now ask you for the sizes of the host and server keys.
For ssh1.0 protocol, key sizes must be between 384 and 2048 bits.
For ssh2.0 protocol, key sizes must be between 768 and 2048 bits.
The size of the host and server keys must differ by at least 128 bits.
Enter the size of host key for ssh1.x protocol [768] :(enter)
Enter the size of server key for ssh1.x protocol [512] :(enter)
Enter the size of host keys for ssh2.0 protocol [768] :(enter)
You have specified these parameters:
    host key size = 768 bits
    server key size = 512 bits
    host key size for ssh2.0 protocol = 768 bits
Is this correct? [yes] (enter)
Setup will now generate the host keys. It will take a minute.
After Setup is finished the SSH server will start automatically.
itsonas1> Tue Apr 14 08:49:22 MST [itsonas1: secureadmin.ssh.setup.success:info]:
SSH setup is done and ssh2 should be enabled. Host keys are stored in
/etc/sshd/ssh_host_key, /etc/sshd/ssh_host_rsa_key, and
/etc/sshd/ssh_host_dsa_key.
```

3. Next, check the status of the ssh protocol's settings. The **options ssh** command will give you the status of the settings (see Example 9-2).

Example 9-2 options ssh command

```
itsonas1*> options ssh
ssh.access                *
ssh.enable                on
ssh.idle.timeout          600
ssh.passwd_auth.enable    on
ssh.port                  22
ssh.pubkey_auth.enable    on
ssh1.enable               off
```

4. If you are satisfied with the settings, run the following command to enable SSH:

```
itsonas1>secureadmin enable ?
```

Here are the different options you can use with the **secureadmin** command:

- **secureadmin setup [-f] [-q] ssh**
- **secureadmin setup [-f] [-q] ssl**
- **secureadmin addcert ssl [<path to CA signed cert>]**
- **secureadmin enable all|ssh|ssh1|ssh2|ssl**
- **secureadmin disable all|ssh|ssh1|ssh2|ssl**
- **secureadmin status**

5. Once we finish the activation of SSH, we run the **secureadmin status** command to be sure SSH is enabled (see Example 9-3).

Example 9-3 secureadmin status command

```
itsonas1> secureadmin status
ssh2    - active
ssh1    - active
ssl     - inactive
```

In our case, the protocol is showing active, but as we saw in Figure 9-65 on page 176, that was not the case. We had to run the **secureadmin setup ssh** command to enable SSH.

Console connection through telnet

Use the Connect Device to Console tool to connect to the storage system console, if the storage system is known to Operations Manager, as shown in Figure 9-69. The storage system must be *connected to a terminal server* for Operations Manager to connect to the storage system console.

Note: Before initiating the console connection, you must set the Console Terminal Server Address in the Edit Settings window.

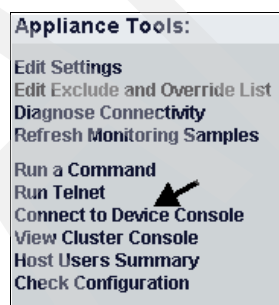


Figure 9-69 Connect to Device Console

Many of the operations you perform from the FilerView of a storage system or from the N series command line can be performed from the Operations Manager user interface. One of those activities is cluster failover. Next, we cover some of the steps to be performed to accomplish this from the Operations Manager interface.

Cluster failover

The following requirements must be met before you can use the cluster console:

- ▶ The Operations Manager license must be installed on your DataFabric Manager Server.
- ▶ An authentication method must be set up for Operations Manager to authenticate to the controller on which takeover and giveback operations are to be performed.

Accessing the cluster

To access the cluster console for an active/active controller, select **Member Details** → **Appliances** → **Tools: View Cluster Console**, as shown in Figure 9-70.

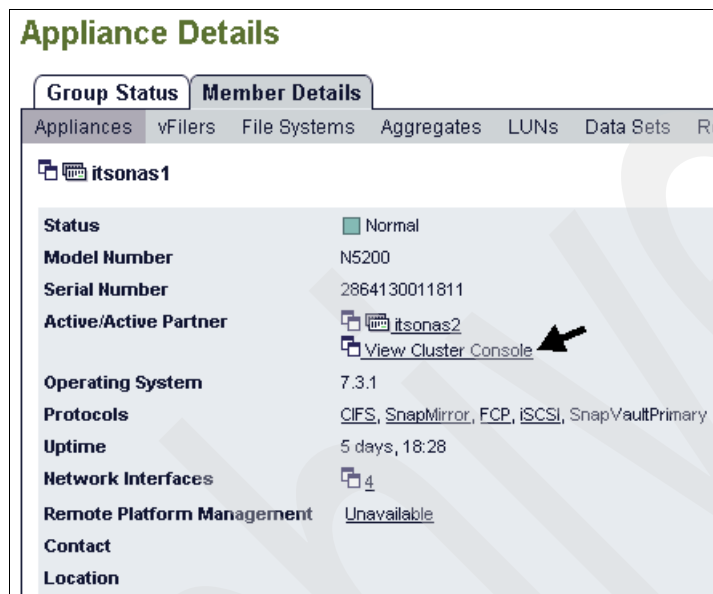


Figure 9-70 View Cluster Console

The Cluster Console window is an automatically refreshing window that displays the nearly real-time status of an active/active configuration, as shown in Figure 9-71.

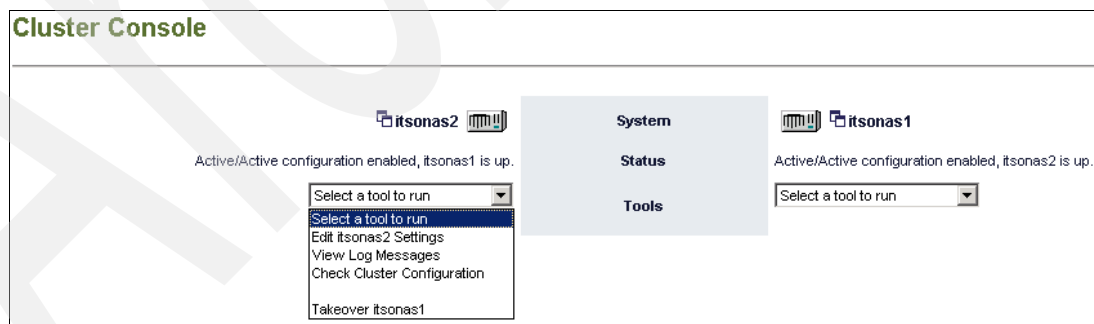


Figure 9-71 Cluster Console window

The Cluster Console window contains the following information:

► **System**

Displays the names of controllers that form an active/active configuration.

The color of the controller icon (next to the controller name) indicates the status of the controller. The following list describes the status that each color indicates:

- Gray: Controller is up and functioning normally.
- Blue: Controller is up, but failover is disabled. Data is available in this condition; however, it is unprotected.
- Red: Controller is down and its services have not been taken over by its partner.
- Pink: Controller hardware is down, but controller's partner has taken over services. Data from the down controller is available in this condition through its partner, but the performance of data service is degraded and the data residing on both controllers is unprotected.
- Yellow: Controller is up and also serving for the partner. Data from both controllers is available, but the performance of data service is degraded and the data residing on both controllers is unprotected.
- White with a ? sign: Controller status is unknown. The controller might be down or not configured correctly to respond to SNMP queries from Operations Manager.

If a controller is in unknown status, use the Diagnose Connectivity tool in the Control Center or run the **dfm host diag** command from the DFM Server CLI to identify the cause of the problem.

– **Status**

Displays the real-time status, as a text message, of controllers in the active/active configuration.

► **Tools**

Provides tools that enable you to edit the settings and perform takeover and giveback operations on the controllers in an active/active configuration. The following tools are available:

– **Edit *system_name* Settings**

Displays the Edit Appliance Settings window, where you can configure the login and password information that Operations Manager uses to authenticate to the controllers.

– **Manage with FilerView**

Invokes the Web-based UI, FilerView, for the controllers. You can use FilerView to make configuration changes on the controllers.

– **View Log Messages**

Displays the syslog messages that are logged on the controllers.

– **Check Cluster Configuration**

Enables you to check and verify the active/active configuration of the controllers. This tool displays the Active/Active Configuration Status window that contains the following information:

- **RSH Enabled:** Reports whether Operations Manager can log in to the controllers that form the active/active configuration.
- **Failover Status:** Reports whether failover is configured and enabled on both controllers.

- **OS Version:** Reports whether the OS versions of the two active/active configurations are the same. If the OS versions are different, the active/active configurations might not interact correctly during failover.
- **Software Licenses:** Reports whether the two controllers have the same list of licensed software. If the sets of licenses are different, some features might not be available after a failover.
- **Options:** Reports whether Data ONTAP options, such as `timed.enable` and `snmp.enable`, that are expected to be the same on both controllers of the active/active configuration are, in fact, the same. If options are different in active/active configurations, the controllers might not behave as expected after a takeover.

To discover all the options that must be or are recommended to be the same on both controllers, see the Data ONTAP man window for the **options** command.

- **Network Configuration:** Reports whether the network interfaces on each controller correctly refer to the partner's interfaces. If not, networking might not work correctly after takeover.

To be able to check the cluster configuration you must have SSH enabled on both heads of the cluster. Refer to “Console connection through SSH” on page 177 for a discussion about setting the SSH protocol on the storage device before running this option. From the appliance detail window, select **View Cluster Console**, as shown in Figure 9-72.

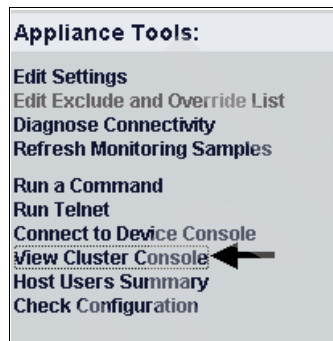


Figure 9-72 Select View Cluster Console

In Figure 9-73, select **Check Cluster Configuration** from the drop-down menu for the heading you want to view. However, the information you receive summarizes the status for both heads of the active/active configuration.

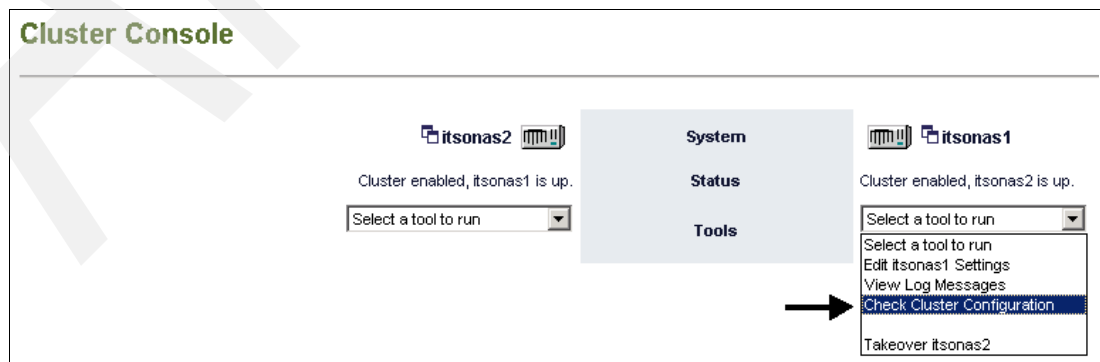


Figure 9-73 Check Cluster Configuration

Finally, you get a summary view of the status of your active/active configuration, as shown in Figure 9-74.

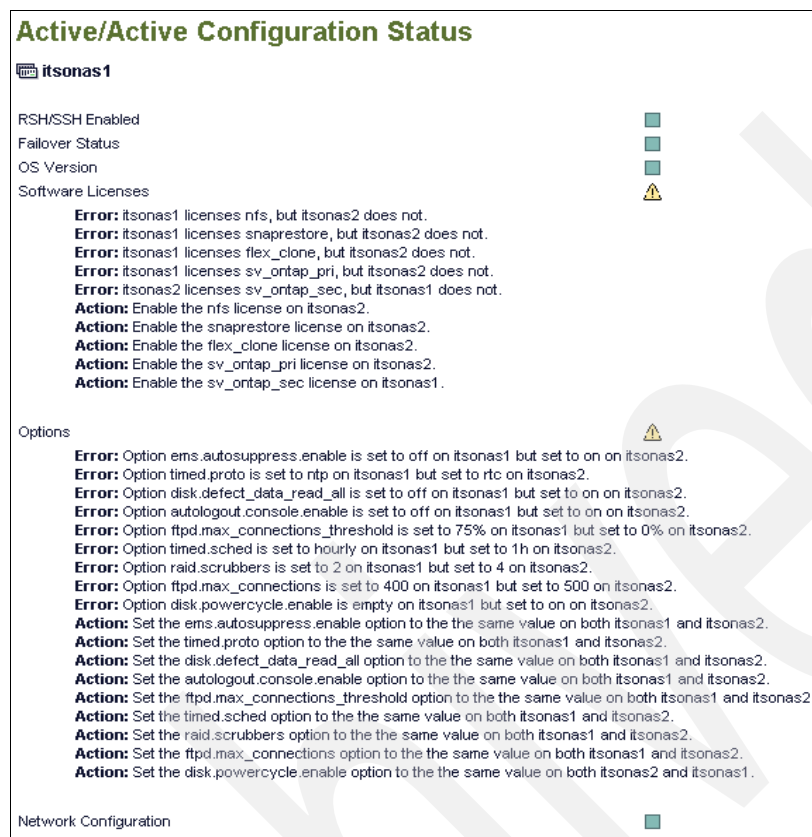


Figure 9-74 Configuration status of cluster

In the Select a tool to run drop-down menu shown in Figure 9-73 on page 183, the following tools are available:

- **Takeover *system_name***
Initiates a takeover of a controller's partner. This tool is available when a controller can take over its partner.
- **Giveback *system_name***
Initiates a giveback of a controller's partner. This tool is available when a controller has taken over its partner.

For more information about how to use the Operations Manager interface to perform takeover and giveback activities for clustered devices, refer to Chapter 12 in *Operations Manager Administration Guide for Use With DataFabric Manager Server 3.7*, GC26-7889.

Remote configuration management

As you monitor your storage systems, you might find that you need to alter the configuration settings of one or more storage systems. Operations Manager provides three methods by which you can remotely configure your storage systems:

- Accessing the storage system CLI
- Accessing FilerView or the Appliance Manager
- Using the Operations Manager multiple-storage system remote configuration feature

In this book, we look at remote configuration through FilerView and Operations Manager. For information about the DFM Server CLI method, refer to Chapter 12 in *Operations Manager Administration Guide for Use With DataFabric Manager Server 3.7*, GC26-7889.

In addition to providing access to the storage system's CLI, Operations Manager enables you to log in to a storage system's management UI, which is FilerView for storage systems. When you invoke FilerView, Operations Manager spawns a new window. Using FilerView, you can edit any or all of a storage system's configuration settings.

You can access FilerView by clicking the icon next to the storage system or vFiler unit name in the details windows for events, storage systems, vFiler units, aggregates, LUNs, qtrees, and volumes.

To access FilerView for a selected storage system or vFiler unit, click the storage system icon next to the storage system or vFiler unit name in the respective details window, as shown in Figure 9-75.

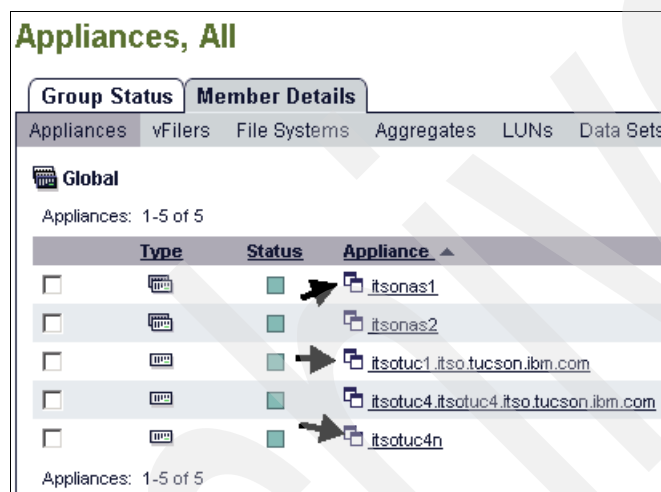


Figure 9-75 Select icons to start FilerView

Once you have authenticated with the storage system, FilerView comes up as a second window, as shown in Figure 9-76.

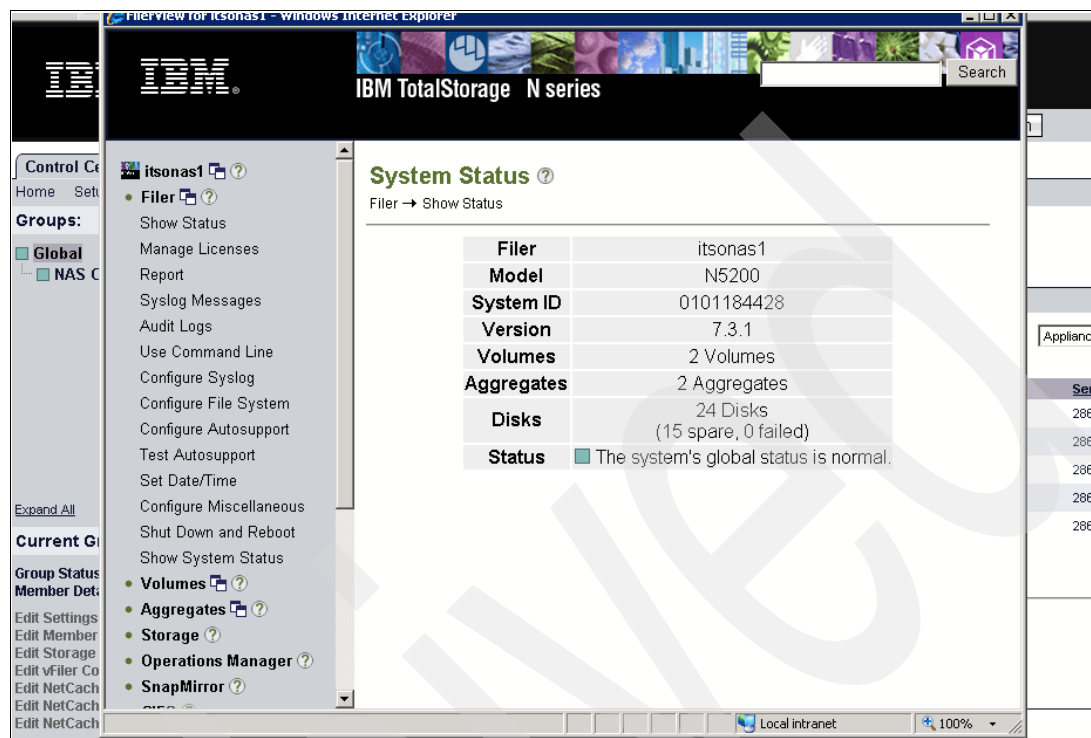


Figure 9-76 FilerView window

From here, you would proceed just as though you were running the FilerView GUI from any browser window.

Remote configuration

The remote configuration feature enables you to remotely configure all storage systems in a group at the same time. By grouping the N series storage systems into special groups called *configuration resource groups* and applying configuration settings to those groups, administrators can remotely configure multiple storage systems in a single action.

You create configuration resource groups using the same method used to create groups. You create an empty group and populate the group with storage systems.

You create empty groups by using the Add a New Group option on the Edit Groups window (select the **Edit Groups** link in the left pane). You populate the group using the Edit Group Membership window (in the left pane, select the group name and select **Edit list** → **Group Membership**).

Configuration files

You can use two methods to create configuration files. If you have an appliance that already has most or all of the settings you desire, you can create a configuration file that copies an storage system's configuration settings. This process is also called *cloning the configuration*. Cloning a configuration is the recommended method of creating a configuration file. If required, you can subsequently edit the cloned file to modify the configuration settings.

Alternatively, you can create an empty configuration file and edit it to contain the desired configuration settings, as shown in Figure 9-77 on page 187.

Storage System Configurations 10 Apr 15:4

Storage System | vFiler | NetCache

Create New Configuration File

Create new configuration file named: pulled from storage system: itsnas1 ▼ Create

Configuration Files

Name	Version	Edit	Last Modified	Select
There are no configuration files.				

Export Configuration Files

☒ To file on the DataFabric Manager server:

☐ To file on your computer

Export

Import Configuration Files

☒ From file on the DataFabric Manager server:

☐ From file on your computer: Browse...

Import

Figure 9-77 Creating and editing existing files

For more information about creating configuration files, refer to the online help.

You can pull configuration files from storage systems by selecting **Management** → **Storage Systems** → **Configuration Files**, as shown in Figure 9-78.

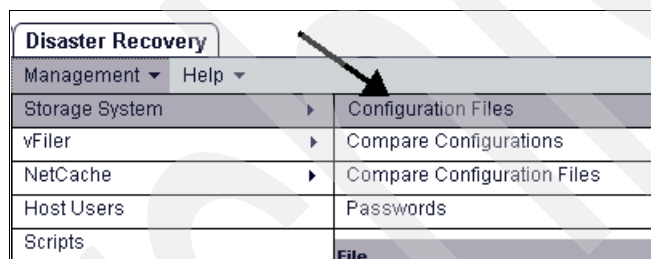


Figure 9-78 Accessing the current configurations files

When you add one or more configuration files to the group, the group becomes a resource configuration group. After the configuration files have been associated with the group, an icon is attached to the group name so that you can identify the group as a configuration resource group, as shown in Figure 9-79 on page 188.

Select **Edit Storage System Configuration**, as shown in Figure 9-79, to reach the configuration window shown in Figure 9-80.



Figure 9-79 Edit Configuration

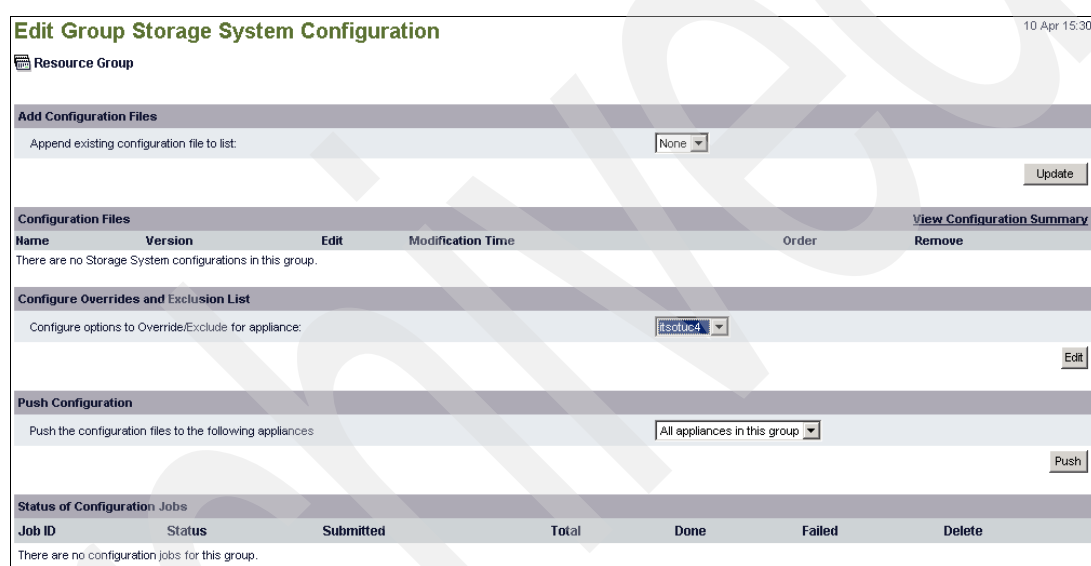


Figure 9-80 Adding, editing, and pushing configuration files

More information about creating and managing multiple files can be found in Chapter 13 in *Operations Manager Administration Guide for Use With DataFabric Manager Server 3.7*, GC26-7889.

LUNs, FCP targets, and SAN hosts

To monitor and manage LUNs, FCP targets, and SAN hosts, the DataFabric Manager Server must first discover them. The DataFabric Manager Server uses SNMP to discover storage systems, but SAN hosts must already have the Operations Manager Host Agent software installed and configured on them before the DataFabric Manager Server can discover them.

After SAN components have been discovered, the DataFabric Manager Server starts collecting pertinent data, for example, which LUNs exist on which storage systems. Data is collected periodically and reported through various Operations Manager reports. (The frequency of data collection depends on the values assigned to the DataFabric Manager Server monitoring intervals.)

The DataFabric Manager Server monitors LUNs, FCP targets, and SAN hosts for a number of predefined conditions and thresholds, such as when the state of an HBA port changes to online or offline or when the traffic on an HBA port exceeds a specified threshold. If a

predefined condition is met or a threshold is exceeded, the DataFabric Manager Server generates and logs an event in its database. These events can be viewed through the Details window of the affected object. Additionally, you can configure the DataFabric Manager Server to send notification about such events (also known as *alarms*) to an e-mail address, a windowr, an SNMP trap host, or a script you write.

In addition to monitoring LUNs (select **Group** → **Member Details** → **LUNs,All**, as shown in Figure 9-81), FCP targets (select **Group** → **Member Details** → **FCP Targets**, as shown in Figure 9-82), and SAN hosts (select **Groups** → **Member Details** → **SAN Hosts,All**, as shown in Figure 9-83 on page 190), you can use the DataFabric Manager Server to manage these components. For example, you can create, delete, or expand a LUN.

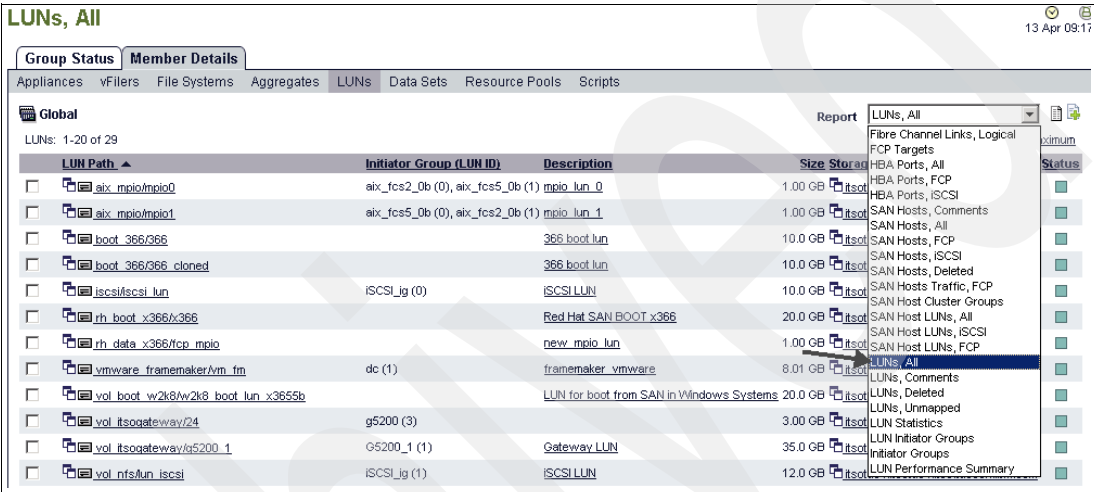


Figure 9-81 Monitoring LUNs

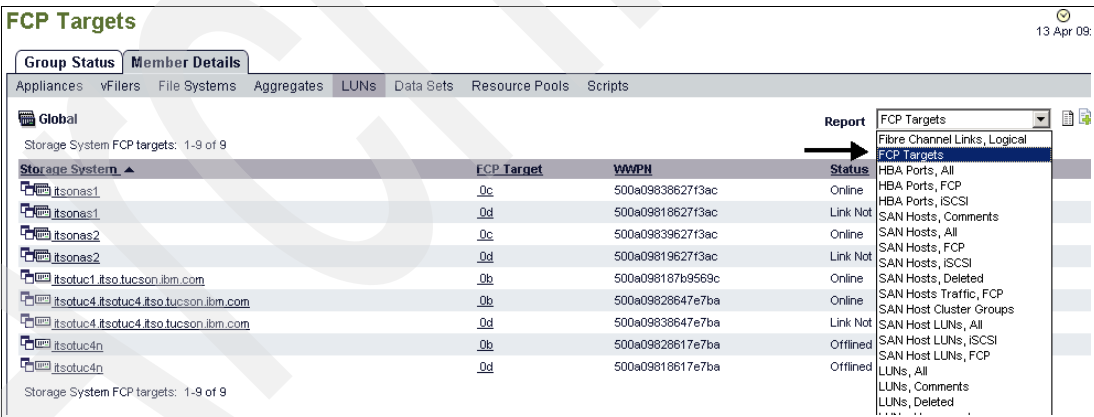


Figure 9-82 Monitoring FCP targets

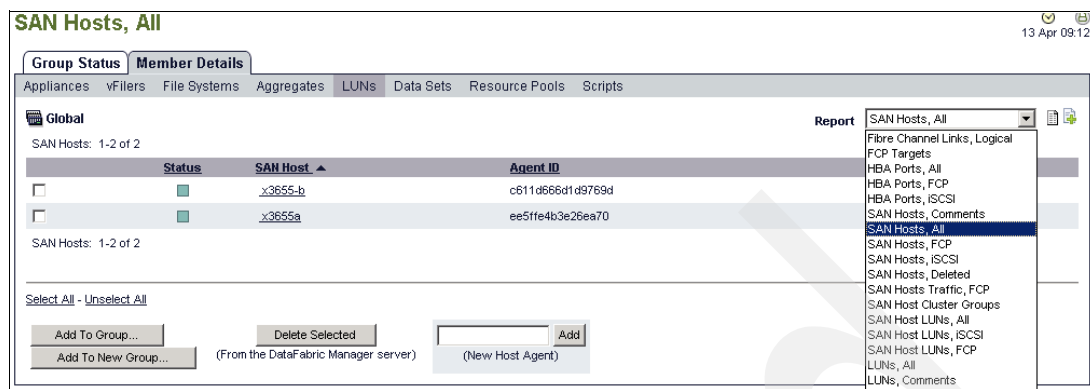


Figure 9-83 Monitoring SAN hosts

The DataFabric Manager Server can automatically discover SAN hosts; however, it does not use SNMP to poll for new hosts. Instead, the special purpose Operations Manager Host Agent software discovers, monitors, and manages SANs on SAN hosts. You must install the Operations Manager Host Agent software on each SAN host that you want to monitor and manage with the DataFabric Manager Server.

After the Operations Manager Host Agent software is installed on a client host and you have installed the DataFabric Manager Server with the Operations Manager license, you can perform a variety of tasks:

- ▶ Monitor basic system information for SAN clients and related devices
- ▶ Perform management functions, such as creating, modifying, or expanding a LUN
- ▶ View detailed LUN information

Operations Manager Host Agent is also used for Storage Resource Management functions, with an File SRM license.

We look at some of these functions in greater detail, but first let us look at some of the prerequisites for SAN hosts. Some of this information is also found in Chapter 5, “Host Agent installation for Windows 2003” on page 79 and Chapter 6, “Host Agent installation for Linux” on page 91.

SAN management prerequisites

The prerequisites for managing SAN hosts with the DataFabric Manager Server include the following:

- ▶ All SAN hosts to be managed by the DataFabric Manager Server must be connected to a TCP/IP network either known to or discoverable by the DataFabric Manager Server. The SAN hosts must be connected to the network through an Ethernet port and must each have a valid IP address.
- ▶ Each SAN host must have the Operations Manager Host Agent software installed on it. The Operations Manager Host Agent software is required for discovering, monitoring, and managing SAN hosts.
- ▶ For LUN management using the Operations Manager server, Windows SAN hosts must have the proper version of SnapDrive® software installed.

LUNs

You can access LUN information from the Control Center Groups window by selecting **LUNs** under the Storage Capacity heading, as shown in Figure 9-84 on page 191.

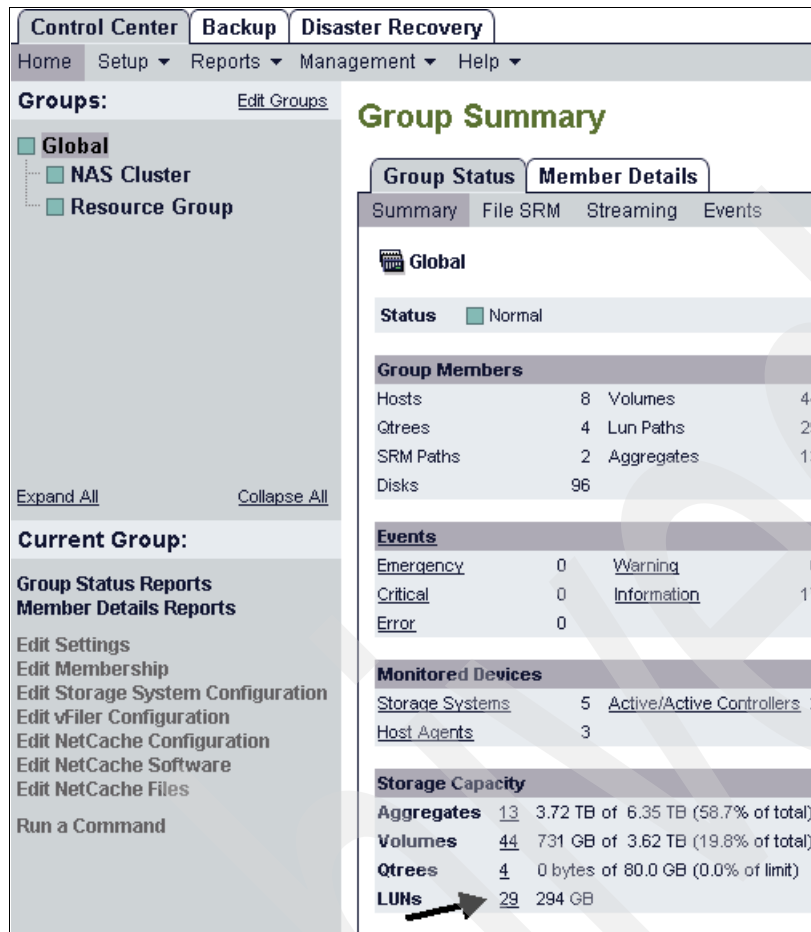


Figure 9-84 Access LUN information from Groups window

You can also see detailed LUN information from the Member Details window by selecting the **LUNs, All** from the drop-down menu, as shown in Figure 9-81 on page 189.

The LUN Details window for a LUN consists of the following information (as shown in Figure 9-85 on page 192):

- ▶ Status of the LUN
- ▶ Storage system on which the LUN exists
- ▶ Volume or qtree on which the LUN exists
- ▶ Size of the LUN
- ▶ Serial number of the LUN
- ▶ Description of the LUN
- ▶ Events associated with the LUN
- ▶ Groups to which the LUN belongs
- ▶ Number of LUNs configured on the storage system on which the LUN exists and a link to a report displaying those LUNs
- ▶ Number of SAN hosts mapped to the LUN and a link to the report displaying those hosts
- ▶ Number of HBA ports that can access this LUN and a link to the report displaying those LUNs

- ▶ Time of the last sample collected and the configured polling interval for the LUN
- ▶ Initiator groups to which the LUN is mapped

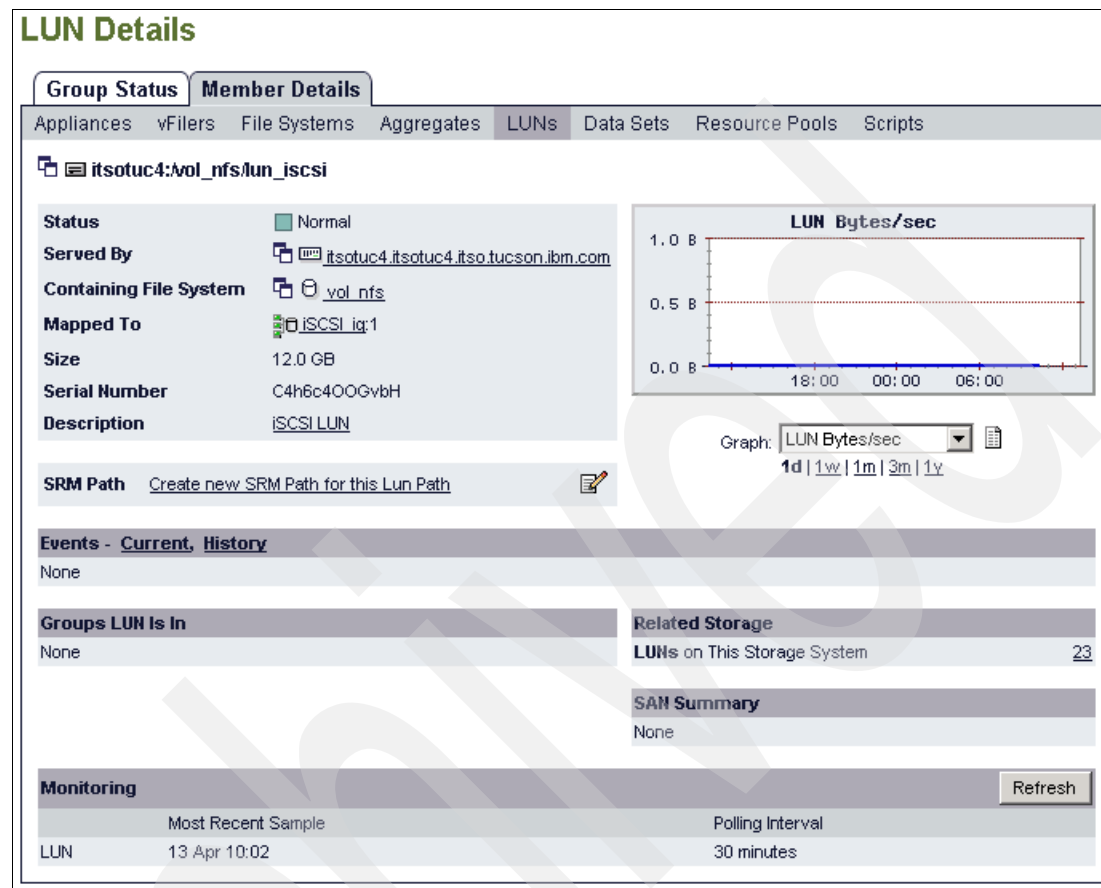


Figure 9-85 LUN Details window

Use the LUN Path Tools links on left side of the LUN Details window to perform the following tasks, as shown in Figure 9-86:

- ▶ **Expand this LUN:** Launches a wizard that helps you expand the LUN.
- ▶ **Destroy this LUN:** Launches a wizard that helps you destroy the LUN.
- ▶ **Refresh Monitoring Samples:** Obtains current monitoring samples from the storage system on which this LUN exists.
- ▶ **Run a Command:** Runs a Data ONTAP command on the storage system on which this LUN exists.

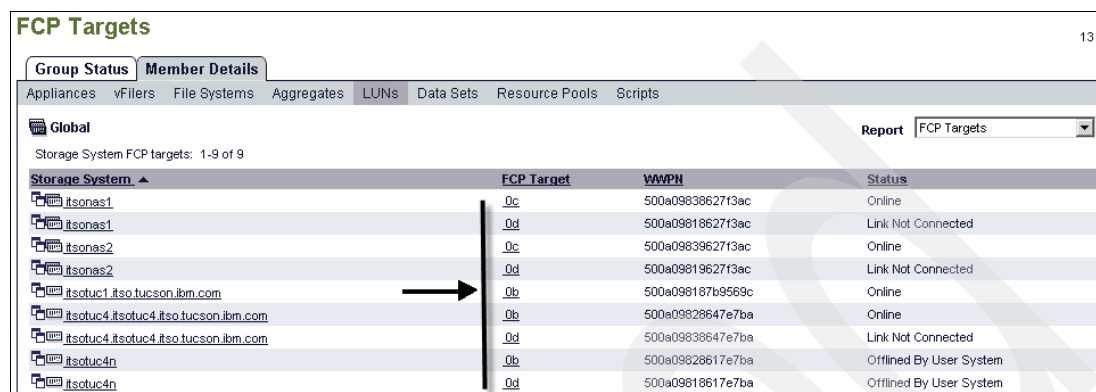
You can also gain access to FilerView from the Details window by clicking next to the Storage System name. From FilerView, you can perform any action on the LUN you would like.



Figure 9-86 LUN Path Tools in lower left of Operations Manager window


FCP Targets

The FCP Targets details window can be accessed by clicking its port number in the FCP Targets Report (as shown in Figure 9-87), which can be accessed by selecting **LUNs** → **FCP Targets**.



FCP Targets			
Group Status		Member Details	
Appliances	vFilers	File Systems	Aggregates
Global			
Storage System FCP targets: 1-9 of 9			
Storage System	FCP Target	WWPN	Status
itsonas1	0c	500a09838627f3ac	Online
itsonas1	0d	500a09818627f3ac	Link Not Connected
itsonas2	0c	500a09839627f3ac	Online
itsonas2	0d	500a09819627f3ac	Link Not Connected
itsonas1:0b	0b	500a098187b9569c	Online
itsonas1:0b	0b	500a09828647e7ba	Online
itsonas1:0d	0d	500a09838647e7ba	Link Not Connected
itsonas1:0d	0d	500a09828617e7ba	Offline By User System
itsonas1:0d	0d	500a09818617e7ba	Offline By User System

Figure 9-87 Select port number for details window



FCP Target Details	
Group Status	
Member Details	
Appliances	vFilers
itsonas1:0c	
Host Containing	itsonas1
Operational Status	Online
Adapter Hardware Version	3.3.24
Adapter Firmware Version	3
Adapter Speed	2 Gbits/s
Fibre Channel Topology	Fabric
Node Name (WWNN)	500a09808627f3ac
Port Name (WWPN)	500a09838627f3ac
Related Devices	
FCP Targets on The Same Storage System	2
HBA Ports accessible to This FCP Target	2
Monitoring	
	Refresh
	Most Recent Sample
	Polling Interval
Ping	13 Apr 11:02
Fibre Channel	13 Apr 11:02

Figure 9-88 FCP Target Details window

The FCP Target details window contains the following information (as shown in Figure 9-88):

- ▶ Name of the storage system on which the target is installed
- ▶ Operational status of the target
- ▶ Adapter hardware and firmware versions
- ▶ Adapter speed
- ▶ FC topology of the target
 - Fabric

- Point-To-Point
- Loop
- Unknown
- ▶ Node name (WWNN) and port name (WWPN) of the target
- ▶ Number of other FCP targets on the storage system on which the target is installed (link to report)
- ▶ Time of the last sample collected and the configured polling interval for the FCP target

Host Agent details

The Host Agent details window displays information reported by the Operations Manager Host Agent on a selected SAN host. You can access the details window for a Operations Manager Host Agent by clicking its name in any of the SAN Host reports. For more information about the SAN host reports, see the online help.

You may also get to the reports window from the main Operations Manager Groups summary window by clicking **Hosts** under Monitored Devices, as shown in Figure 9-89.

The screenshot shows the IBM System Storage N series Operation Manager interface. The top navigation bar includes 'Control Center', 'Backup', and 'Disaster Recovery'. Below this is a menu bar with 'Home', 'Setup', 'Reports', 'Management', and 'Help'. The left sidebar shows a tree view of groups: 'Global' (selected), 'NAS Cluster', and 'Resource Group'. The main content area is titled 'Group Summary' and has two tabs: 'Group Status' and 'Member Details'. The 'Group Status' tab is active, showing a 'Summary' sub-tab. The status is 'Normal'. Below this, there's a 'Group Members' table with columns for component name and count. The 'Events' table shows counts for various event types. The 'Monitored Devices' table at the bottom lists 'Storage Systems' (5), 'Active/Active Controllers' (2), and 'Host Agents' (3). A black arrow points to the 'Host Agents' entry.

Group Summary			
Group Status		Member Details	
Summary			
Global			
Status	Normal		
Group Members			
Hosts	8	Volumes	44
Qtrees	4	Lun Paths	29
SRM Paths	2	Aggregates	13
Disks	96		
Events			
Emergency	0	Warning	0
Critical	0	Information	17
Error	0		
Monitored Devices			
Storage Systems	5	Active/Active Controllers	2
Host Agents	3		

Figure 9-89 Access to Host Agents details window

When you click **Host Agents**, you will be presented with the list of host agents being monitored in your Operations Manager infrastructure, as shown in Figure 9-90 on page 195.



Figure 9-90 Select the agent for details

The Details window for a host agent on a SAN host contains the following information (as shown in Figure 9-91 on page 196):

- ▶ Status of the SAN host and the time since the host has been up
- ▶ The operating system and the Operations Manager Host Agent version, as well as protocols and features running on the SAN host
- ▶ The MSCS configuration information about the SAN host, if any, such as the cluster name, cluster partner, and cluster groups to which the SAN host belongs
- ▶ The events that occurred on this SAN host
- ▶ The devices related to the SAN host, such as the storage systems accessible from the SAN host
- ▶ Graphs of information, such as the HBA port traffic per second or the HBA port frames for different time intervals
- ▶ Time of the last sample collected and the configured polling interval for the SAN host.

Host Agent Details

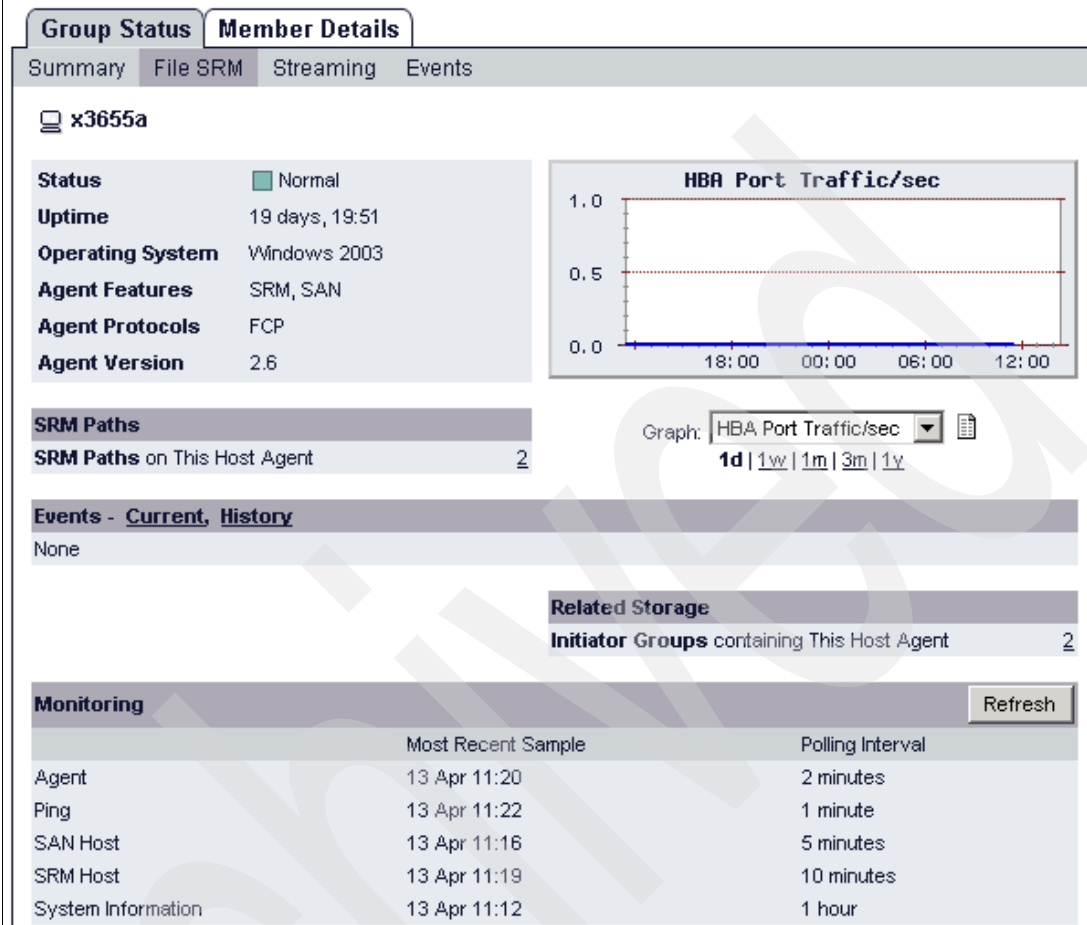


Figure 9-91 Host Agent Details window

The Host Tools list on the Host Agent Details window enables you to perform the following tasks (as shown in Figure 9-92 on page 197):

- ▶ **Edit Settings:** Displays the Edit Host Agent Settings window where you configure login, password, administrative transport, port information for the SAN host, and the user name and password for CIFS access in Operations Manager.
- ▶ **Create a LUN:** Takes you to a LUN creation window that helps you create a LUN.
- ▶ **Diagnose Connectivity:** Automates connectivity troubleshooting.
- ▶ **Refresh Monitoring Samples:** Obtains current monitoring samples from the SAN host.
- ▶ **Manage Host Agent:** Allows you to edit settings for the host agent, including monitoring and management of API passwords and the HTTP and HTTPS ports, enabling remote upgrading, and specifying a filewalk log path.




Figure 9-92 Tasks to perform on the host agent

You can group LUNs, storage systems, or SAN hosts for easier management and to apply access control, as shown in Figure 9-93.

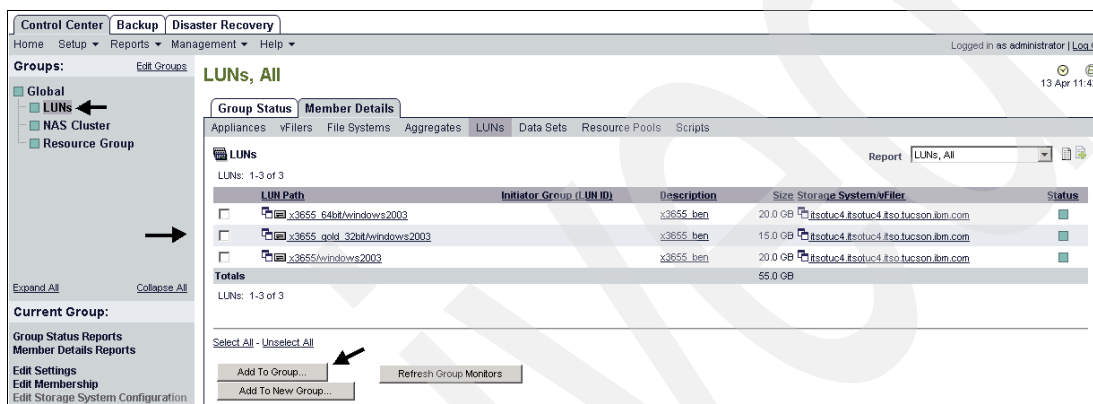


Figure 9-93 Example of grouping LUNs

Both storage systems and SAN hosts are considered storage systems for the purpose of creating groups. Therefore, when you create a group of storage systems or SAN hosts, the type of the created group is “Appliance resource group.” You can also add storage systems or SAN hosts to an Appliance resource group.

However, a group containing LUNs cannot contain any other type of object. When you create a group of LUNs, the created group is “LUN resource group.”

You cannot group HBA ports or FCP targets.

From the Administrators window (select **Setup** → **Administrative Users** to access it), you can allow an administrator access for managing all your SAN hosts and LUNs. The GlobalSAN role allows an administrator to create, expand, and destroy LUNs.

To create a new administrator with management capabilities for your SAN hosts and LUNs, fill in the fields for the Add a New Administrator option on the Administrators window. Select GlobalSAN from the Roles list, as shown in Figure 9-94.

There are other actions you can perform on SAN objects as well. For complete information about all of the monitoring and managing functions of LUNs, FCP targets, and SAN hosts, refer to Chapter 10 in *Operations Manager Administration Guide for Use With DataFabric Manager Server 3.7*, GC26-7889.

Add Roles

Administrative Users **Roles** **Details**

Add a New Role

Role Name

Description

Owner Role

Capabilities

Resource	Operations	View
Add Capabilities...		

Inherit Capabilities from other Roles

GlobalSAN
GlobalSDConfig
GlobalSDDataProtection
GlobalSDDataProtectionAndRestore
GlobalSDFullControl
GlobalSDSnapshot
GlobalSDStorage
GlobalSRM
GlobalWrite

>>
<<

Figure 9-94 Setting up a SAN administrator

Performance Advisor operation and configuration

Performance Advisor is a Java-based client application that provides you with a single location to view comprehensive performance information about storage systems and vFiler units.

Performance Advisor runs in the IBM System Storage N series Management Console. You can use the performance information to achieve the following objectives:

- ▶ Monitor storage systems or vFiler units for usage levels and optimal functioning by accessing historical and real-time data for all monitored devices recognized by the DataFabric Manager Server.
- ▶ Identify potential blockages in the data infrastructure.
- ▶ Perform short-term trend analysis for the data infrastructure.

10.1 Performance Advisor overview

The Performance Advisor application relies on a topology that includes the N series Management Console (which provides the interface to Performance Advisor), DataFabric Manager Server (which includes a *performance-monitoring server*), and one or more storage systems or vFile units monitored by the DataFabric Manager Server. This architecture is shown in Figure 10-1.

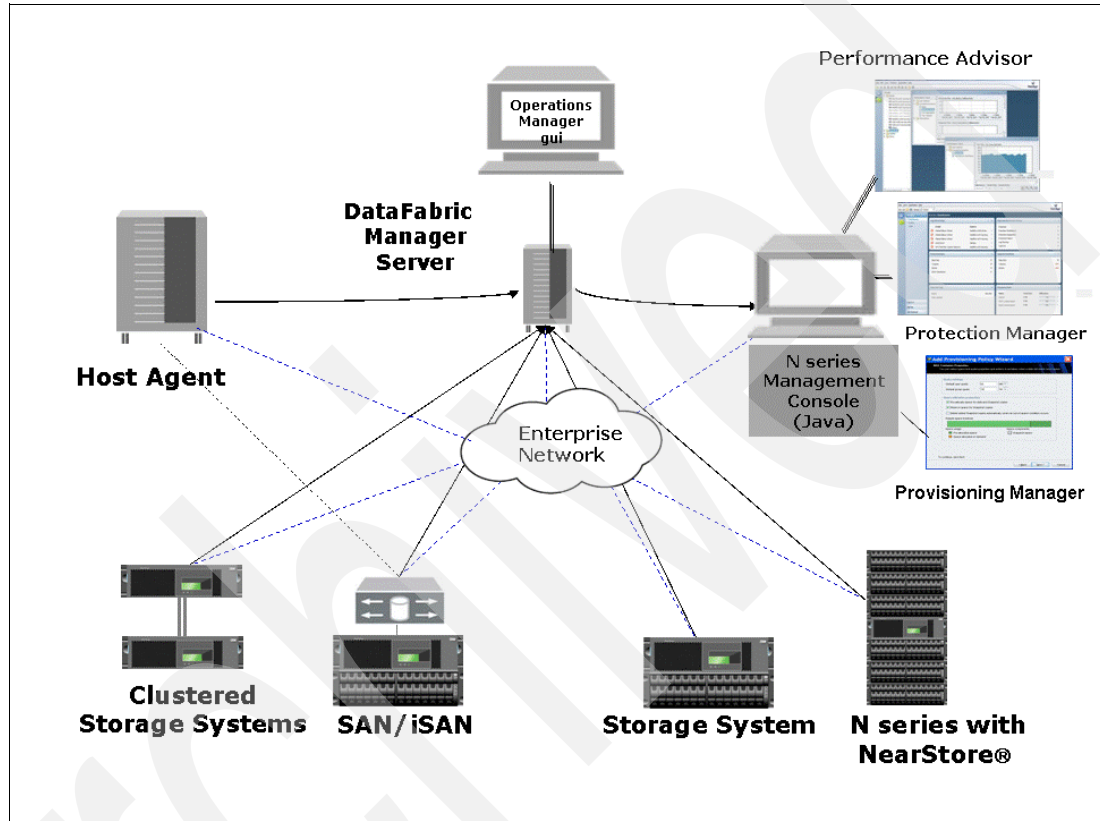


Figure 10-1 Architecture

These components interact in the following ways:

- ▶ The performance-monitoring server collects and stores performance data from one or more storage systems or vFile units in the network. The monitoring server is installed and enabled by default as part of the DataFabric Manager Server installation.
- ▶ Performance Advisor reads the performance data stored on the monitoring server and generates performance-related views and charts. Performance Advisor is installed on a Windows or Linux system separate from the system on which the DataFabric Manager Server is installed.

Note: Performance Advisor is available only at the host on which N series Management Console is installed. It cannot be viewed remotely through a browser.

- ▶ One or more storage systems or vFile units monitored by the DataFabric Manager Server.

To minimize CPU utilization, Performance Advisor does not collect real-time data for any view that includes a bar graph. Although performance monitoring server supports simultaneous connections to multiple Performance Advisor instances, connection to many instances of Performance Advisor simultaneously can result in slower collection of data.

Performance Advisor takes its administrator infrastructure from Operations Manager. You use the administrator roles in Performance Advisor to control access to performance data in Operations Manager.

The administrator roles determine the availability of menu options in the application. Dimmed options in the application interface indicate that the options are not available to you for the selected storage system or vFiler unit.

Administrator roles in Performance Advisor

Performance Advisor takes its administrator infrastructure from Operations Manager. You use the administrator roles in Performance Advisor to control access to performance data in Operations Manager.

The administrator roles determine the availability of menu options in the application. Dimmed options in the application interface indicate that the options are not available to you for the selected storage system or vFiler unit.

The following roles help you to work with performance views in the interface:

- ▶ The GlobalRead role allows you to view all existing performance views.
- ▶ The GlobalWrite and GlobalDelete roles allow you to create, edit, and delete custom views within any group.
- ▶ The GlobalPerfManagement role allows you to manage views, event thresholds, and alarms, apart from viewing performance information in the application. You can use this role to work with custom views.

Note: You can execute all performance monitoring tasks if you have, at minimum, the GlobalRead, GlobalWrite, and GlobalPerfManagement roles.

In the next few windows, you will see what changing a couple of roles will do. Currently, we have an Administrator named, “Roman”. From the Administrators window in Operations Manager, you can see that “Roman” has no rights (Figure 10-2) When Roman logs into the N series Management Console, he has no access, as shown in Figure 10-3.

Administrator Name	Roles	Email	Pager	View	Edit	Delete
Everyone				view	edit	
X3655A\Administrator	GlobalFullCo...			view	edit	<input type="checkbox"/>
X3655A\Roman				view	edit	<input type="checkbox"/>

Delete Selected

Figure 10-2 Admin user has no roles

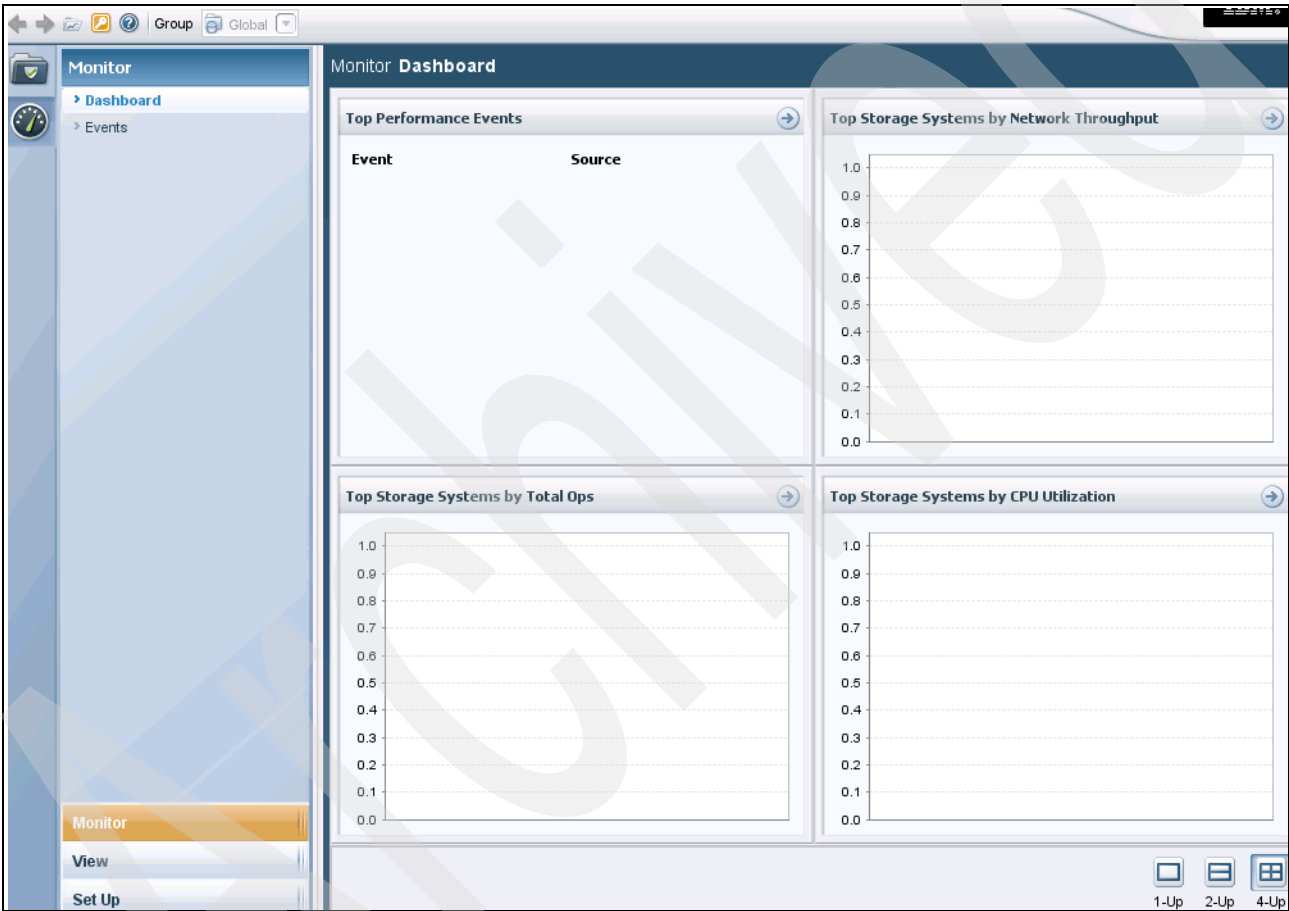


Figure 10-3 User logged in with no rights

Once we add a privilege to the “Roman” administrator, he can now view the Performance Advisor windows, as shown in Figure 10-4.

Administrator Name	Roles	Email	Pager	View	Edit	Delete
Everyone				view	edit	
X3655A\Administrator	GlobalFullCo...			view	edit	<input type="checkbox"/>
X3655A\Roman	GlobalRead			view	edit	<input type="checkbox"/>

Delete Selected

Figure 10-4 User given global read access

The user “Roman” is given access to monitor Performance Advisor, as shown in Figure 10-5.

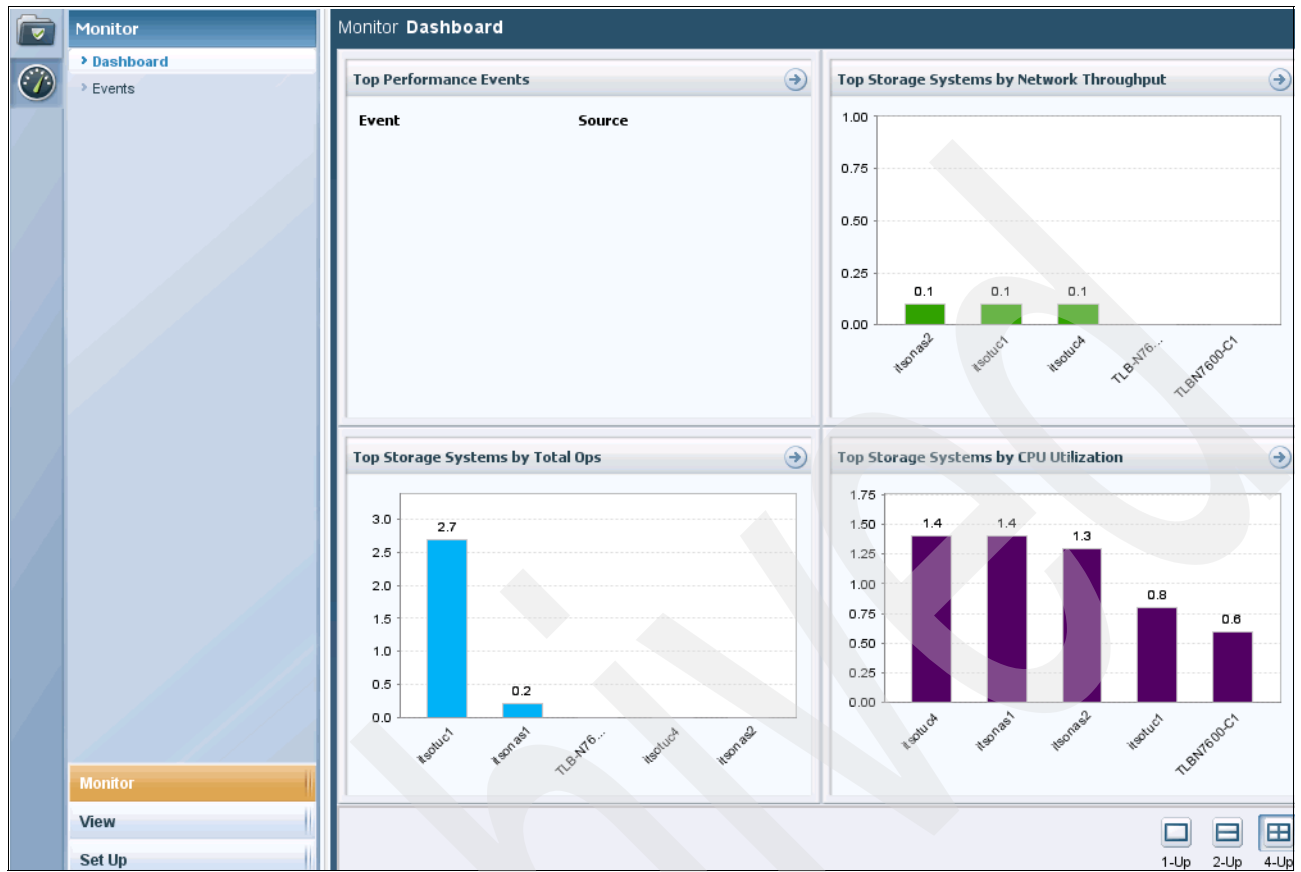


Figure 10-5 User now has read access to Performance Advisor

When we try to create a custom view (which we cover later in “Custom views” on page 228), most of the process can be done, but when we try to assign this view to a system or systems, we do not have access to any storage devices, as shown in Figure 10-6. Only if we are given the other two global rights, such as GlobalWrite or GlobalPerfManagement, will we be able to complete the process.

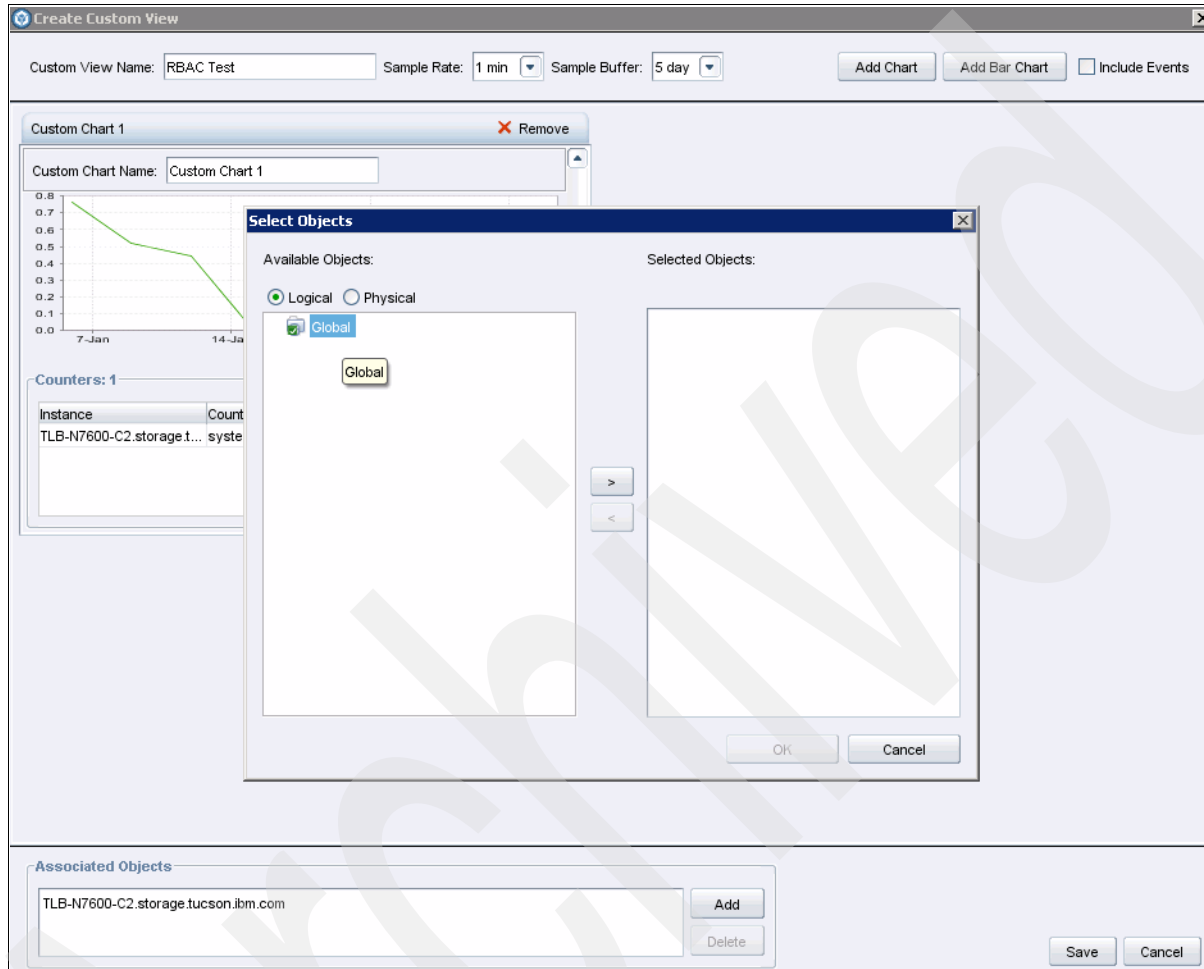


Figure 10-6 Cannot create custom view because of RBAC

10.2 N series Management Console

N series Management Console is the client platform for IBM management software applications. N series Management Console runs on a Windows or Linux system separate from the server on which Operations Manager is installed. See Chapter 7, “N series Management Console installation on Linux” on page 99 and Chapter 8, “N series Management Console installation for Windows” on page 107 for more detailed information.

N series Management Console allows storage, application, and server administrators to perform daily tasks without having to switch between separate user interfaces.

In this chapter, our emphasis is focused on the Performance Advisor portion of the Management Console. Figure 10-7 on page 205 shows the Dashboard when the Management Console is focused on Performance Advisor. The speedometer icon in the

upper left corner of the window is used to focus the console on the Performance Advisor windows.

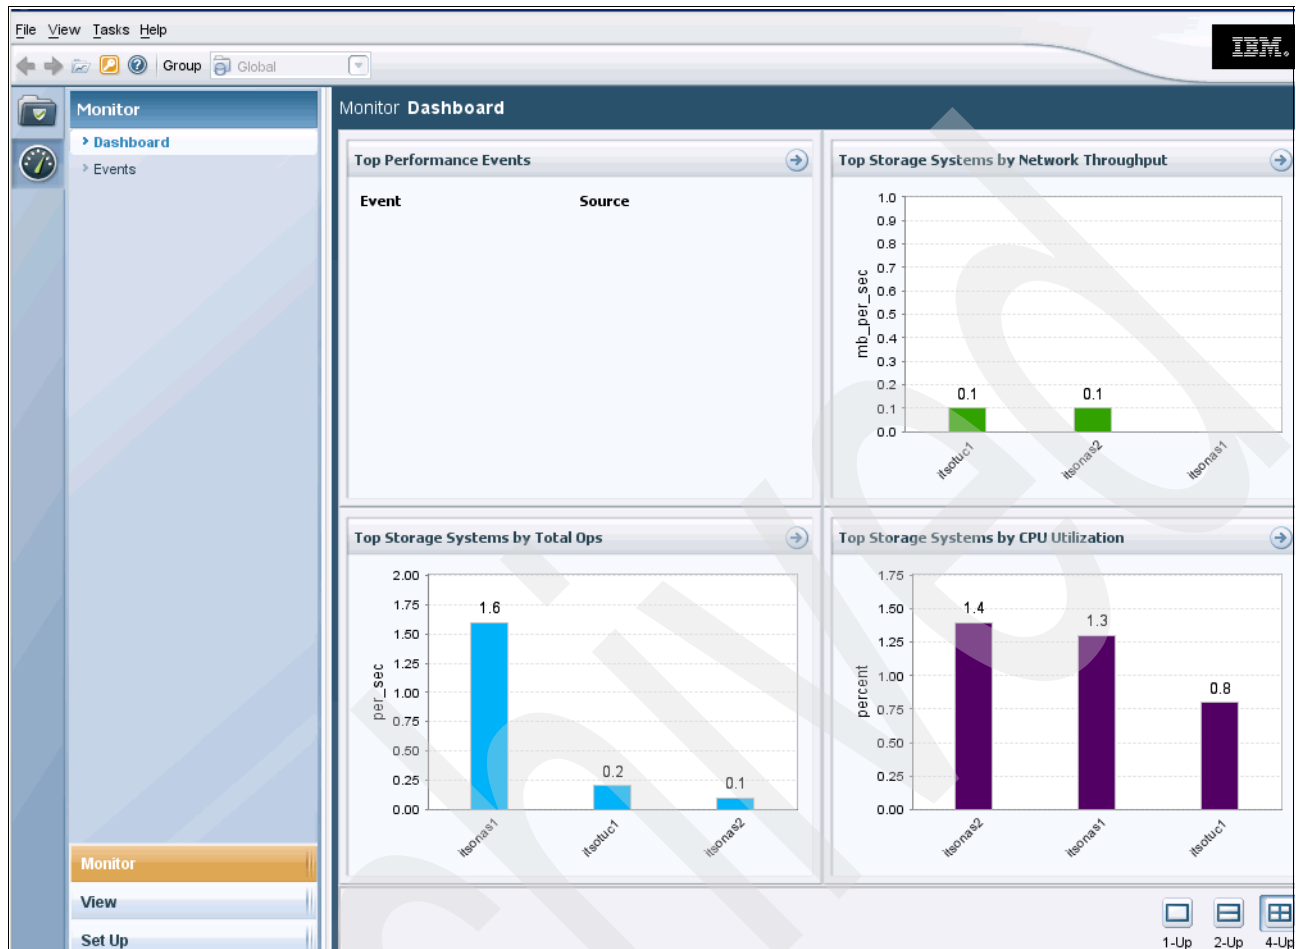


Figure 10-7 Performance Advisor Dashboard

Dashboard

Refer to Figure 10-7 as we describe each pane.

► Top Performance Events

This Dashboard pane lists a table of the most recent five highest severity events, ordered by time. The first column in the table gives a brief description of the event, while the second column describes the source of the event.

When you click >>, the Monitor Events window is displayed. Clicking the source of the event displays the Summary View window of the object in the navigation pane.

► Top Storage Systems by Network Throughput

This Dashboard pane displays a bar chart of the top five storage systems in Operations Manager, ordered by the highest network throughput. The vertical axis displays the megabytes of throughput per second. The number above each bar chart displays the exact value of the throughput per second from that system. The horizontal axis displays the name of the storage system.

When you click the bar chart, the Summary View window of the object is displayed in the navigation pane, as shown in Figure 10-8. The view expands to show more detail, as shown in Figure 10-9.




Figure 10-8 Click bar for expanded view

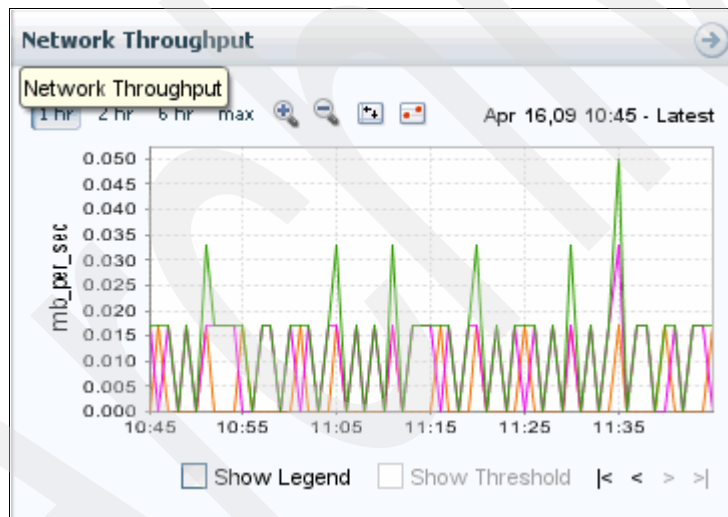


Figure 10-9 Expanded view

► Top Storage Systems by Total Ops (Figure 10-10)

This Dashboard pane displays a bar chart of the top five storage systems in Operations Manager, ordered by the highest total operations per second. The vertical axis displays the total operations per second for that storage system. The horizontal axis displays the storage system name.

When you click a bar chart, the Summary View window of the object is displayed in the navigation pane, as shown in Figure 10-11.




Figure 10-10 Click bar for expanded view




Figure 10-11 Expanded view of Ops

► Top Storage Systems by CPU Utilization (Figure 10-12)

This Dashboard pane displays a bar chart of the top five Operations Manager storage systems, sorted by the highest average CPU utilization. The number above each bar chart displays the exact value of the throughput per second from that system. The vertical axis displays the percentage CPU usage of the storage system. The horizontal axis displays the name of the storage system.

When you click the bar chart, the Summary View window of the object is displayed in the navigation pane, as shown in Figure 10-13.




Figure 10-12 Click bar for expanded view




Figure 10-13 Expanded view of CPU utilization

Although these panes are the default pane you see when bringing Performance Advisor up for the first time, you have the option of selecting a different default view. Refer to “Changing the default view” on page 220 for more information.

Performance Advisor

Performance Advisor provides a single location from which you can view comprehensive storage system and MultiStore® vFiler unit performance information and perform short-trend analysis. Performance Advisor also helps you identify bottlenecks and potential bottlenecks in the data infrastructure.

Performance Advisor is automatically enabled with the Operations Manager Core license. For more information, see the *Performance Advisor Administration Guide for Use with DataFabric Manager Server*, GC26-7897.

Objects and counters

Performance Advisor uses a variety of objects and counters to provide performance data collected from the various storage systems in your network.

Performance objects

Performance objects are collections of counters for specific subsystem types. The names of the objects correspond to the subsystem names that are displayed in the navigation pane of Performance Advisor (Figure 10-14 on page 210 and Figure 10-15 on page 211).

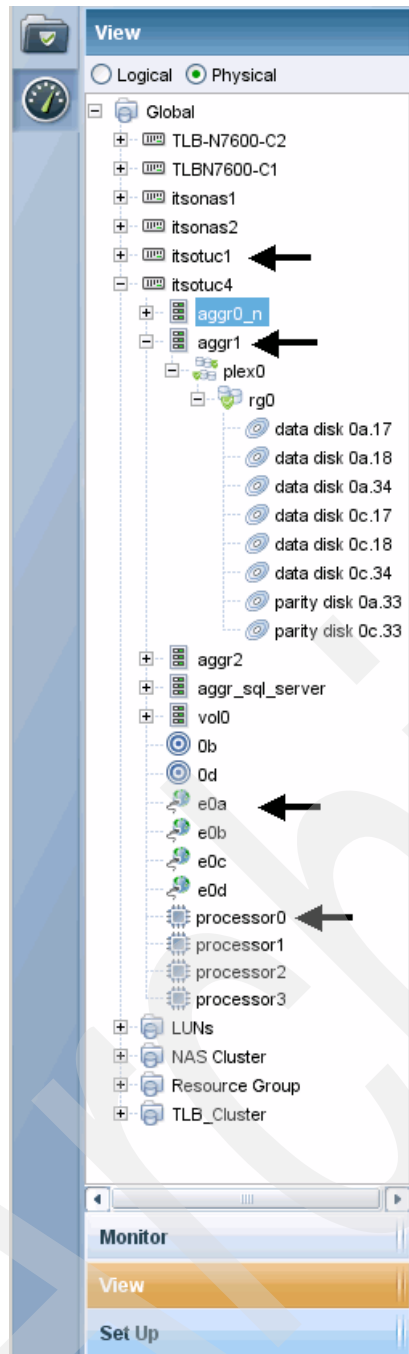


Figure 10-14 Objects - Physical View

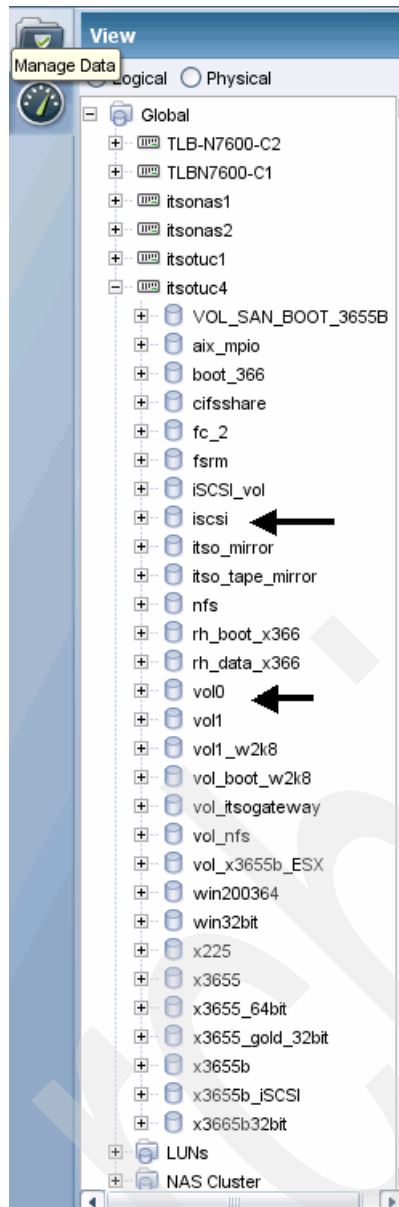


Figure 10-15 Objects - Logical View

In Figure 10-14 on page 210 and Figure 10-15, you can see that a performance object can be any number of things. Aggregates, LUNs, processors, adapters, volumes, disks, and so on, can be performance objects. Refer to the *Performance Advisor Administration Guide for Use with DataFabric Manager Server*, GC26-7897 for a complete list of performance objects.

Performance counter

A performance counter is a statistical measurement of activity on a storage system or storage subsystem.

The Operations Manager performance monitoring server collects the data from the monitored storage systems or vFiler units. Performance Advisor uses this data to generate a performance chart or part of a performance chart.

For example, the application might use the volume's total_ops counter of a storage system to generate a chart that displays total operations per second on that volume on a minute-to-minute basis, as shown in Figure 10-16.

The server collects the data for a counter from the time the counter is added to a view. Once the view is deleted, data collection for the counter stops. When you add a counter to a view, data collection begins with the counter in that view.

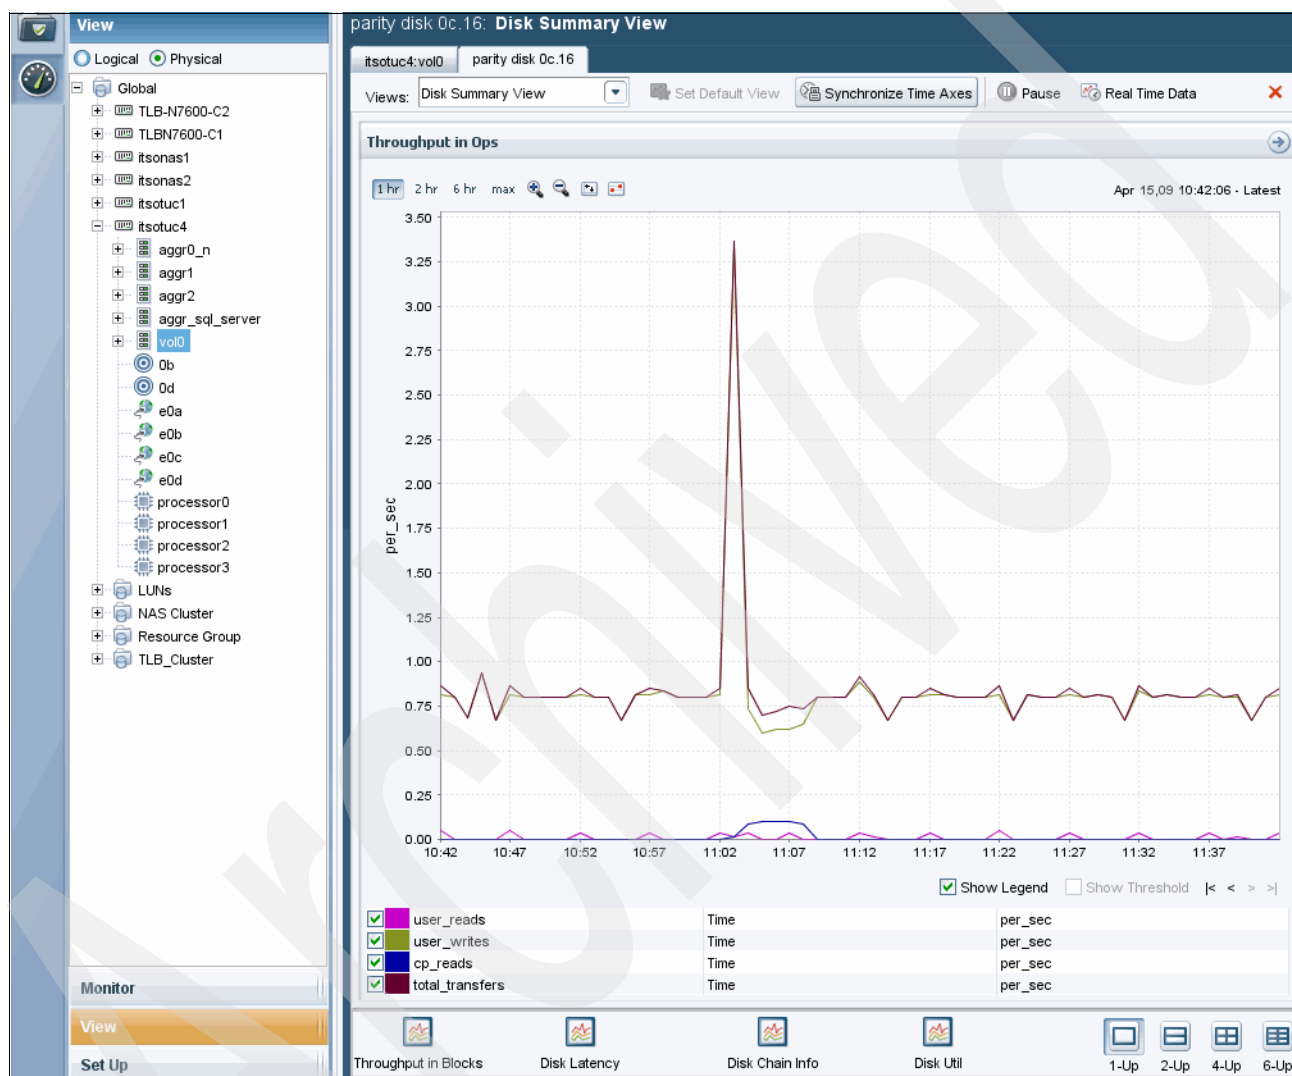


Figure 10-16 Volume total operations per second example

In Figure 10-16 on page 212, if you want to get a minute by minute view, you can highlight the spike with your cursor and the view will expand to the desired view, as shown in Figure 10-17 and Figure 10-18 on page 214.




Figure 10-17 Drag cursor over spike

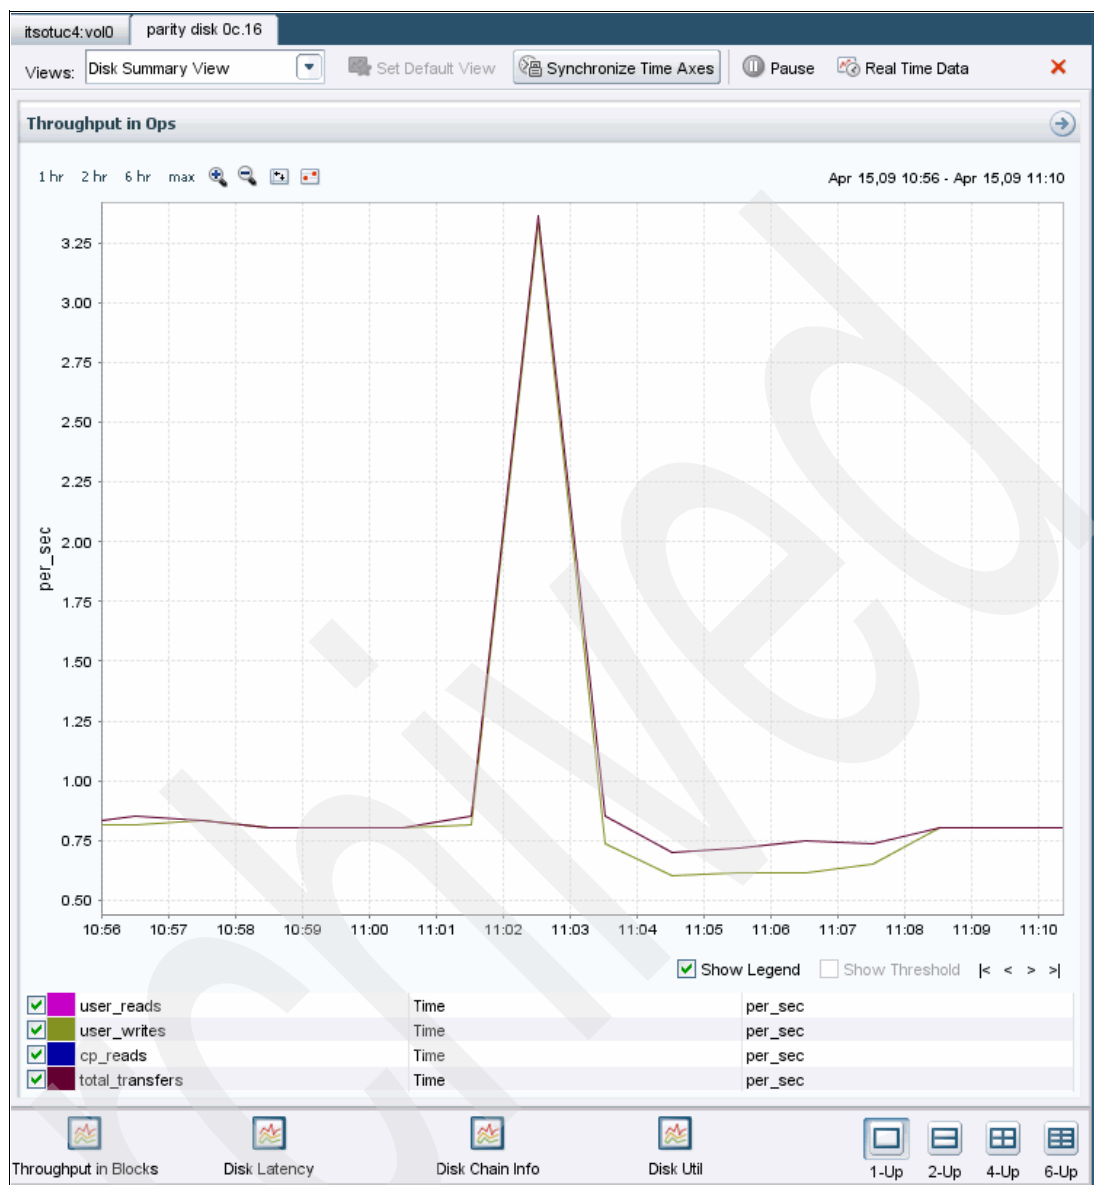


Figure 10-18 Graph expands to show minute by minute

There are counters for many objects in the N series storage systems, but if a particular function is not licensed, you will receive no data. For example, every storage system has an HTTP_ops counter. If HTTP is not licensed, the counter will contain zero data. The same goes for iSCSI and others, that is, no license, no counters. Refer to the *Performance Advisor Administration Guide for Use with DataFabric Manager Server*, GC26-7897 for a list of counters included in your storage system.

10.3 Setting up Performance Advisor

Before you use Performance Advisor, you must install N series Management Console and ensure that the prerequisites are met.

The prerequisites for Performance Advisor are as follows:

1. Install N series Management Console. Refer to Chapter 7, “N series Management Console installation on Linux” on page 99, which has detailed steps for installing the N series Management Console.
2. After you have installed the Management Console, enable the Performance Monitoring server by selecting, in the Control Center tab, **Setup** → **Options** → **Performance Advisor** (Figure 10-19).

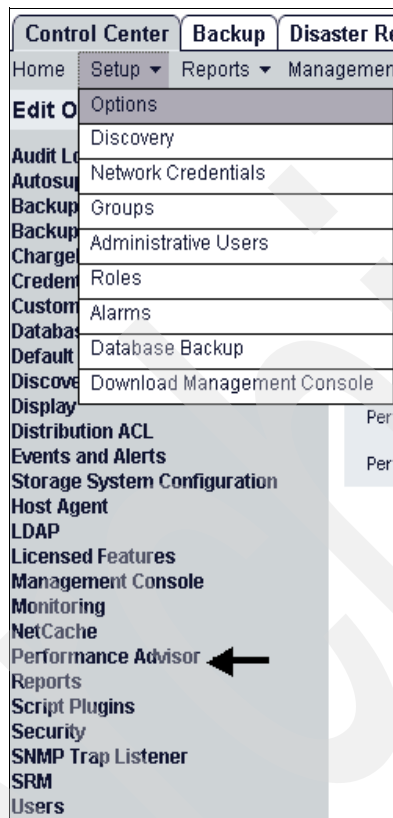


Figure 10-19 Select Performance Advisor from the Options menu

After you select Performance Advisor, the window shown in Figure 10-20 appears. The DataFabric Manager Server enables the Performance Monitoring server by default.

Options

Performance Advisor

Server HTTP Port: 8088

Server HTTP Enabled: Yes

Server HTTPS Port: 8488

Server HTTPS Enabled: Yes

Performance Advisor Enabled: Enabled

Performance Advisor Transport: httpOnly

Figure 10-20 Performance Advisor Enabled window

3. Install Data ONTAP V7.1 or later on all storage systems and the hosts of all vFiler units for which you want to monitor performance data.
4. Expose the storage systems and vFiler units for which you want to monitor performance data to Operations Manager.

By exposing, this refers to being sure that all systems to be monitored are either on the same IP network as the DataFabric Manager Server or that storage systems on other networks have been added to Operations Manager for monitoring. See Chapter 9, “Configuring Operations Manager” on page 127 for adding additional networks to monitor. Figure 10-21 shows the option window for adding additional networks.

Networks To Discover 07 Apr 10:40

Add a New Network

Network Address (e.g. 172.24.1.0, 172.24.1/24):

Network Mask:

Back Add

Network Address	Network Mask	Hop Count	Last Searched	Edit	Delete
9.11.218.0	255.255.255.0	0	07 Apr 10:32	edit	<input type="checkbox"/>

Delete Network Delete Network and Hosts

Figure 10-21 Making other networks available for monitoring

10.4 Working with the N series Management Console interface

There are two ways to start the N series Management Console interface:

- ▶ Select **Start** → **All Programs** → **IBM** → **N series Management Console**, as shown in Figure 10-22.



Figure 10-22 Starting from menu selections

- ▶ Alternatively, you can start the N series Management Console on Windows by clicking the N series Management Console icon on the desktop, as shown in Figure 10-23.



Figure 10-23 N series Management Console icon

To stop Performance Manager, select **File** → **Exit**, as shown in Figure 10-24.




Figure 10-24 Stopping Performance Advisor

Viewing performance data

You can view the performance data of an object by selecting the object in the navigation tree. From the navigation tree, select **View**. A tree of items (storage systems, LUNs, and groups that you have created within Operations Manager) will be displayed, as shown in Figure 10-25.

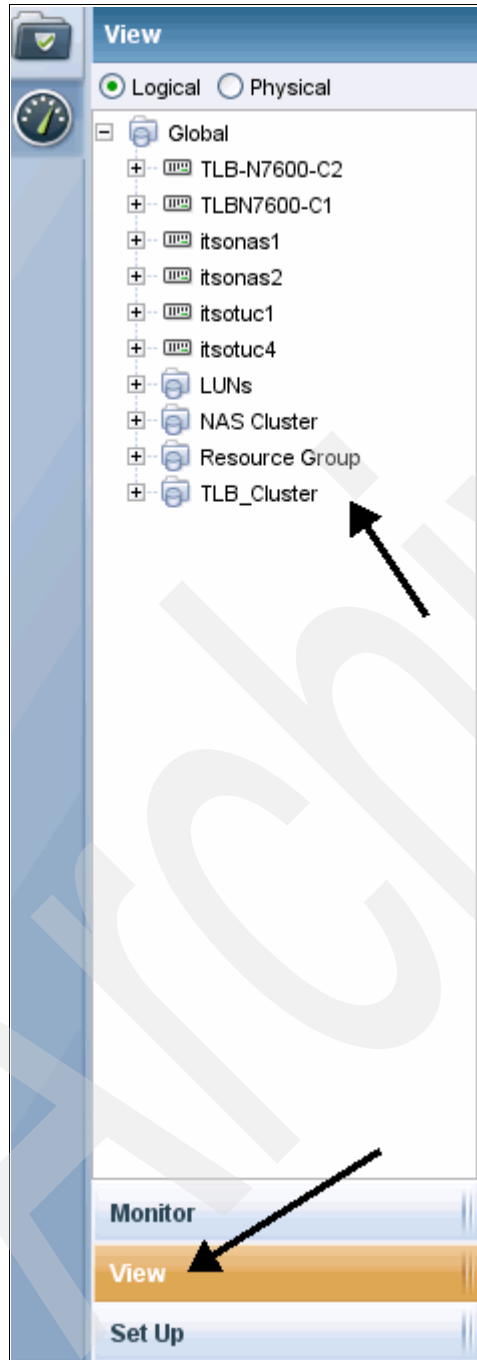


Figure 10-25 Viewing performance data

From this view, you can click the + next to any of the items to expand the view. By expanding the tree on the left and selecting an item from the branch, you can view the performance data in the right hand pane, as shown in Figure 10-26 on page 219. Remember, there may not be

data for every element. As mentioned before, there may not be any data to collect from that item, or the license may not be installed to support that element.

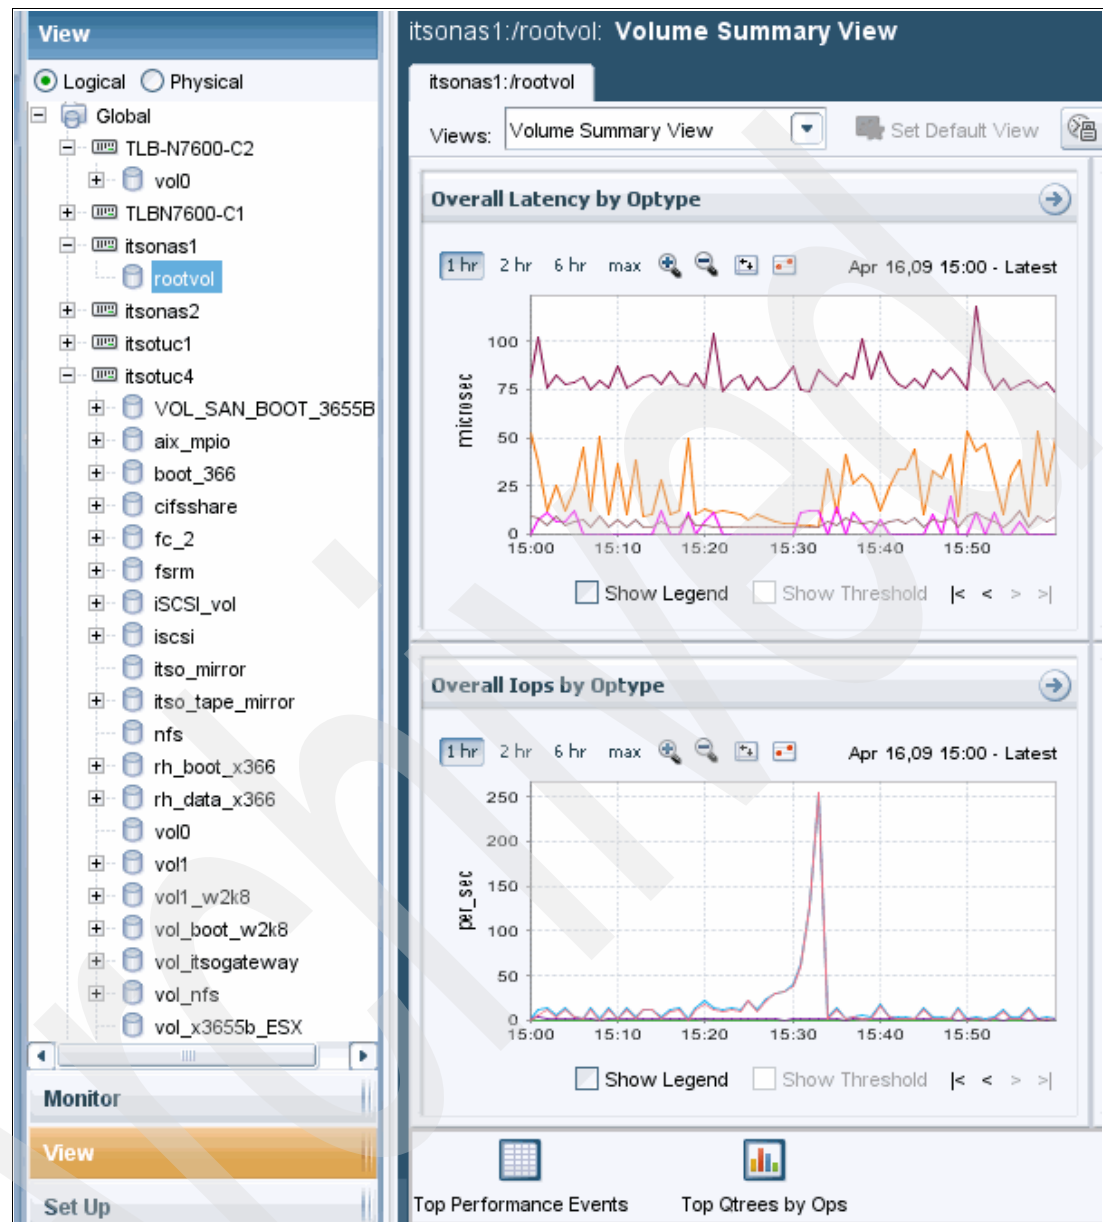


Figure 10-26 Selecting an item in the left pane displays performance data in the right pane

Changing the default view

The initial, default view is the Group Summary View of Performance Advisor. You can change the default view to another view. For example, you want to change the default view to Group-Top Objects instead. To do this task, you must do the following steps:

1. Select **Views** from the main dashboard and select the drop-down menu from the Views menu, as shown in Figure 10-27.

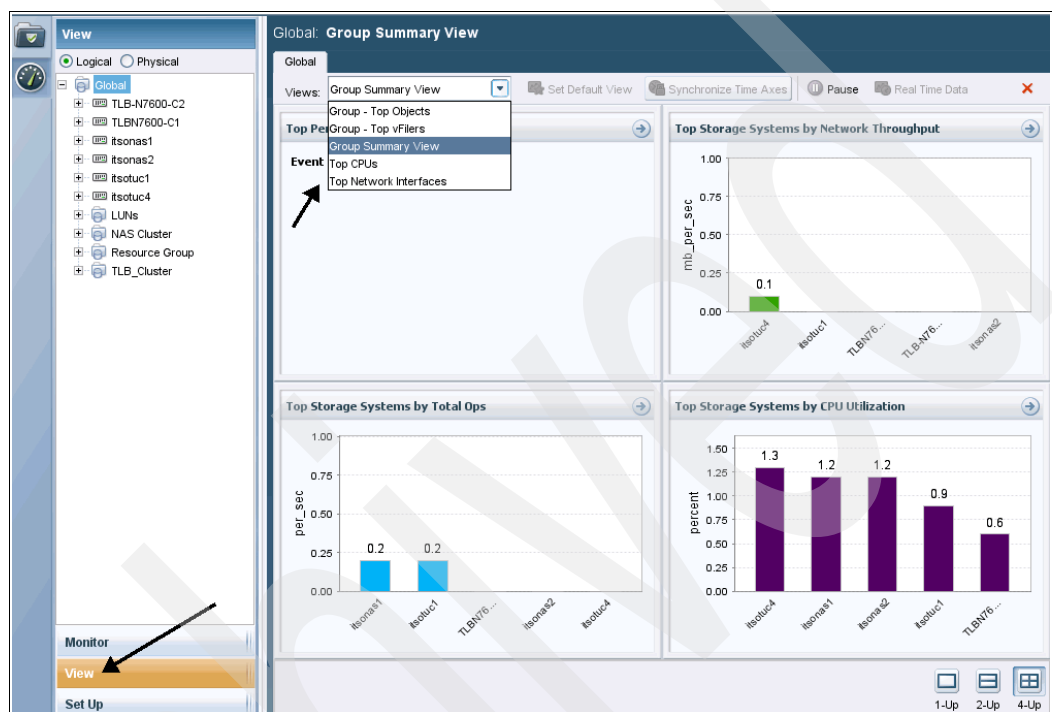


Figure 10-27 Select View and the view you want to see from the Views drop-down menu

2. When you select the **Top Objects** menu item, the window changes and the option to set this window as the default becomes active. When you select **Set Default View**, this view becomes the default, as shown in Figure 10-28 on page 221.

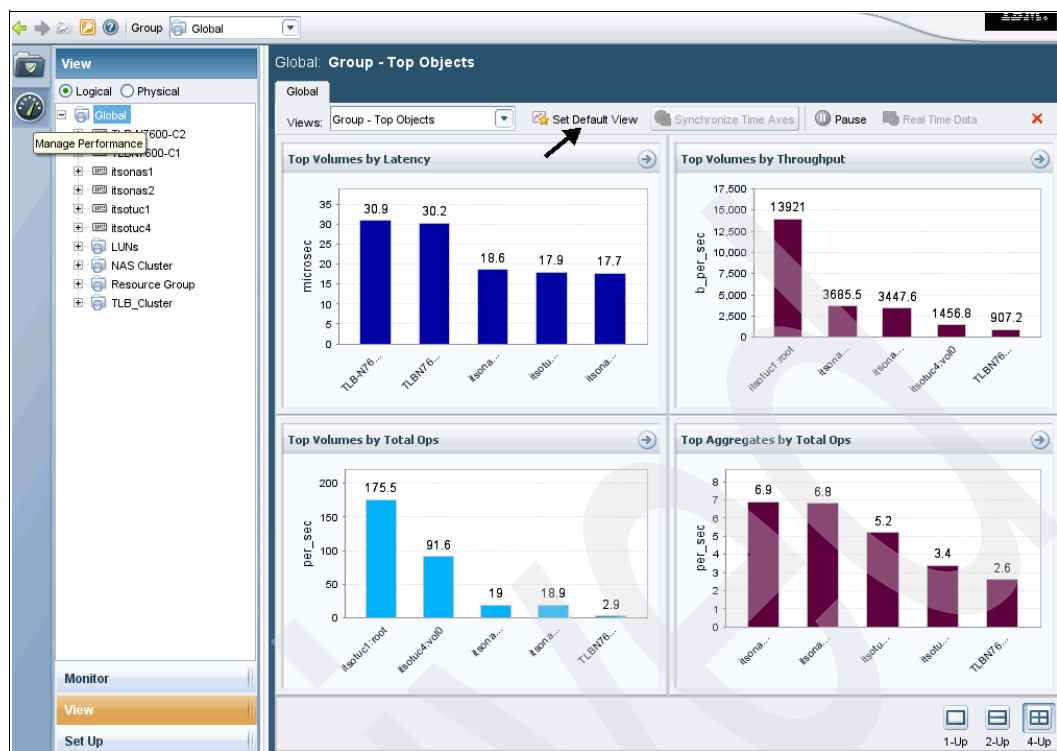


Figure 10-28 Select Set Default View window

The default view can also be selected for each logical and physical device. The default view for each device is the Summary view. However, each storage device shown in the left hand pane has its own group of views.

Any of these views can be set as the default so that when you select to view this storage device, your default view will be the first view displayed, as shown in Figure 10-29.

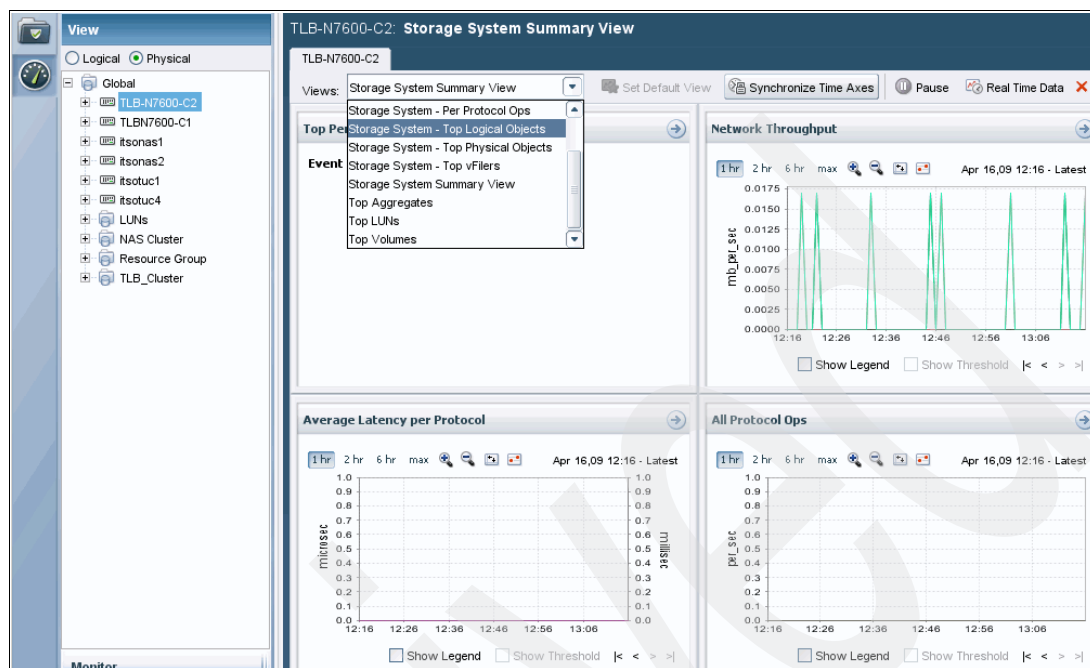


Figure 10-29 Selecting default views for storage systems

Refer to the *Performance Advisor Administration Guide for Use with DataFabric Manager Server*, GC26-7897 for a list of views that are available for each level of the tree.

Tabbed views

There may be a time when you want to compare data from different storage devices. The tabbed view will allow you to compare, almost side by side, similar information from separate storage devices. For example, let us say that we want to see performance comparisons for aggr0 on an N series N7600 cluster. With the tabbed view, we can open both devices in a tab and compare the data.

In this view, we have opened both aggr0 disks and can even synchronize the time to compare the information. (Notice the Synchronize Time Axis button under the tabs.) You can synchronize the axes when you want the time axes of all line charts in a view to be in sync, which helps in detecting performance blockage by interpreting data through views. This view shows 9-13.126L0 (LUN 0 for Cluster 1). The next tab shows 9.13.126L1 (LUN 0 for Cluster 2). By moving from one tab to the other, we can compare percent utilization for this particular time frame. Refer to Figure 10-30 for more details.

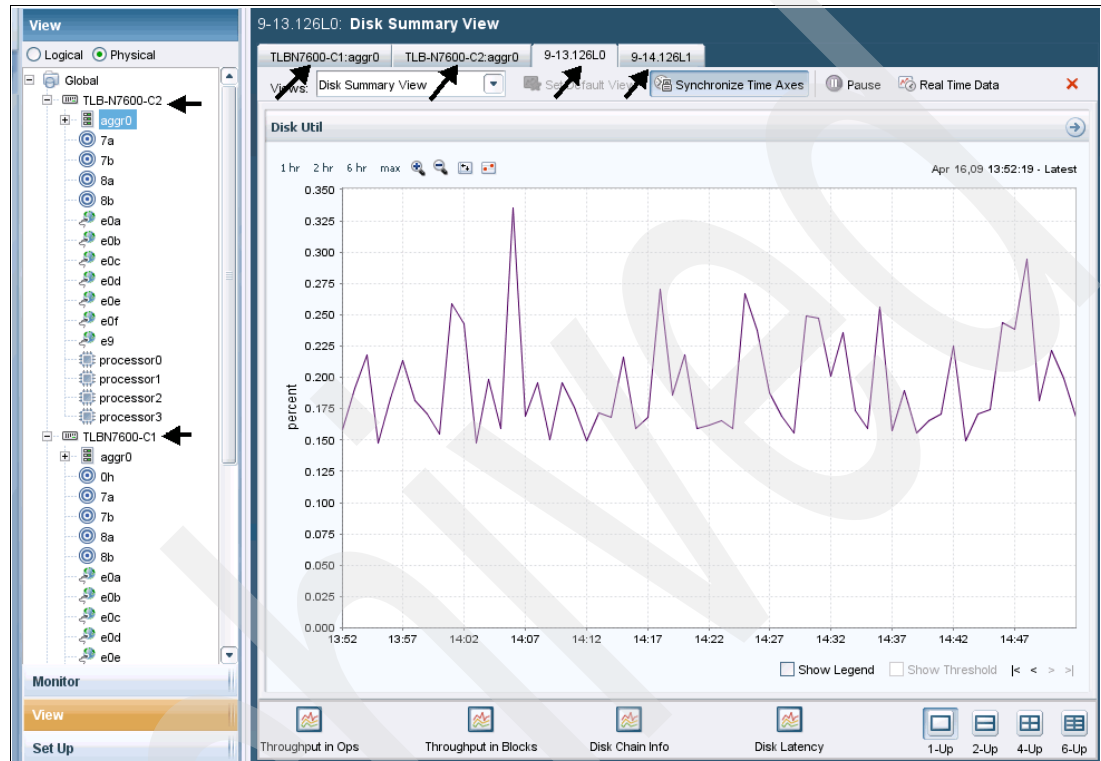


Figure 10-30 Tabbed view comparison

In the next view, shown in Figure 10-31, notice the icons at the bottom of the window. By running your mouse over these icons, your mouse turns into a plus. You can then grab these icons and drag them into the graph area. When this is done, the window changes to show the information that the icon indicated.

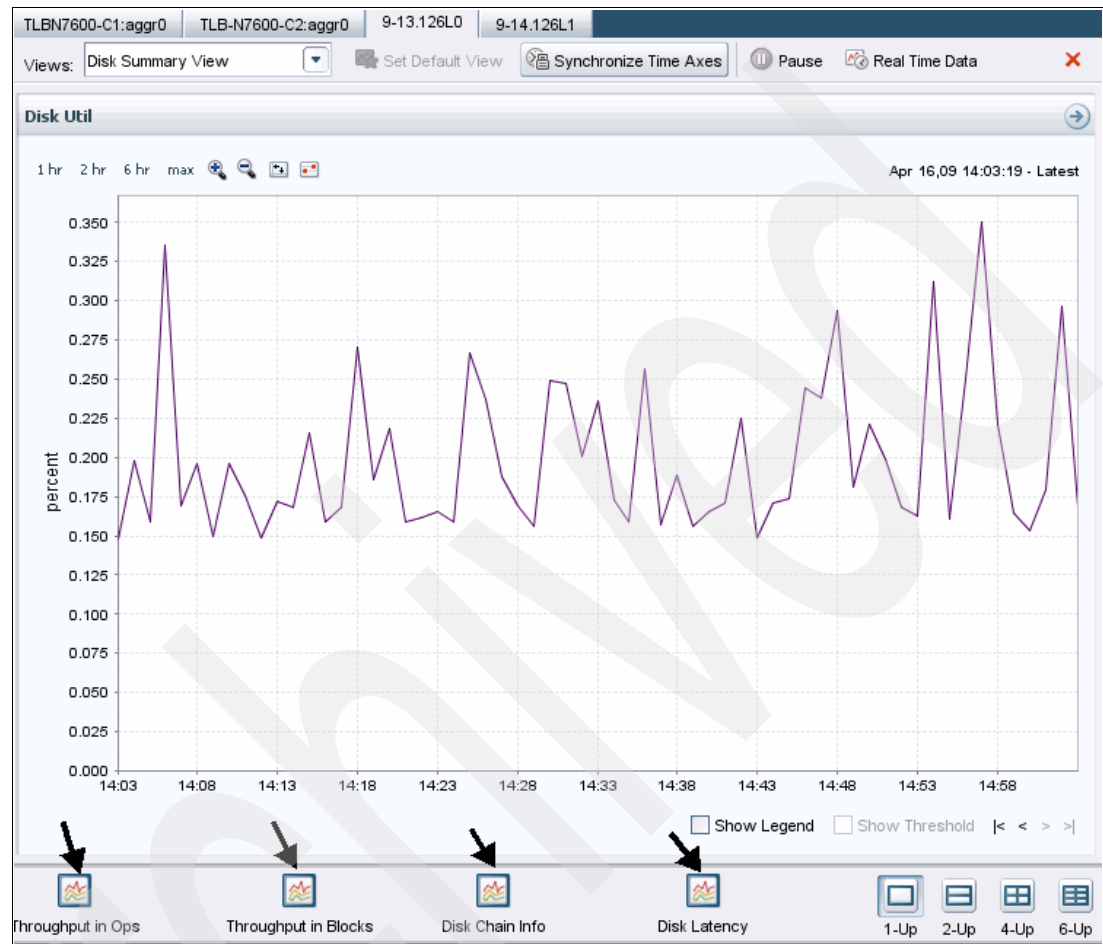


Figure 10-31 Select additional icons for other information

When you drag and drop the icon into the graph area, the new data appears. We select the “Throughput in Blocks” icon, as shown in Figure 10-32.




Figure 10-32 Drag and drop the icon

Another useful item is the legend. When you select the legend, the graph expands to show the various functions that make up this graph, as shown in Figure 10-33.

Notice also the marks in the bottom right corner of Figure 10-32 on page 225. These will allow you to move forward and backward in time on the timeline.

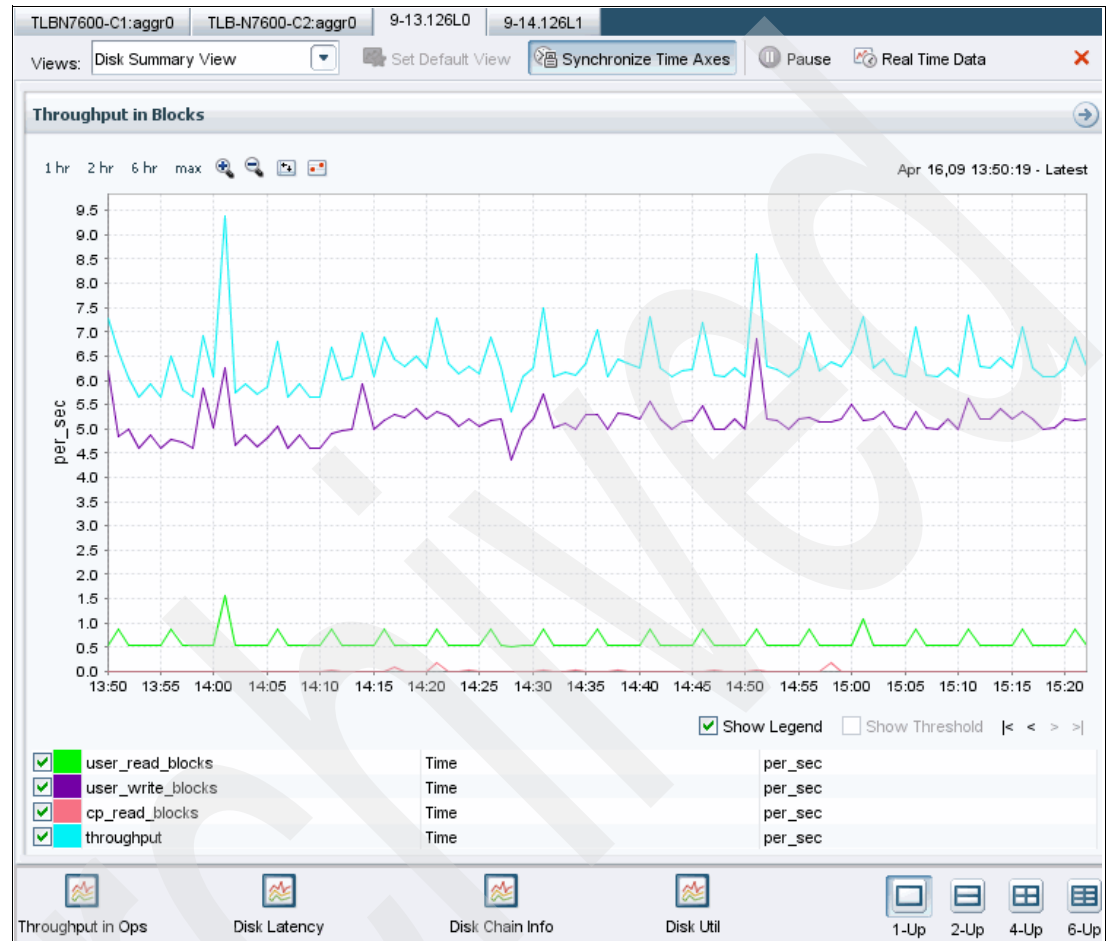


Figure 10-33 Graph with legend

Notice the icons in the lower right hand corner of the display. We can create a view where all of the items on this graph can be displayed at one time. By selecting the 6-up view, we can see all of the data at one time. However, it makes for a very complicated chart, as shown in Figure 10-34 on page 227.

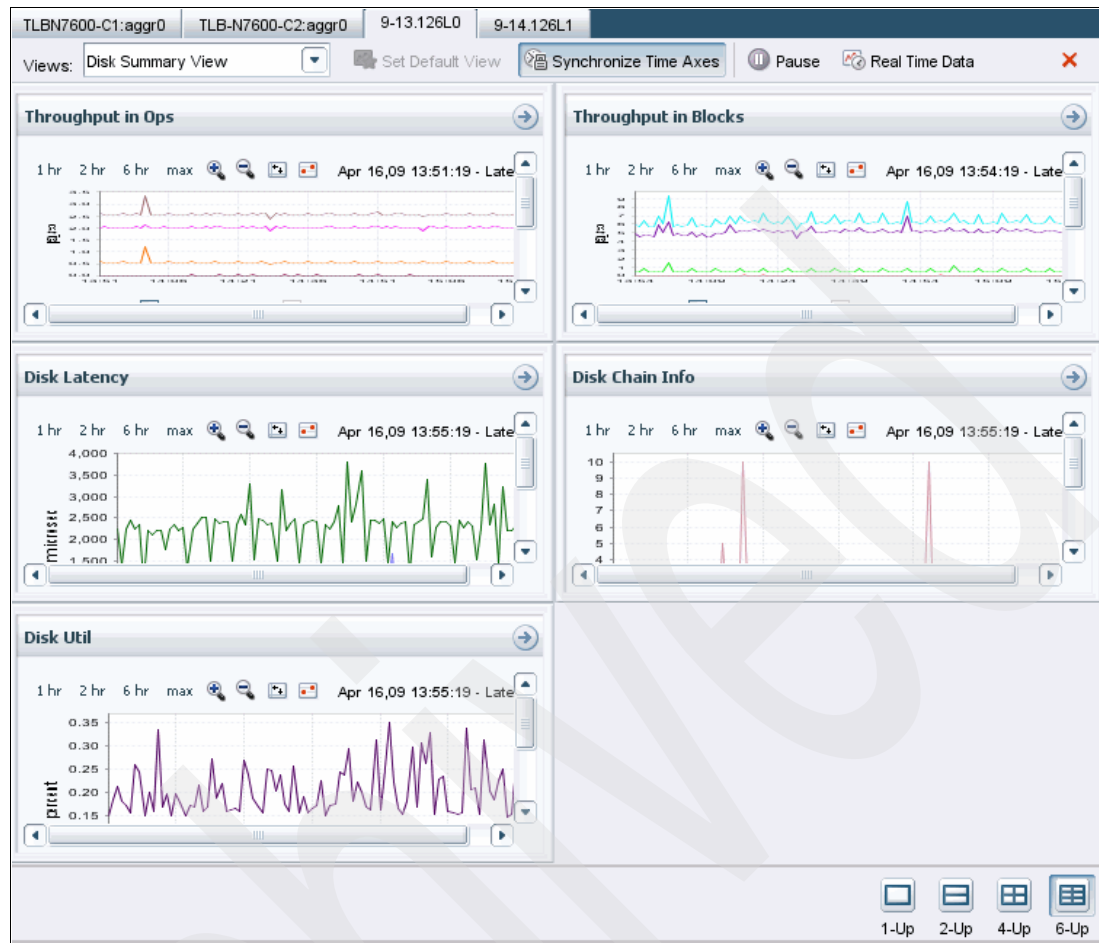


Figure 10-34 All graphs view

Although we have been discussing the tabbed view, these features are available in many of the other views as well. The functions discussed here are global in nature and can be found in most views used in Performance Advisor.

If you want to print the graph, place your cursor over the graph to be printed, right-click it, and select **Print**. (Be sure you have a printer setup first before trying this function.) Refer to Figure 10-35 for more details.

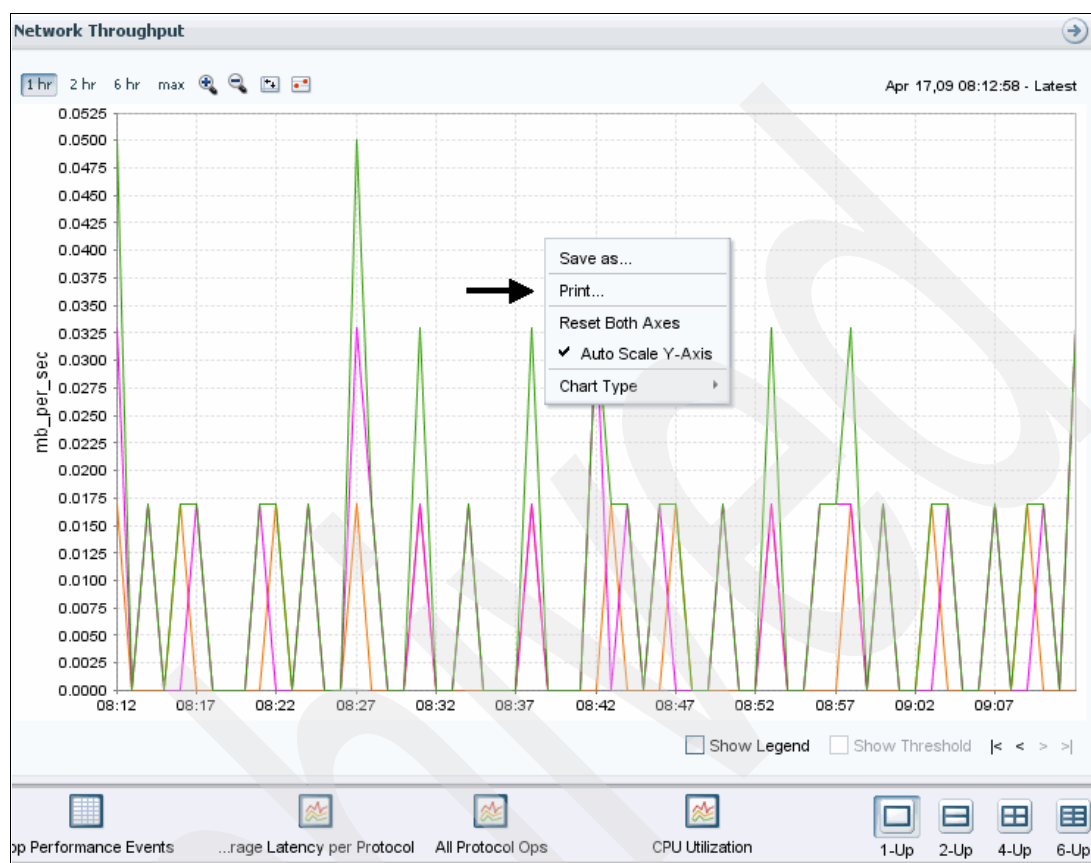


Figure 10-35 Printing the graph

Custom views

Up to this point, we have been discussing how Performance Advisor appears by default and some of the options you can perform to customize your displays. However, there may come a time when the defaults may not meet your needs or you want a certain view that Performance Advisor does not have. That is when you can turn to the customizing function.

Custom views expand the ways that the performance client can display data, enabling you to present data in combinations or styles that are not possible with the default views. For example, you might want to monitor a single counter across several storage systems or vFiler units, or you might find the counter combinations in the default views to be useful, but you would prefer a different type of chart.

You must create a custom view if you want to access real-time data, display counters in combinations that are not offered by default views, use chart types different than the default views, and use different views to compare the performance of the same counter on different objects.

If you use custom views, there are no restrictions on which counters you can combine in a single view. Additionally, you can include multiple counters in a single chart and multiple charts in a single view.

When creating custom views, it is important to remember to group like information together. Although you can combine any counters in a single chart, keep in mind that if you combine counters with different units, the chart multiplies vertical ("y") axis with the respective unit labels of the counters. For example, combining the disk_data_read counter (KB per second) with the processor_busy counter (percentage) results in a chart with multiple vertical ("y") axis, where each axis the respective unit labels of the counters.

Let us take a look at creating a custom view. Figure 10-36 shows our starting point for creating custom views. Do these steps:

1. Select **Setup** and **Custom Views** in the left hand pane of the Performance Advisor window.

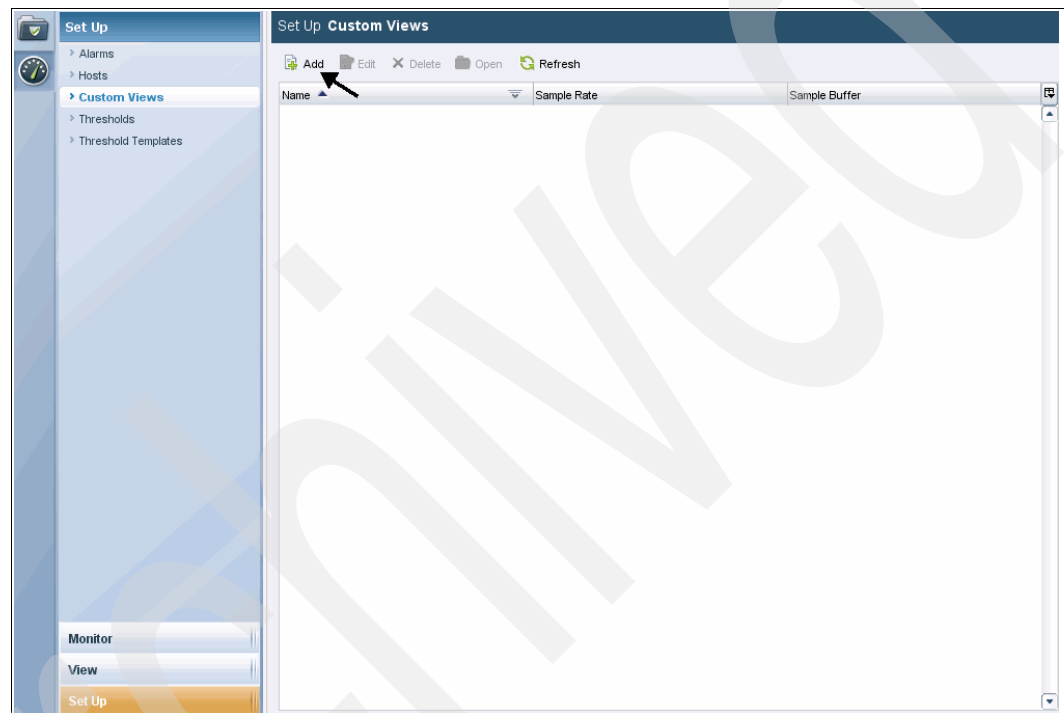


Figure 10-36 Starting point for custom views

2. Click **Add** and you will come to the customization window, as shown in Figure 10-37.

Create Custom View

Custom View Name: Sample Rate: 1 min Sample Buffer: 5 day ☐ Include Events

Custom Chart 1

Custom Chart Name:

Counters

Instance	Counter	Legend Color

Associated Objects

Figure 10-37 Custom chart creation window

3. You need to supply the name of the chart, sample rate, duration, and then click the **Add** button to supply the counters for the chart, as shown in Figure 10-38.

The screenshot shows the 'Create Custom View' dialog box. At the top, there are fields for 'Custom View Name' (ITSO Sample), 'Sample Rate' (1 min), and 'Sample Buffer' (5 day). There are buttons for 'Add Chart', 'Add Bar Chart', and a checkbox for 'Include Events'. Below this is a section for 'ITSO Sample_1 Network' with a 'Remove' button. Underneath is a 'Custom Chart Name' field containing 'ITSO Sample_1 Network'. Below that is a 'Counters' section with a table that has columns 'Instance', 'Counter', and 'Legend Color'. The table is empty, and there are 'Add' and 'Delete' buttons next to it. An arrow points to the 'Add' button. At the bottom is an 'Associated Objects' section with an empty field and 'Add' and 'Delete' buttons. Finally, there are 'Save' and 'Cancel' buttons at the bottom right.

Figure 10-38 Fill in the name of the chart

In our example, we look at the inbound, outbound, and total network load of our busiest storage system, itsotuc4. In Figure 10-39 you can see that there are many other options we could choose. Refer to the *Performance Advisor Administration Guide for Use with DataFabric Manager Server*, GC26-7897 for a complete list of counters provided by Operations Manager.

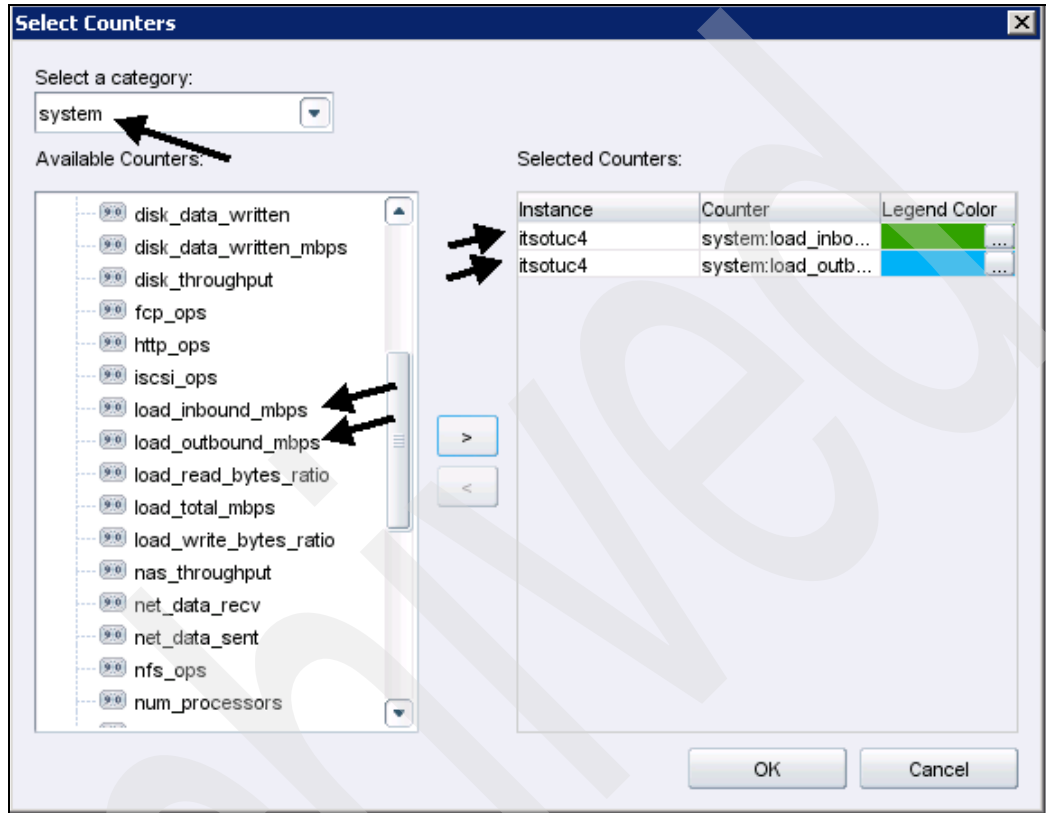


Figure 10-39 Select the counters you want to be a part of this chart and click OK

4. Once you have made your selections, click **OK** and you see an example of your chart. From here, you can add other counters or just click **Save** to keep the chart, as shown in Figure 10-40.

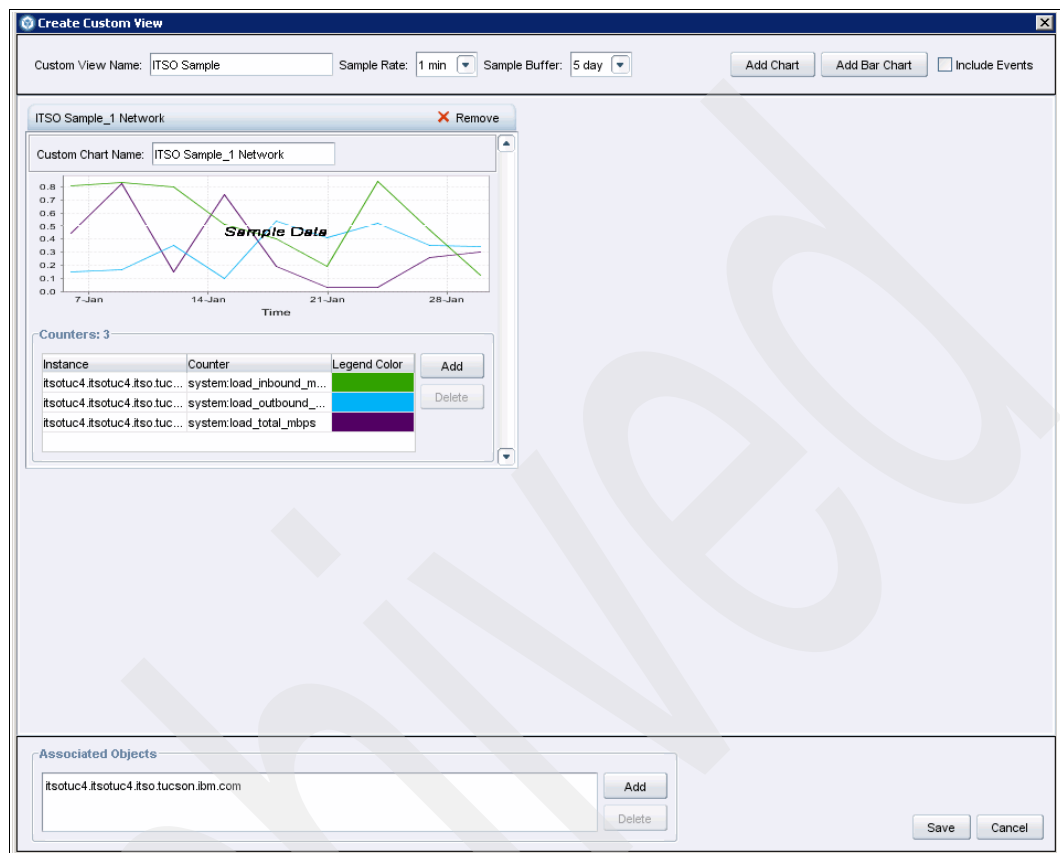


Figure 10-40 Your custom chart

- Once you click **Save**, the chart becomes a part of Performance Advisor and can be selected for viewing at any time, as shown in Figure 10-41.

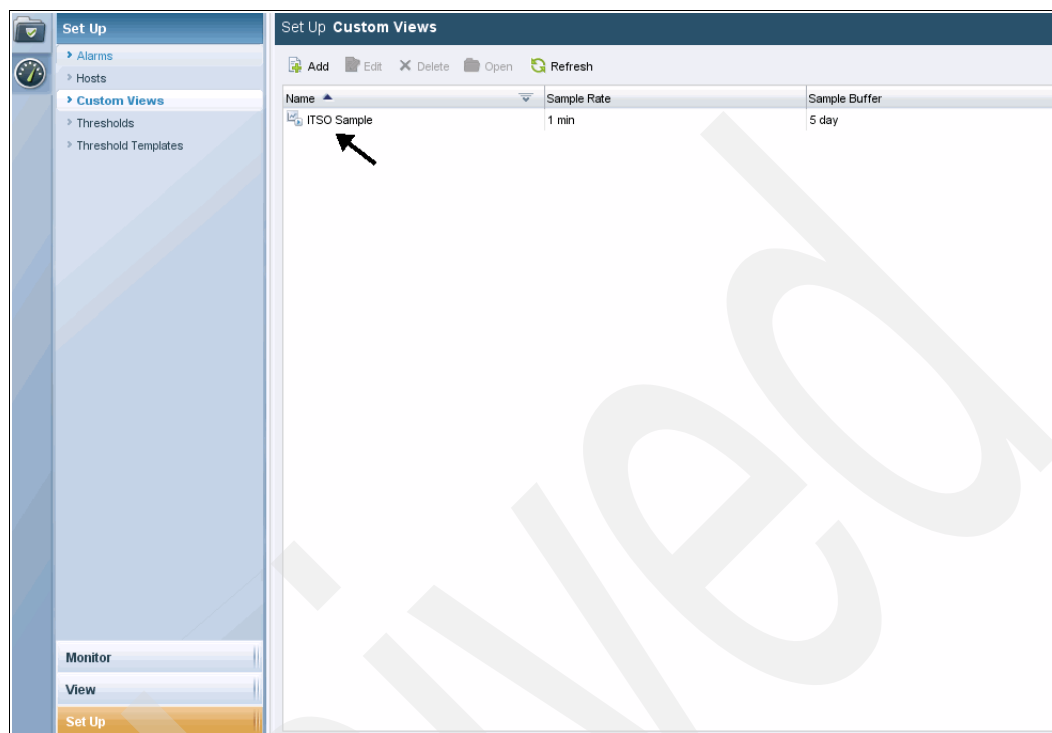


Figure 10-41 Custom charts can be selected for viewing at any time

- Once you have saved the chart, it can be selected from the View section of the Performance Advisor. This chart and any others that you create become a part of the menu selections for the devices you are monitoring. In this case, we created this chart to view network information for the storage system itsotuc4. In the selection menu for charts to view for this system, our ITSO Sample is available, as shown in Figure 10-42.

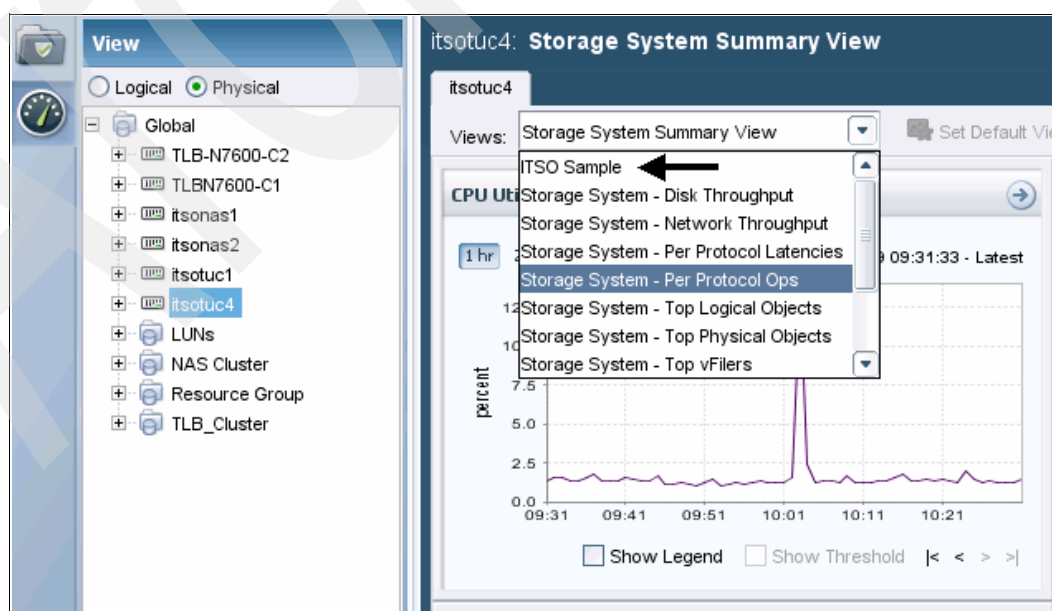


Figure 10-42 Our chart can be selected from the list of views for this storage system

Real Time Data view

For this example, we use the custom network chart we created in Figure 10-40 on page 233.

First, we need to look at some items on the chart. When we click the **Real Time Data** button, the title of the chart now shows “(Real Time)” next to the chart name, as shown in Figure 10-43.

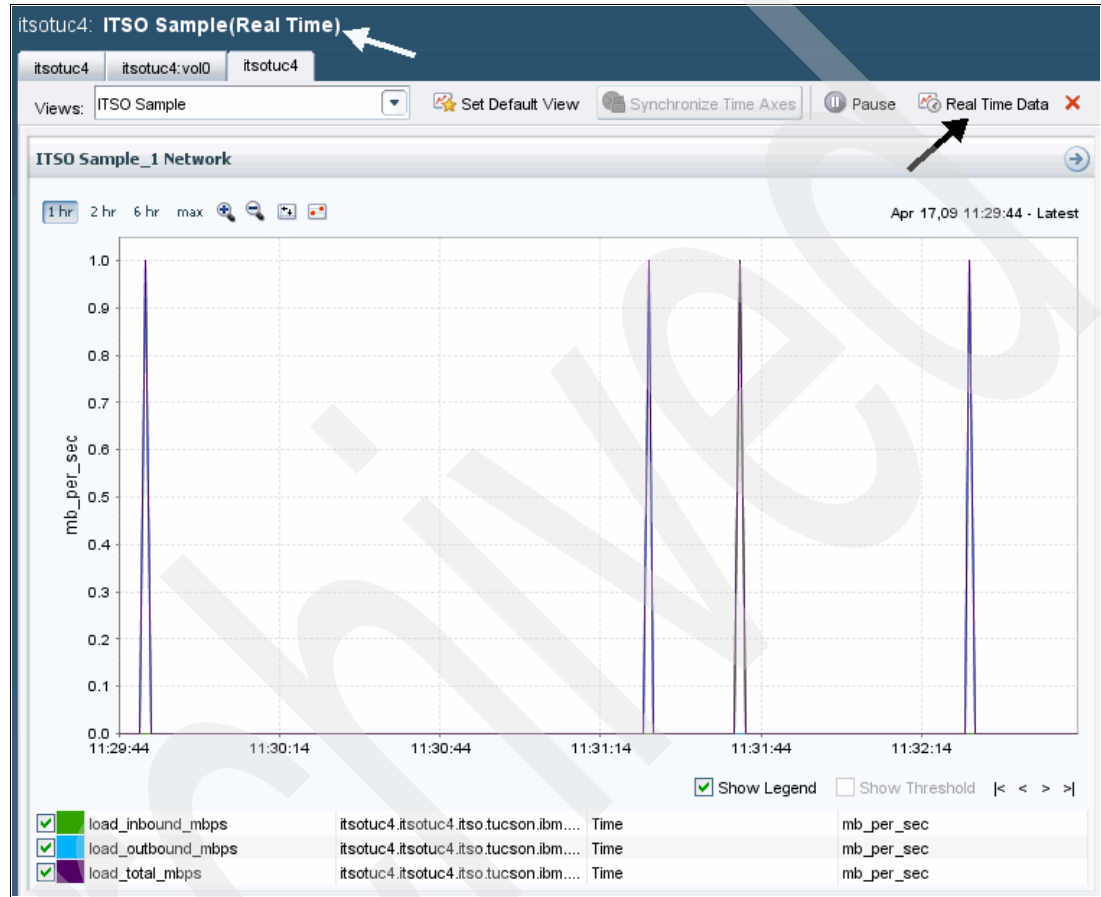


Figure 10-43 Real Time Data

When you click the **Real Time Data** button, you will get a pop-up window that will allow you to customize your sample rate and buffer time, as shown in Figure 10-44.

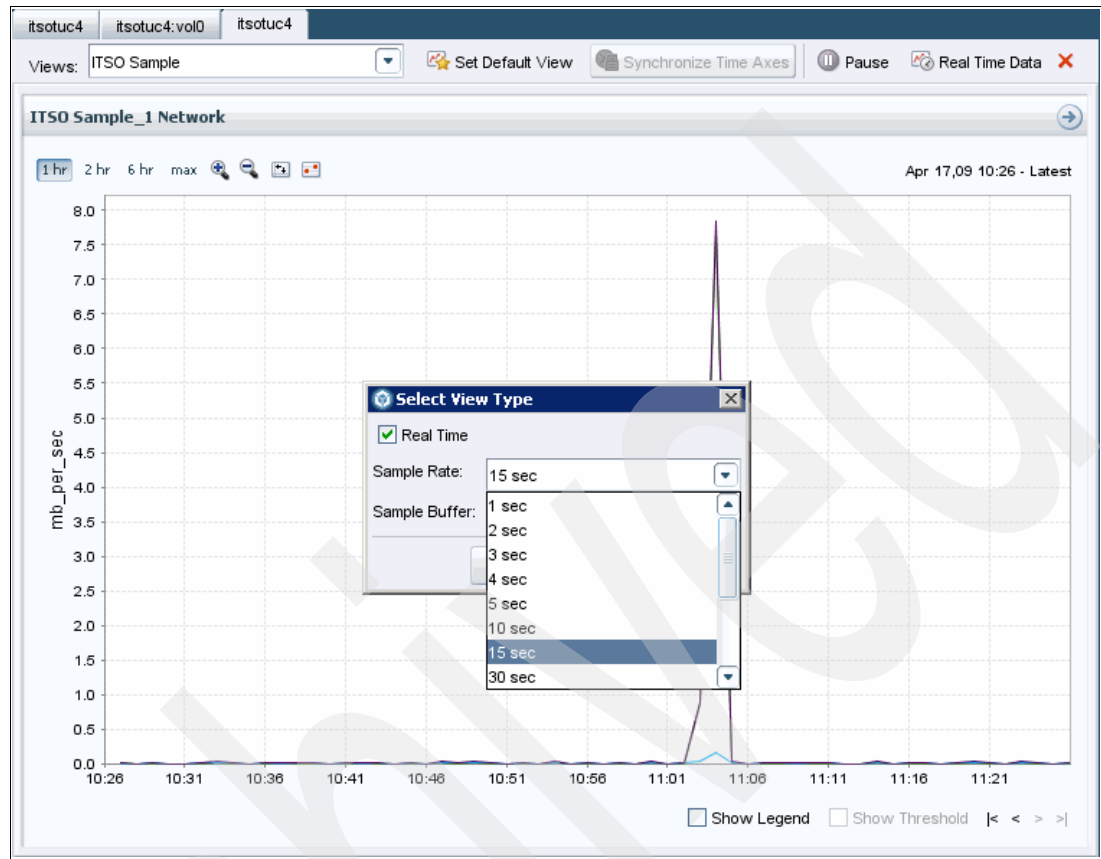


Figure 10-44 Customize Sample Rate

You will be able to set your sample rate as well as the time frame you want to take the samples, as shown in Figure 10-45 on page 237. Just realize the more frequent the sample, the more of a load your are putting on your network.

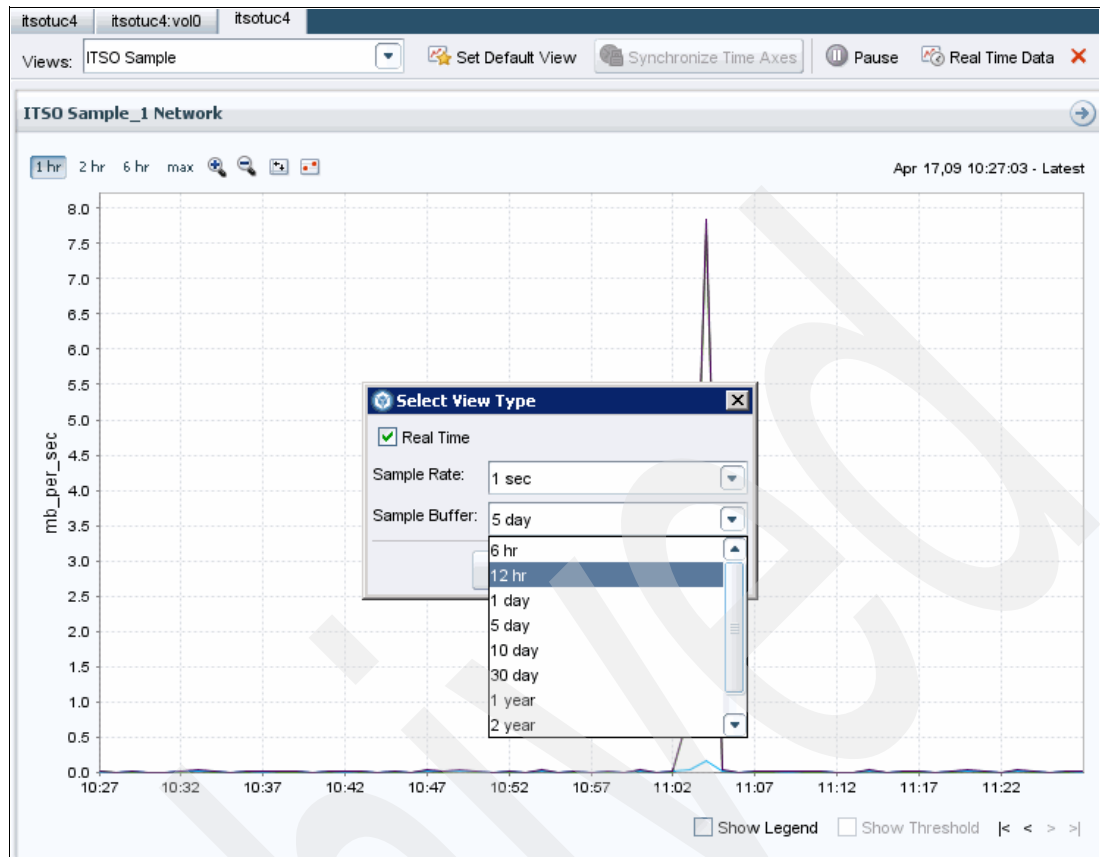


Figure 10-45 Select Sample Buffer

Performance Advisor will warn you about the possible performance impact for what you are about to do, as shown in Figure 10-46.




Figure 10-46 Warning of performance impact

Once you select **OK**, the collection of data will begin and continue for the time frame specified, as shown in Figure 10-47.

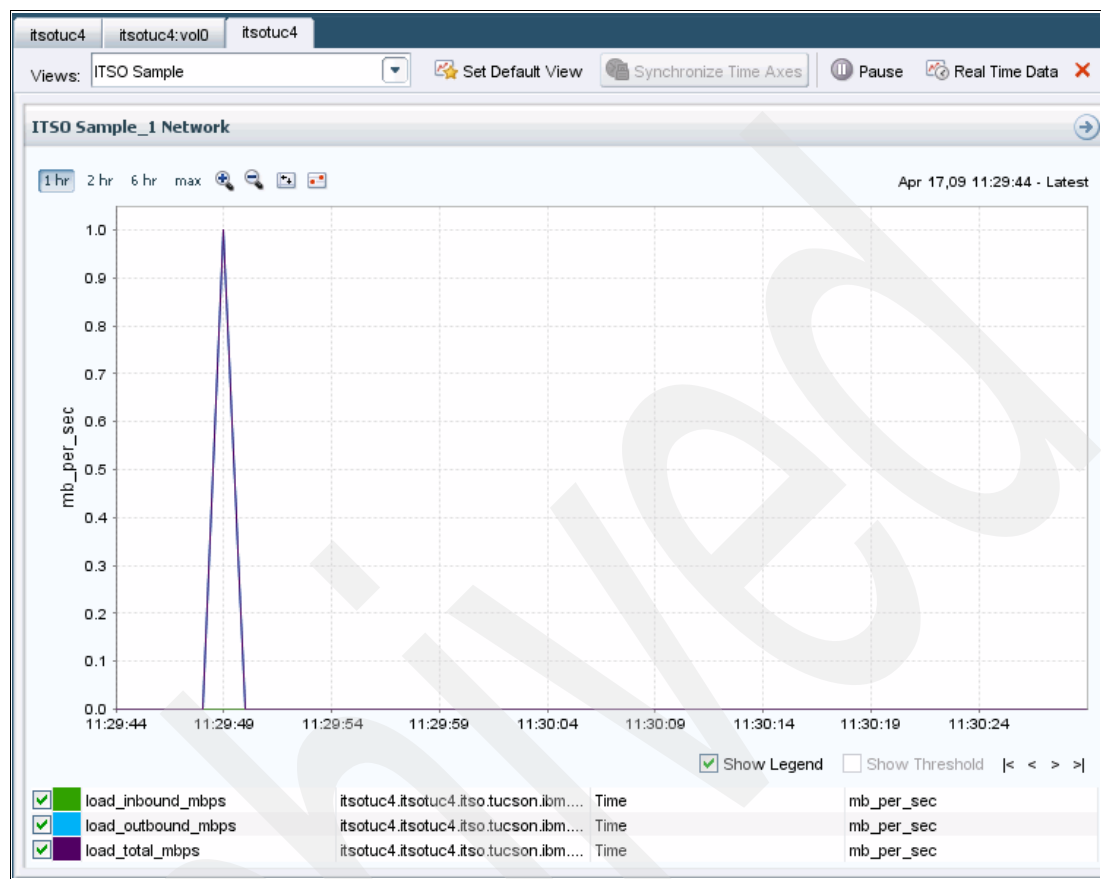


Figure 10-47 Real time collection begins

10.5 Events, alarms, and thresholds

In this section, we cover events from Operations Manager, alarms that can be configured, and thresholds that can be modified to enable you to monitor Operations Manager using the N series Management Console. We will only cover the Performance Advisor in this chapter. Refer to Chapter 13, “Protecting your data with Protection Manager” on page 313 to learn more about the N series Management Console and how it is used for those licensed features.

Events

Events are generated automatically when a predefined condition occurs (for example, Management Station License Nearly Expired) or when an object crosses a threshold (for example, Management Station Node Limit Reached). *Event messages* inform you when specific events occur. All events are assigned a severity type and are automatically logged in the Events window .

You can configure alarms to send notifications automatically when specific events or severity types occur, as shown in Figure 10-48 on page 239.

Severity type	Description
Normal	A previous abnormal condition for the event source returned to a normal state and the event source is operating within the desired thresholds. To view events with this severity type, you select the All option.
Information	The event is a normal occurrence—no action is required.
Warning	The event source experienced an occurrence that you should be aware of. Events of this severity do not cause service disruption and corrective action might not be required.
Error	The event source is still performing; however, corrective action is required to avoid service disruption.
Critical	A problem occurred that might lead to service disruption if corrective action is not taken immediately.
Emergency	The event source unexpectedly stopped performing and experienced unrecoverable data loss. You must take corrective action immediately to avoid extended downtime.
Unknown	The event source is in an unknown state. To view events with this severity type, you select the All option.

Figure 10-48 Event severity types

Note: Performance Advisor uses only the Normal and Error events. All events will be listed under the Operations Manager Events tab.

Operations Manager has many events already configured. See Appendix A in *Operations Manager Administration Guide for Use With DataFabric Manager Server 3.7*, GC26-7889 for a list of events and severity levels. All events that occur within Operations Manager are reflected within the N series Management Console under System Events. You can see some of them in Figure 10-50 on page 240. However, be aware that the events reflected in the Performance Advisor Events window are only of the Error level.

All events are actually reflected in the N series Management Console from Operations Manager, as shown in Figure 10-49 on page 240. All alarms that are set in Operations Manager are also reflected here. Therefore, you can create the alarms for the events you want to track in either location.

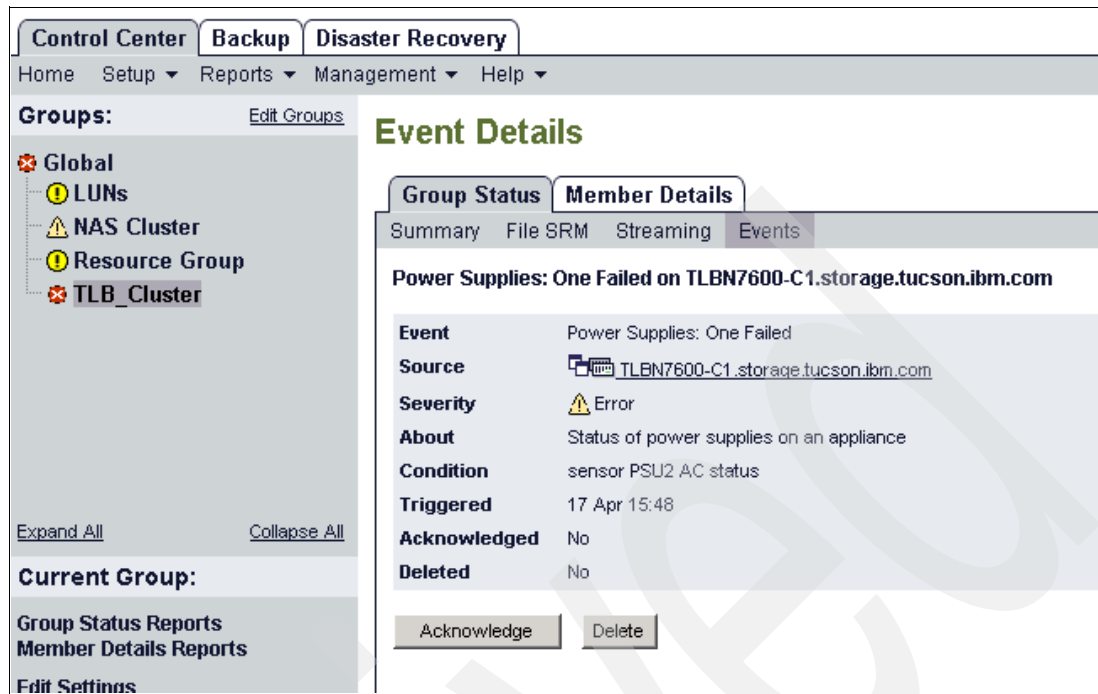


Figure 10-49 Event as seen in Operations Manager

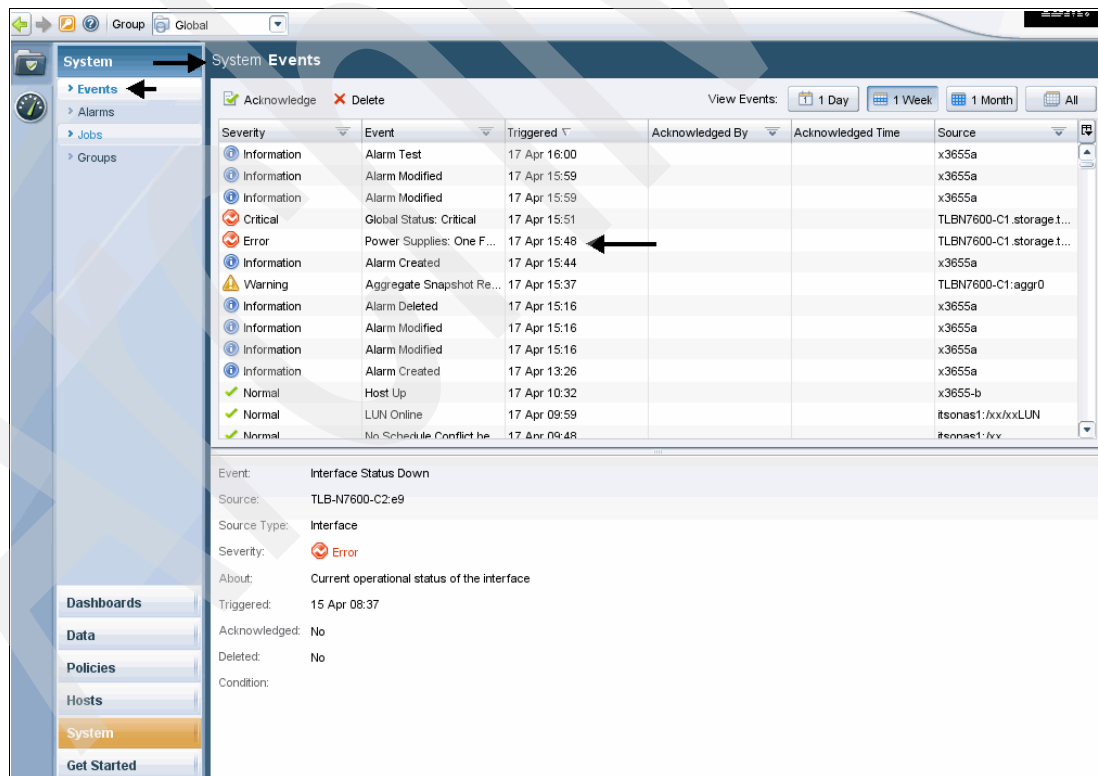


Figure 10-50 Events seen in N series Management Console

Alarms are configured notifications that are sent whenever a specific event or an event of a specific severity type occurs, and are not necessarily related to a specific user. Alarms are used to monitor and manage data sets and resources as a whole.

Alarms are *not* the events themselves, only the *notification* of events. Alarms can be configured within Operations Manager or within the N series Management Console itself.

For example, let us assume that we want to be informed if there is a power failure that affects our N series storage system that has controlled access. Although Operations Manager is monitoring the power, it will not inform us if anything is wrong if we do not set up an alarm. So, in the N series Management Console (or within Operations Manager; the alarm will work from either location), we set up an alarm to notify us if we lose power. We could have remote access to the Operations Manager GUI and view any events that are taking place within our infrastructure, but for this example, we only want to know if power is disrupted.

Also, be aware that there are two places within the N series Management Console that you can create alarms. You can view the alarm from the System Events window, as shown in Figure 10-50 on page 240. You can also set up alarms in the Performance Advisor view.

Creating alarms in the N series Management Console

Begin the alarm setup by selecting the **Set Up** button at the bottom of the left pane. When the Set Up Alarm window appears, select **Add**, as shown in Figure 10-51.

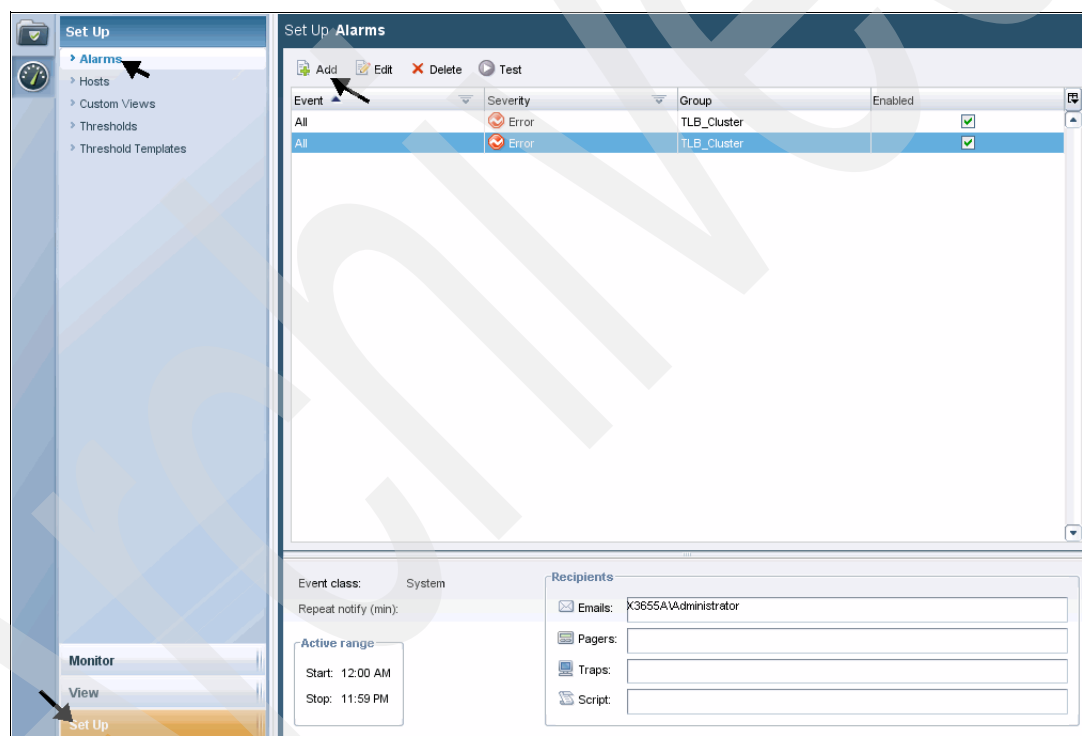


Figure 10-51 Set Up Alarm window

Once you click **Add**, the New Alarm Wizard will open, as shown in Figure 10-52. From here, you will be guided through a few windows to create a new alarm. Click **Next** to begin the setup.



Figure 10-52 New Alarm Wizard

From the Group Selection window, you can pick the group you want associated with this alarm, as shown in Figure 10-53. You can even pick “Global” if you want the alarm to apply to all storage systems.

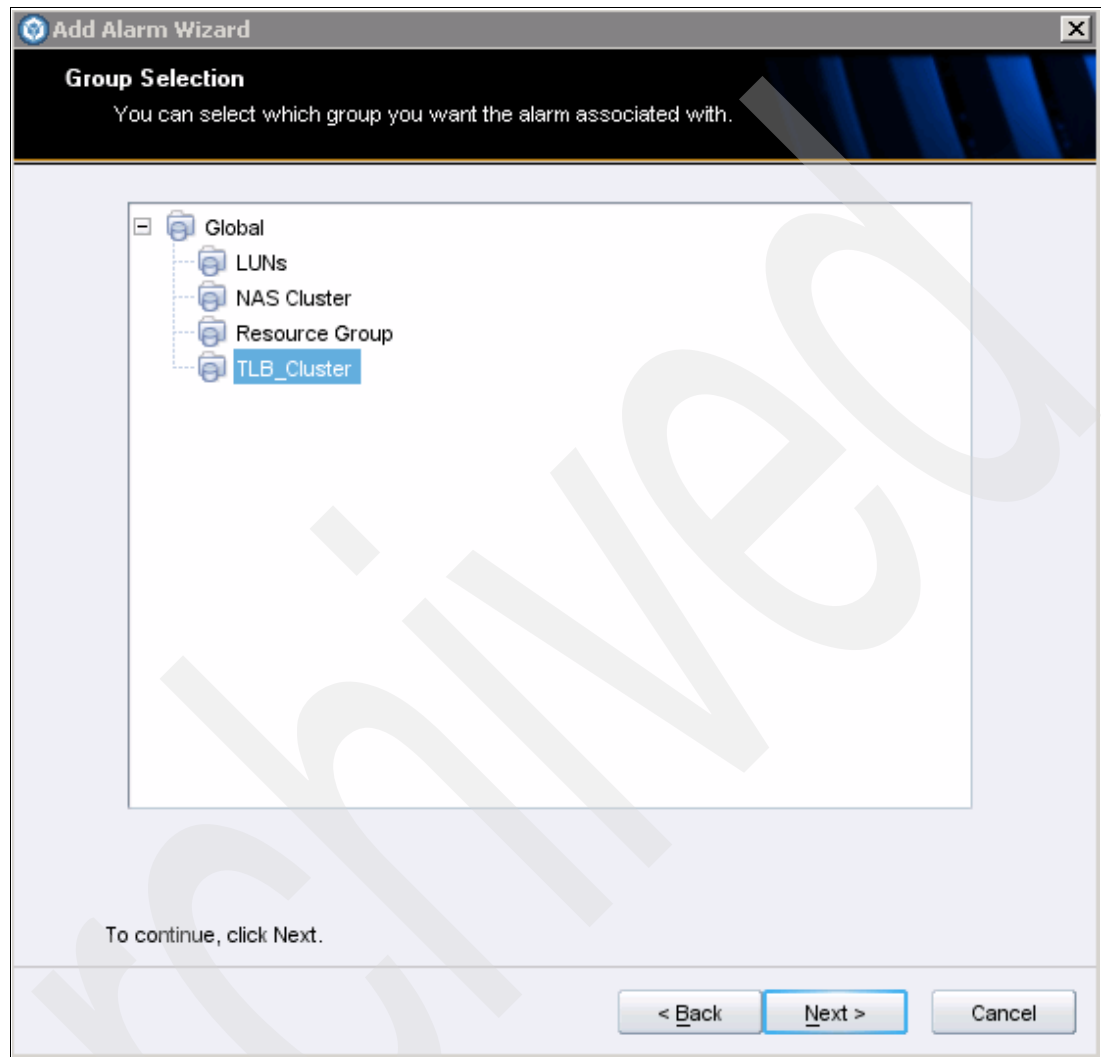


Figure 10-53 Group Selection window

In the window shown in Figure 10-54, you can either create the alarm by event name or severity.

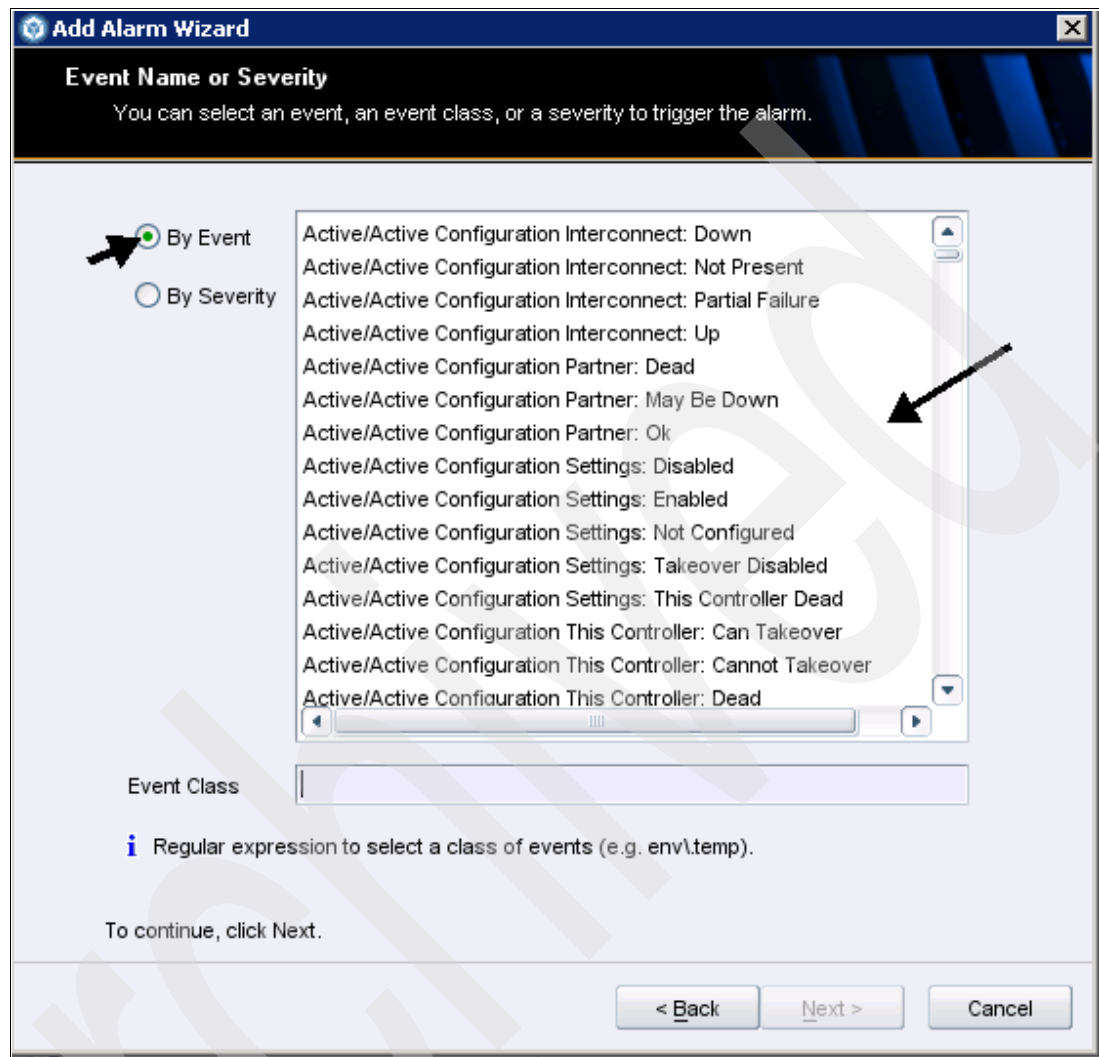


Figure 10-54 Select alarm by event name or severity

You can also set the alarm by severity. Note the Event Class field; here you place a regular expression to select the class of events, as shown in Figure 10-55.

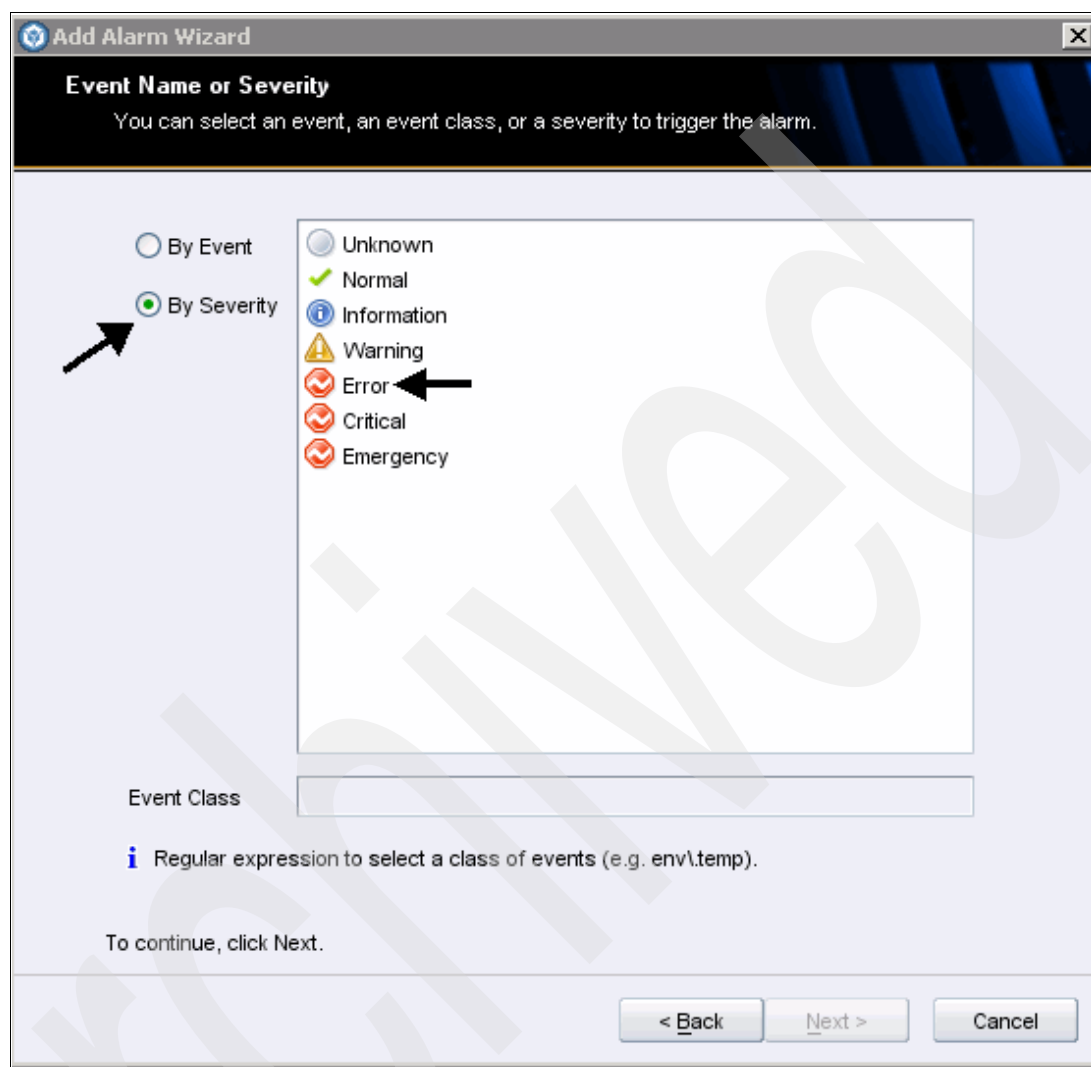
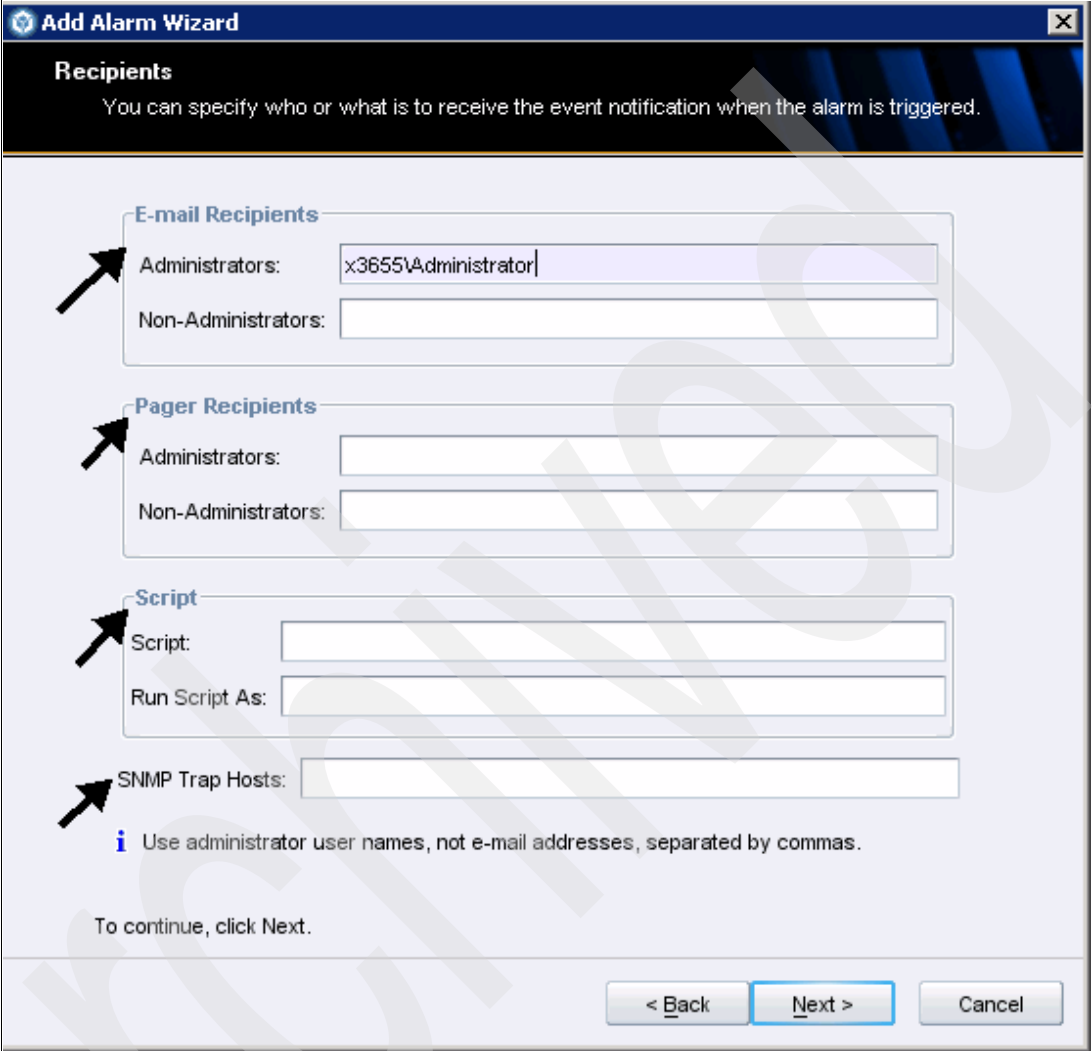


Figure 10-55 Select by severity

You use *regular expressions* to specify event classes. Regular expressions (regex) are familiar to those who have worked extensively with UNIX command-line tools. Regular expressions are rules, or pattern descriptions, that typically use the word "matches" in the expression. They are built by combining expressions to form text strings, such as "hello", which matches "hello" or "hihellothere". Regular expressions are used to compare patterns in text strings until a match occurs. When a match occurs, the function returns true. If no match occurs, the function returns false. Actions are executed only when a pattern-matching expression is true. Refer to the help information in Operations Manager under "Alarms" for more information about regular expressions.

From the Recipients window, shown in Figure 10-56, you can arrange to send an e-mail, page, or even run a script to accomplish a task based on a particular alert. You can also send the alarm to a SNMP host.



The image shows a screenshot of the 'Add Alarm Wizard' window, specifically the 'Recipients' step. The window has a title bar with a gear icon and the text 'Add Alarm Wizard'. Below the title bar, the word 'Recipients' is displayed in a bold font, followed by the instruction: 'You can specify who or what is to receive the event notification when the alarm is triggered.' The main area of the window is divided into four sections, each with a blue header and a text input field. The first section is 'E-mail Recipients', with an arrow pointing to the 'Administrators' field which contains the text 'x3655\Administrator'. The second section is 'Pager Recipients', with an arrow pointing to the 'Administrators' field. The third section is 'Script', with an arrow pointing to the 'Script' field. The fourth section is 'SNMP Trap Hosts', with an arrow pointing to the 'SNMP Trap Hosts' field. Below these sections, there is a blue information icon followed by the text: 'Use administrator user names, not e-mail addresses, separated by commas.' At the bottom of the window, there is a message: 'To continue, click Next.' and three buttons: '< Back', 'Next >', and 'Cancel'.

Add Alarm Wizard

Recipients
You can specify who or what is to receive the event notification when the alarm is triggered.

E-mail Recipients
Administrators: x3655\Administrator
Non-Administrators:

Pager Recipients
Administrators:
Non-Administrators:

Script
Script:
Run Script As:

SNMP Trap Hosts:

i Use administrator user names, not e-mail addresses, separated by commas.

To continue, click Next.

< Back Next > Cancel

Figure 10-56 Recipients of the alert

In Figure 10-57 you are able to decide when the alarm is active. For example, you can have a notification any time of the day or just between certain hours. Also, you can check the **Repeat Notify** check box if you want the alert sent until it is acknowledged.

Add Alarm Wizard

Details

You can specify when the alarm is active and whether you want the notification repeated.

Active Range

Start: 12:00 AM

Stop: 11:59 PM

☒ Repeat Notify 30 Minutes

To continue, click Next.

< Back Next > Cancel

Figure 10-57 When is the alarm active

You will be presented with a summary window, as shown in Figure 10-58. If you are satisfied, click **Finish** to activate the alarm.



Figure 10-58 Summary window

You will now see the alarm in the Alarms window of the N series Management Console, as shown in Figure 10-59.




Figure 10-59 Alarm displays in the Set Up Alarms window

As covered in Chapter 9, “Configuring Operations Manager” on page 127, all of these windows are presented as one window. You can create the same alarms there. In the N series Management Console, it took seven window to set up an alarm. You can do the same thing in Operations Manager in just one window, as shown in Figure 10-60 and “Thresholds” on page 250.

The screenshot shows the 'Alarms' window with the title 'Alarms' and a timestamp '18 Apr 13:09'. The main section is 'Add an Alarm (Advanced)'. It contains three main sections: 'Conditions', 'Recipients', and 'Other Options'.

Conditions:

- Group: ☒ Global
 - ☐ LUNs
 - ☐ NAS Cluster
 - ☐ Resource Group
 - ☐ TLB_Cluster
- Select Event By: ☒ Severity (dropdown: Critical or Worse), ☐ Event Name (dropdown: 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:
- Time To:
- Repeat Notify:
- Repeat Interval:
- Disable:

At the bottom right, there are links for 'Simple Version' and an 'Add' button.

Figure 10-60 Operations Manager alarm one window setup

Thresholds


Thresholds are limits that are set on counters and associated with a set of objects. A threshold breach for a specified interval of time raises events and alarms. They are used to identify potential blockages in the performance of storage systems.

You can use thresholds to determine the point at which Performance Advisor should generate an event. This point depends on the threshold interval and the counter value you set. If you set a threshold on the counter such as the average latency counter on any volume, then an event occurs when the threshold breaches a particular counter value for an interval.

You cannot apply thresholds in combination with the following pairs of objects:

- ▶ Network interface and targets
- ▶ Aggregates and vFiler units
- ▶ vFiler units and disks.
- ▶ Disks and LUNs
- ▶ Disks and qtrees
- ▶ Disks and volumes

There are several thresholds that have already been set in Operations Manager. From the Control Center window, select **Setup** → **Options** → **Default Thresholds** to see a list of thresholds that have already been set, as shown in Figure 10-61. You can change or modify these thresholds.

Options 

Default Thresholds

Host CPU Too Busy Threshold (%)	95
Host CPU Busy Threshold Interval	15 minutes
Aggregate Full Threshold (%)	90
Aggregate Nearly Full Threshold (%)	80
Aggregate Full Threshold Interval	0 seconds
Aggregate Overcommitted Threshold (%)	100
Aggregate Nearly Overcommitted Threshold (%)	95
Aggregate Snapshot Reserve Full Threshold (%)	90
Aggregate Snapshot Reserve Nearly Full Threshold (%)	80
Volume Full Threshold (%)	90
Volume Nearly Full Threshold (%)	80
Volume Full Threshold Interval	0 seconds
Volume Quota Overcommitted Threshold (%)	100
Volume Quota Nearly Overcommitted Threshold (%)	95
Volume Growth Event Minimum Change (%)	1
Volume Snap Reserve Full Threshold(%)	90
Volume No First Snapshot Threshold(%)	90
Volume Nearly No First Snapshot Threshold(%)	80
Volume Space Reserve Depleted Threshold(%)	90
Volume Space Reserve Nearly Depleted Threshold(%)	80

Figure 10-61 Thresholds in Operations Manager

Within Performance Advisor, you can create other objects in which to set thresholds. In order to start the process, select **Setup → Thresholds**. You will see the window shown in Figure 10-62.

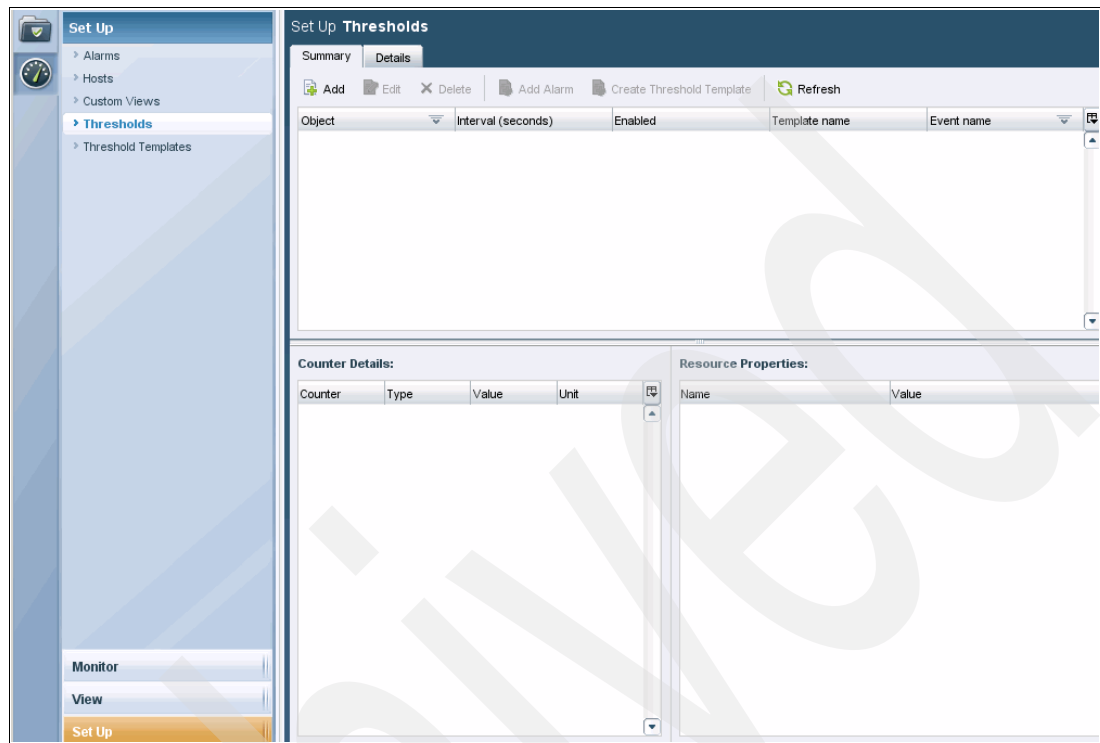


Figure 10-62 Setup Thresholds window

If you are familiar with N series products and comfortable with setting up thresholds, you can use the Add function in this window (we will discuss the feature in the following paragraphs). However, if you would like to see an exhaustive list of counters you can set thresholds for, select the **Details** tab, which will bring up the window shown in Figure 10-63 on page 253.

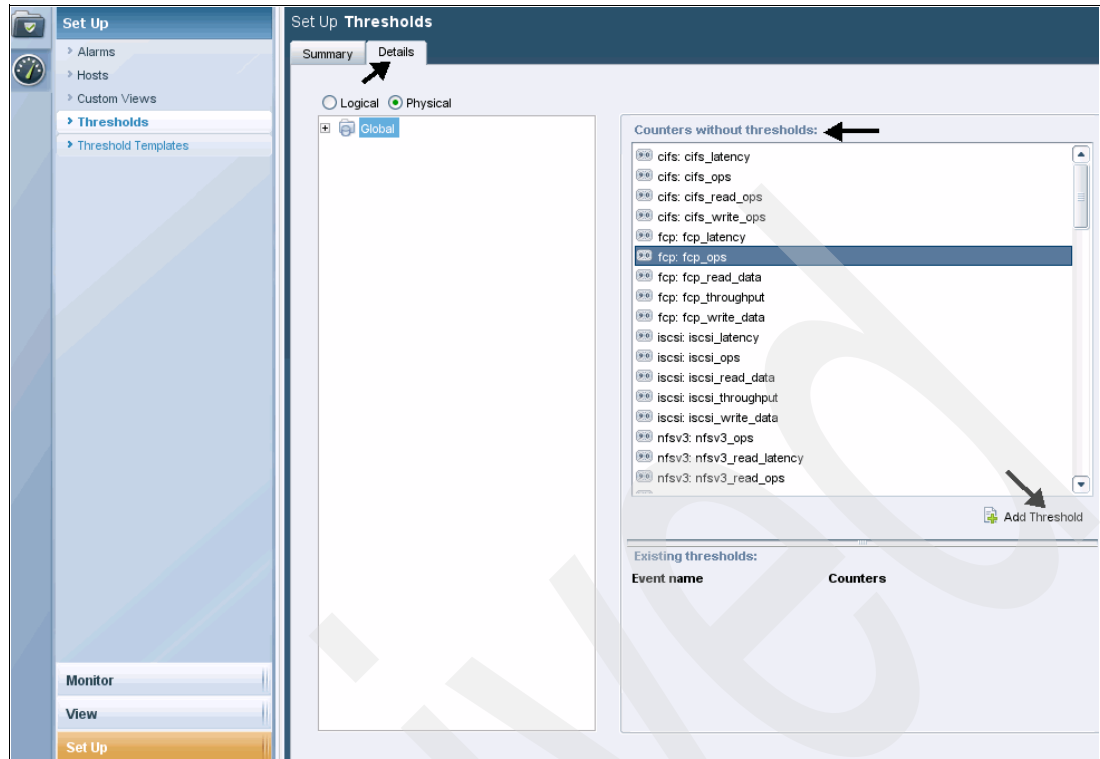
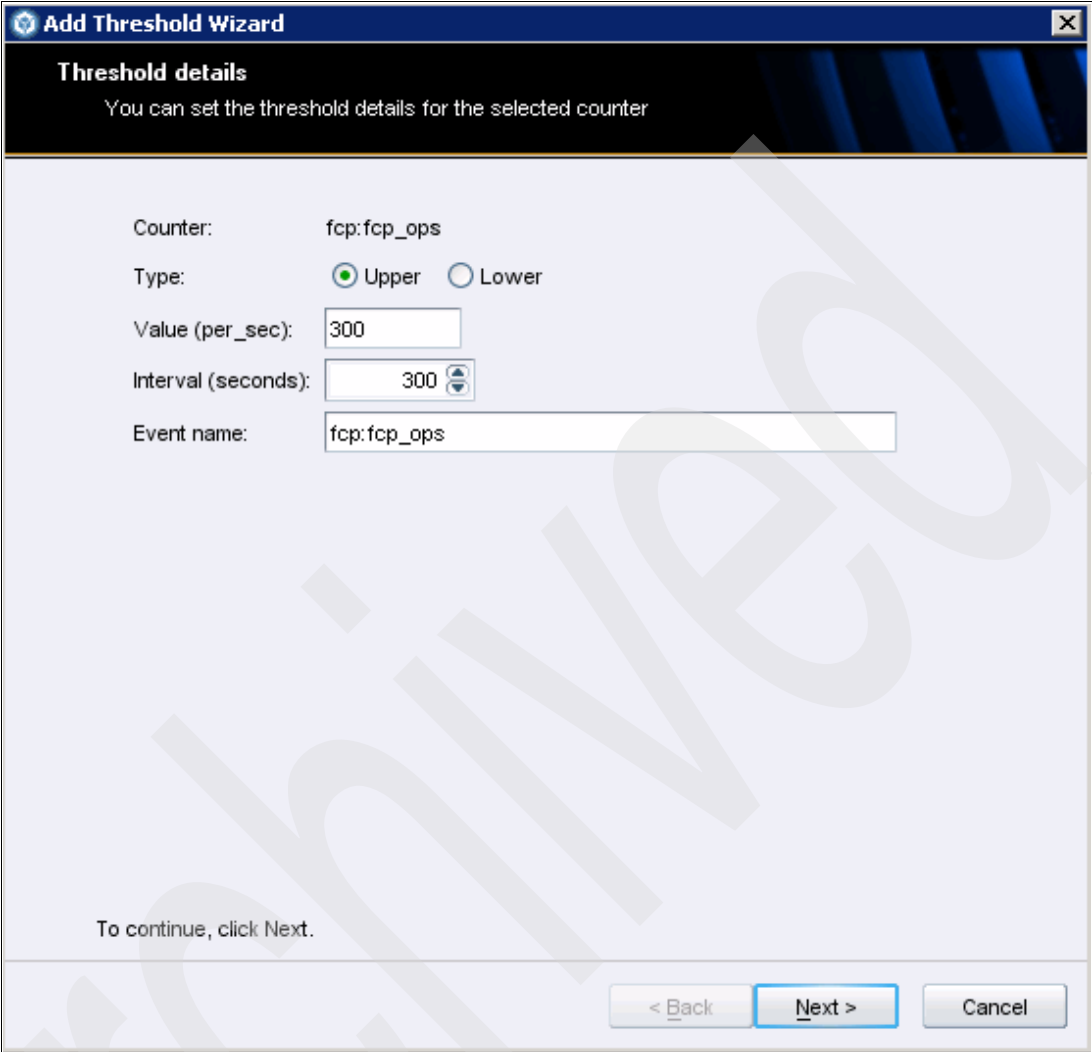


Figure 10-63 Thresholds details window

In this window, you see a list of counters without thresholds. One good thing about this window is that it gives the name of the counter for which you can set a threshold. This is different from the Add Threshold wizard in that you have to provide the counter name for which you want to set a threshold. If you do not know the name of the counter, it is difficult to set the threshold.

In this window, select the **Add Threshold** button in the lower right of the window.

The window shown in Figure 10-64 appears.



The image shows a Windows-style dialog box titled "Add Threshold Wizard" with a close button (X) in the top right corner. Below the title bar is a dark blue header area with the text "Threshold details" and a subtitle "You can set the threshold details for the selected counter". The main area of the dialog is light blue and contains the following fields:

- Counter: fcp:fcp_ops
- Type: ☒ Upper ☐ Lower
- Value (per_sec): 300
- Interval (seconds): 300 (with a spin button)
- Event name: fcp:fcp_ops

At the bottom left, it says "To continue, click Next." At the bottom right, there are three buttons: "< Back", "Next >" (which is highlighted with a blue border), and "Cancel".

Figure 10-64 Select upper or lower limit for the threshold

In this window, you are able to select an upper or lower limit for your threshold. When you have filled in the value, click **Next**.

You will be given a summary window, as shown in Figure 10-65. If you are satisfied, click **Finish**.

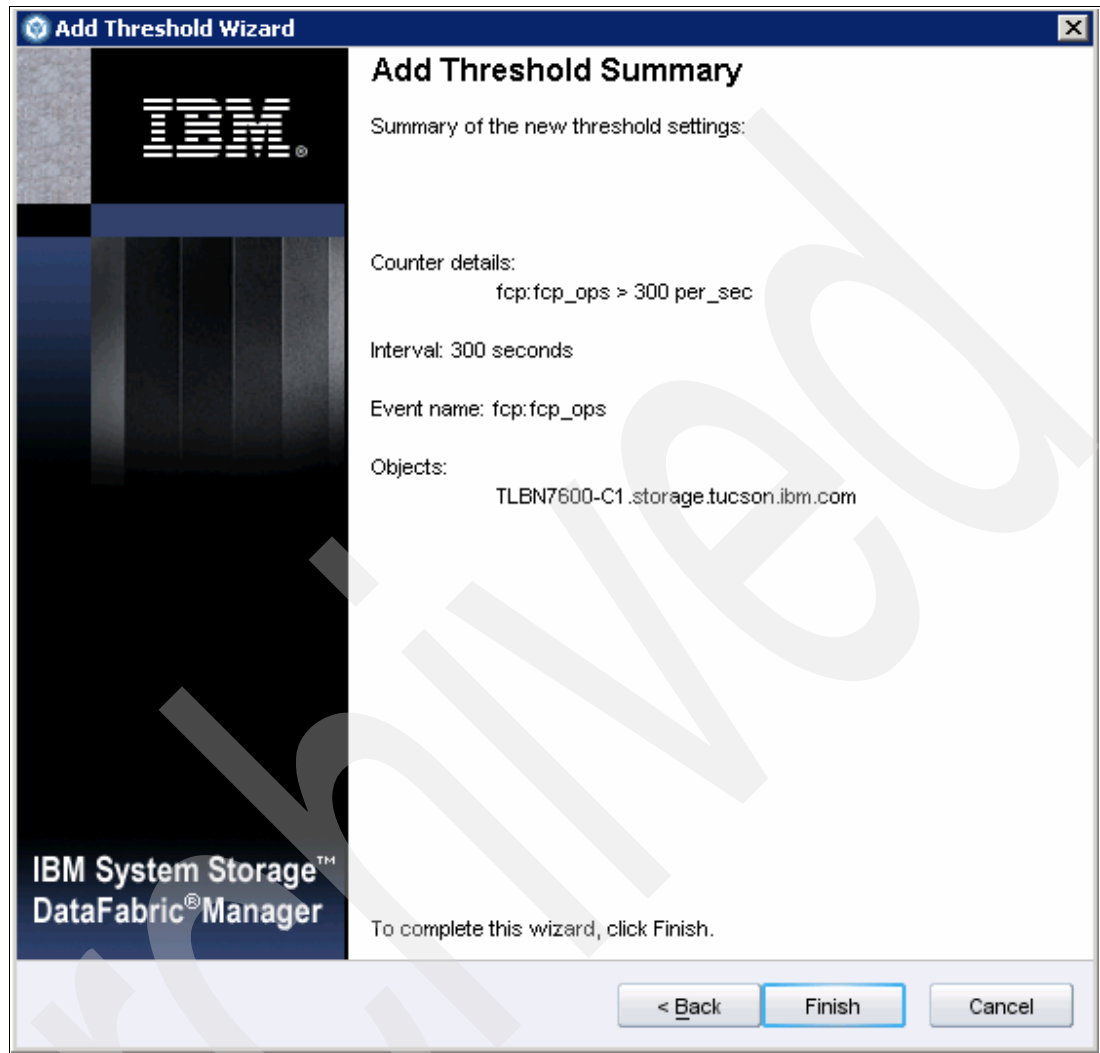


Figure 10-65 Summary window for adding a threshold

Your new threshold is now set, as shown in Figure 10-66.

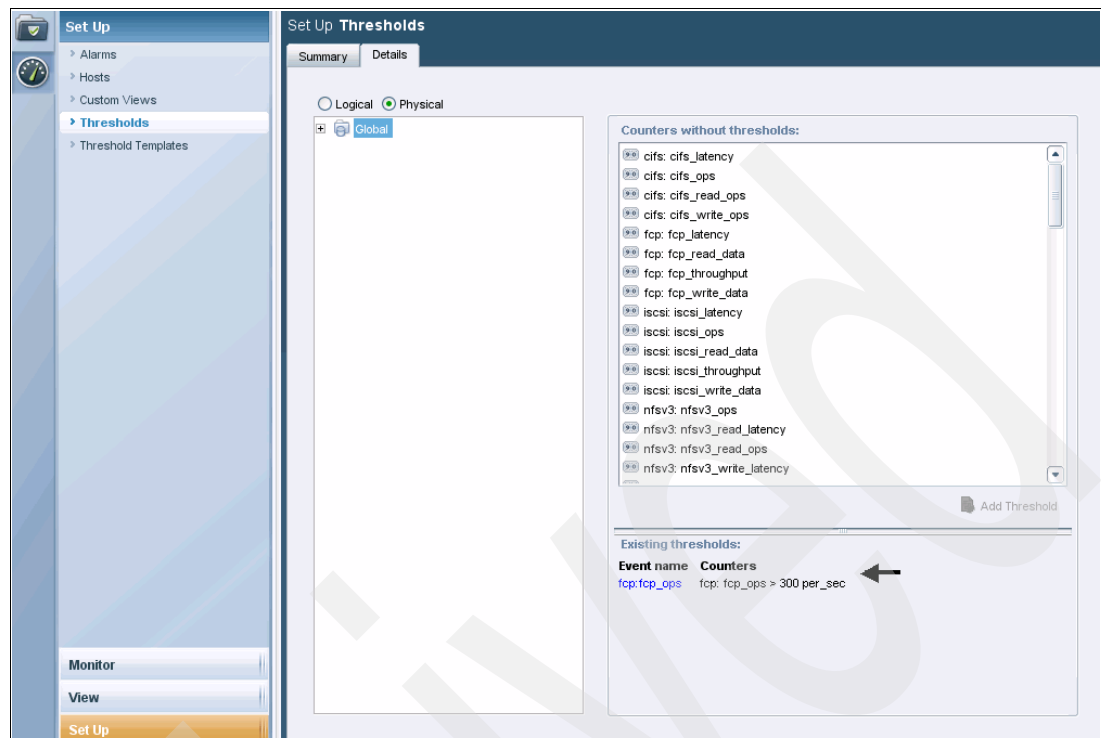


Figure 10-66 New threshold added

Thresholds can be set at the system, aggregate, volume, and disk level. As you step through the logical and physical selections of the tree, the counters you can monitor vary. For example, object types that represent the physical resources in a storage system, such as disks, aggregates, memory, network interfaces, resource pools, and RAID groups, are known as physical objects.

Figure 10-67 gives you an example of the counters available at the physical level. In this particular example, we can set up thresholds based on aggregates, Ethernet ports, fiber HBA, or any “physical” object we want to monitor.

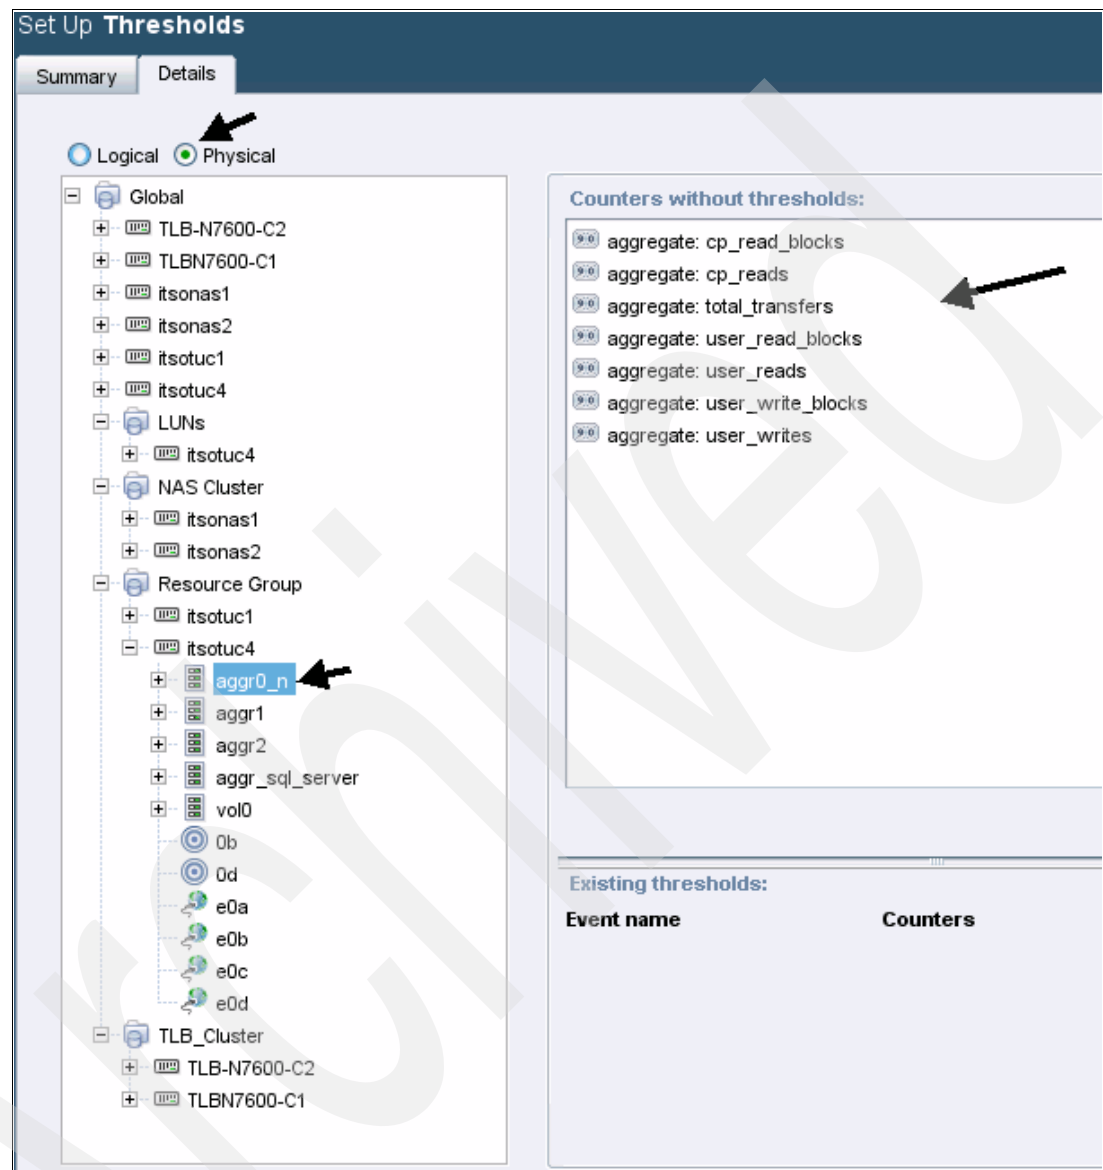


Figure 10-67 Physical level counters example

On the other hand, logical objects are object types that represent storage containers, such as volume, qtree, LUNs, vFiler units, and data sets, as shown in Figure 10-68.

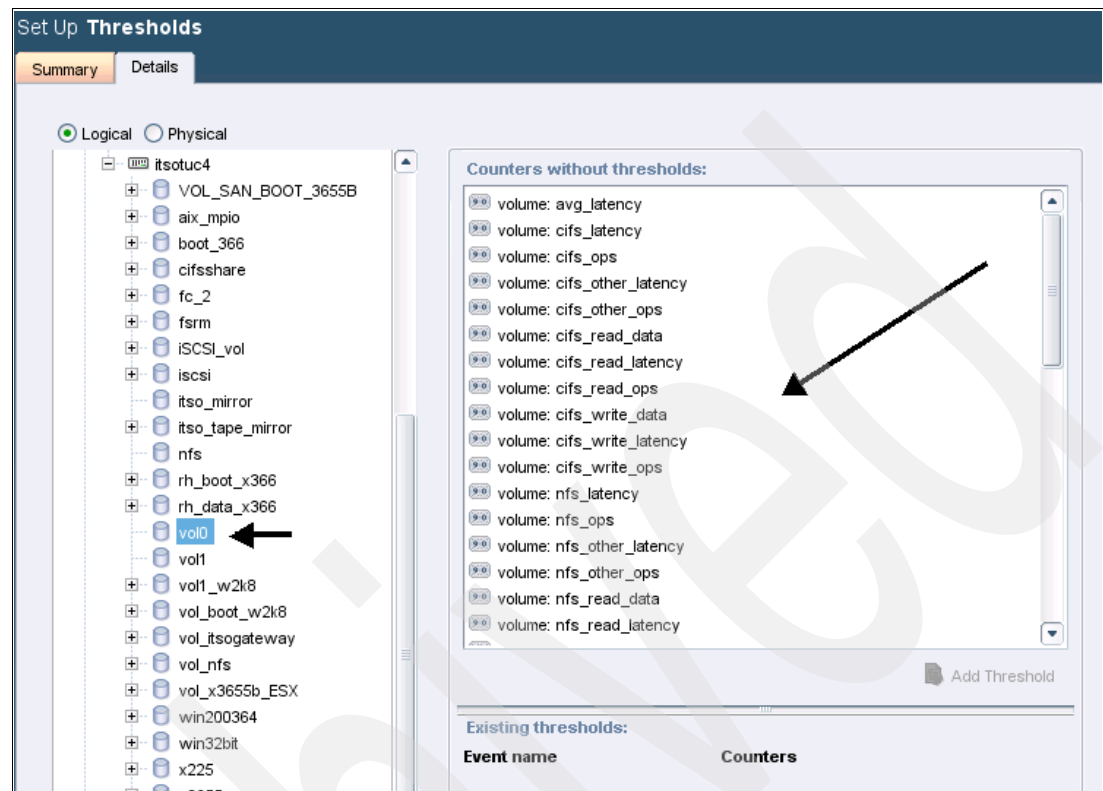


Figure 10-68 Logical level counters example

As you can see from these examples, Operations Manager is capable of monitoring a vast array of counters. It is possible to set up a threshold for each one.

Adding a threshold from the Summary view is a bit different. You need to know the name of the counter you want to set the threshold for in order to create the threshold. It is easier to set the threshold in the detail view, but we will discuss how to add a threshold from the summary window.

Note the arrows in Figure 10-69. Not only can you add a threshold from this window, but you can also edit thresholds, create templates, delete thresholds, and set alarms based on a threshold breach.

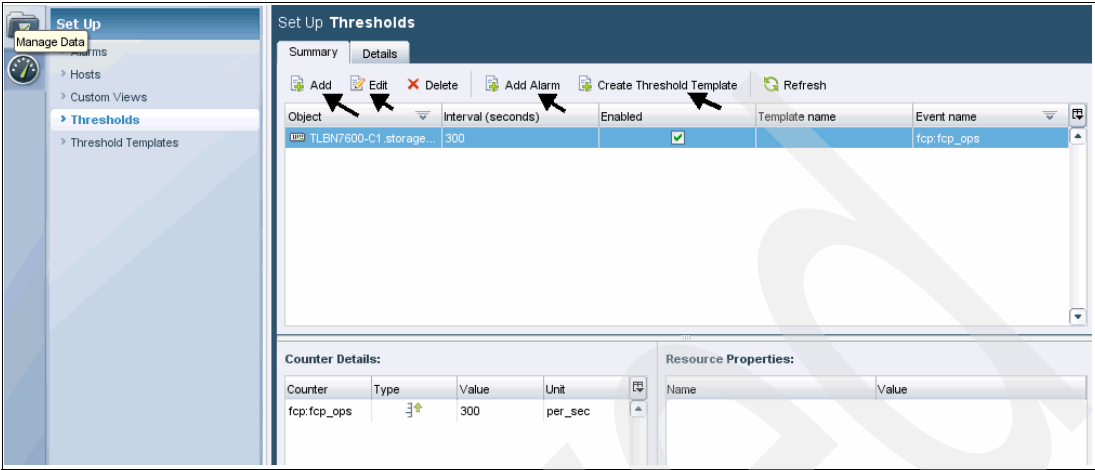


Figure 10-69 Threshold summary window

First, let us discuss the steps to create a threshold. As we noted when we created an alarm, the N series Management Console uses wizards to create some of its options. When you click Add in the window shown in Figure 10-69 on page 259, you will be presented with the wizard window shown in Figure 10-70. Click **Next**.

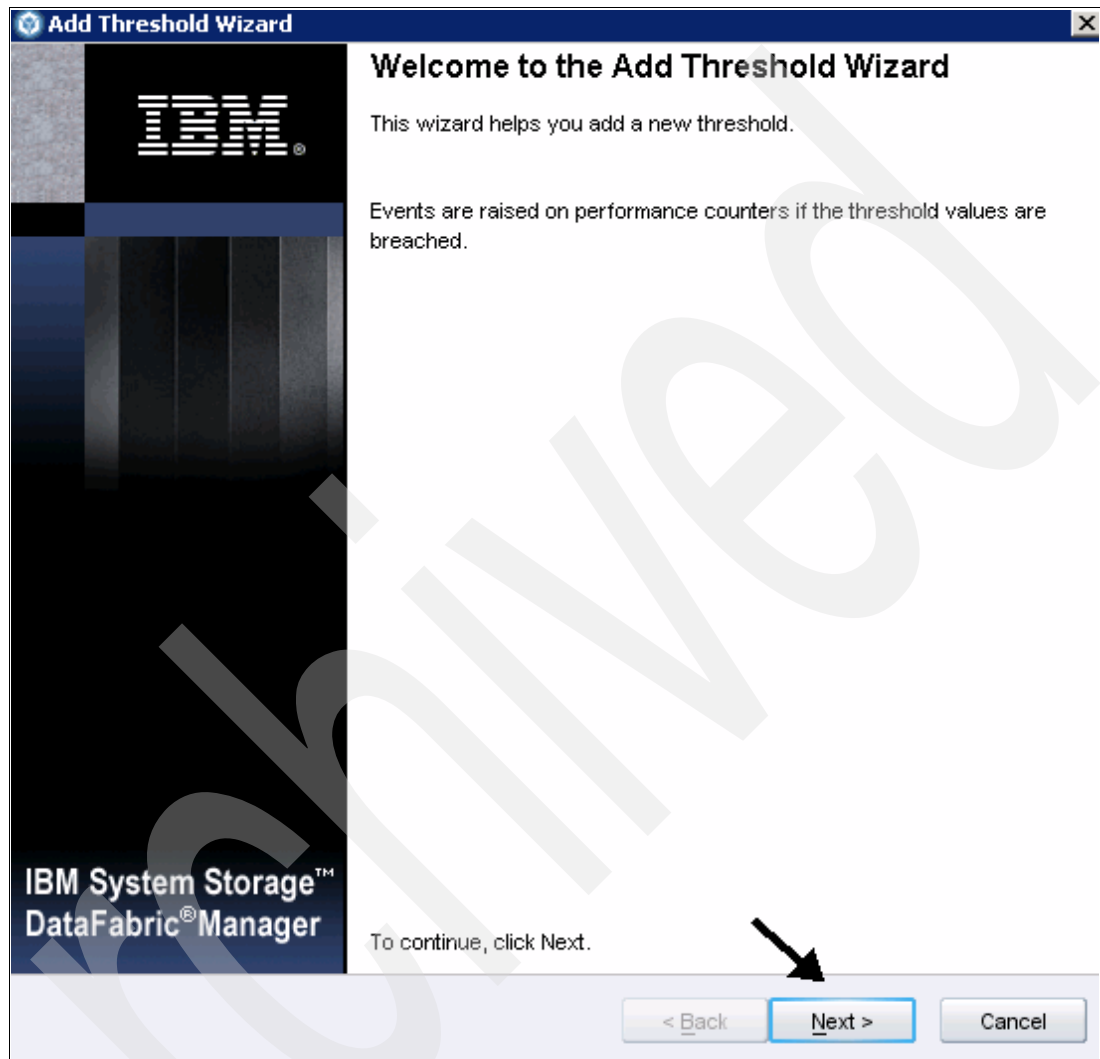


Figure 10-70 Add Threshold wizard

The first thing you will notice is that you have the ability to select multiple items at one time, as shown in Figure 10-71. When this is done, a threshold breach is not reported until all of the threshold limits have been met.

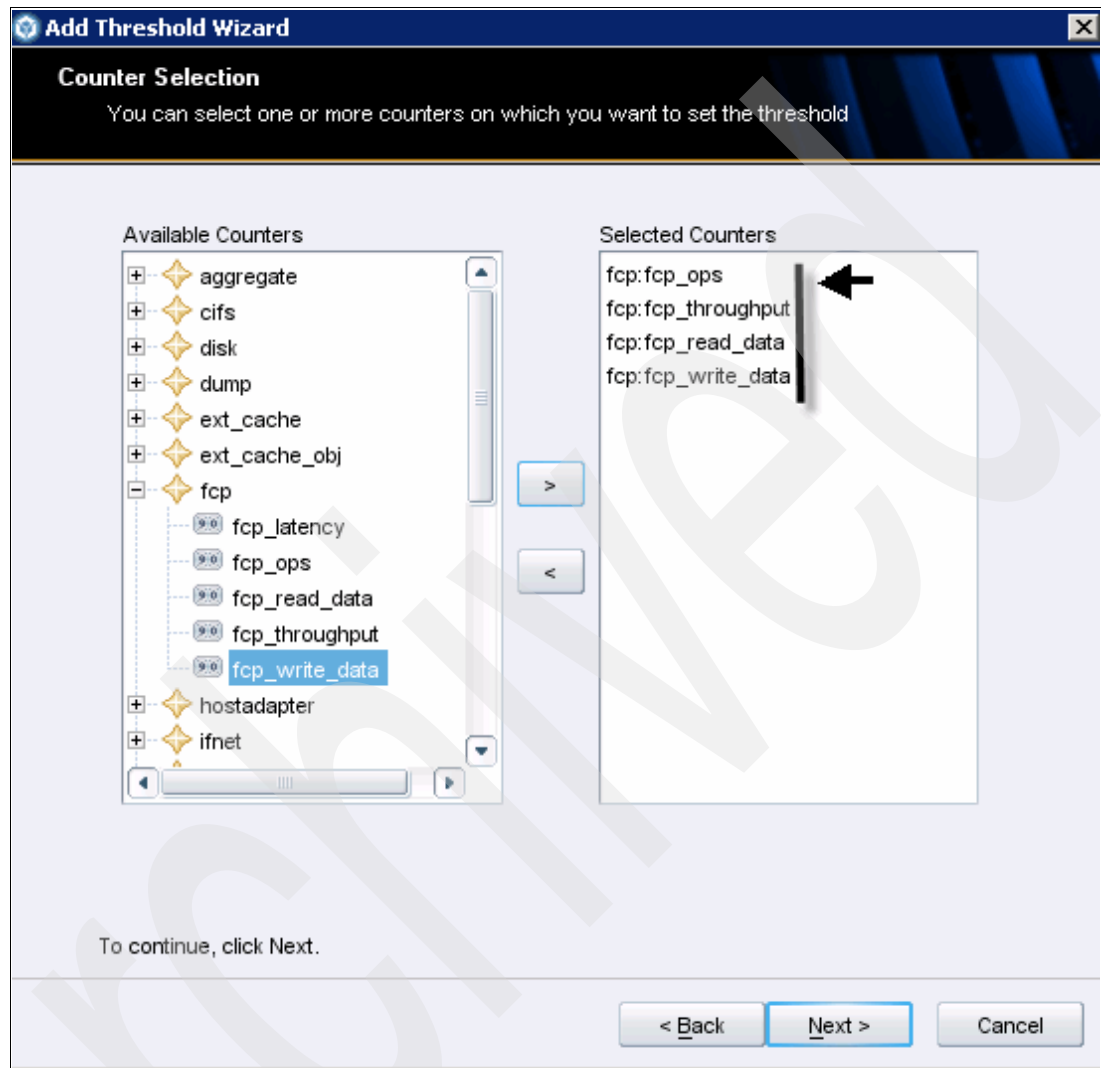


Figure 10-71 Multiple counters can be set

Once you have decided what you want to monitor, you can select the storage system to which you want this threshold to apply, as shown in Figure 10-72.

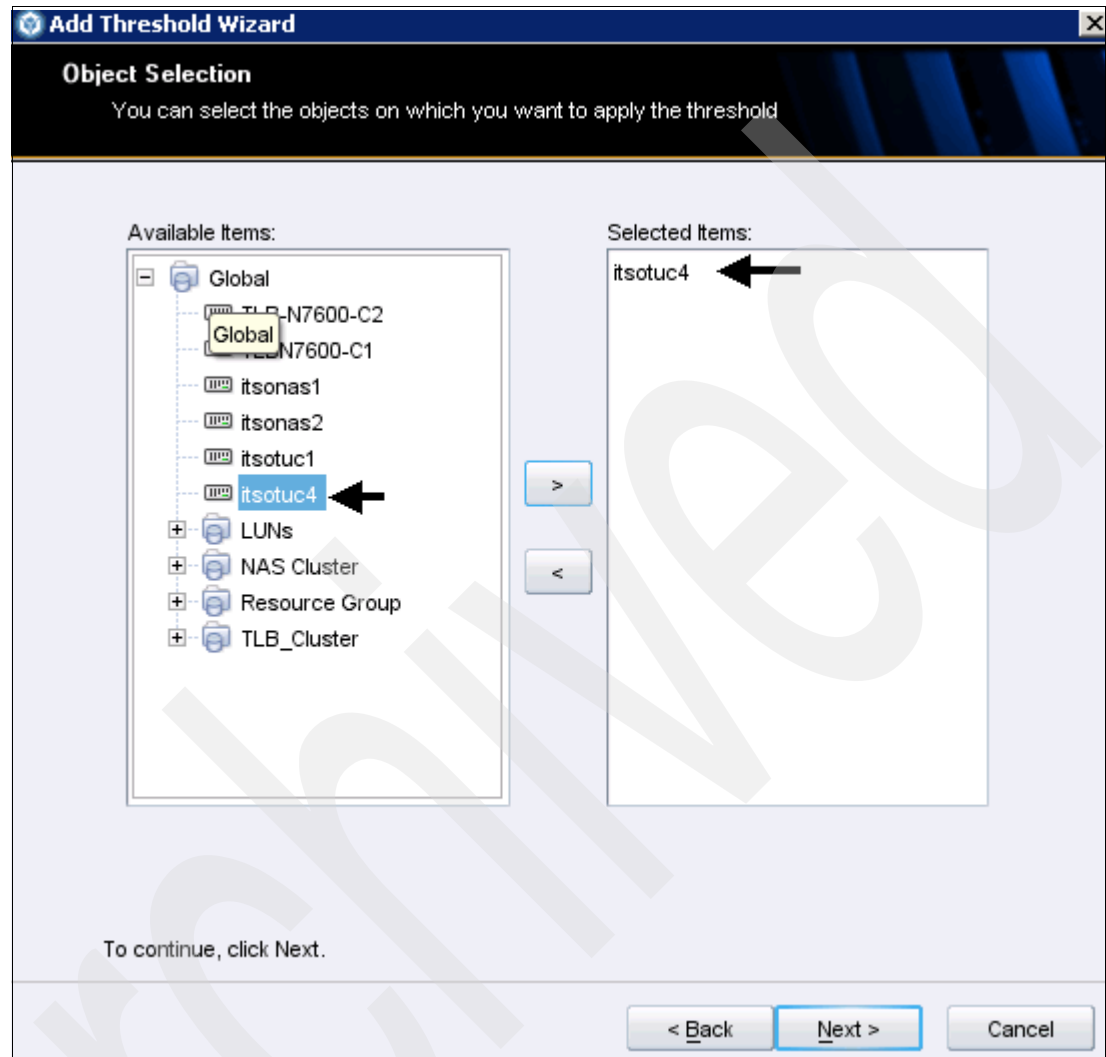


Figure 10-72 Select storage system or object you want to monitor

Once you have selected the system you want to monitor, you can set the thresholds you would like for the items selected. Here we used arbitrary figures. These are in no way suggestions; we just put numbers in for reference only.

In Figure 10-73, you can decide on the high or low threshold levels and the units you want sampled. Our case is based on per-second. You also have to provide a name for this threshold. When you are satisfied with this threshold, click **Next** to continue.

The screenshot shows a configuration window for setting threshold levels. At the top, there is a label "Interval(seconds):" followed by a text box containing the value "300" and a small up/down arrow icon. Below this is a table with four columns: "Counter", "Type", "Value", and "Unit". The table contains four rows of data. To the right of the "Type" column, there are three green up arrows and one black arrow pointing to the "Value" column. Below the table, there is a label "Event name:" followed by a text box containing the value "fcp:data" and a black arrow pointing to it. At the bottom of the window, there is a small text label "To continue, click Next".

Counter	Type	Value	Unit
fcp:fcp_ops		300	per_sec
fcp:fcp_throughput		300	b_per_sec
fcp:fcp_read_data		300	b_per_sec
fcp:fcp_write_data		300	b_per_sec

Event name: fcp:data

To continue, click Next

Figure 10-73 Threshold levels

Figure 10-74 shows a summary window. Click **Finish** if you are satisfied, or you can click **Back** to make changes..

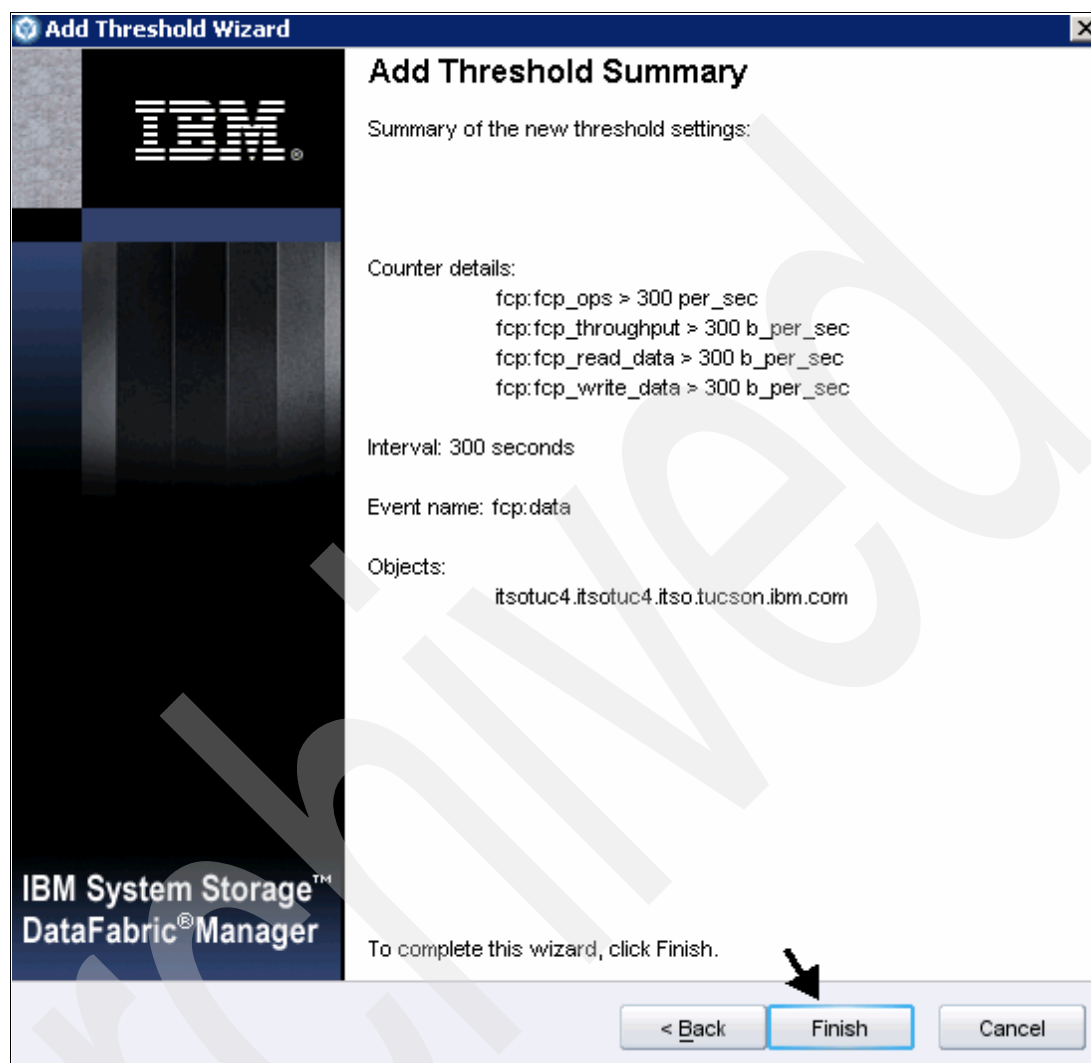


Figure 10-74 Add Threshold Summary window

When the threshold creation is complete, the threshold will appear in the N series Management Console window, as shown in Figure 10-75 on page 265.

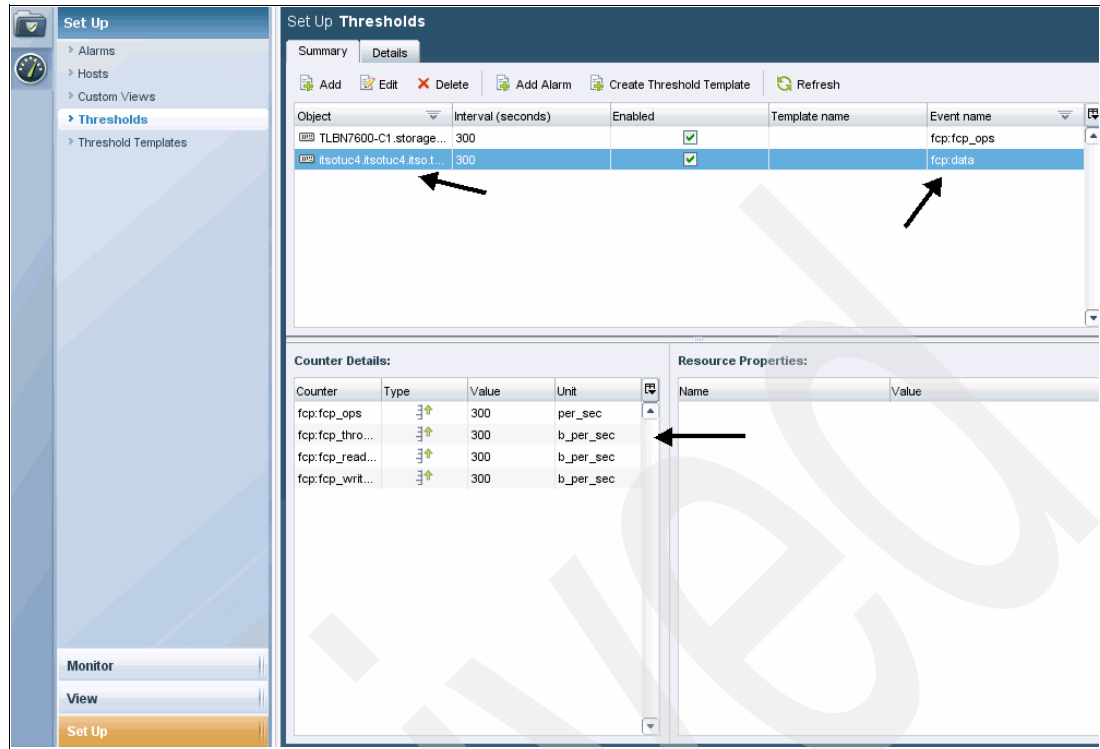


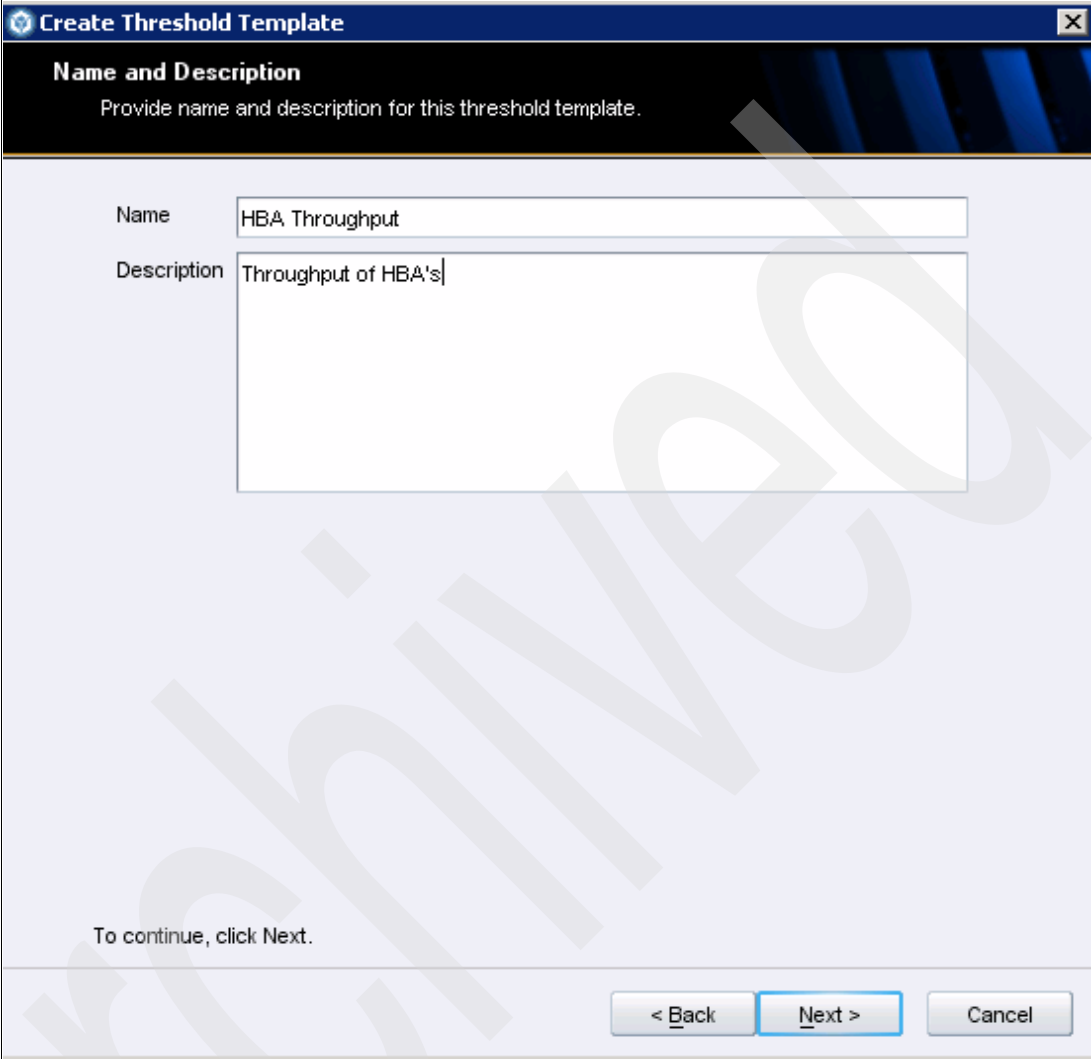
Figure 10-75 New threshold appears

Once a threshold has been created, you can edit it, delete it, create a template out of it, or set an alarm for notification if the threshold is breached.

A threshold template is a set of thresholds. You can apply a threshold template to one or more applicable objects in a storage system.

An object can be assigned to a template only if there is at least one applicable threshold in the template. For example, you cannot assign a volume object if volume threshold was not present in the template.

To create a template, select **Create Threshold Template** from the top of the window. The window shown in Figure 10-76 will appear.



Create Threshold Template

Name and Description
Provide name and description for this threshold template.

Name: HBA Throughput

Description: Throughput of HBA's

To continue, click Next.

< Back Next > Cancel

Figure 10-76 Create Threshold Template Name and Description

You have the opportunity to add thresholds to the template you are creating in addition to the one you have selected, or you can edit the values you have already set in your threshold, as shown in Figure 10-77 on page 267.

Add
 Delete

Object	Counter	Type	Value	Unit
fcp	fcp_ops	upper	300	per_sec

Event name: fcp:fcp_ops

Filter based on properties:

Add
 Delete

Name	Value
------	-------

Figure 10-77 Add or edit items to your template

Click **Next**, and you can review your template. If you are satisfied, click **Finish**. The template will appear in the Threshold Templates view of the N series Management Console, as shown in Figure 10-78.

IBM N Series Management Console: Manage Performance - administrator on x3655a

File View Tasks Help

Group Global

Set Up

- Alarms
- Hosts
- Custom Views
- Thresholds
- Threshold Templates

Set Up Threshold Templates

Add
 Edit
 Copy
 Delete
 Objects
 Import
 Export
 Refresh

Name	Description
HBA Throughput	Throughput of HBA's

Events:

fcp:fcp_ops

Counters:

Threshold interval(seconds): 300

Counter	Type	Value	Unit
fcp:fcp_ops		3000	per_sec

Figure 10-78 Threshold Templates view

You can associate any alarm with any threshold using the same view that you used to create alarms in N series Management Console. To do this task, select **Alarms** within Performance Advisor, as shown in Figure 10-79. Refer to “Creating alarms in the N series Management Console” on page 241” for more information.

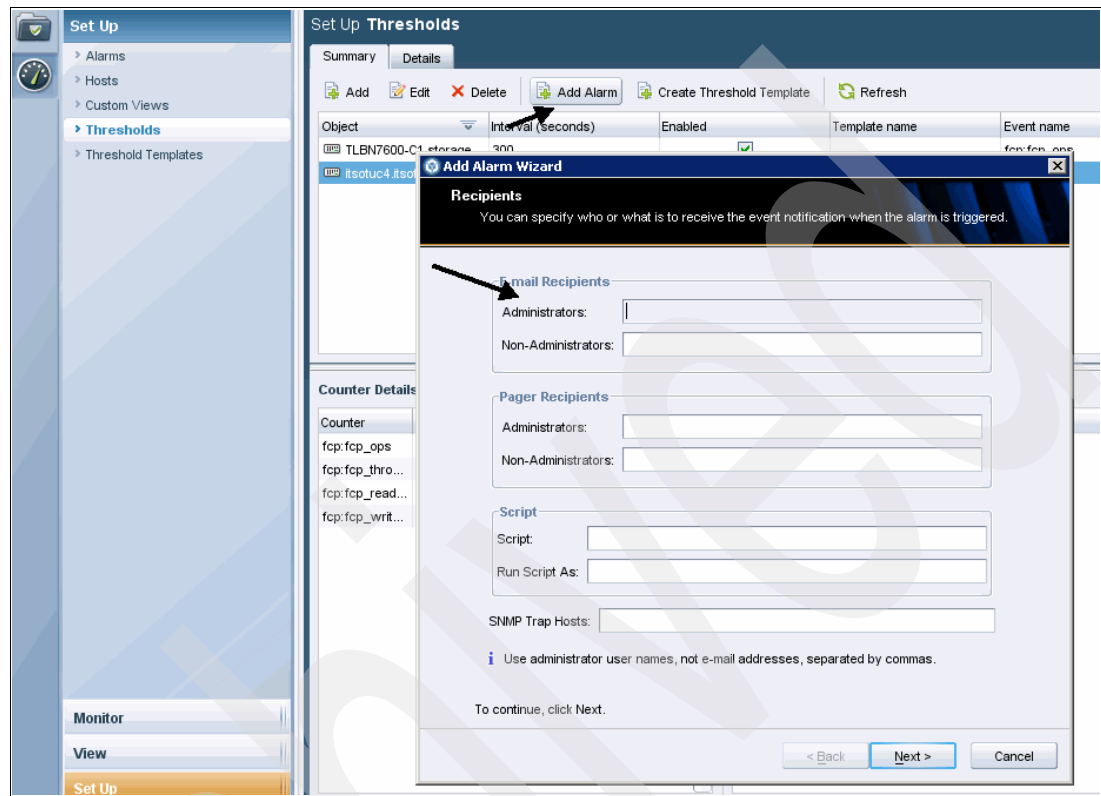


Figure 10-79 Adding an alarm to a threshold

In this chapter, we have shown you many of the features and functions of the Performance Advisor portion of the N series Management Console. This is not an exhaustive list or representation of all that can be seen and done with Performance Advisor. We highly recommend that you refer to *Performance Advisor Administration Guide for Use with DataFabric Manager Server*, GC26-7897 for all of the features and functions of this product.



File Storage Resource Manager

This chapter describes and gives examples about File Storage Resource Manager (FRSM) usage. Refer to Chapter 3, “Installing Operations Manager: Windows 2003 64-bit operating system” on page 37 for more information about FRSM. You can also find detailed information about FRSM in “*Operations Manager Administration Guide for Use With DataFabric Manager Server 3.7*, GC26-7889.

FRSM is part of the Storage Management function of Operations Manager, as shown in Figure 11-1).

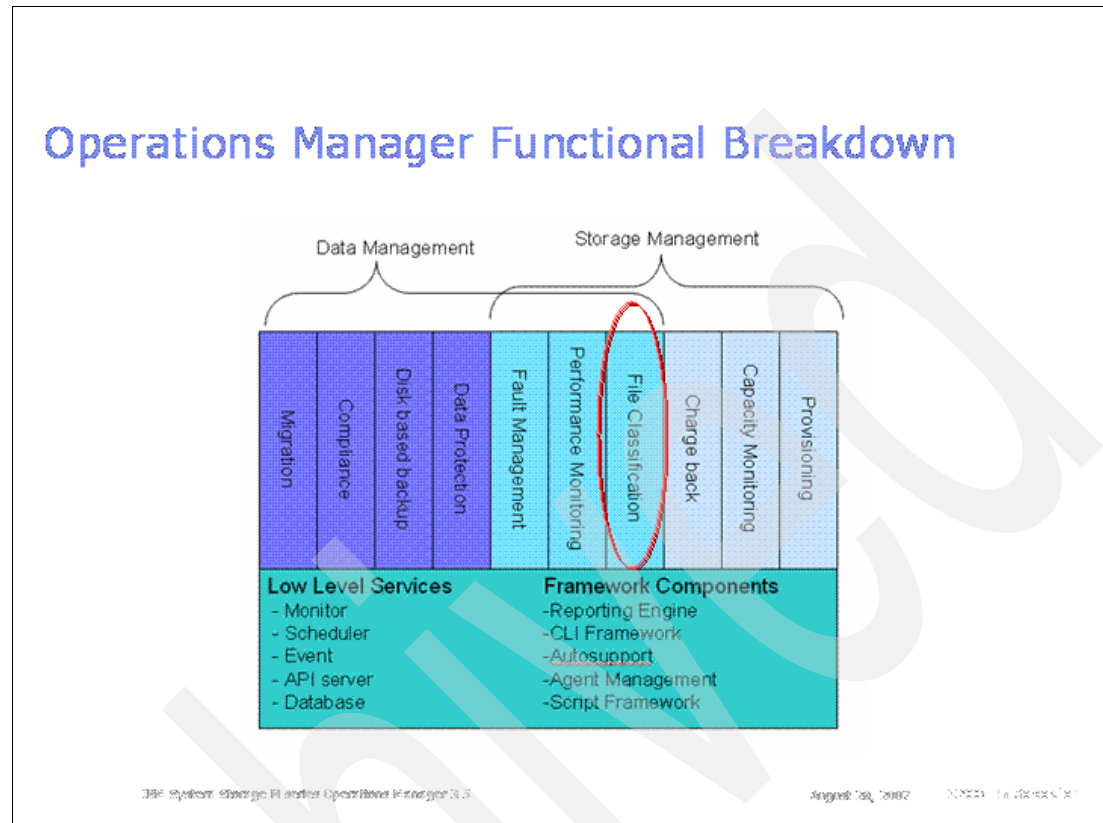


Figure 11-1 Operations Manager functions

Our environment includes host agents for Windows Server® 2003 32-bit and 64-bit, as well as Linux. We found that there were no differences in the way the information was obtained or displayed from FSRM regardless of the operating system.

11.1 Tracking file system usage and capacity information

In order to use FSRM, you must first install the Operations Manager Host Agent for the host you are wanting to monitor. The Host Agent will install on various operating systems, including Windows, Solaris, Linux, and VMWare. Refer to the *Operations Manager Host Agent 2.6 Installation and Administration Guide*, GC26-7894 for specific information about supported operating systems.

Operations Manager Host Agent is software that resides on a Windows, Linux, or Solaris host. It collects information, such as OS name, version, HBA information, and file system meta data, and then sends that information back to the DataFabric Manager Server. Users can create reports of the collected information by using the Operations Manager interface or the DataFabric Manager DFM Server CLI. To enable a target host to communicate with the DataFabric Manager Server, install and configure the Operations Manager Host Agent software on that host. After the DataFabric Manager Server discovers that instance of Operations Manager Host Agent, no further configuration is required. Operations Manager Host Agent does not initiate any management actions on the Windows, Linux, or Solaris host; it is strictly a passive agent. It acts only on requests from external management applications, such as the DataFabric Manager Server.

You need Operations Manager Host Agent *only* if you want to monitor SAN hosts or FSRM-generated file system data through Operations Manager.

After you install Operations Manager Host Agent on a host, you can use Operations Manager to perform a variety of SAN and FSRM functions.

- ▶ **SAN capabilities:** Using Operations Manager Host Agent and Operations Manager, you can perform the following SAN tasks:
 - Monitor basic system information for the SAN hosts.
 - View detailed HBA and LUN information.
- ▶ **FSRM capabilities:** Using Operations Manager Host Agent and Operations Manager, you can perform the following FSRM tasks:
 - Collect storage usage data at the file and directory level.
 - Identify a variety of file-related information, for example, largest files, oldest files, or space consumed per file type.

Note: Although host agents can be used both to monitor SAN hosts and to provide file system meta data, this chapter only describes the setup for File SRM.

11.2 About FSRM

Use File Storage Resource Manager (FSRM) to gather file system meta data and generate reports about different characteristics of that meta data. The DataFabric Manager Server interacts with the Host Agent software residing on remote Windows, Solaris, or Linux workstations or servers (called *hosts*) to recursively examine the directory structures (*paths*) you have specified.

For example, if you suspect that certain file types are consuming excessive storage space on your storage systems, you would deploy one or more host agents and then configure to walk a path, which might be mounted on top of a Data ONTAP LUN, volume, or qtree. FSRM can be configured to generate reports periodically detailing which files are consuming the most space, which files are old or have been accessed recently, and the types of files (.doc, .gif, .mp3, and so on) on the file system. You can then make an intelligent choice about how to efficiently use your existing space.

For example, the FSRM Summary window will give you a snapshot of some of this information, as shown in Figure 11-2.

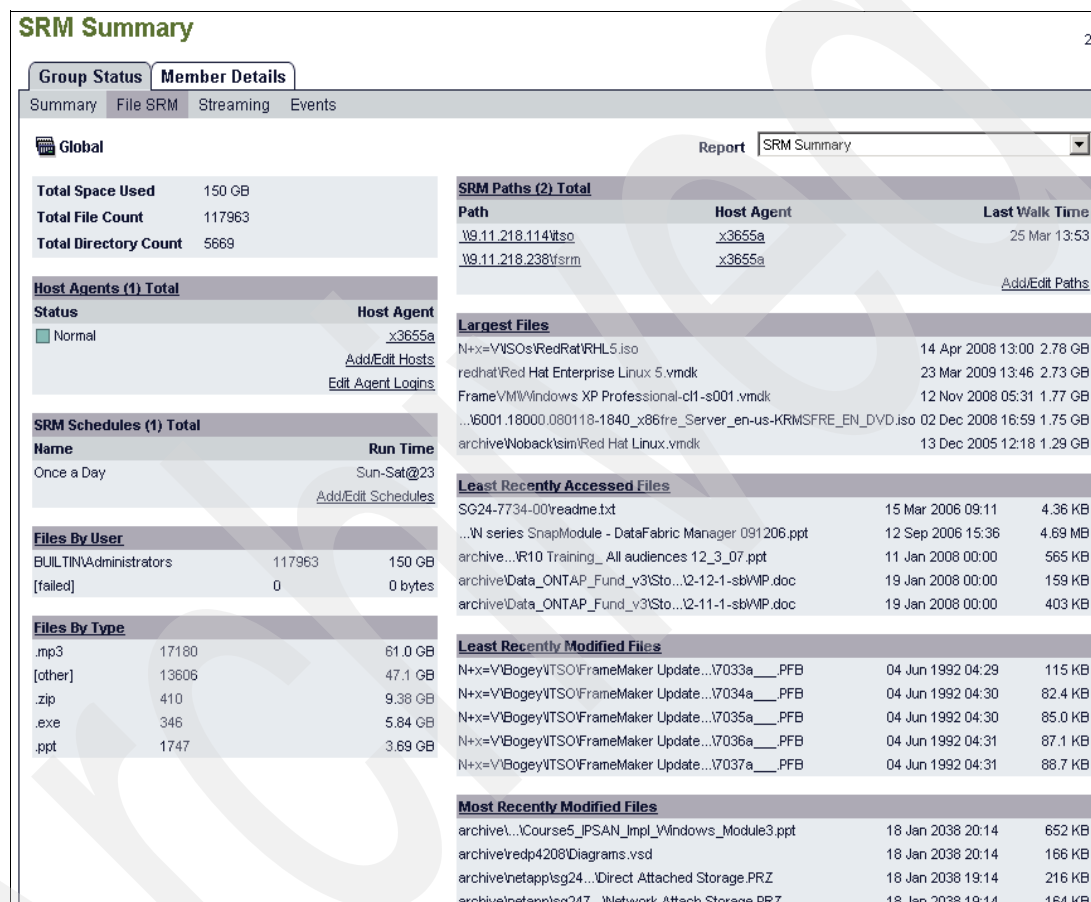


Figure 11-2 FSRM example of various meta data information

From the SRM Summary window, you can select the drop-down menu and get more detail about other FSRM information. For example, you can see where the largest files are, what the oldest files are, what kind of files are taking up the most space, and so on, by selecting one of the options under this menu, as shown in Figure 11-3 on page 273.

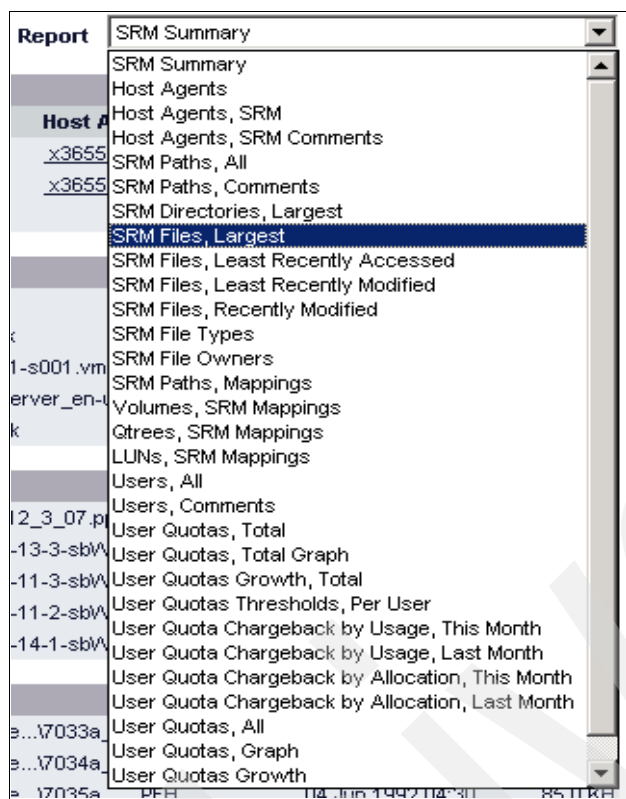


Figure 11-3 SRM menu selections

Note: The File SRM tab in the Operations Manager includes other storage monitoring utilities, for example, chargeback and quota reporting. However, this chapter uses the term *FSRM* to refer specifically to the file system meta data collection functionality.

11.3 How FSRM works

The DataFabric Manager Server monitors directory paths that are visible to the host agent. Therefore, if you want to enable FSRM monitoring of storage systems, the remote host must mount a share using NFS or CIFS, or the host must use a LUN on the storage system.

Note: Operations Manager cannot obtain FSRM data for files located in volumes that are not exported by CIFS or NFS. Host agents can also gather FSRM data about other file system paths that are not on a storage system, for example, local disk or third-party storage systems.

11.4 Prerequisites

Table 11-1 lists the prerequisites for using File SRM.

Table 11-1 Prerequisites

Prerequisite	Description
File SRM license	You must have a valid File SRM license installed on your Operations Manager. If you do not have an File SRM license, contact your sales representative. After you install the File SRM license, the Quotas tab in the Operations Manager is renamed “File SRM” and all of the FSRM features become available.
Connection to TCP/IP network	All FSRM hosts must be connected to a TCP/IP network that is either known to or discoverable by the Operations Manager. The hosts must be connected to the network through an Ethernet port and must have a valid IP address.
Host Agent software	Each workstation from which FSRM data is collected must have the Host Agent installed. Version 2.6 or later is recommended.
Visible directory paths	All directory paths to be monitored must be visible to the host agent. For example, to enable FSRM monitoring, the host agent must mount a system share (volume or qtree) using NFS or CIFS, or the host agent must use a LUN on the system.

To set up and configure FSRM, complete the following steps:

1. Identify FSRM host agents.
2. Add new host agents manually if they have not been discovered.
3. Set up Host Agent administration access on the hosts to be monitored. You can verify the host administration access by checking the SRM Summary window
4. Add paths.
5. Set up path-walk schedules.

Note: To use host agents for FSRM, you must set your login to admin.

11.5 Quick reference for FSRM tasks

The following types of tasks are included in this quick reference for Operations Manager:

- ▶ Where to find the SRM summary window
- ▶ Host Agent management tasks
- ▶ Path management tasks
- ▶ SRM report tasks“

Where to find the SRM Summary window

You can access the SRM Summary window from the Group Status tab by clicking the SRM link.

11.5.1 Host Agent management tasks

This section discusses the most common Host Agent management tasks.

Add a new host agent

To add a new host agent, do these steps:

1. Click the **Group Status** tab.
2. Click the **File SRM** menu, select **SRM Summary** from the Report drop-down list, and then select the **Add/Edit Hosts** link, as shown in Figure 11-4.

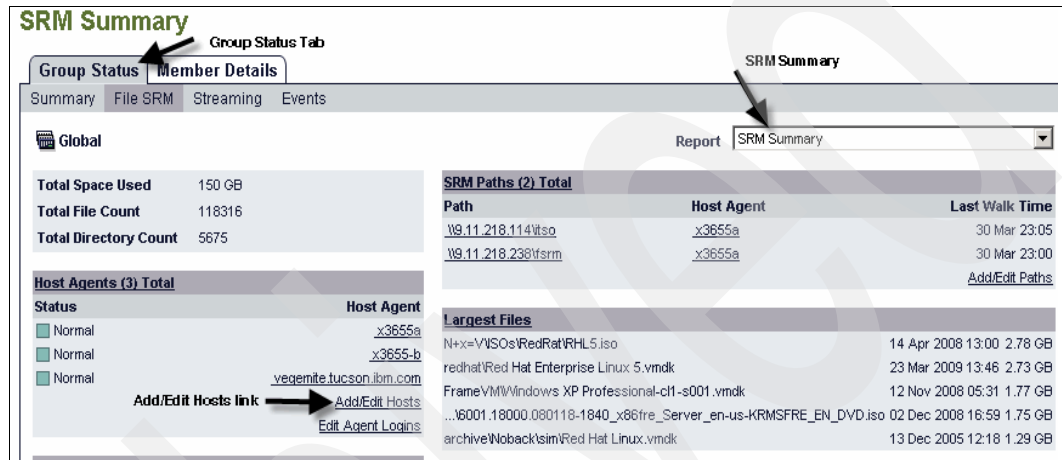


Figure 11-4 SRM Summary window

Edit host agent settings (including passwords)

To edit a host agent's settings, including passwords, do either one of the following:

1. Click the **Group Status** tab, click the **File SRM** menu, select **SRM Summary** from the Report drop-down list, select the **Add/Edit Hosts** link, select the host agent name, and click **Edit**, as shown in Figure 11-5.

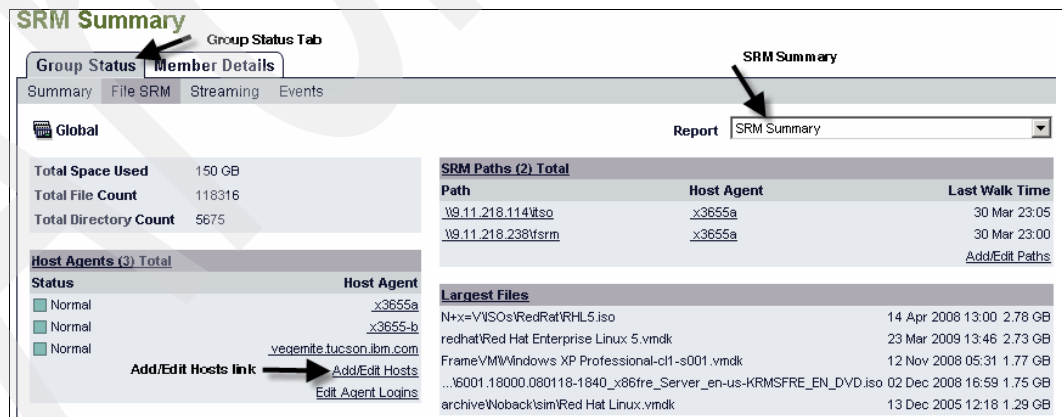


Figure 11-5 Edit host agent settings

2. Click the **Group Status** tab, click the **File SRM** menu, select **SRM Summary** from the Report drop-down list, select the **Add/Edit Hosts** link, and click Host Name.

Figure 11-6 shows the window where you would add the new host agent.

Figure 11-6 Add or Edit Host Agents window

Edit host agent properties

To edit the host agent's properties (for example, monitoring and management of API passwords, the HTTP port, and HTTPS settings), do these steps:

1. Click the **Group Status** tab, click the **File SRM** menu, select **SRM Summary** from the Report drop-down list, and select the **Add/Edit Hosts** link, as shown in Figure 11-7.

Status	Host Agent
Normal	x3655a
Normal	x3655-b
Normal	vegemite.tucson.ibm.com

Path	Host Agent	Last Walk Time
\\9.11.218.114\\iso	x3655a	30 Mar 23:05
\\9.11.218.238\\srm	x3655a	30 Mar 23:00

File Name	Size	Last Walk Time
N+x=VISOs\\RedHat\\RHL5.iso	2.78 GB	14 Apr 2008 13:00
redhat\\Red Hat Enterprise Linux 5.vmdk	2.73 GB	23 Mar 2009 13:46
FrameVM\\Windows XP Professional-cl1-s001.vmdk	1.77 GB	12 Nov 2008 05:31
...\\6001.18000.080118-1840_x86fre_Server_en-us-KRMSFRE_EN_DVD.iso	1.75 GB	02 Dec 2008 16:59
archive\\noback\\sim\\Red Hat Linux.vmdk	1.29 GB	13 Dec 2005 12:18

Figure 11-7 SRM Summary window

2. In the Add or Edit Host Agents window, shown in Figure 11-8 on page 277, select the host agent name and click **Edit**. The Edit Host Agent Settings window for that host agent should appear, as shown in Figure 11-9 on page 277; change the properties as needed.

Add or Edit Host Agents

Add a New Host Agent

Host Name:

Login: (Global De

Password:

Administration Transport: (Global De

Administration Port: (leave empty to use default)

Host Agent Name	Edit
vegemite	Edit
x3655-b	Edit Host Name Edit
x3655a	Edit

Figure 11-8 Add or Edit Host Agents window

Edit Host Agent Settings

Edit x3655a

Primary IP Address: 9.11.218.144

Login: admin (Monitoring and Management

Password: (leave empty to use default)

Administration Transport: (Global Default)

Administration Port: (leave empty to use default) 4092

Agent CIFS Account: (user name for CIFS access, leave empty to use default) Administrator

Agent CIFS Password: (password for CIFS access, leave empty to use default)

Owner Email:

Owner Name:

Resource Tag:

Figure 11-9 Edit Host Agents Settings window

Configure host agent administration access

To configure the host agent's administration access, do either of the following:

1. Go to the Edit Agent Logins window by clicking the **Group Status** tab, clicking the **File SRM** menu, selecting **SRM Summary** from the Report drop-down list, and selecting the Edit Agent Logins, as shown in Figure 11-10. You should get the window shown in Figure 11-11. Make the necessary changes.

Figure 11-10 SRM Summary window

Figure 11-11 Edit Logins for Host Agents window

2. Create a global default using the Host Agent Options window. You can access this window by selecting, under the Control Center tab, **Setup** → **Options**, and then clicking the **Host Agent** link, as shown in Figure 11-12 on page 279. The window appears as shown in Figure 11-13 on page 279. Make any necessary changes.

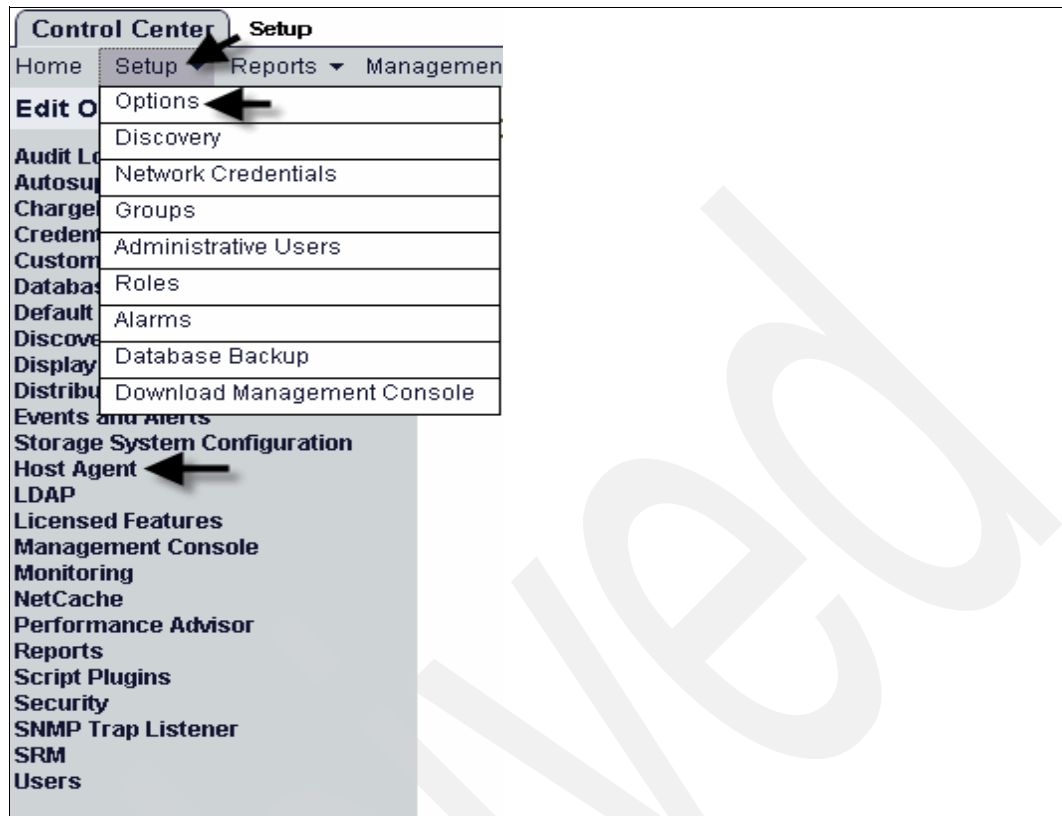


Figure 11-12 Edit Options window

The screenshot shows the 'Options' window with the 'Host Agent' section selected. The section contains the following fields and values:

Field	Value
Host Agent Login	admin (Monitoring and Management)
Host Agent Monitoring Password
Host Agent Management Password
Host Agent Administration Transport	HTTP
Host Agent Administration Port	4092
Host Agent CIFS Account (default account to use for CIFS filewalks)	
Host Agent CIFS Password (default password to use for CIFS filewalks)

Figure 11-13 Options window

Group SRM hosts

To group SRM hosts, go to the Host Agents, SRM view by clicking the **Group Status** tab, clicking the **File SRM** menu, selecting **SRM Summary** from the Report drop-down list, and selecting **Host Agents (3) Total**, as shown in Figure 11-14. In Figure 11-15, select the hosts to group and click **Add to Group**.

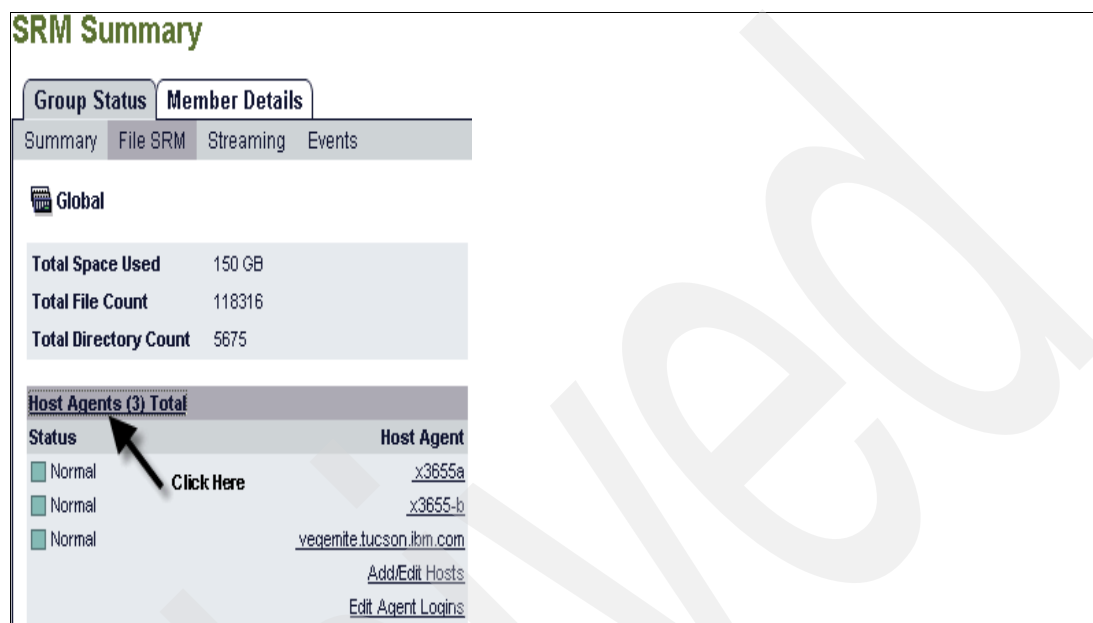


Figure 11-14 SRM Summary window

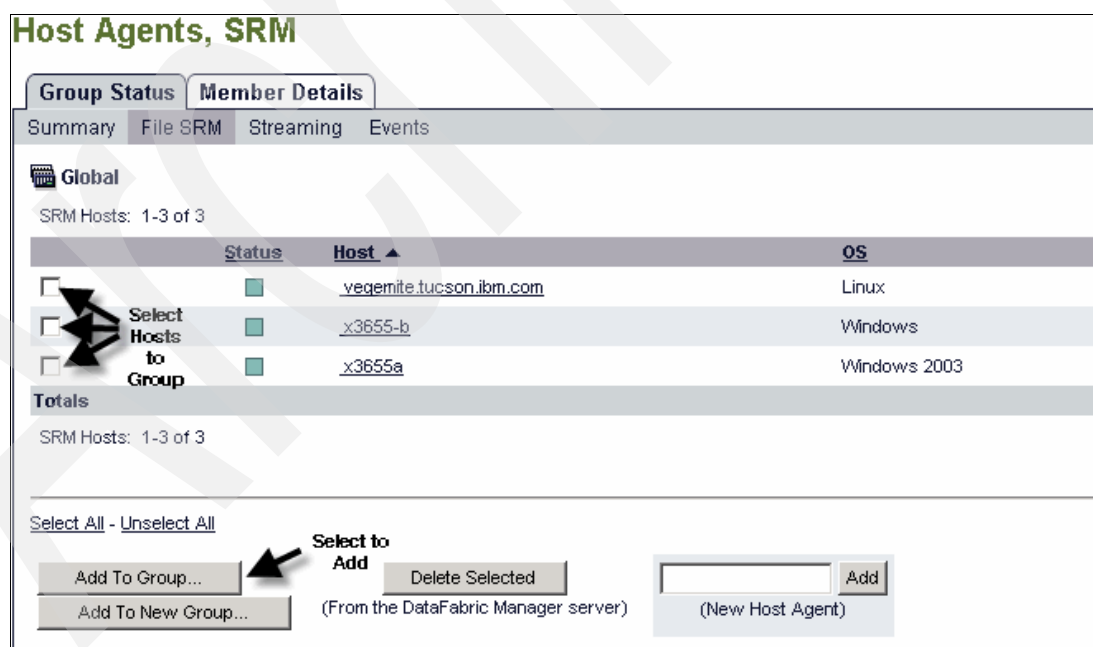


Figure 11-15 Add to Group

List available host agents

To list the available host agents, go to the Add or Edit Host Agents window by clicking the **Group Status** tab, clicking the **File SRM** menu, selecting **SRM Summary** from the Report drop-down list, and selecting the **Add/Edit Hosts** link, as shown in Figure 11-16.

SRM Summary

Group Status Tab: Group Status, Member Details

Summary | File SRM | Streaming | Events

Global

Total Space Used: 150 GB
Total File Count: 118316
Total Directory Count: 5675

Host Agents (3) Total

Status	Host Agent
Normal	x3655a
Normal	x3655-b
Normal	vegemite.tucson.ibm.com

Add/Edit Hosts link: Add/Edit Hosts, Edit Agent Logins

Report: SRM Summary

SRM Paths (2) Total

Path	Host Agent	Last Walk Time
\\9.11.218.114\\iso	x3655a	30 Mar 23:05
\\9.11.218.238\\fsrm	x3655a	30 Mar 23:00

[Add/Edit Paths](#)

Largest Files

File Name	Last Walk Time	Size
N\\x=VISOs\\RedHat\\RHL5.iso	14 Apr 2008 13:00	2.78 GB
redhat\\Red Hat Enterprise Linux: 5.vmdk	23 Mar 2009 13:46	2.73 GB
FrameVM\\Windows XP Professional-ch1-s001.vmdk	12 Nov 2008 05:31	1.77 GB
...\\6001.18000.080118-1840_x86fre_Server_en-us-KRMSFRE_EN_DVD.iso	02 Dec 2008 16:59	1.75 GB
archive\\Noback\\sim\\Red Hat Linux.vmdk	13 Dec 2005 12:18	1.29 GB

Figure 11-16 SRM Summary window

Review host agent global settings

To review a host agent's global settings, go to the Host Agent Options window by selecting, in the Control Center tab, **Setup** → **Options**, and then selecting the **Host Agent** link, as shown in Figure 11-17. The window shown in Figure 11-18 on page 282 should appear.

Control Center: Setup

Home | Setup | Reports | Management

Edit Options

Audit Log

Autosup

Charge

Credent

Custom

Databas

Default

Discover

Display

Distribu

Events and Alerts

Storage System Configuration

Host Agent

LDAP

Licensed Features

Management Console

Monitoring

NetCache

Performance Advisor

Reports

Script Plugins

Security

SNMP Trap Listener

SRM

Users

Figure 11-17 Edit Options window

Options

Host Agent

Host Agent Login

admin (Monitoring and Management)

Host Agent Monitoring Password

.....

Host Agent Management Password

.....

Host Agent Administration Transport

HTTP

Host Agent Administration Port

4092

Host Agent CIFS Account

(default account to use for CIFS filewalks)

Host Agent CIFS Password

(default password to use for CIFS filewalks)

.....

Figure 11-18 Options window

Obtain SRM Host Agent information

To obtain SRM Host Agent information, go to the Host Agent Details window by clicking the **Group Status** tab, clicking the **File SRM** menu, selecting **Host Agent, SRM** from the Report drop-down list, and selecting the host agent name, as shown in Figure 11-19.

Host Agents, SRM

Group Status

Member Details

Select Host Agent, SRM

Summary

File SRM

Streaming

Events

Global

SRM Hosts: 1-3 of 3

Report

Host Agents, SRM

	Status	Host	OS	SRM Paths	Total KBytes	Total Files
<input type="checkbox"/>		vegemite.tucson.ibm.com	Linux	0		
<input type="checkbox"/>		x3655-b	Windows	0		
<input type="checkbox"/>		x3655a	Windows 2003	2	150 GB	118326
Totals				2	150 GB	118326

Figure 11-19 Host Agents, SRM window

Change the SRM Host Agent monitoring interval

To change the SRM Host Agent monitoring interval, go to the Monitoring Options window by selecting, in the Control Center tab, **Setup** → **Options**, click the **Monitoring** link, and then click the **Agent monitoring interval** option, as shown in Figure 11-20. Make any necessary changes.

Control Center

Home

Setup

Reports

Management

Help

Edit Options

Discovery

Network Credentials

Groups

Administrative Users

Roles

Alarms

Database Backup

Download Management Console

Events and Alerts

Storage System Configuration

Host Agent

LDAP

Licensed Features

Management Console

Monitoring

NetCache

Performance Advisor

Reports

Options

Monitoring

Agent Monitoring Interval	2 minutes
File SRM Monitoring Interval	5 minutes
Cluster Failover Monitoring Interval	5 minutes
Config Conformance Monitoring Interval	4 hours
CPU Monitoring Interval	5 minutes
Disk Free Space Monitoring Interval	30 minutes
Disk Monitoring Interval	4 hours
Environmental Monitoring Interval	5 minutes
File System Monitoring Interval	15 minutes

Figure 11-20 Agent monitoring interval

Disable host agent discovery

To disable host agent discovery, go to the Host Agent Discovery option in the Options window by selecting, in the Control Center tab, **Setup** → **Options**, and then click the **Discovery** link, as shown in Figure 11-21. Make any necessary changes.

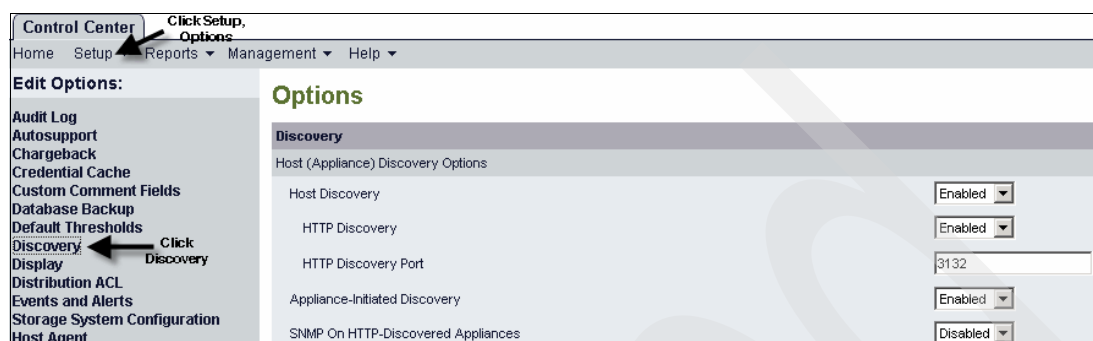


Figure 11-21 Discovery

Delete a host agent

To delete a host agent, do these steps:

1. Go to the Add or Edit Host Agents window by clicking, in the Group Status tab, **File SRM**, select **SRM Summary** in the Report drop-down menu, and then click **Host Agents (3) Total**, as shown in Figure 11-22.

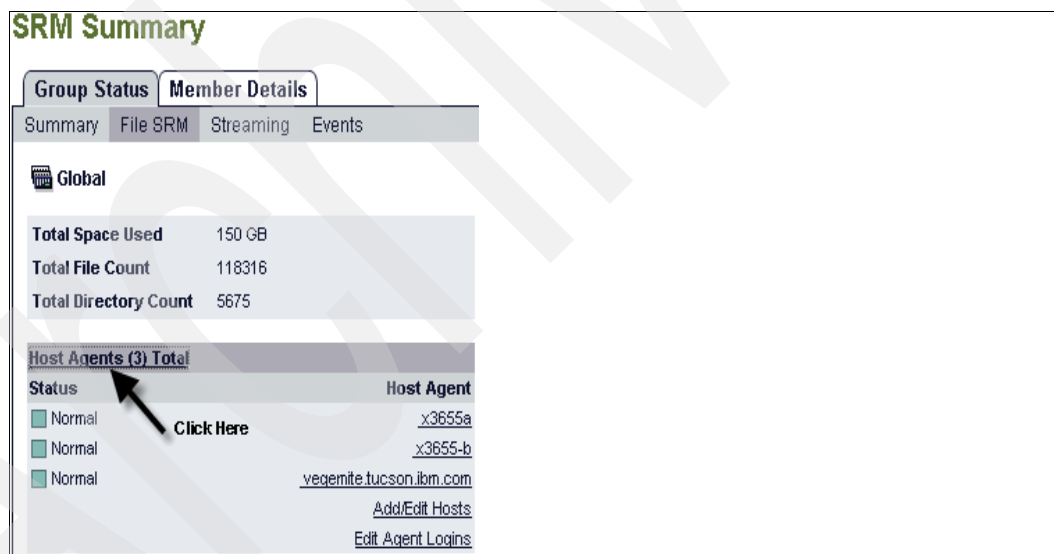


Figure 11-22 SRM Summary window

2. In the Host Agents, SRM window, check the check boxes of the host agents that you want to delete and click **Delete Selected**, as shown in Figure 11-23.

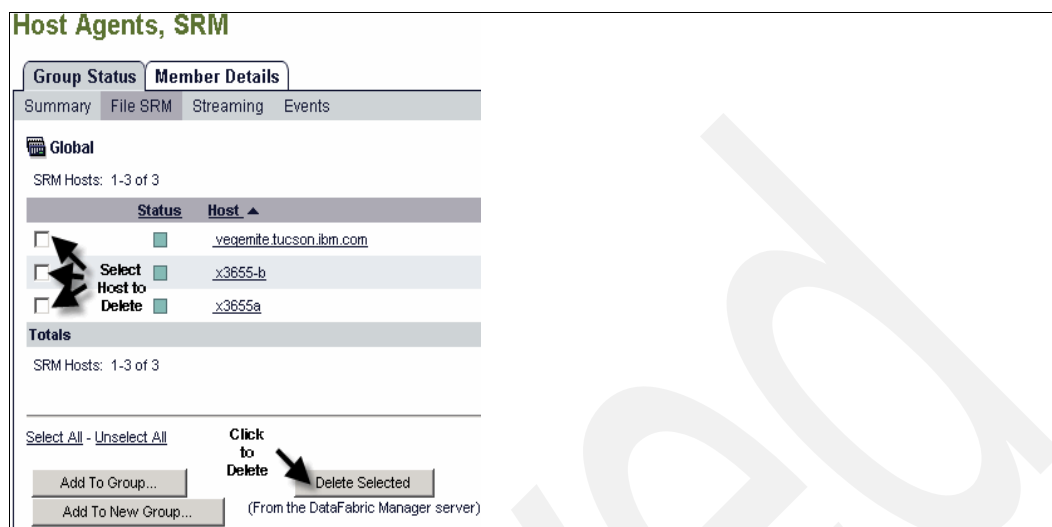


Figure 11-23 Host Agents, SRM window

11.5.2 Path management tasks

This section lists the common path management tasks and the location of the Operations Manager user interface windows that enables you to complete them.

Automapping requirements

You can automatically create a new path for an object using the “Create new SRM Path for this object” link in the details window for the object. To use this feature, you must ensure the following:

- ▶ The host agent is set up and properly configured.
- ▶ The host agent’s passwords match those set in Operations Manager.
- ▶ The host agent has access to the volume, qtree, or LUN:
 - If the host agent is a Windows host, you must ensure that the CIFS passwords match.
 - If the object is a LUN on a Windows host, SnapDrive must be installed and the LUN must be managed by SnapDrive.
 - If the host agent is a UNIX host, then the volume or qtree must be NFS mounted.
 - If the object is a LUN on a UNIX host, the LUN must be formatted and mounted directly into the file system (volume managers are not supported).
- ▶ The Host Agent Login and Management Password are set correctly.

You can also manually map SRM paths to volumes, and LUNs.

Add paths

To add paths, you can either use the automapping feature or manually add an SRM path.

To use the automapping feature, use the **Create new SRM Path for this object** link, which can be found by selecting, in the Member Details tab, **File Systems** or **LUNs**, and then selecting the correct file system, as shown in Figure 11-24 on page 285. In the window shown in Figure 11-25, click **Create new SRM Path for this Volume**. Make any necessary changes.

File Systems, All

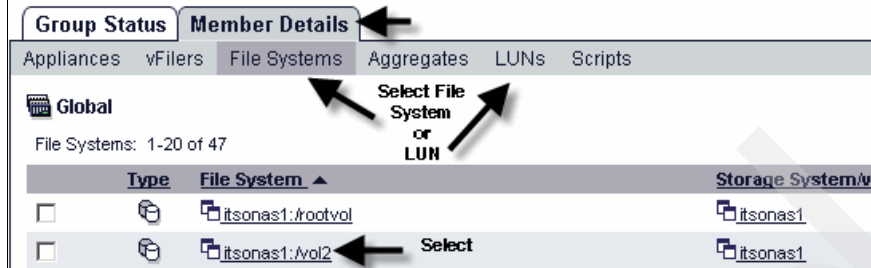


Figure 11-24 File Systems, All window

Volume Details

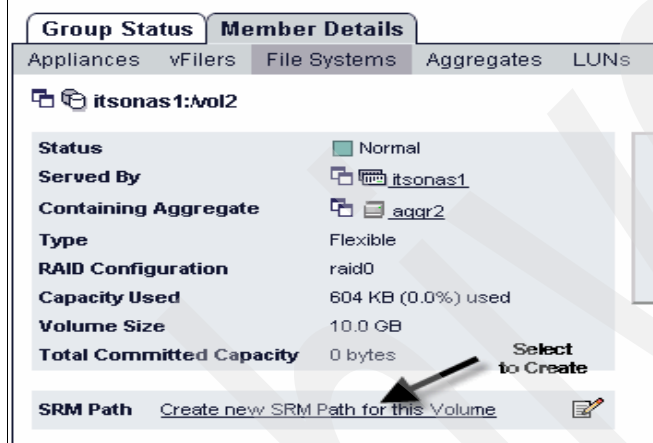


Figure 11-25 Volume Details window

To manually add an SRM path, go to the Add or Edit SRM Paths window, which can be found by clicking, in the Group Status tab, **File SRM**, selecting **SRM Summary** in the Report drop-down menu, and then clicking the **Add/Edit Paths** link, as shown in Figure 11-26.




Figure 11-26 SRM Summary

All directory paths to be monitored must be visible to the host agent. Before setting up FSRM paths and schedules, you must enable administrative access to your host agents.

About path walks

A path walk is the process of recursively examining a directory path for file level statistics. Path walks are scheduled using Operations Manager and executed by the Host Agent software. The Host Agent software scans all subdirectories of the specified directory path and gathers per-file and per-directory data.

SRM path walk recommendations

SRM path walks can consume considerable resources on the SRM host agent and on Operations Manager. Therefore, schedule your SRM path walks to occur during off-peak hours. Also, do not schedule multiple, simultaneous SRM path walks on the same SRM host agent.

Create path walk schedules

To create path walk schedules, go to the Edit SRM Path Walk Schedules window, which can be found by selecting, in the Group Status tab, **File SRM**, selecting **SRM Summary** from the Report drop-down menu, and then clicking the **Add/Edit Schedules** link, as shown in Figure 11-27. Make any necessary changes.

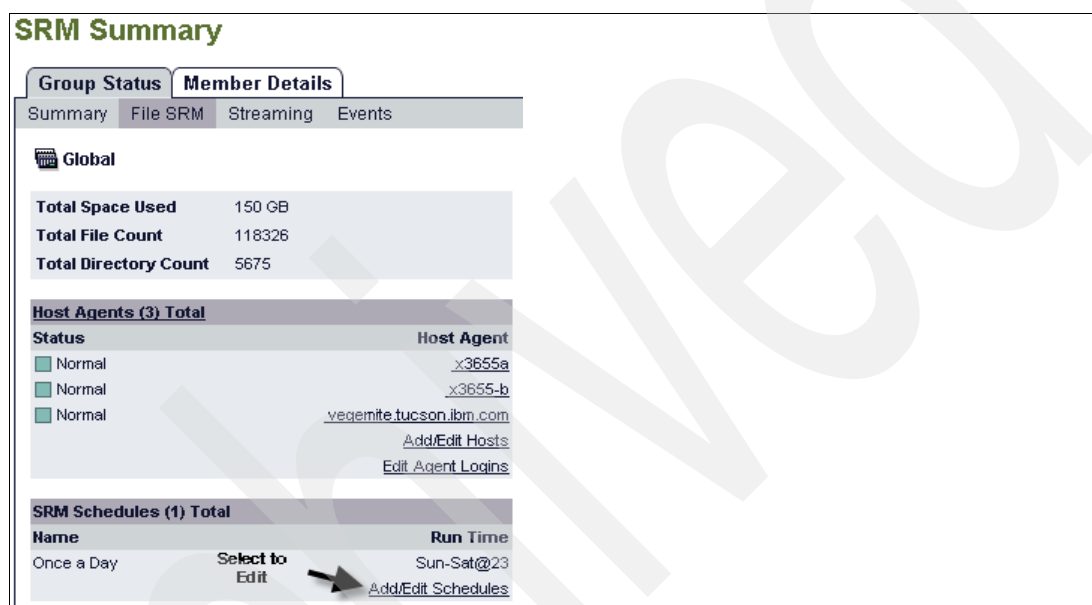


Figure 11-27 SRM Summary window

Specify path walk times

To specify path walk times, do these steps:

1. Select **File SRM** in the Group Status tab, select **SRM Summary** in the Report drop-down menu, and then select the **Add/Edit Schedules** link. The window shown in Figure 11-28 should appear.



Figure 11-28 Edit SRM Path Walk Schedules window

2. Select the schedule name (in our case, **Once a Day**). The window shown in Figure 11-29 on page 287 should appear.

SRM Path Walk Schedule Details

Schedule: Once a Day

Schedule Name:

Once a Day

☐ Use as Default for New SRM Paths

Used By These SRM Paths:

[x3655a\W9.11.218.114\its](#)
[x3655a\W9.11.218.238\fsrm](#)

Walk Times:

Sun-Sat @ 23

Add Walk Times

Select to Add

Figure 11-29 SRM Path Walk Schedule Details window

3. Select the **Add Walk Times** link. The window shown in should appear. Make any necessary changes in this window.

SRM Path Walk Schedule Times

Edit Schedule: Once a Day

Schedule Name:

Once a Day

☐ Use as Default for New SRM Paths

Used By These SRM Paths:

[x3655a\W9.11.218.114\its](#)
[x3655a\W9.11.218.238\fsrm](#)

Walk Times:

All Days

Weekdays

Weekends

Sunday

Monday

Tuesday

Wednesday

Thursday

Friday

Saturday

at

00:00

01:00

02:00

03:00

04:00

05:00

06:00

07:00

08:00

09:00

10:00

11:00

12:00

13:00

14:00

15:00

16:00

17:00

18:00

19:00

20:00

21:00

22:00

23:00

Figure 11-30 SRM Path Walk Schedule Times window

Manually start or stop an SRM path walk

To manually start or stop an SRM path walk, go to the SRM Path Details window by selecting the SRM path name, as shown in Figure 11-31.

SRM Paths (2) Total			
Path		Host Agent	Last Walk Time
W9.11.218.114\its	Select a Path	x3655a	31 Mar 23:05
W9.11.218.238\fsrm		x3655a	31 Mar 23:00
Add/Edit Paths			

Figure 11-31 SRM Paths (2) Total window

Chapter 11. File Storage Resource Manager 287

Click **Start**, as shown in Figure 11-32. If an SRM path walk is in progress, the Start button changes to a Stop button.



Figure 11-32 SRM Path Details window

Review SRM path details

to review SRM path details, go to the SRM Path Details window by clicking the SRM path name.

Edit SRM paths

To edit SRM paths, go to the Add or Edit SRM Paths window by selecting, in the Group Status tab, **File SRM**, clicking **SRM Summary** in the Report drop-down menu, and clicking the **Add/Edit Paths** link, as shown in Figure 11-33. Make any necessary changes in the window that appears.

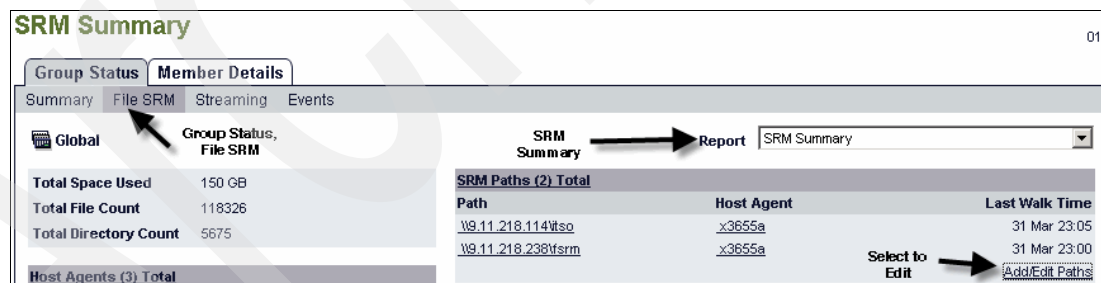


Figure 11-33 SRM Summary window

Review SRM path walk schedule details

To review SRM path walk schedule details, go to the SRM Path Walk Schedule Details window by clicking the schedule name.

Delete SRM paths

To delete SRM paths, go to the Add or Edit SRM Paths window by selecting, in the Group Status tab, **File SRM**, clicking **SRM Summary** in the Report drop-down menu, and clicking the **Add/Edit Paths** link. The window shown in Figure 11-34 on page 289 should appear. In this window, check the check boxes for the paths you want to delete and click **Delete Selected**.

Add or Edit SRM Paths 01 Apr 10:54

Add a New SRM Path

SRM Host: [Add Host](#)

Path:

Schedule: [Add A Schedule](#)

[Add SRM Path](#)

Path	Host	Last Walk Time	Schedule	Edit	Delete
W9.11.218.114\lso	x3655a	31 Mar 23:05	Once a Day	Edit	<input type="checkbox"/>
W9.11.218.238\lfrsm	x3655a	31 Mar 23:00	Once a Day	Edit	<input type="checkbox"/>

[Delete Selected](#)

Figure 11-34 Add or Edit SRM Paths window

Delete SRM path walk schedules

To delete SRM path walk schedules, go to the Edit SRM Path Walk Schedules window by selecting, in the Group Status tab, **File SRM**, clicking **SRM Summary** in the Report drop-down menu, and clicking the **Add/Edit Schedules** link. The window shown in Figure 11-35 should appear. In this window, check the check boxes of the schedules you want to delete and click **Delete Selected**.

Edit SRM Path Walk Schedules 01 Apr 10:58

Add a New Schedule

Schedule Name:

Schedule Template: [Add](#)

Schedule Name	Walk Times	Used By	Delete
Once a Day	Sun-Sat@23	W9.11.218.114\lso W9.11.218.238\lfrsm	<input type="checkbox"/>

[Delete Selected](#)

Figure 11-35 Edit SRM Path Walk Schedules window

Viewing File SRM reports

Operations Manager provides three levels of file system statistics:

- ▶ Consolidated data gathered from all paths
- ▶ SRM path-specific data
 - This is a summary of the data for all directories in the specified path.
- ▶ Directory-level data
 - This contains the data for the specified directory only.

If the Host Agent is installed on UNIX, File SRM feature tracks only the files having a file extension that exactly matches the file type specification in Operations Manager. For example, files that end in .JPG will not match the .jpg file type if the host agent is on UNIX, even though they would match if the agent was on Windows. Running the host agent on Windows avoids this problem.

SRM report tasks

This section shows the common SRM report tasks and the location of the Operations Manager user interface window that enables you to complete them.

View SRM reports

To view SRM reports, go to the report's window by selecting, in the Group Status tab, **File SRM**, clicking **SRM Summary** in the Report drop-down menu, and then selecting the report's name in the Report drop-down list, as shown in Figure 11-36.

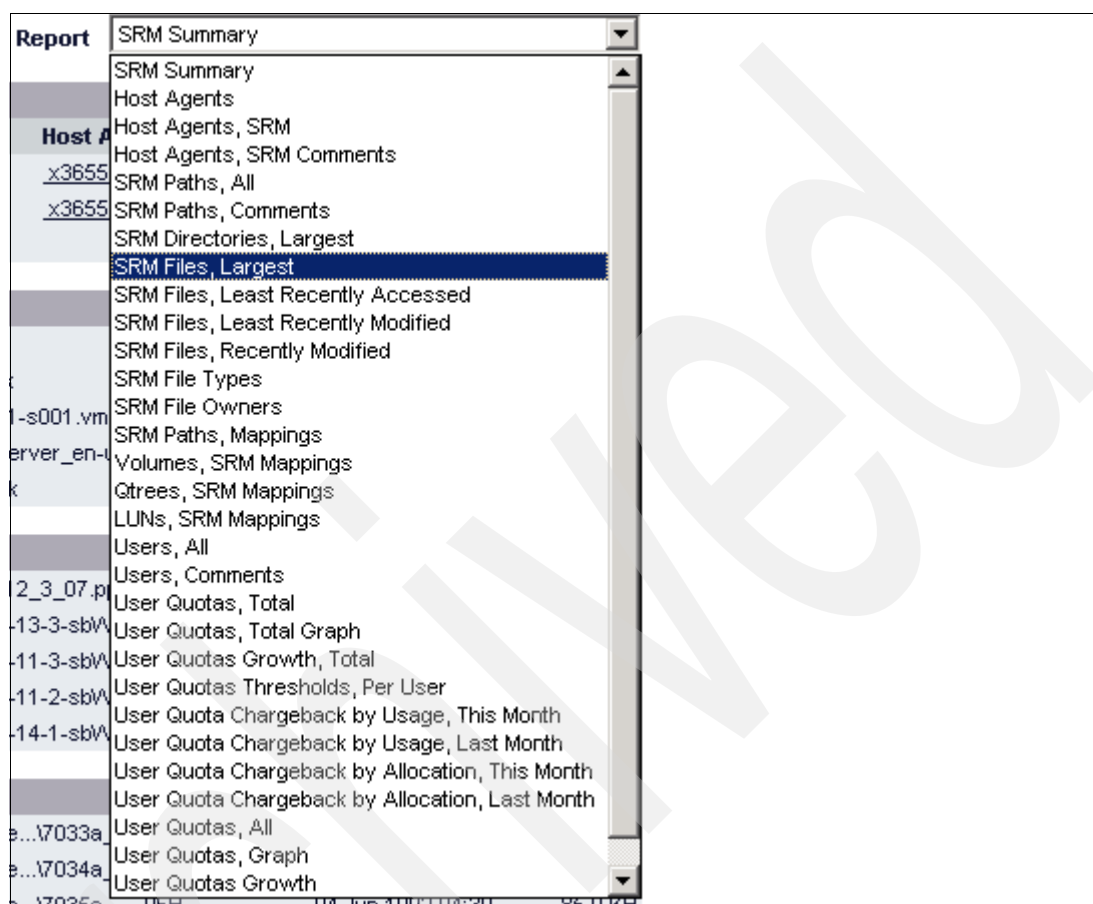


Figure 11-36 Report window

11.6 Identifying SRM host agents

If you have installed a File SRM license on the workstation, Operations Manager automatically discovers all hosts that it can communicate with. Communication between the host agents and Operations Manager takes place over HTTP or HTTPS (port 4092 or port 4093, respectively).

Identifying SRM host agents

To determine which SRM host agents have been discovered, select, in the SRM Summary window, **Host Agents, SRM** from the Report drop-down list. The Host Agents, SRM window is displayed (Figure 11-37 on page 291).

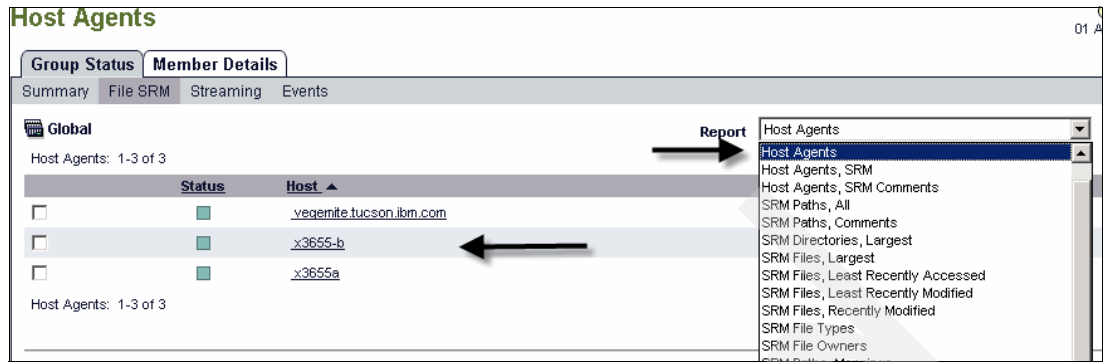


Figure 11-37 Host Agents window

Disabling host agent discovery

If you do not want Operations Manager to automatically discover host agents, disable the Host Agent Discovery by completing these steps:

1. From any window, select **Options** from the Setup drop-down menu. The Options window is displayed, as shown in Figure 11-38.

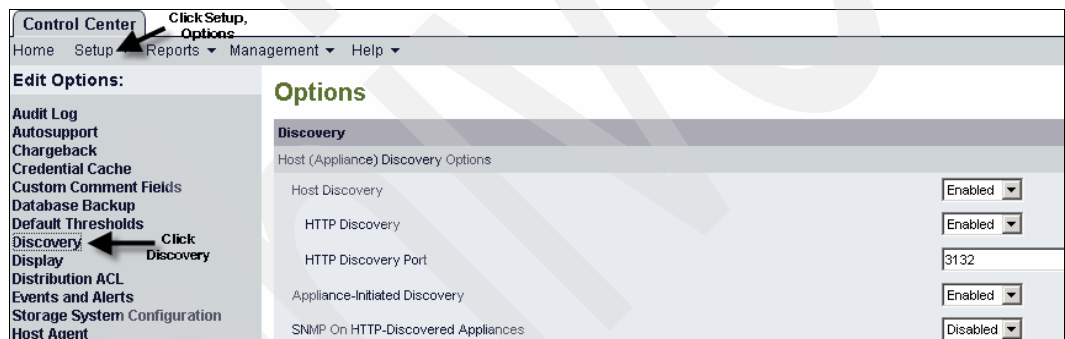


Figure 11-38 Options window

2. Click **Discovery** from the Edit Options pane on the left. The Discovery options window is displayed.
3. In the (Host) Appliance Discovery Options section, disable the Host Agent Discovery option.

11.7 Managing host agents

As a reminder, host agents are not needed on every server or workstation on your network. It is only required for servers you want to monitor for either SAN HBA information or file meta data information provided by CIFS or NFS shares.

Adding host agents

To add host agents, complete these steps:

1. From the SRM Summary window, click the **Add/Edit Hosts** link in the Host Agents Total section. The Add or Edit Host Agents window is displayed, as shown in Figure 11-39.

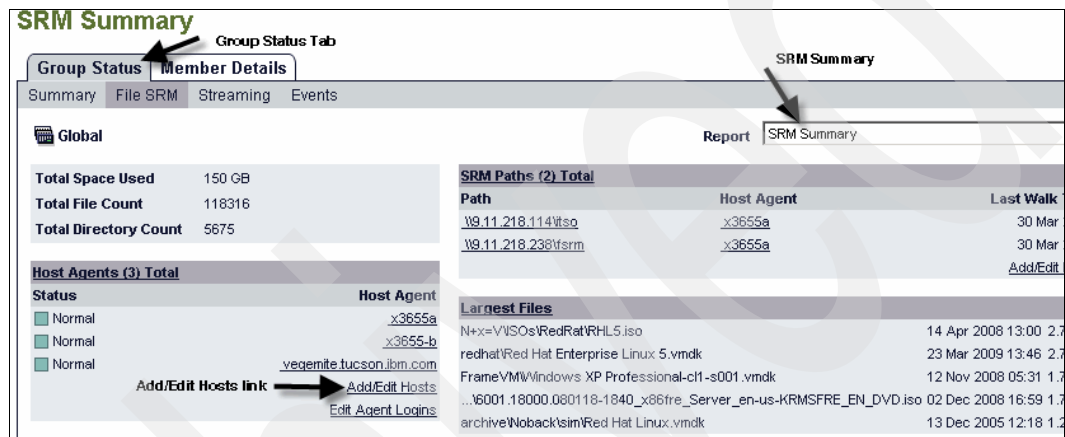


Figure 11-39 SRM Summary window

2. Enter the required information in the Add a New Host Agent section and click **Add**, as shown in Figure 11-40.




Figure 11-40 Add or Edit Host Agents window

Editing settings for a single host agent

To change the password or other settings for a single host agent and override those specified in the global Host Agent Options, complete these steps:

1. From the SRM Summary window, click the host agent name in the Host Agents Total section, as shown in Figure 11-41 on page 293. The Host Agent Details window is displayed.

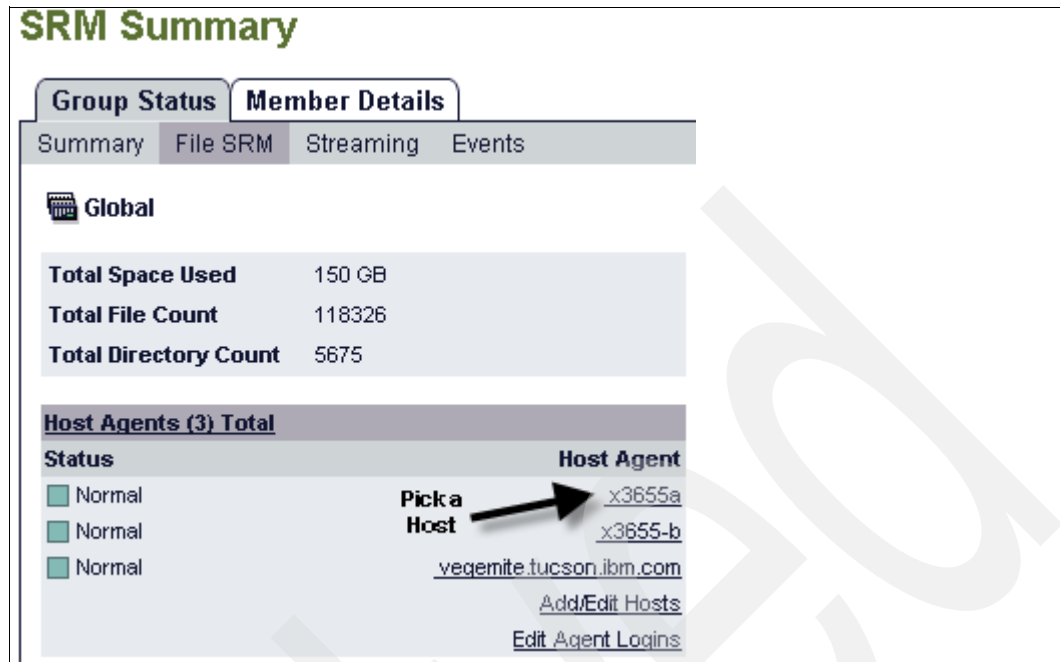


Figure 11-41 SRM Summary window

- Click the **Edit Settings** link in the Appliance Tools list (at the lower left of Operations Manager), as shown in Figure 11-42. The Edit Host Agent Settings window is displayed, as shown in Figure 11-43 on page 294. Modify the fields, as needed, and click **Update**.

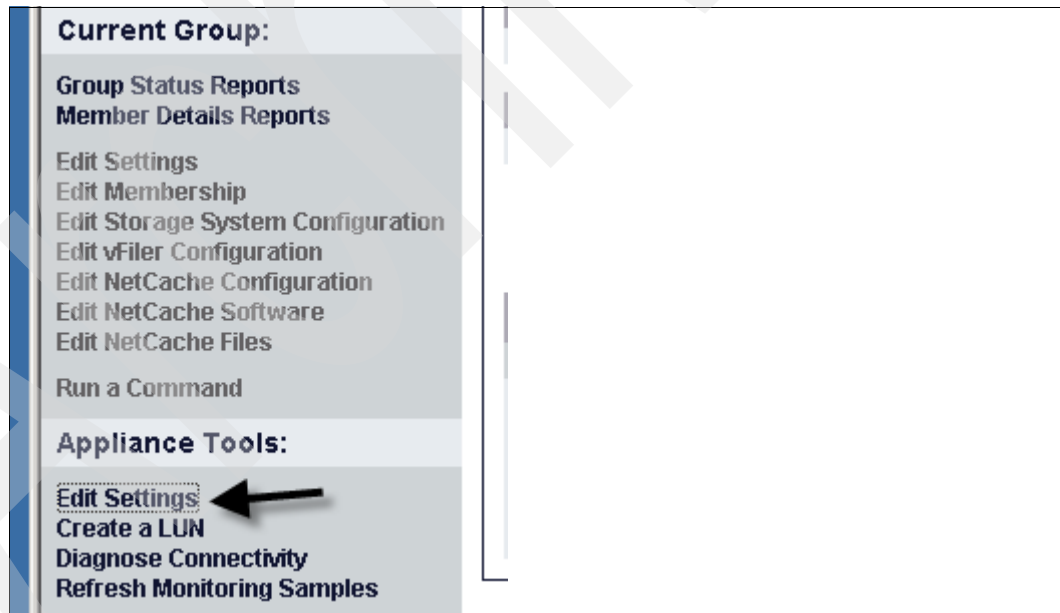


Figure 11-42 Edit Settings

Edit Host Agent Settings

Edit x3655a

Primary IP Address	9.11.218.144
Login	admin (Monitoring and Management)
Password (leave empty to use default)
Administration Transport	(Global Default) ▾
Administration Port (leave empty to use default)	4092
Agent CIFS Account (user name for CIFS access, leave empty to use default)	Administrator
Agent CIFS Password (password for CIFS access, leave empty to use default)
Owner Email	
Owner Name	
Resource Tag	

Figure 11-43 Edit Host Agent Settings window

Editing settings for multiple host agents

To change the password or other settings for multiple host agents, complete these steps:

1. From the SRM Summary window, click the **Edit Agent Logins** link in the Host Agents Total section, as shown in Figure 11-44. The Edit Logins for Host Agents window is displayed.

SRM Summary

Group Status Tab

Group Status | Member Details

Summary | File SRM | Streaming | Events

Global

Report SRM Summary

Total Space Used 150 GB Total File Count 118316 Total Directory Count 5675	SRM Paths (2) Total <table> <tr> <th>Path</th> <th>Host Agent</th> <th>Last Walk</th> </tr> <tr> <td>\9.11.218.114\iso</td> <td>x3655a</td> <td>30 Mar</td> </tr> <tr> <td>\9.11.218.238\fsrm</td> <td>x3655a</td> <td>30 Mar</td> </tr> </table>	Path	Host Agent	Last Walk	\9.11.218.114\iso	x3655a	30 Mar	\9.11.218.238\fsrm	x3655a	30 Mar
Path	Host Agent	Last Walk								
\9.11.218.114\iso	x3655a	30 Mar								
\9.11.218.238\fsrm	x3655a	30 Mar								

Host Agents (3) Total

Status	Host Agent
Normal	x3655a
Normal	x3655-b
Normal	vegemite.tucson.ibm.com

Add/Edit Hosts link → [Add/Edit Hosts](#)
[Edit Agent Logins](#)

Largest Files

File Name	Size	Last Modified
N+=V\ISOs\RedHat\RHLS.iso	14 Apr 2008 13:00	2.7
redhat\Red Hat Enterprise Linux 5.vmdk	23 Mar 2009 13:46	2.7
FrameVM\Windows XP Professional-cl1-s001.vmdk	12 Nov 2008 05:31	1.7
...6001.18000.080118-1840_x86fre_Server_en-us-KRMSFRE_EN_DVD.iso	02 Dec 2008 16:59	1.7
archive\Woback\sim\Red Hat Linux.vmdk	13 Dec 2005 12:18	1.2

Figure 11-44 SRM Summary window

2. Click the boxes of the host agents you want to modify, as shown in Figure 11-45 on page 295.

Edit Logins for Host Agents 01 Apr 11:50

Edit Agent Logins

Login: (Global Default) ▼

Password (leave empty to use default):

Agent CIFS Account (user name for CIFS access, leave empty to use default):

Agent CIFS Password (password for CIFS access, leave empty to use default):

Change login settings for multiple agents by selecting the checkboxes below.

Host Agent	Agent OS	Login	CIFS Account
<input type="checkbox"/> vegemite.tucson.ibm.com	Linux	🟡 (default)	N/A (default)
<input type="checkbox"/> x3655-b	Windows	🟡 (default)	🔑 (default)
<input type="checkbox"/> x3655a	Windows 2003	🔑 admin	🔑 Administrator

Select All - Unselect All

Update

Figure 11-45 Edit Logins for Host Agents window

3. Modify the fields, as needed and click **Update**.

Settings specified in the Edit Logins for Host Agents window override the settings specified in the global Host Agent Options section on the Options window.

Editing host agent monitoring intervals

To change the global monitoring interval for host agents, complete these steps:

1. From any window, select **Options** from the Setup menu. The Options window is displayed.
2. Select **Monitoring** from the Edit Options list (in the left pane), as shown in Figure 11-46.

Control Center

Home Setup Reports Management Help

Edit Options

- Options
- Discovery
- Network Credentials
- Groups
- Administrative Users
- Roles
- Alarms
- Database Backup
- Download Management Console
- Storage System Configuration
- Host Agent
- LDAP
- Licensed Features
- Management Console
- Monitoring
- NetCache
- Performance Advisor
- Reports

Monitoring

Host Monitoring Interval	2 minutes
File Monitoring Interval	5 minutes
Cluster Failover Monitoring Interval	5 minutes
Config Conformance Monitoring Interval	4 hours
CPU Monitoring Interval	5 minutes
Disk Free Space Monitoring Interval	30 minutes
Disk Monitoring Interval	4 hours
Environmental Monitoring Interval	5 minutes
File System Monitoring Interval	15 minutes

Figure 11-46 Options window

3. Enter (or modify) a value for SRM Host Monitoring Interval and click **Update**.

Deleting host agents

To delete one or more host agents, complete the following steps.

1. From the SRM Summary window, click the **Add/Edit Hosts** link in the Host Agents Total section. The Add or Edit Host Agents window is displayed.

2. Click the check box(es) of the host agent(s) you want to delete and then click **Delete Selected**, as shown in Figure 11-47.

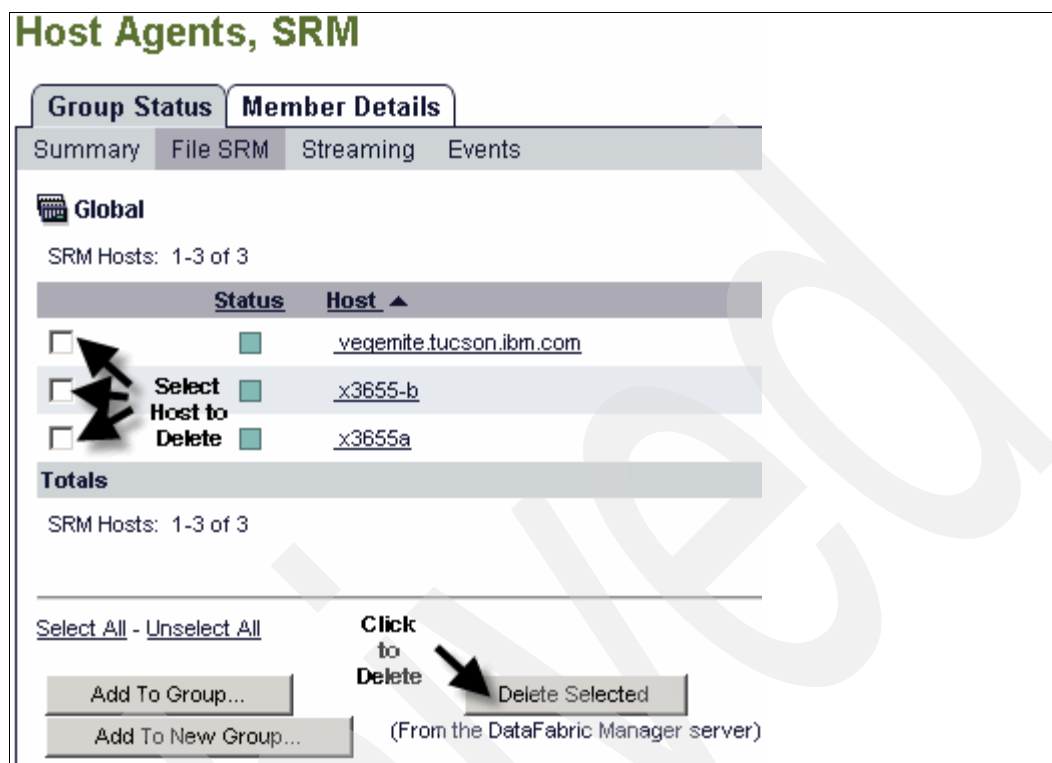


Figure 11-47 Host Agents, SRM window

11.8 Configuring host agent administration access

An important step in configuring host agent access is password administration.

About the Host Agent software passwords

Host agents have two user name and password pairs: one pair for monitoring and one pair for administration tasks.

► Monitoring tasks

The default host agent user name and password allows monitoring only; you cannot perform FSRM functions. The values are as follows:

- User name=guest
- Password=public

Any sessions initiated by the Operations Manager using this user name and password are limited to basic monitoring operations. If you later decide to change the guest password on the host agent, you must also set the same user name and password in Operations Manager, using the Host Agent Monitoring Password option on the Options window by selecting **Setup menu** → **Options** → **Host Agent**, as shown in Figure 11-48 on page 297.



Figure 11-48 Where to find and change monitoring passwords

To enable management, the password must be changed from *public*, as shown in Figure 11-49.

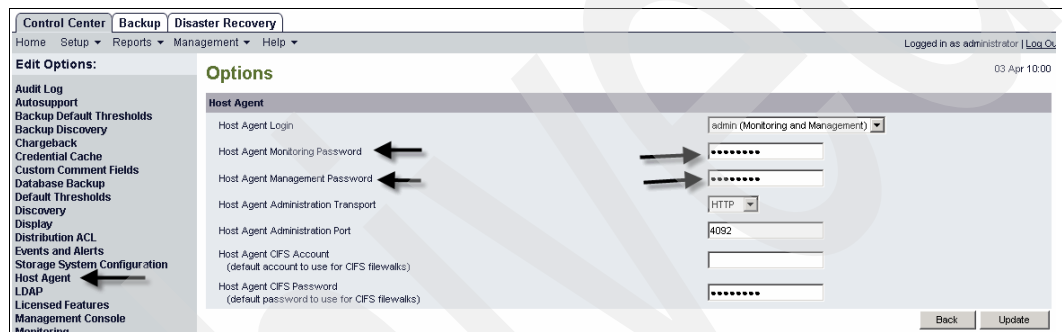


Figure 11-49 Change password from *public*

► Administration tasks

The administration user name and password allows read-write access and is required for FSRM functions. The default values are as follows:

- User name=admin
- Password=userspecified

You specify the password in the Host Agent's configuration UI (<http://name-of-agent:4092/>). This user name and password allows full access to the Host Agent. After setting the administration user name and password in the Host Agent, you must also set the same user name and password in Operations Manager on the Options window by selecting **Setup** → **Options** → **Host Agent**.

See the online help for detailed information about setting passwords.

Administration settings you must configure

You must enable administration access to your host agents before you can use the FSRM feature to gather statistical data.

Global options apply to all affected devices that do not have individual settings specified for them. For example, the Host Agent Login option applies to all host agents. The host agent access and communication options are globally set for all storage systems using the values specified in the Host Agent Options section on the Options window by selecting Setup > Options, as shown in Figure 11-48 on page 297.

Default values are initially supplied for these options. However, you should review and change the default values as necessary.

To enable administrative access, the passwords set in Operations Manager must match those set for the Host Agent software. Table 11-2 describes the options you must set to enable administration access.

Table 11-2 Administrative access

Access type	Operations Manager options	Host Agent software option
Monitoring only	Host Agent Login=guest Host Agent Monitoring Password	Monitoring API Password
Management (required for FSRM)	Management Host Agent Login=admin Host Agent Management Password= <i>your-administration-password</i>	Management API Password

Enabling administration access for one or more host agents

To enable administration access for one or more selected host agents, do these steps:

1. From the SRM Summary window, click **Edit Agent Logins** in the Host Agents Total section. The Edit Logins for Host Agents window is displayed.
2. Select the check boxes of the host agents for which you want to enable administration access, as shown in Figure 11-50.
3. Modify the fields as needed and then click **Update**.

Host Agent	Agent OS	Login	CIFS Account
<input type="checkbox"/> vegemite.tucson.ibm.com	Linux	! (default)	N/A (default)
<input checked="" type="checkbox"/> x3655-b	Windows	! (default)	! (default)
<input checked="" type="checkbox"/> x3655a	Windows 2003	! admin	! Administrator

Figure 11-50 Edit Logins for Host Agents window

Enabling administration access globally for all host agents

To enable administration access globally for all host agents, do these steps:

1. From any Summary window, select **Options** from the Setup drop-down menu. The Options window is displayed.
2. Select **Host Agent** from the Edit Options list (in the left pane).
3. Enter (or modify) the required information and click Update, as shown in Figure 11-51.

The screenshot shows the 'Options' window with the 'Host Agent' tab selected. The left pane lists various options, and the right pane contains the configuration fields. The 'Host Agent Login' field is a dropdown menu showing 'admin (Monitoring and Management)'. The 'Host Agent Monitoring Password' and 'Host Agent Management Password' fields are masked with dots. The 'Host Agent Administration Transport' field is a dropdown menu showing 'HTTP'. The 'Host Agent Administration Port' field is a text box containing '4092'. The 'Host Agent CIFS Account' and 'Host Agent CIFS Password' fields are also masked with dots. At the bottom right, there are 'Back' and 'Update' buttons. An arrow points to the 'Update' button.

Figure 11-51 Options window

This option changes *all* host agent login names and passwords, unless the host agent has a different login name or password already specified for it. For example, if an administrator has specified a password other than Global Default in the Password field of the Edit Host Agent Settings window (select **appliance name** → **Tools list** → **Edit Settings** to get to this window), changing the global password option does not change the storage system's password.

11.9 Managing FSRM search paths

In order to monitor file meta data information, you will need to define storage paths to the data you want monitored.

About SRM paths

File SRM paths define the location in the file system that is to be indexed for data. FSRM paths have the following properties:

- ▶ They must be defined for a specific host.
- ▶ They can be walked or monitored by multiple host agents.
- ▶ They can be grouped like any other storage object.
- ▶ They can be mapped (linked) to volumes, qtrees, and LUNs.

The FSRM path walk feature can cause performance degradation. However, you can schedule your path walks to occur during low-use or non-business hours.

Adding paths

To add SRM paths, do these steps:

1. From the SRM Summary window, click the **Add/Edit Paths** link in the SRM Paths Total section. The Add or Edit SRM Paths window is displayed.
2. From the SRM Host drop-down list in the Add a New SRM Path section, select the name of the host agent that you want to monitor.

3. Enter a path name, select a schedule, and then click **Add SRM Path**, as shown in Figure 11-52.

Figure 11-52 Add or Edit SRM Paths window

Valid path formats

The following examples show valid path entries:

```
host:/u/earth/work
host:/usr/local/bin
host:/engineering/toi
host:C:\Program Files
```

For CIFS, you must specify the path as a UNC path, as follows:

```
host:\\storage system\share\dir
```

Path names for CIFS

For CIFS systems, always use Universal Naming Convention (UNC) path names.

In Windows operating systems, the UNC format is as follows:

```
\\servername\sharename\path\filename
```

The SRM feature does not convert mapped drives to UNC path names. For example, suppose that drive H: on the system host5 is mapped to the following path name:

```
\\abc\users\jones
```

The path entry host5:H:\ fails because the FSRM feature cannot determine what drive H: is mapped to. The following path entry is correct:

```
host5:\\abc\users\jones
```

Quoting conventions for specifying the paths from the DFM Server CLI

Windows requires that you use double quotation marks to enclose all strings that contain spaces. For example:

```
C:\dfm srm path add "inchon:C:\Program Files"
```

UNIX requires that you double all backslashes, unless the argument is enclosed in double quotation marks. This convention is also true for spaces in file names. For example:

```
$ dfm srm path add inchon:C:\\Program\ Files
$ dfm srm path add "inchon:C:\Program Files"
$ dfm srm path add oscar:/usr/local
$ dfm srm path add oscar:/usr/home
```

Viewing file level details for a path

To view file level details about an SRM path, click, in the SRM Summary window, a path name in the SRM Paths Total section, as shown in Figure 11-53 on page 301. The SRM Paths Details window is displayed, as shown in Figure 11-54 on page 301.

- Click the **Browse Directories** link in the SRM Path Tools list (at the lower left of Operations Manager), as shown in Figure 11-56. The Directory Details window is displayed.

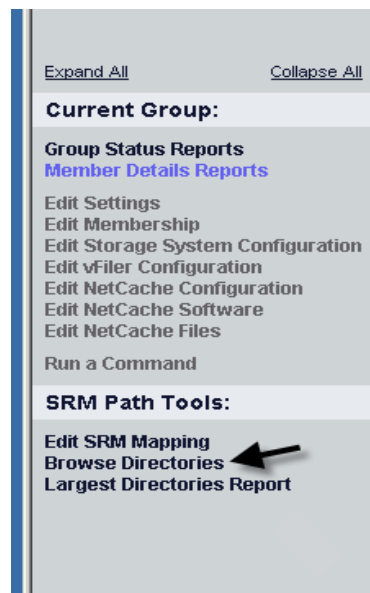


Figure 11-56 Browse Directories

- To view an expanded view of directory information that includes a listing of files by type and by user, click the **Extended Details** link (at the upper right corner of the File SRM tab window), as shown in Figure 11-57. The window shown in Figure 11-58 on page 303 should appear, showing more details about the directory.

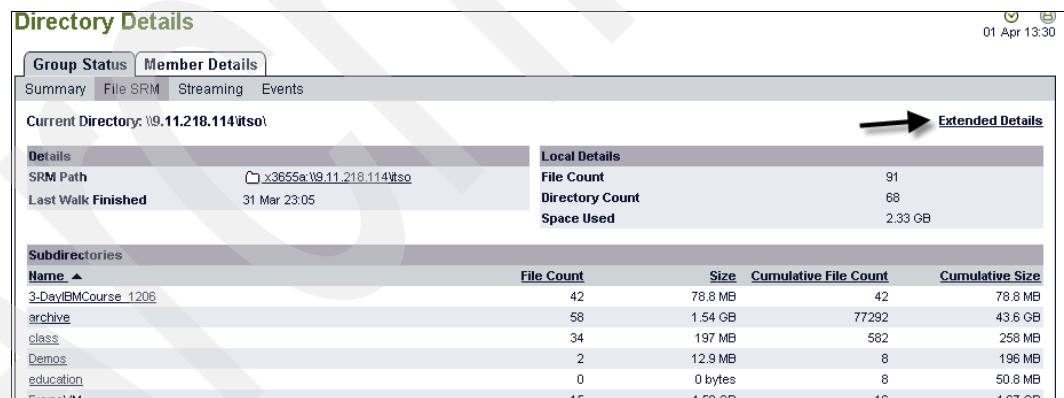


Figure 11-57 Directory Details window

Files By Type			Local Files By Type		
File Type	File Count	Size	File Type	File Count	Size
.mp3	17180	61.0 GB	.mp3	1	941 MB
[other]	13644	47.1 GB	[other]	20	828 MB
.zip	412	9.43 GB	.exe	7	379 MB
.exe	346	5.84 GB	.ppt	27	236 MB
.ppt	1747	3.69 GB	.doc	1	1.82 MB
.jpg	6380	3.53 GB	.jpg	19	1.31 MB

Figure 11-58 Directory Details expanded

Restricting access to file system data

To restrict access to private or sensitive file system information, remove the GlobalSRM role from the access privileges in the Administrators window. You can get to this window by selecting **Setup** and then the **Administrative Users** link, as shown in Figure 11-59.

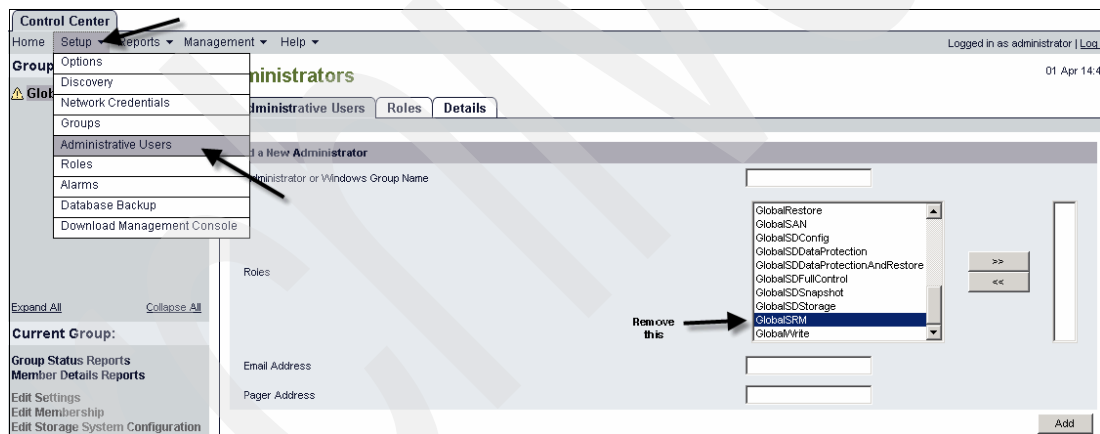


Figure 11-59 Restricting file system access

11.10 Scenario: identifying the oldest files in a storage network

Here we discuss how to identify the oldest files in a storage network.

Description

To locate the oldest files residing in a storage network so that the files can be archived to a near-line system. To find the oldest files, the administrator uses File SRM.

Summary of tasks

This process consists of the following high-level tasks:

- Check FSRM prerequisites.
- Verify administrative access.
- Verify host agent communication.

- ▶ Create a new group. Group the host agents in a logical way. For example, group the engineering host agents together if you want to search for them separately.
- ▶ Add a FSRM path.
- ▶ Add a schedule.
- ▶ Group the FSRM paths.
- ▶ View a report listing the oldest files.

Check FSRM prerequisites

Before using the FSRM feature for the first time, the administrator verifies that all prerequisites are met by referring to the “Difference between capacity reports and file system statistics” topic in the online help.

Verify administrative access

To verify administrative access with the administrator, do these steps:

1. Select **Options** from the Setup drop-down menu.
2. Click **Host Agent** in the Edit Options section.
3. Change the Host Agent Login to Admin.
4. Verify that the Host Agent Management Password is set.
5. Click **Update** to apply the changes.
6. Click the **Home** link to return to the Control Center.
7. Click the **File SRM** tab to return to the SRM Summary window.

Refer to Figure 11-48 on page 297 and Figure 11-49 on page 297 for more information.

Verify host agent communication

To verify that Operations Manager can communicate with the host agents, as the administrator, do these steps:

1. Click the **File SRM** tab and then select **SRM Summary** from the Report drop-down list.
2. Check the list of host agents to view the status. If the status is Unknown, the host agent login settings might not be properly configured.
3. If the status of one or more of the storage systems is Unknown, click the **Edit Agent Logins** link.
4. Select the host agents for engineering that the administrator wants to communicate with.
5. Edit the login or password information.
6. Click **Update**.
7. Click the **File SRM** tab to return to the SRM Summary window.

Create a new group

The administrator wants to find the oldest files in its engineering department first. To find these files, the administrator groups the host agents in the engineering domain together. To create a new group of hosts, as the administrator, do these steps:

1. From the SRM Summary window, select **Host Agents, SRM** from the Report drop-down list.
2. Select the check box to the left of each host agent in the engineering domain.
3. From the buttons at the bottom of the window, click **Add To New Group**.

4. When prompted, enter the name Eng and click **Add**. Operations Manager refreshes.
5. Select **SRM Summary** from the Report drop-down list.

Add a FSRM path

To add an FSRM path, as the administrator, do these steps:

1. From the SRM Summary window, click **Add/Edit Paths**.
2. Select a host from the SRM Host drop-down list.
3. Enter the path to be searched in the Path field.
4. Click the **Add A Schedule** link.

Add a schedule

To create the path-walk schedule, as the administrator, do these steps:

1. Click the **Add/Edit Schedules** link in the SRM Summary window (File SRM tab).
2. In the Add a New Schedule section, enter a meaningful name for the schedule.
3. In the Schedule Template list, select a schedule or select **None**.
4. Click **Add**.
5. In the SRM Path Walk Schedule Times window, select the days and times to start the SRM path walks.
6. Click **Update**.
7. Click the **Home** link to navigate back to the main window.
8. Click the **File SRM** tab.
9. From the SRM Summary window, click **Add/Edit Paths**.
10. In the Add a New SRM Path section, select a host agent to associate the new schedule with.
11. In the Schedule field, select the schedule name the administrator just created.
12. Click **Add SRM Path**.

Group the FSRM paths

The company administrator wants to see consolidated data, so the administrator chooses to group the SRM paths. To group the SRM paths, as the administrator, do these steps:

1. Click the **File SRM** tab.
2. Select **SRM Paths, All** from the Report drop-down list.
3. Select the check box to the left of each SRM path that the administrator wants to group.
4. From the buttons at the bottom of the window, click **New Group**.
5. When prompted, enter a name for the group.
6. Operations Manager adds the new group and refreshes.

View a report listing the oldest files

To view a report listing the oldest files in the SRM path, as the administrator, do these steps:

1. Click the **Home** link.
2. Click the **File SRM** tab.
3. Select the engineering group in the Groups section at the left side of the tab window.
4. Select the report **SRM Files, Least Recently Accessed** from the Report drop-down list.
5. Review the data.



Protection Manager setup

This chapter will help you prepare and install Protection Manager.

12.1 Host prerequisites

Table 12-1 details the hardware and software prerequisites for installing IBM System Storage N series DFM Server on a Linux system on physical hardware or in a VMware ESX guest image. Protection Manager is a feature that is built into Operations Manager and is activated by a license, so the requirements for Protection Manager are identical to that of DFM Server.

Table 12-1 Installation requirements for Operations Manager on Linux hosts

Linux workstation or server	
Hardware requirements	Software requirements
Intel-based PC with a single 2 GHz CPU (Xeon or Pentium 4) 4 GB of free disk space minimum, 8 GB recommended 1 GB of memory minimum	Red Hat Enterprise Linux AS 4 (Update 3 or later) for x86, 32-bit and 64-bit Red Hat Enterprise Linux Advanced Platform 5 for x86, 32-bit and 64-bit SUSE Linux Enterprise Server 9 (Service Pack 2 or later) for x86, 32-bit and 64-bit SUSE Linux Enterprise Server 10 for x86, 32-bit and 64-bit
Linux servers on VMware ESX V3.0.1 server or later	
Hardware requirements	Software requirements
Intel-based PC with a single 2 GHz CPU (Xeon or Pentium 4) 4 GB of free disk space minimum, 8 GB recommended 1 GB of memory minimum Single dedicated network interface	Red Hat Enterprise Linux AS 4 (Update 3 or later) for x86, 32-bit and 64-bit

Note: DataFabric Manager (DFM) server V3.7 is not supported on Windows NT 4.0, Windows 2000, Windows XP, Solaris 8, or distributions of Linux not listed in Table 12-1.

Operations Manager V3.7 does not support VMware VMotion and VMware High Availability features.

These requirements are for a Operations Manager installation with only basic system monitoring enabled. If you enable additional features and monitor additional objects, a more powerful platform is probably required. Examples of objects and features that might require a more powerful platform include additional storage systems, qtrees, user quotas, and use of the Storage Resource Management, Performance Advisor, Business Continuance Option, Provisioning Manager, or Protection Manager features.

It is important that you run Operations Manager on a system that is running no other applications. Running other applications that will take away CPU, I/O, and memory bandwidth from Operations Manager may cause Operations Manager to be unable to carry out its tasks properly or reliably.

12.2 License requirements

You must have a valid DataFabric Manager Server license key to complete the Operations Manager installation. You can access your license key at:

<http://www.ibm.com/storage/nas/>

After you have accessed the site, follow the options provided to access the license keys.

After you complete the installation, you can enter additional license keys in the Options window in Operations Manager. You can install (or upgrade to) Operations Manager V3.7 using the core license key.

However, if you do not have the core license key, you need the following licenses to monitor and manage your storage systems:

- ▶ DataFabric Manager Server license
- ▶ Additive license

DataFabric Manager Server license

The DataFabric Manager Server license is the server license with a unique serial number that tracks the number of Operations Manager installations. You must have this license to enable features. The node count is one.

Additive license

The additive license is an additional license with a unique serial number that is used to increase the node count and enable the features.

After you have installed the prerequisite licenses for base Operations Manager, you then need to install a license for Protection Manager. You can do this from the command prompt, as we demonstrate in 12.3, “Installing the license” on page 310.

As we mentioned previously, Protection Manager feature has the benefits shown in Table 12-2.

Table 12-2 Protection Manager benefits

To Use	Install this product	Which enables these features
Protection Manager	<ul style="list-style-type: none">▶ DataFabric Manager Server license▶ Operations Manager license▶ Protection Manager license▶ NetSeries Management Console	<ul style="list-style-type: none">▶ Automated policy-based data protection for NAS and SAN storage systems▶ SnapVault, Open Systems SnapVault, and SnapMirror management▶ Policy conformance checking and alerting Monitoring<ul style="list-style-type: none">– Reports– Storage usage and availability, such as qtrees, volumes, aggregates, LUNs, and disks– Storage systems– vFiler units– NetCache appliances– Real-time streaming events

12.3 Installing the license

DFM Server has a very powerful DFM Server CLI called **dfm**. From a terminal window, you can execute the **dfm** command to view a list of actions you can carry out.

One of the commands is **dfm licenses**. From a terminal window, we can execute the **dfm licenses help** command to list the syntax of the command, as shown in Example 12-1.

Example 12-1 dfm licenses help command

```
[root@vegemite ~]#  
[root@vegemite ~]# dfm licenses help
```

NAME
 license -- enable or disable licensed features

SYNOPSIS
 dfm license disable <feature> ... | <license-key> ...
 dfm license install <license-key> ...
 dfm license list [<feature> ...]

DESCRIPTION
 The license command manages the list of licensed features, enabling them (with license codes) or disabling them.

 License codes are 14-character strings which enable product features.

```
[root@vegemite ~]#
```

We can execute the **dfm licenses list** command to list the licenses currently installed, as shown in Example 12-2.

Example 12-2 Using the dfm command to list installed licenses

```
[root@vegemite ~]# dfm licenses list
```

Feature	License Code	Serial #	Nodes	Expiration
core	rayehkru coulic	demo	250	18 May 2009
srm	ZUCJFHSFJXJWQE	demo	250	22 Jun 2009

```
[root@vegemite ~]#
```

We can install the Protection Manager License by executing the **dfm licenses install** command and then verify the installation by executing the **dfm licenses list** command again, as shown in Example 12-3 on page 311.

Example 12-3 Installing the Protection Manager license

```
[root@vegemite ~]#  
[root@vegemite ~]# dfm licenses install trzmnkru coulic  
Enabled Protection Manager license.  
  
[root@vegemite ~]#  
[root@vegemite ~]#  
[root@vegemite ~]# dfm licenses list
```

Feature	License Code	Serial #	Nodes	Expiration
core	rayehkru coulic	demo	250	18 May 2009
srm	ZUCJFHSFJXJWQE	demo	250	22 Jun 2009
dataprotection	trzmnkru coulic	demo	250	18 May 2009

```
[root@vegemite ~]#
```

Note: The evaluation licenses we use have a time limit. If you wish to evaluate Protection Manager or DFM Server, contact your IBM Sales Representative and request evaluation licenses.

If you wish to remove a license, you will need to disable it, as shown in Example 12-4.

Example 12-4 Removing or disabling a license

```
root@vegemite ~]#  
[root@vegemite ~]# dfm licenses disable ZUCJFHSFJXJWQE  
Disabled File SRM Option license.  
[root@vegemite ~]#  
[root@vegemite ~]# dfm licenses list
```

Feature	License Code	Serial #	Nodes	Expiration
core	rayehkru coulic	demo	250	18 May 2009
dataprotection	trzmnkru coulic	demo	250	18 May 2009

```
[root@vegemite ~]#
```

12.4 Running N series Management Console

If you had N series Management Console running while you installed the license, you need to restart it. When N series Management Console starts up and you have pointed it to your DFM Server, you will now see new navigation panes for Protection Manager, as shown in Figure 12-1.

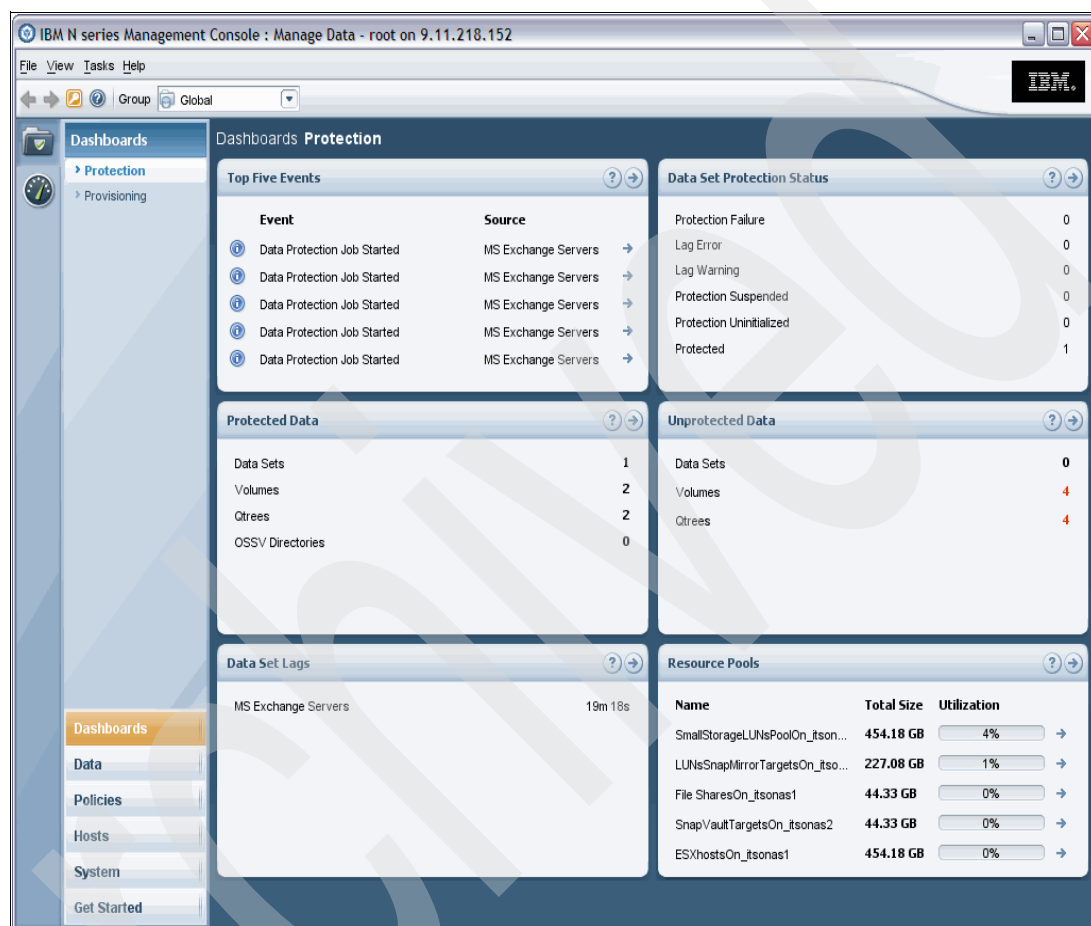


Figure 12-1 Protection Manager Dashboard as seen from N series Management Console



Protecting your data with Protection Manager

This chapter introduces you to N series Protection Manager, provide information about its setup and installation, and give examples of common tasks.

13.1 Introduction to Protection Manager concepts

Protection Manager helps you manage your backup and mirror relationships and perform failover operations easily and efficiently by eliminating repetitive tasks and automating some tasks. Typically, data and resource management is time consuming because it involves manual analysis and management of storage capacity, network bandwidth, schedules, retention policies, and other infrastructure variables. Protection Manager simplifies this work by employing configuration policies, convenient wizards, and automated verification of certain aspects of the data protection configuration. It lets you launch a backup, restore, or failover operation with a single click.

Protection Manager can perform the following actions:

- ▶ Use policies to manage primary data, storage, and backup and mirror relationships.
- ▶ Manage local and remote backups and mirror copies.
- ▶ Provision secondary storage for backups and mirrored copies based on policies you assign.
- ▶ Enable disaster recovery capability if you install the licensed disaster recovery option.
- ▶ Automatically validate your backup and disaster recovery configuration with a conformance checker.

Anyone using Protection Manager should be familiar with general data protection and disaster recovery concepts. Protection Manager uses Data ONTAP data protection technologies, such as Snapshot copies, SnapVault, Open Systems SnapVault, and SnapMirror.

Figure 13-1 on page 315 presents a conceptual overview of how we can use Protection Manager to protect data. In this scenario, we assume that resource pools are used not only to designate backup and SnapMirror targets, but can also be used to designate production data.

Data sets are used to group production data resource pools into logical groups that represent, for example, application servers, database servers, file and print servers, or groups of Web servers.

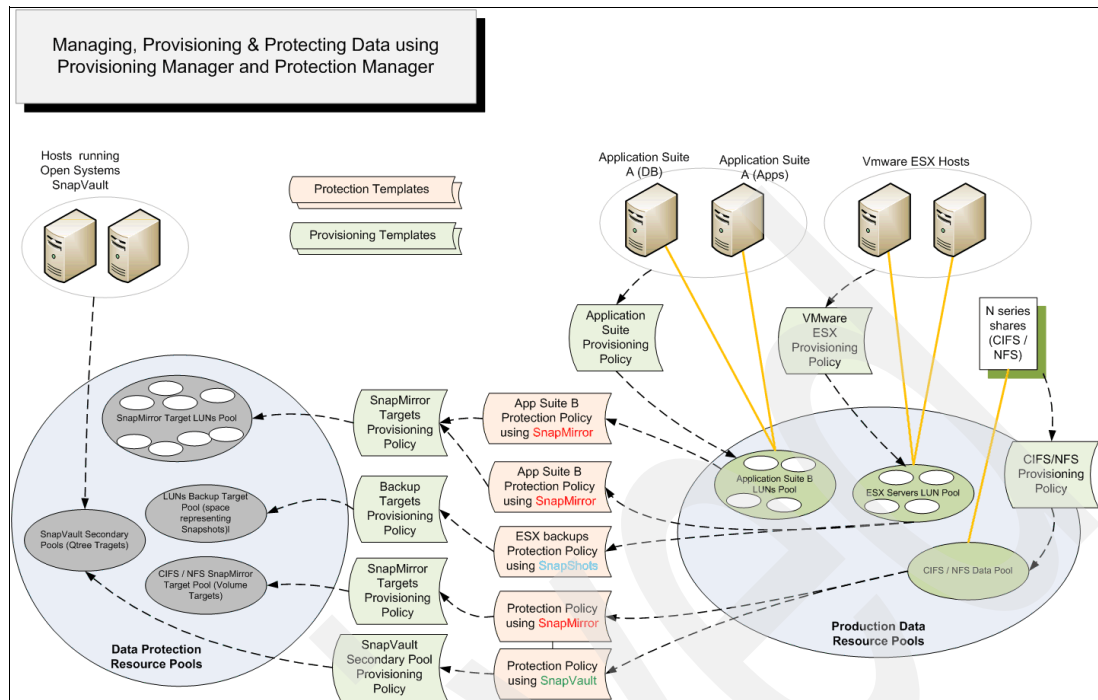


Figure 13-1 Conceptual overview of data protection using resource pools and provisioning policies and protection policies to enforce outcomes

Protection Manager V3.7.1 and earlier does not support the following features:

- ▶ Synchronous or semi synchronous mirroring using SnapMirror
- ▶ Automatically increasing the size of traditional volumes
- ▶ Backing up a volume to a qtree on a destination system
- ▶ Non-ASCII characters for CIFS data on vFile units, that is, unicode names are not yet supported.

13.1.1 Data sets

In the simplest terms, a data set is a *collection* of user data that you manage as a single unit, plus all the replicas of that data (Figure 13-2). The data is identified by the volume, qtree, or directory on the same storage system in which it is located. Because you manage a data set as a single unit, its members should have common management requirements. In Protection Manager, members of a data set should share the same data protection requirements. In Provisioning Manager, each node in a data set might not necessarily share the same provisioning requirements, but all members of each node should share the same protection requirements. Therefore, you will need to strike a balance with how you wish to designate datasets for protection and provisioning.

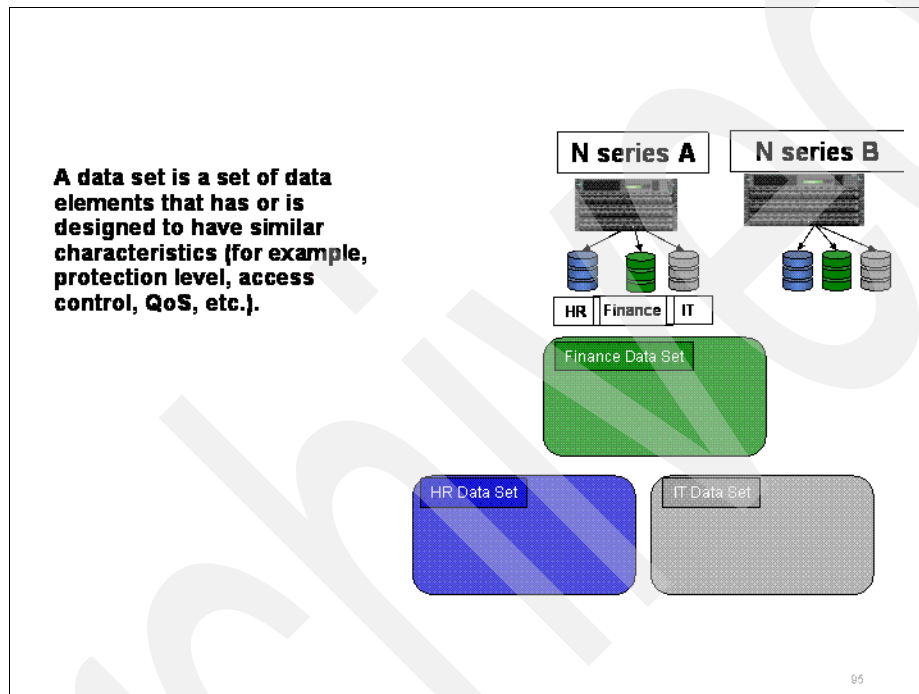


Figure 13-2 Data set

For example, different types of data supporting the same application would probably share the same protection requirements of the application. You want to group that data in the same data set, even if the data were stored in different volumes or qtrees. By configuring protection appropriate for the application and applying that protection to the data set, you apply it to all the data set members. So, a three tier application comprising a Web tier, application tier, and database tier may, for example, be grouped together as one data set and have the same protection *requirements* (such as backup first before mirroring to another storage array). Ultimately, your policy should reflect the business recovery requirements for the application that your organization has specified.

For provisioning, if the data set had a protection policy that created primary, backup, and mirror nodes, your provisioning requirements for each node might be different. You might want to provision the primary node on high-availability storage but provision the mirror node on less expensive, secondary storage.

13.1.2 Protection policies

A protection policy is a set of *rules* that specify the intended management of data set members. You can apply the same policy to multiple data sets, leveraging your configuration of the policy across the data sets.

A protection policy itself is not effective until it is applied to one or more data sets. At the time you apply a policy to a data set, Protection Manager will interrogate the data set and create the necessary scripts to help it enforce the rules set in the policy. Some of these scripts are initialization scripts, designed to set up the relationships that help enforce the protection policy. Other scripts are associated with enforcing the policies with schedules defined within your policy. The status of these jobs (current state, success failure, and so on) will be reflected in the Events and Jobs queues in N series Management Console as well as on the Protection Manager Dashboard. Some of these policies are run within Operations Manager server while others are set up to run automatically within the N series storage system itself.

Note: It is therefore, very important to make sure that the plug-ins used by DFM Server to monitor and manage your N series storage systems are of the correct version. Periodic review of your current Data ONTAP patch levels, DFM Server patch levels, and DFM plug-in versions are necessary in order to make sure DFM Server interfaces correctly with your N series storage systems. These plug-ins should not be confused with Host Agent plug-ins.

Note: If you update a policy, the update is propagated across all the data sets to which the policy is applied.

From the Protection Manager modules within N series console, you can use the following policies to quickly implement changes across an entire organization.

Protection policies

Protection policies define *which* data used for backups (Snapshots) are created on the primary storage, *when* to transfer the copies through a SnapMirror to a secondary or tertiary location and *how much* throttle to apply to the data being transferred at scheduled times. Protection policy settings also define *how long* to retain copies at each backup location and the warning and error thresholds for *lag time*. You cannot override a policy for specific members of a data set; if some members of a data set require different policy parameters, you need to move those members to a different data set.

Note: In Operations Manager, Protection Manager is not able to set up and monitor Synchronous SnapMirror relationships or Semi-Synchronous SnapMirror relationships. Unless otherwise stated, assume that all references to SnapMirror in this chapter will use periodic and scheduled replication of data to facilitate mirroring.

Types of data protection

This information describes the three types of data protection (local backup, remote backup, and mirror) that Protection Manager provides.

The Protection Manager protection policies enable you to provide a combination of three types of data protection:

- ▶ Local backup
- ▶ Remote backup
- ▶ Mirror

Local backup

Local backup protection (also referred to as Snapshot copy protection) is the periodic capture of the active data on an IBM storage system in backup images and the storage of those images on that same system. If active data on the local system is accidentally deleted or corrupted, it can quickly be restored with the most recent image stored locally from the last local backup job.

Local backup operations are typically employed on the primary storage systems, where data is being actively updated and where, in event of accidental data loss, data restoration from the last hour or two might be required.

Local backup protection is based on IBM System Storage N series Snapshot copy technology.

Note: Local backup protection is not available for data sets that include Open Systems SnapVault (OSSV) directories.

Remote backup

Remote backup protection is the periodic capture and copying of active data from a source storage system to a remote secondary or tertiary storage system. If data in the source storage system is lost and unrecoverable from its local backup (for example, if the source system is damaged), then data can still be quickly restored from the remote backup site.

Remote backup operations are employed from primary to secondary storage and from secondary to tertiary storage in circumstances where secure storage of backup data at a remote site might be required.

If a disaster recovery license is installed, Protection Manager enables you to supplement remote backup operations with additional failover instructions that transfer primary storage function to a secondary storage site if disaster or mishap disables or destroys the original primary storage site.

Remote backup protection is based on IBM Storage System N series SnapVault technology and IBM Qtree SnapMirror technology.

SnapMirror

SnapMirror protection is the periodic exact mirroring of all volume data (both active and protected) from an N series source storage system to an N series destination storage system (see Figure 13-3 on page 319). If data in the source storage system is lost or made unavailable (for example, if the source system is damaged), then that same data can quickly be made available from the destination mirror site.

Mirror operations are employed from primary to secondary storage and from secondary to tertiary storage, in circumstances where there is secure mirroring of that data, and in event of breakdown at the source site, quick availability of that data from a second site might be required.

If a disaster recovery license is installed, Protection Manager enables you to supplement mirror operations with additional failover instructions that transfer primary storage function to a secondary storage site if disaster or mishap disables or destroys the original primary storage site.

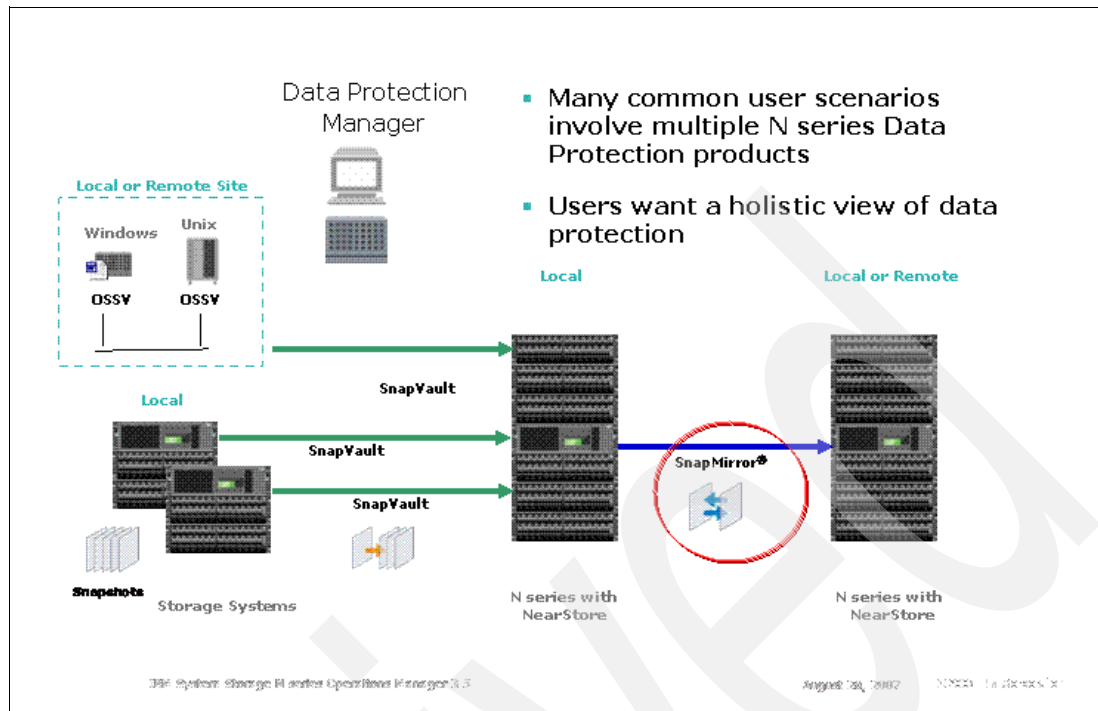


Figure 13-3 SnapMirror

Mirror protection is based on IBM System Storage N series Volume SnapMirror technology.

Note: Mirror operations that are scheduled and applied to a storage system through Protection Manager replace and cancel any other SnapMirror operations or jobs that were configured on that storage system locally by other means. Backup operations that are scheduled and applied to a storage system through Protection Manager run in addition to any other SnapVault operations that were configured on that node locally by other means.

Note: The lag time is the difference in age between the source data and the backup or mirror targets. If the lag time exceeds the duration between replications, your data availability will not meet your Recovery Point Objectives (RPO) or data age objectives. This is the mechanism where you ensure that business recovery Service Level Agreements (SLAs) are met. You can set alerts to notify you if the SLAs will not be met, enabling you to take corrective action.

Disaster recovery policies

A protection policy that supports failover from a primary storage system to a secondary storage system is considered to be capable of disaster recovery. The disaster recovery node is always directly connected to the primary node through the Fibre Channel SAN or IP network and its storage is made available after a disaster. Protection is possible by regularly replicating changes to data in production to the DR storage system. Special DR scripts are run by Operations Manager and are configured through the Protection Manager DR component to help complete the failover steps. This component is only made available through an appropriate license. You may have to restart the N series Management Console to see the new features after you add the Protection Manager DR license.

Provisioning policies

A provisioning policy defines how you want to have storage provisioned, exported, and managed, and what your space requirements are. For example, a provisioning policy might specify that when a storage container reaches the Nearly Full or Full threshold, an event message is sent or the size of the volume is increased (which provides more space for all the qtrees in the volume). A provisioning policy applies to all volumes, qtrees, or LUNs in a data set node. Policies also support error conditions (such as insufficient free space on the aggregate) to alert you if a policy cannot be enforced.

A policy is applied against an entire data set. You *cannot* assign different provisioning policies to *individual members* within a data set. However, with the protection license, and if the data set has a mirror or backup node, you can create and assign a different policy that defines provisioning and storage management on the mirror or backup node.

13.1.3 Resource pools

A *resource pool* is a collection of physical storage (comprising storage systems or aggregates) from which new volumes or LUNs can be provisioned to contain data (see Figure 13-4 on page 321). If you assign a storage system to a resource pool, all aggregates on that storage system become available for provisioning. Existing volumes do not get deleted. New volumes will be created to service a provisioning policy. Any unused physical resource in a resource pool is potentially eligible for provisioning. You can organize physical resources into resource pools by location, performance, or other important factors. The protection and provisioning applications apply IBM best practices in provisioning storage, so automatically provisioned volumes or LUNs meet the necessary requirements for compatible software version, licensing, and available space.

In Protection Manager, you typically assign a resource pool to the backup and mirror destinations of a data set. The protection application can then automatically provision volumes out of the physical resources in the resource pool to contain backups and mirror copies. To prevent conflicts, physical storage assigned to one resource pool cannot be assigned to a second resource pool. With Protection Manager and Provisioning Manager, you can use resource pools to fulfill requests for storage space for the primary or secondary data of a data set. By applying a provisioning policy to a data set node, the provisioning application applies the resiliency characteristics and space settings in the policy to automatically select the resources needed to fulfill a provisioning request.

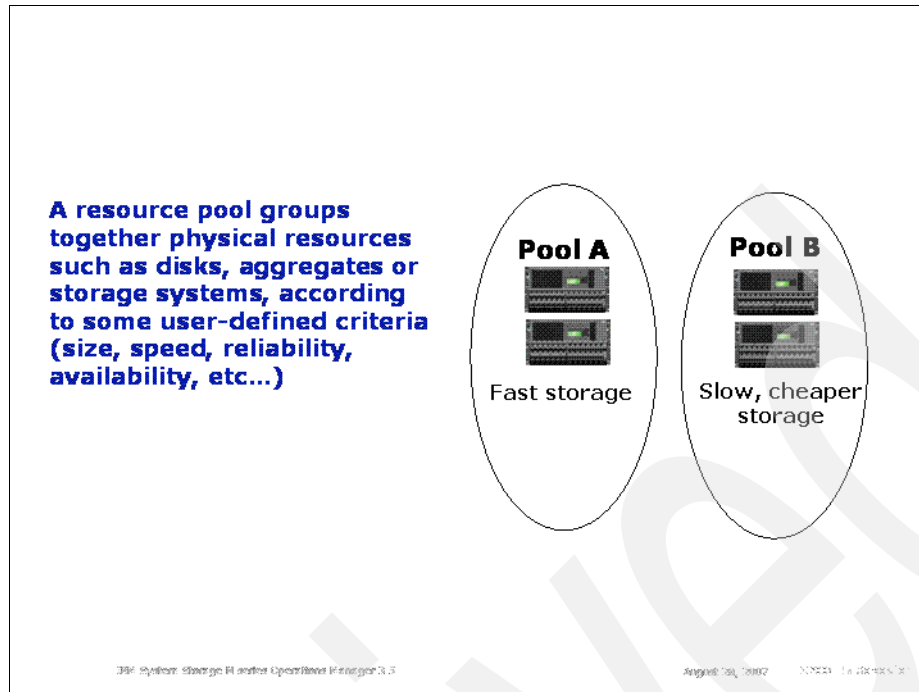


Figure 13-4 Resource pools

Note: In Operations Manager, you can now use Provisioning Policies to provision data volumes in a resource pool for actual production use. This is a new feature and works well when provisioning new data sets. Care must be taken when attempting to provision on imported aggregates, as there are certain restrictions that must be adhered to.

13.1.4 Protection Manager Dashboard

Protection Manager is accessed through the N series Management Console. This console make the Protection Manager components visible once you have added the necessary Protection Manager license to the DFM Server through the Operations Manager interface. You may need to restart the N series Management Console in order for it to recognize that you have the necessary licenses for Protection Manager.

When you launch the N series Management Console, you will see a new icon on the left side navigation pane that looks like a folder. Clicking this icon brings you to the Protection Manager Dashboard, as shown in Figure 13-5 on page 323.

The Protection Manager Dashboard provides a one-stop summary window of all information relevant to protecting your data. The main panes are:

- Top Five Events summary pane

This pane provided a list of the top five events that have occurred within your managed environment. The arrow to the right of each of the top five events is a shortcut that will take you the specific event information in the System Events window.

- ▶ Protected Data summary pane

This pane provides an overview of the number of data sets, volumes, qtrees, and OSSV directories that are currently protected.

- ▶ Data Set Protection Status summary pane

This pane provides a quick and effective overview of any pending problems with any of the data currently being protected.

- ▶ Unprotected Data summary pane

This pane provides you with a quick overview of any data currently in an unprotected state. The Data Sets row represents data sets that have been created but not yet protected. The Volumes and Qtrees rows represents data that has not yet been grouped into data sets or have been assigned a protection policy.

- ▶ Data Set Lags summary pane

This pane provides you with information of the lag (how long the backups are out of sync with the source) for a particular data set. As long as the lag time is less than the backup frequency time, the status is considered to be good. If the lag time exceeds the backup frequency assigned to a data set, then the protection status of that data set at risk and SLAs, if any, are not being met.

- ▶ Resource Pools summary pane

This pane provides an overview of your resource pools and how much storage is currently utilized.

Note: The arrow in the top right corner of each pane shown in Figure 13-5 on page 323 is a shortcut to the corresponding detailed pane.

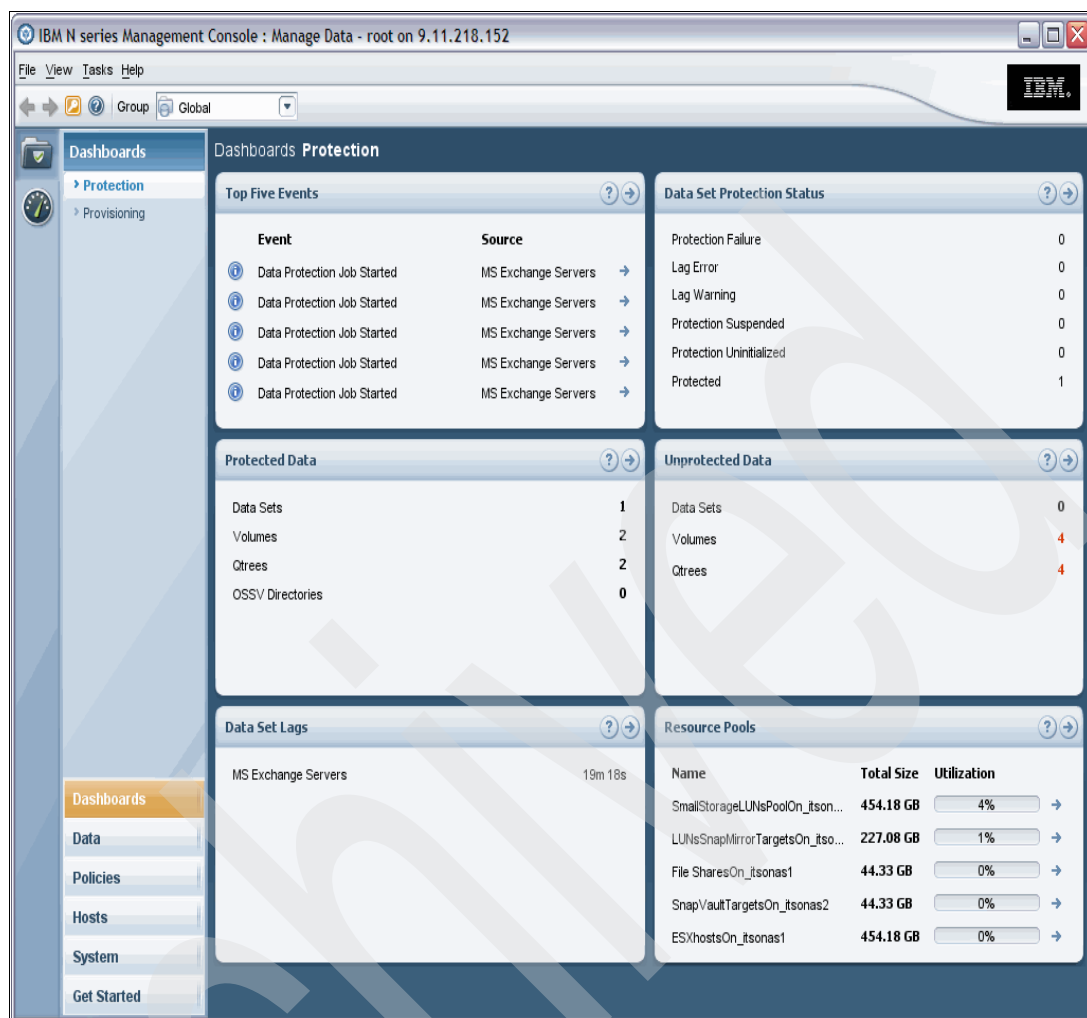


Figure 13-5 Protection Manager Dashboard window

13.1.5 Protection of discovered data

Protection Manager relies on a DFM Server database for information about storage systems and hosts. The DFM Server periodically polls the state of the registered storage systems and hosts. Protection Manager is able to tell you about storage systems and hosts and the protection status of their data. If the data is unprotected, the application reports the data in the dashboard as unprotected. From the Unprotected Data pane, you can drill-down to find detailed information that helps you decide whether to back up and mirror the data.

Note: To add a new storage system to Protection Manager, first configure the DFM Server to discover and manage it. You can do this by using the “Add Storage System” wizard in the Storage systems submenu of the Hosts section of the N series Management Console. Protection Manager will then notice the storage system and report on it.

If new volumes are created in an aggregate that is already protected, that protection is extended automatically to them. Likewise, for new Open Systems SnapVault clients that replicate to new volumes on aggregates in Storage systems or on a vFiler unit that are currently protected, the protection is automatically extended to these new volumes.

For example, if you create a FlexVol® volume on a storage system that is a member of a protected data set, the protection configured for that data set is applied automatically to the data in the new FlexVol volume. Protection Manager creates backup and mirror relationships for the new data, as defined in the policy applied to the data set, and provisions storage on the destination systems for copies of the new data.

Figure 13-6 shows two N series storage systems, *itsonas1* and *itsonas2*, which have resources that are listed as unprotected. Itsonas1's aggregates are displayed in Figure 13-6. Operations Manager will consider these resources as protected once you add these aggregates (or the entire storage system) to a current or new data set that is protected with a protection policy. In our scenario, we attempt to create a new data set for the Microsoft Exchange aggregate and assign a protection policy to protect these volumes and LUNs within this aggregate.

Figure 13-6 is the summary window for Protection Manager. You can see in the pane associated with unprotected data that there is storage that is currently unprotected.

Note: If you create or modify aggregates or volumes through Filerview or the text console on an N series storage system through the command-line interface (CLI), you may not see the changes reflected in Operations Manager or Protection Manager until the next refresh cycle of the DFM Server's monitoring agents. The default refresh cycle is 15 minutes and this can be changed through the Operations Manager interface; however, note that increasing the refresh frequency will, in turn, increase the performance load on the DFM Server.

Note: If you modify aggregate names or volume names through Filerview or the N series CLI, you will see the changes reflected in Operations Manager or Protection Manager at the next refresh cycle. Operations Manager (and Protection Manager) does not become confused when aggregate names, volume names, and LUN names are changed through Filerview or the N series CLI.

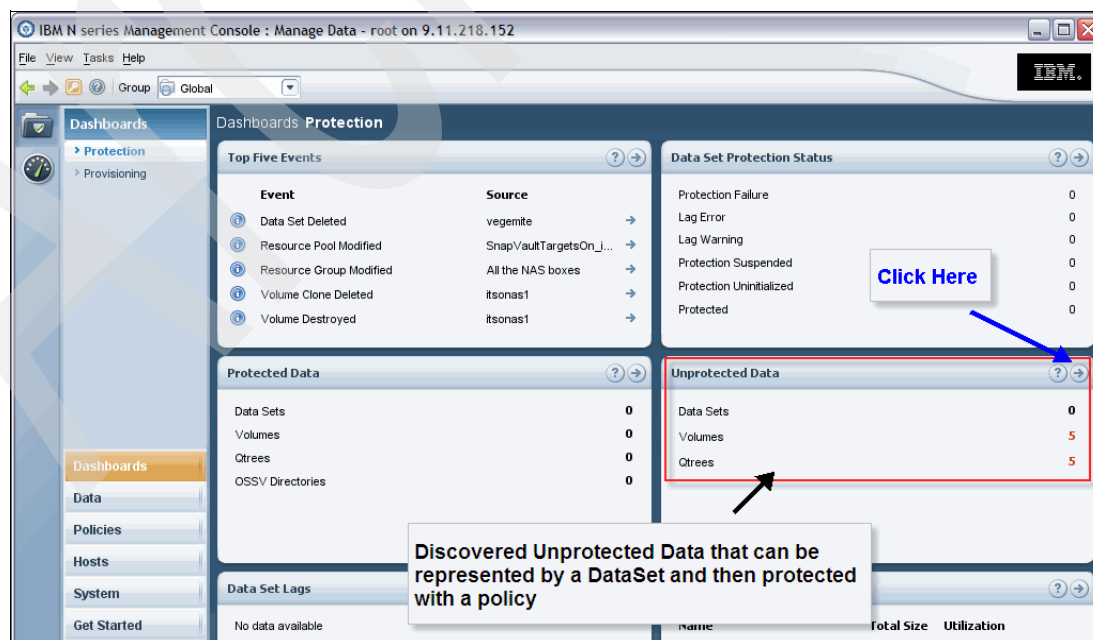


Figure 13-6 Overview of unprotected data

Figure 13-7 Shows the Unprotected Data window. In this window, you can see the two storage systems. The highlighted system, *itsonas1*, has a number of aggregates and volumes that need to be protected.

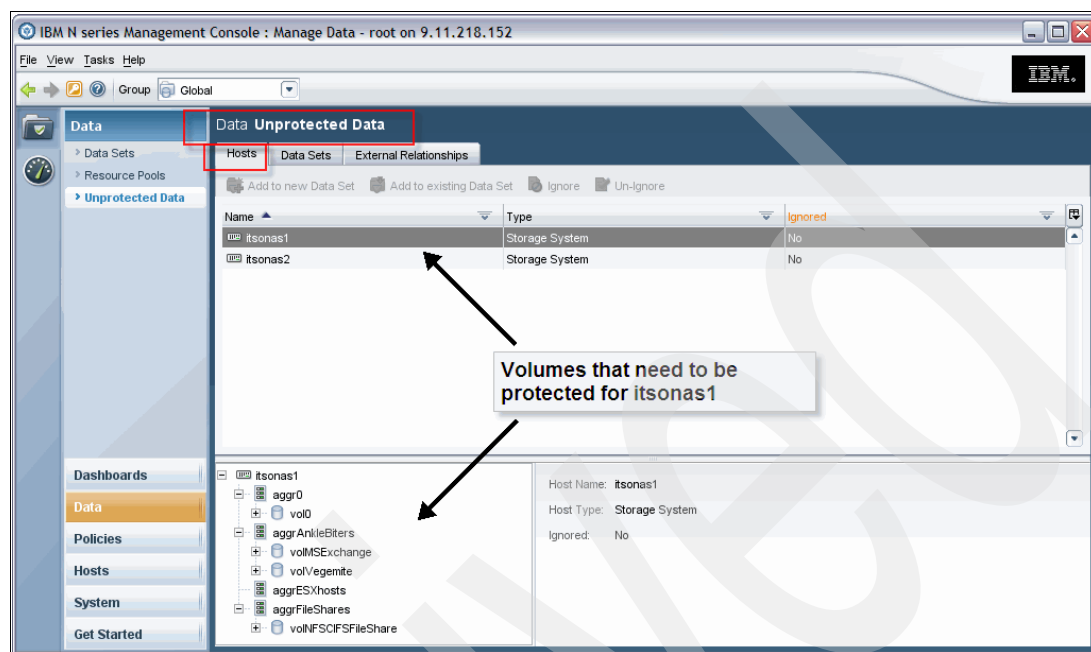


Figure 13-7 Report of unprotected storage systems and their data

13.2 Restoring data

The primary aim for protecting data is to be able to restore that data in the event there is some form of data loss. An protection strategy is regarded as effective if the following are true:

- ▶ The data sets are well defined and correctly encompass all relevant data associated with the business application or data resource that needs protection.
- ▶ The applied Protection Policy correctly reflects the business recovery outcomes desired.
- ▶ The appropriate data protection agents (if any) on hosts accessing this data are correctly configured and are able to be manipulated or driven by protection scripts.
- ▶ Protection scripts (if any) are written to properly capture the data in such as way as to guarantee a successful restore or recovery.
- ▶ Alerting mechanisms are well defined and are able to promptly relay to the relevant staff, information about error or warning events.
- ▶ Resources used for protecting data sets are adequately provisioned and able to cope with various unexpected changes and interruptions to a replication activity, including timely alerts of impending shortage of available storage resources as backup targets.
- ▶ There is a simple and effective restoration mechanism, designed to minimize recovery times.

While the scope for most of these items are outside the scope of this IBM Redbooks publication, we are able to demonstrate the steps associated with restoring data using Protection Manager. This will be covered in 13.6, "Demonstration of restoring data" on page 397.

13.3 Setting up the environment for demonstration

In this section, we set up an environment to demonstrate the use of Protection Manager. We define a data set from unprotected data and employ a protection policy against it. This, in turn, will require a resource pool and a provisioning policy for data protection use.

We will create a data set to represent an existing Microsoft Exchange environment.

We will use a protection policy to backup and mirror the source LUNs (on *itsonas1*) to another storage system (*itsonas2*).

To begin with, we use pre-created protection policies, provisioning policies and resource pools in order to maintain the conceptual flow of the demonstration. For details about how to create resource pools or provisioning policies, refer to Chapter 14, “Provisioning Manager setup” on page 437. We will, later in this chapter, review the steps we took to create the protection policy, provisioning policy, and resource pool for this demonstration.

We demonstrate the process of adding an Open Systems SnapVault (OSSV) host to a DFM Server, creating a data set to protect its data, and assigning a protection policy to protect the data.

Finally, in this demonstration, we restore data from a backup.

See 13.3.1, “Overview of the protection policies templates” on page 326 for an overview of existing protection policies.

See 13.3.2, “Creating a new protection policy” on page 332 for details about how we create a custom protection policy based on the existing (template) policies.

See 13.7.2, “Creating the provisioning policies” on page 414 for a demonstration of the steps we took to create one of the provisioning policies used here.

Tip: In a customer environment, it is likely that the customer has defined storage according to usage, availability, and performance tiers. In this case, you may decide to create resource pools to match these definitions and group all LUN based aggregates into these pools and make use of resource labels that represent application suites to distinguish between them.

A protection policy is applied against all the members of a data set. You may have application suites that depend on multiple storage tiers (and resource pools). You can have individual provisioning policies for provisioning storage from these pools.

Care must be taken in setting up your environment in order to ensure you get the organizational data protection outcomes you are looking for. This is a standard requirement that is applicable regardless of the storage technology and backup mechanisms you use and is not unique to N series technology.

13.3.1 Overview of the protection policies templates

Protection Manager has a number of predefined policies that can be used directly or as templates. We can create new policies that are based on these templates by highlighting a template and then selecting the **Copy** button. Protection Manager will immediately copy the template and prefix the text with “Copy of...”. You can then select the copied template and click the **Edit** button to edit and rename the new policy.

Figure 13-8 shows the list of templates. We now select a template and walk through its various properties. In this example, we select the template “Chain of two Mirrors”, which essentially uses SnapMirror to copy source data to a mirror target and then, at a later time, mirror that target to another target. This policy is useful for situations where frequent backups need to be made of the production (source) storage to an intermediary point for rapid restore or recovery, and less frequent copies need to be made for DR or archive purposes.

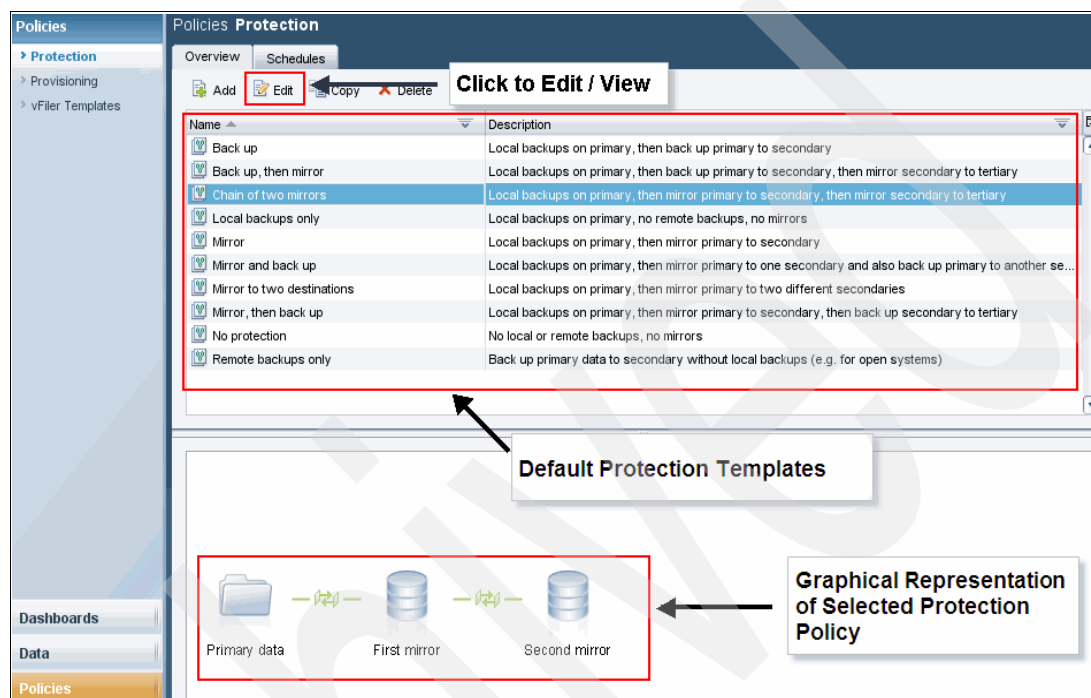


Figure 13-8 Overview of protection policy templates window

The properties window shown in Figure 13-9 shows the name and description of the template. You can change these fields to anything you like. When creating a new policy that is based on this template, change these fields on the copied policy to reflect the specific use you have planned for this policy.

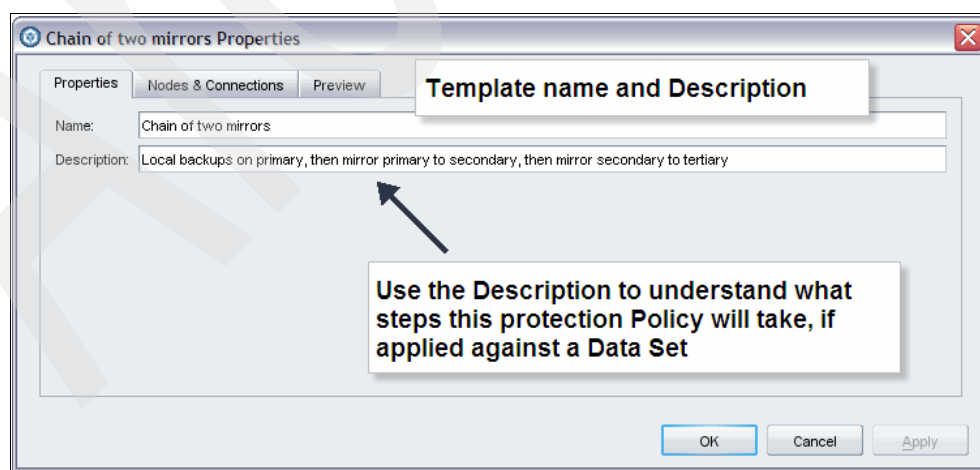


Figure 13-9 Properties tab within a protection policy

Figure 13-10 shows the Nodes & Connections tab. This tab is where most of the policy actions are defined. In some policies, the workflow steps here can be more simplistic. In this policy, the workflow makes a mirror copy of the source data at a user defined schedule. This schedule is executed relative to the time zone the storage system is located in. It then mirrors the copy to another destination at another user defined schedule.

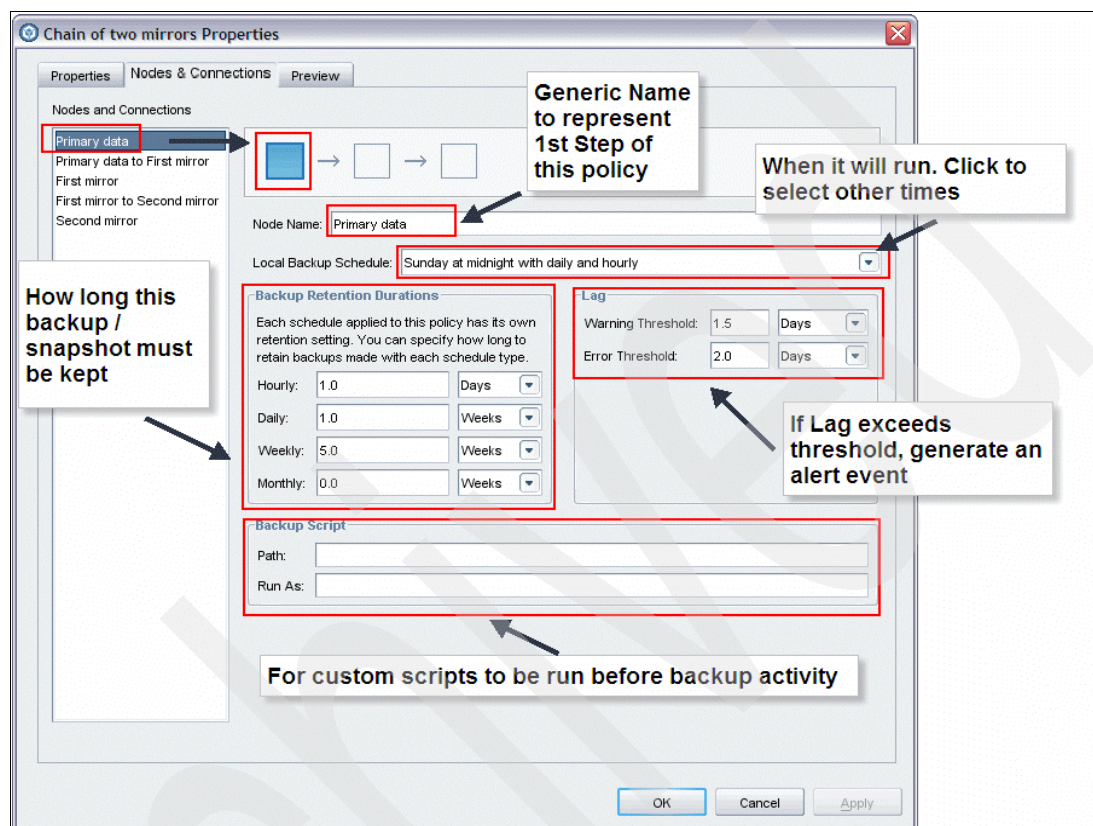


Figure 13-10 Primary data in Nodes & Conditions tab within a protection policy

The first step in the workflow allows you to define what needs to be done with Snapshots of the source data and when. It also allows you to determine how long to keep the copy, which is usually according to the retention policy your business requires for that data set. The lag threshold area allows you define acceptable lag warning and error thresholds that satisfy your business protection criteria. If you have custom applications that need specific actions to be carried out prior to the backup, then you may provide a backup script. An example would be a script that takes a database into hot-backup mode so that data can be flushed to disk prior to the snapshot being taken.

Figure 13-11 on page 329 shows the next step in the workflow. Here we can define the frequency of the mirror as well as how much network bandwidth the mirror should take while replicating. Throttling the bandwidth is very useful when using remote connections that are not dedicated to SnapMirror and need to be shared.

Note: Setting the throttle bandwidth too low or using a network bandwidth that is too low for successful replication within the specified frequency will cause the replication policy to be ineffective and may affect your business data protection SLAs.

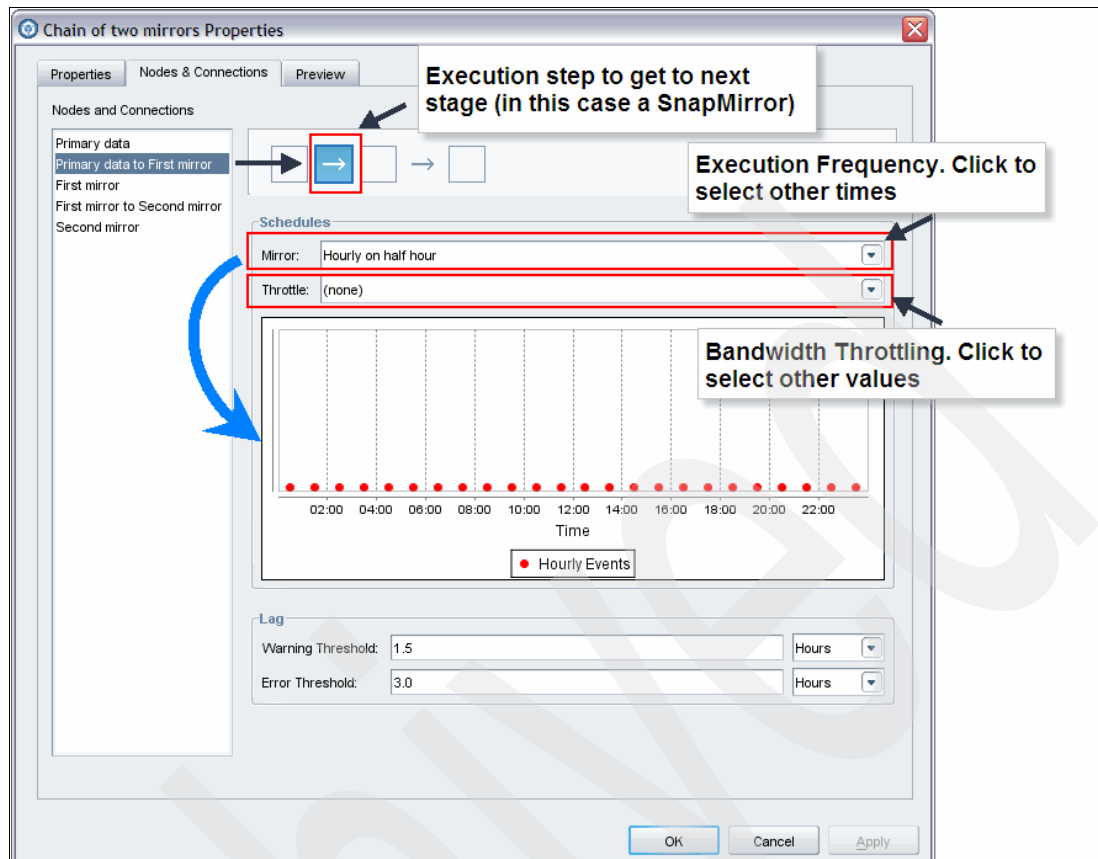


Figure 13-11 Execution step on primary data within a protection policy

An effective replication frequency is determined by:

1. Assessing the amount of data that needs to be replicated at each cycle by understanding the rate of change to the source data set
2. Assessing the time taken to complete the replication given the available bandwidth for the task
3. Allowing adequate extra time to ensure that unforeseen issues that may extend the replication duration do not impact on or collide with the next replication period
4. Ensuring that adequate extra time is also available to accommodate any other activities that also need to complete before the next replication cycle, such as Snapshots for backups
5. Ensuring that the start time of each replication cycle does not clash with any other protection activities that may be running on the same data set

Note: The execution of the policy occurs relative to the time zone where the data set is located, as defined in the data set's definition.

Note: A scheduled snapshot on a SnapMirror source will fail if a currently running SnapMirror operation on that source has not yet completed. If this happens, updates to the source data may not be backed (through Snapshot) up until the next cycle. This may not necessarily affect SnapMirror operations unless the SnapMirror is configured to use the specific Snapshot created.

Figure 13-12 shows the first mirror target for this two stage replication. In this scenario, we may envisage the mirror target being located near the data source and with adequate bandwidth between the two storage systems.

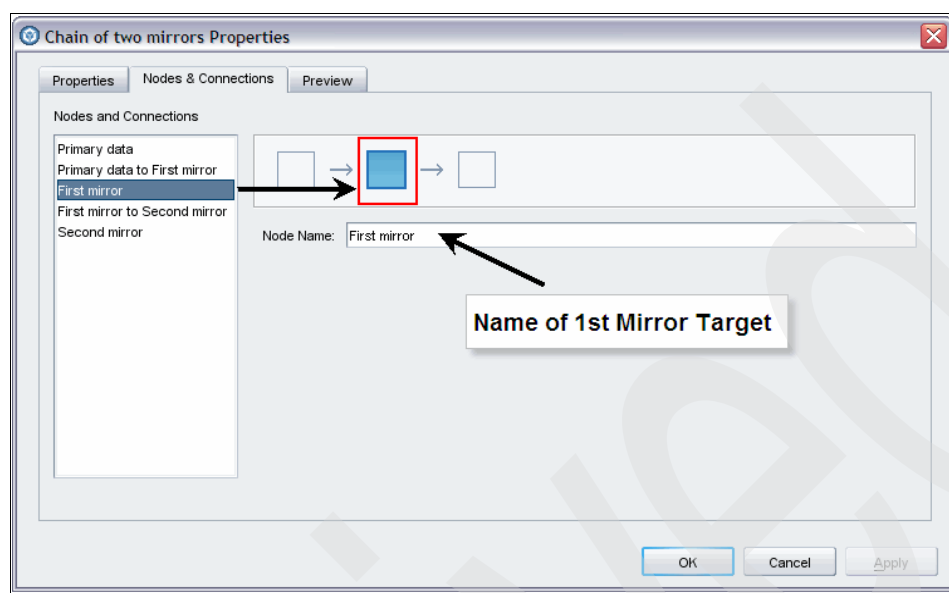


Figure 13-12 First mirror way point for the protection policy

Note: A policy is not effective unless it is applied against a data set. When you apply a policy against a data set, a new wizard will launch to gather specific information, such as the data set name, target resource pools for first mirror, target resource pools for second mirror, and so on. See 13.4.2, “Creating the data set” on page 334 for details about applying a protection policy against a data set.

Figure 13-13 on page 331 shows the next replication schedule. In this scenario, we may envisage the second mirror site to be at a remote location and that it requires an hourly backup that is out of phase by 30 minutes from the previous mirror workflow step.

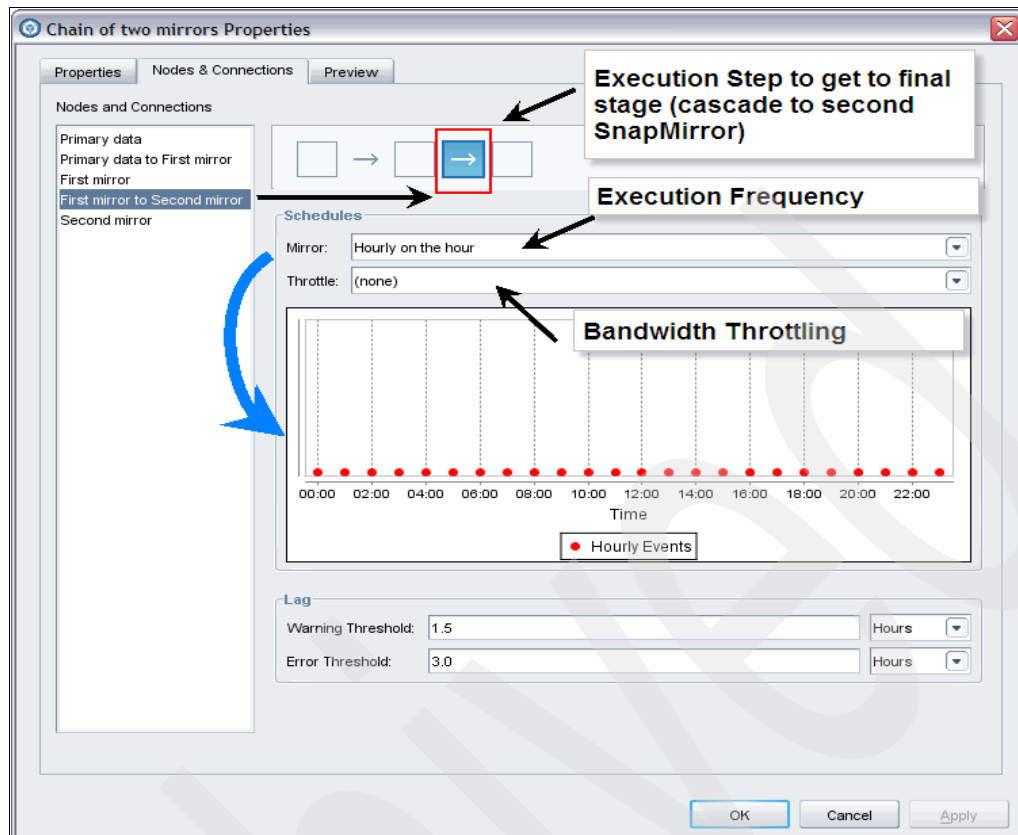


Figure 13-13 Execution step against first way point within the protection policy

Figure 13-14 shows the final stage of the protection policy wherein we define the name of the second mirror. These names (first mirror and second mirror) are used by Protection Manager to assist in creating unique relationship names that are also humanly readable when viewing from FilerView.

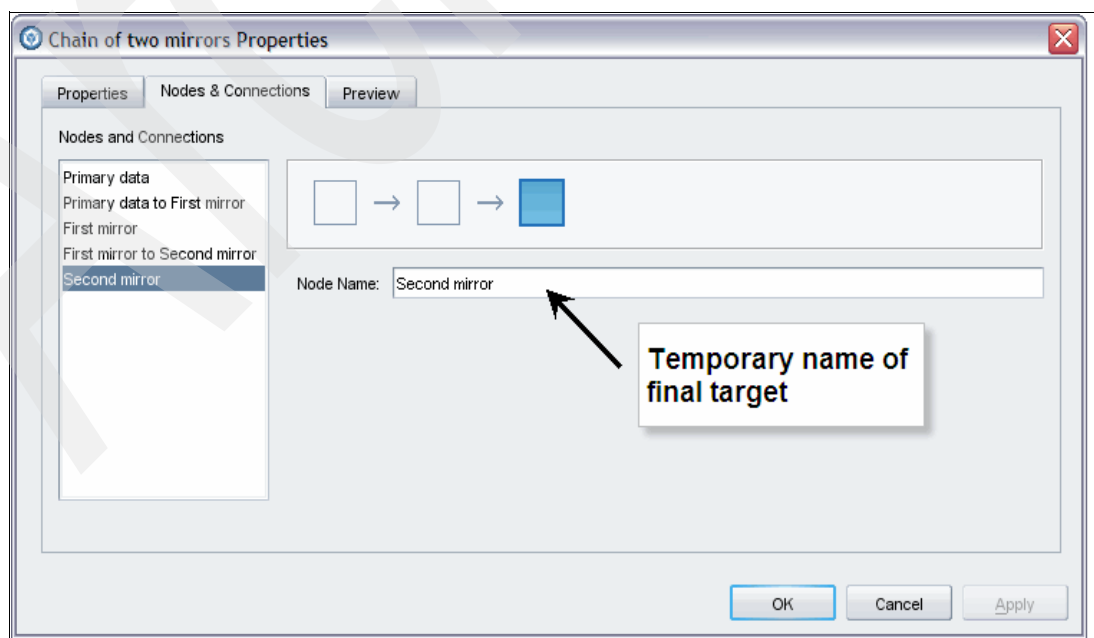


Figure 13-14 Final way point within the protection policy

Figure 13-15 shows the outcomes of any changes proposed and a recommendation of corrective action (if any is needed).

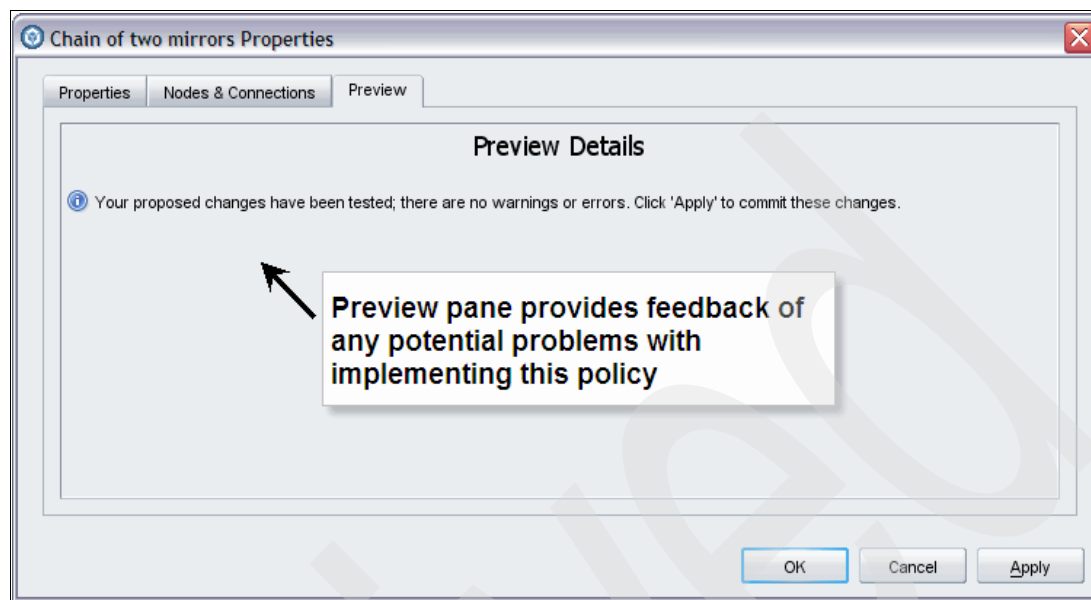


Figure 13-15 Review of changes tab within a protection policy

13.3.2 Creating a new protection policy

As stated before, we can create new policies that are based on these templates by highlighting a template and then selecting the **Copy** button. Protection Manager will immediately copy the template and prefix the text with “Copy of...”. You can then select the copied template and click the **Edit** button to edit and rename the new policy.

13.4 Demonstration of protecting data with a data set

In this section, we create a data set to represent a Microsoft Exchange Server volume. There are a number of steps to successfully protecting data.

1. Locate an unprotected data source (this could be a storage system, aggregate, or volume).
2. Create a data set with the located data source as a member.
3. Assign an appropriate protection policy to the data set.
4. Provision the resources pools in order to enforce the protection policy.
5. Monitor the Protection Manager Dashboard for details on the health of the protected Data set.

Note: The “Protect” Command located on the data set window makes it possible to launch a special backup on demand, such as prior to the beginning of a change activity.

13.4.1 Identifying unprotected data

The simplest way to identify unprotected data is to use the Protection Manager Dashboard. This Dashboard has a section that presents an overview of unprotected data. Use the arrow in the top right corner of this pane to get to the Unprotected Data summary window, as shown in Figure 13-16.

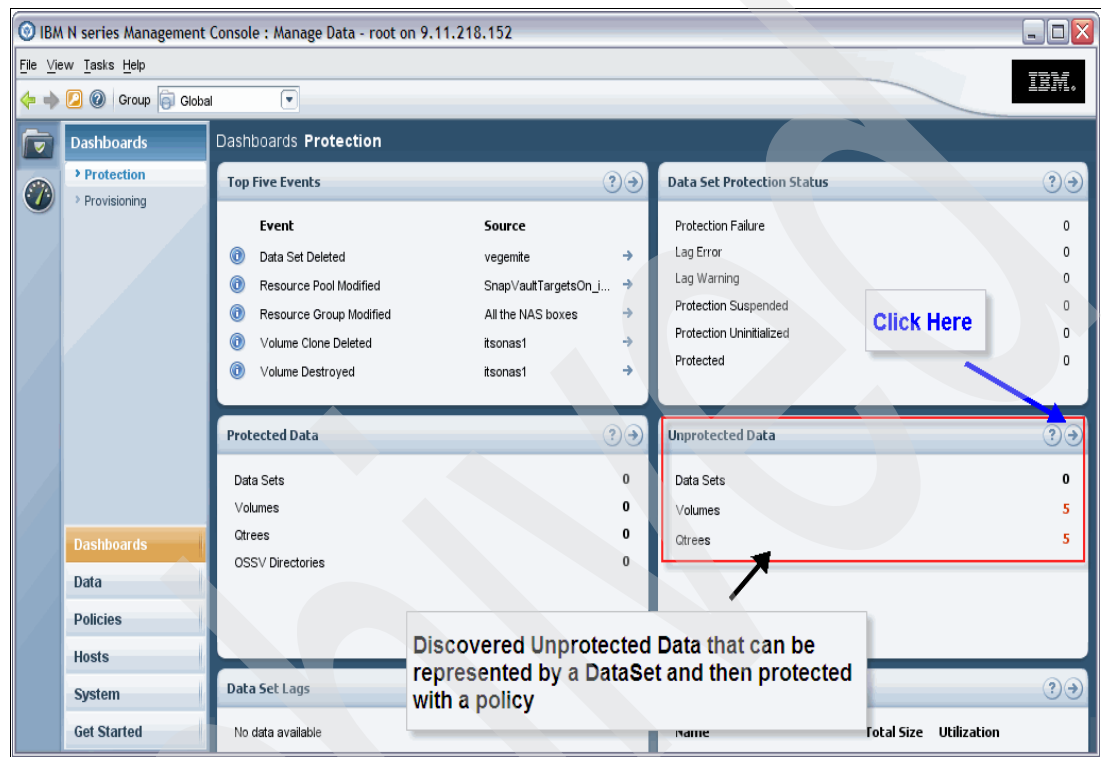


Figure 13-16 Protection Manager Dashboard showing unprotected data

13.4.2 Creating the data set

When you click the arrow pointed to by the blue marker, as shown in in Figure 13-16 on page 333, you will arrive at the Hosts tab for Unprotected Data, as shown in Figure 13-17. From here we are able to identify resources that need protection. We select the *volExchange* volume on *itsonas1* and add it to a new data set.

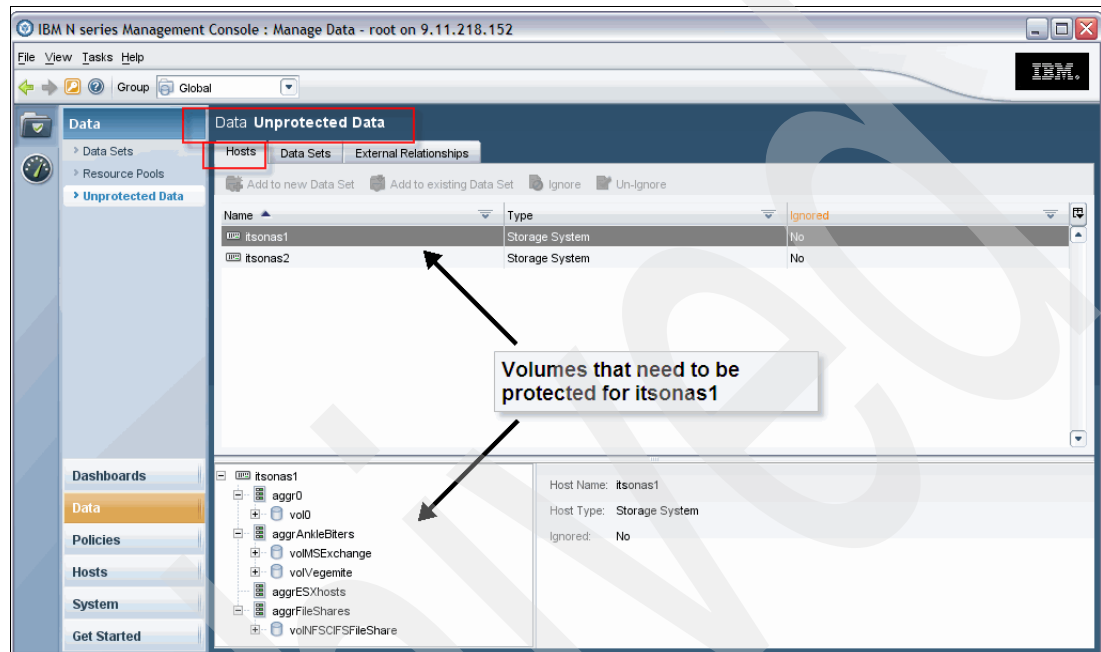


Figure 13-17 Identifying unprotected data

To add a volume to a new data set, first select the volume or volumes and then click the **Add to new Data Set** button, as shown in Figure 13-18 on page 335.

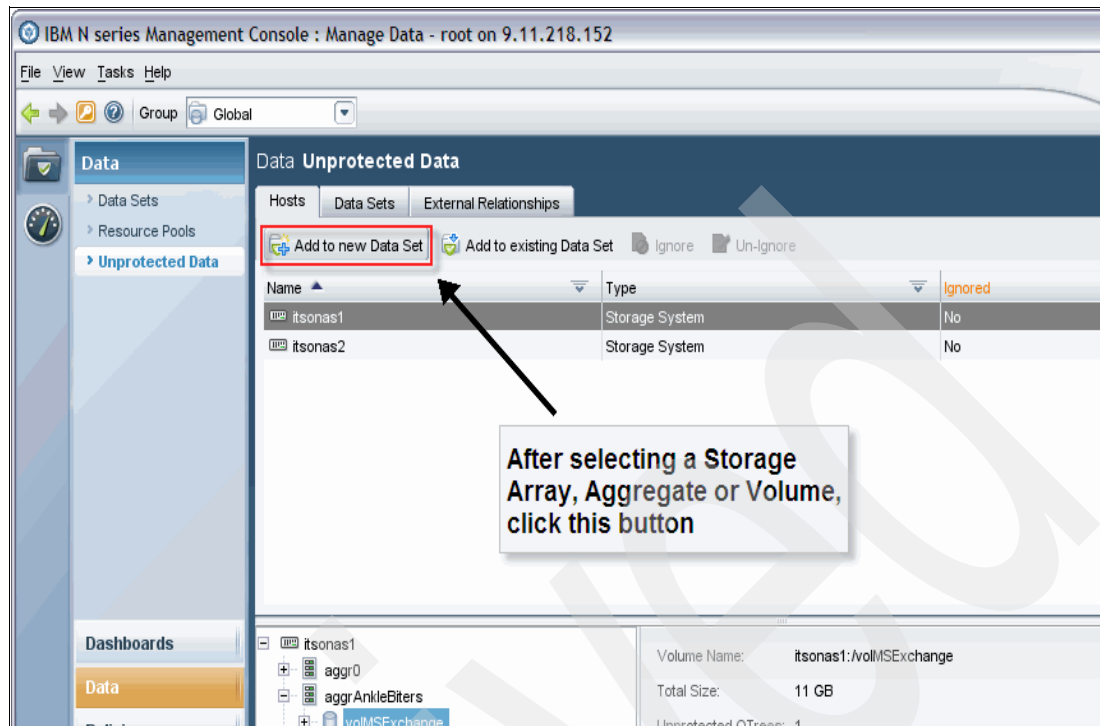


Figure 13-18 Adding unprotected data to a new data set

This will launch the Add Data Set Wizard, which will walk you through the steps for creating a new data set.

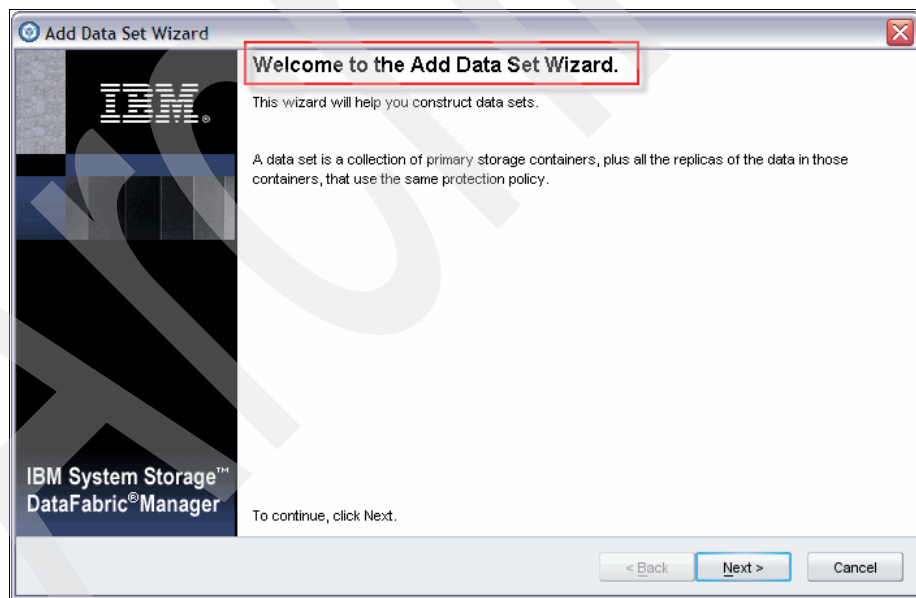


Figure 13-19 Add Data Set Wizard welcome window

A data set is uniquely identified by its name. It is helpful to ensure that the name correctly reflects the use of this data set. It is also important to indicate to Protection Manager the time zone in which the data set resides. This helps Protection Manager ensure that scheduled events like protection activities occur at the correct times relative to the time zone of the data set being protected.

Figure 13-20 shows the details we entered to create the “MS® Exchange Servers” data set.

Note: While the Description, Owner, and Admin fields are optional, you are strongly advised to make use of these fields to self document the function of this data set.

Add Data Set Wizard

General Properties
You must provide a name for the new data set. Other properties are optional.

Name: MS Exchange Servers
Description: All MS Exchange servers
Owner: IT Admin
Contact: cactus@tso.tucson.ibm.com
Time Zone: America/Panama
America/Pangnirtung
America/Paramaribo
America/Phoenix
America/Port-au-Prince

To continue, click Next.

< Back Next > Cancel

Figure 13-20 Uniquely identify the data set with a name and time zone

A data set needs to be added to a group. This feature makes it possible to collectively operate data sets for that group. You will not be able to create a new group in this window. If you need a new group, create it first before you launch the wizard.

Add Data Set Wizard

Group
You may select the group to which the new data set will be added

Global
All the NAS boxes
Linux Hosts
Windows Hosts

To continue, click Next.

< Back Next > Cancel

Figure 13-21 Add the data to a group

As we are attempting to create a new data set to protect existing data, we will select **Assign Resources Manually**, as shown in Figure 13-22.

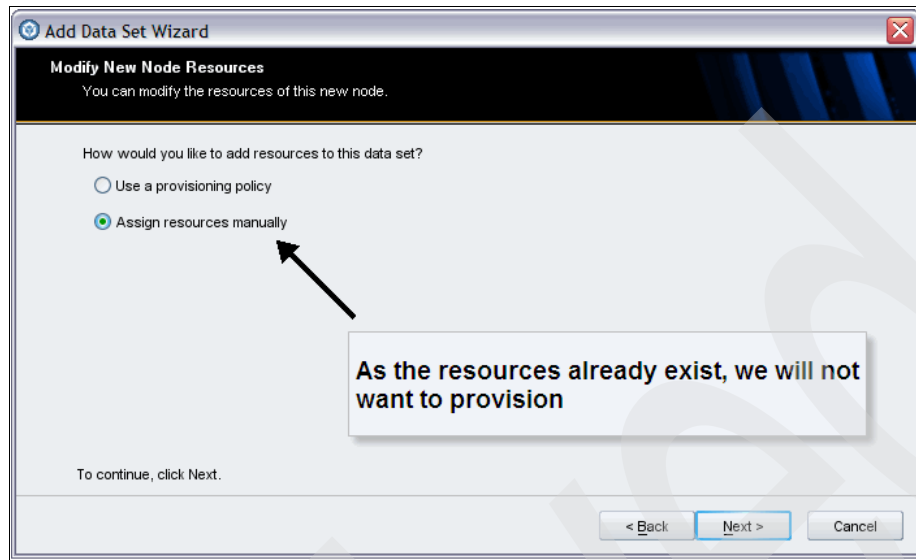


Figure 13-22 Opting to add resources manually

If there are additional resources you wish to add to this data set, you can do so in the next window, as shown in Figure 13-23.

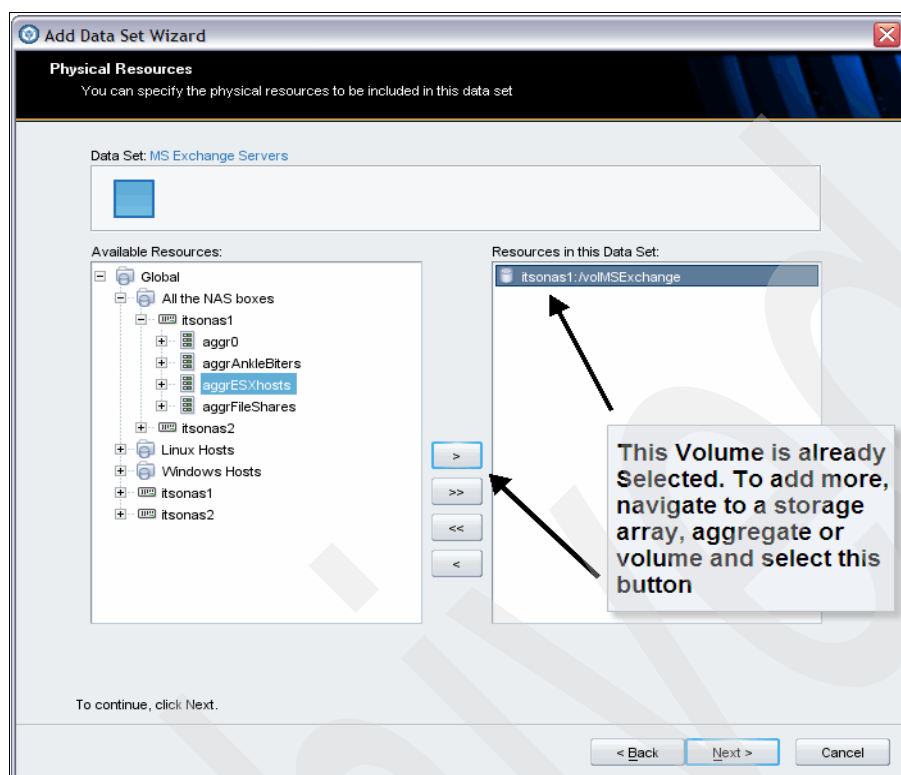


Figure 13-23 Adding resources manually

A data set can contain resources from multiple N series storage systems.

Finally, the Preview window reviews what you wish to do and will inform you if it perceives any warnings or errors. If no warnings or errors are found, you will get a successful result, as shown in Figure 13-24.

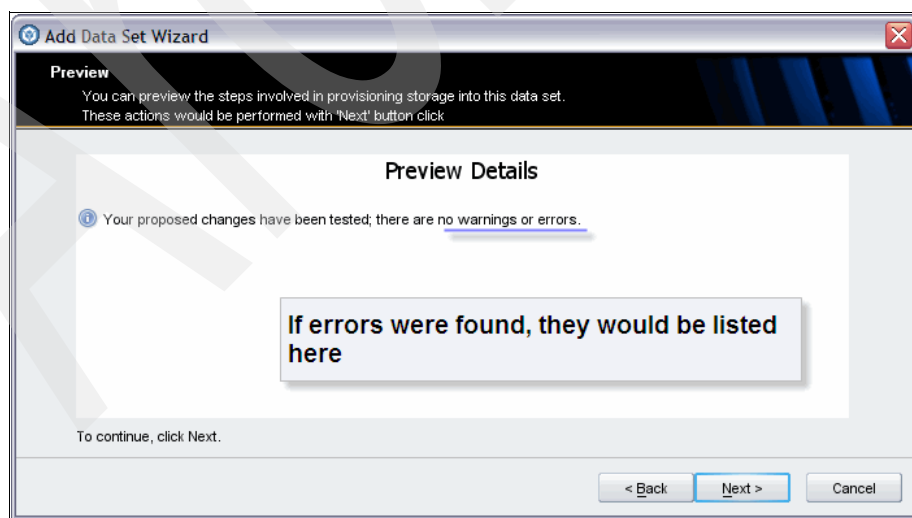


Figure 13-24 Previewing details of tasks to be executed

The final window actually summarizes the tasks that the Protection Manager Wizard will carry out, as shown in Figure 13-25.

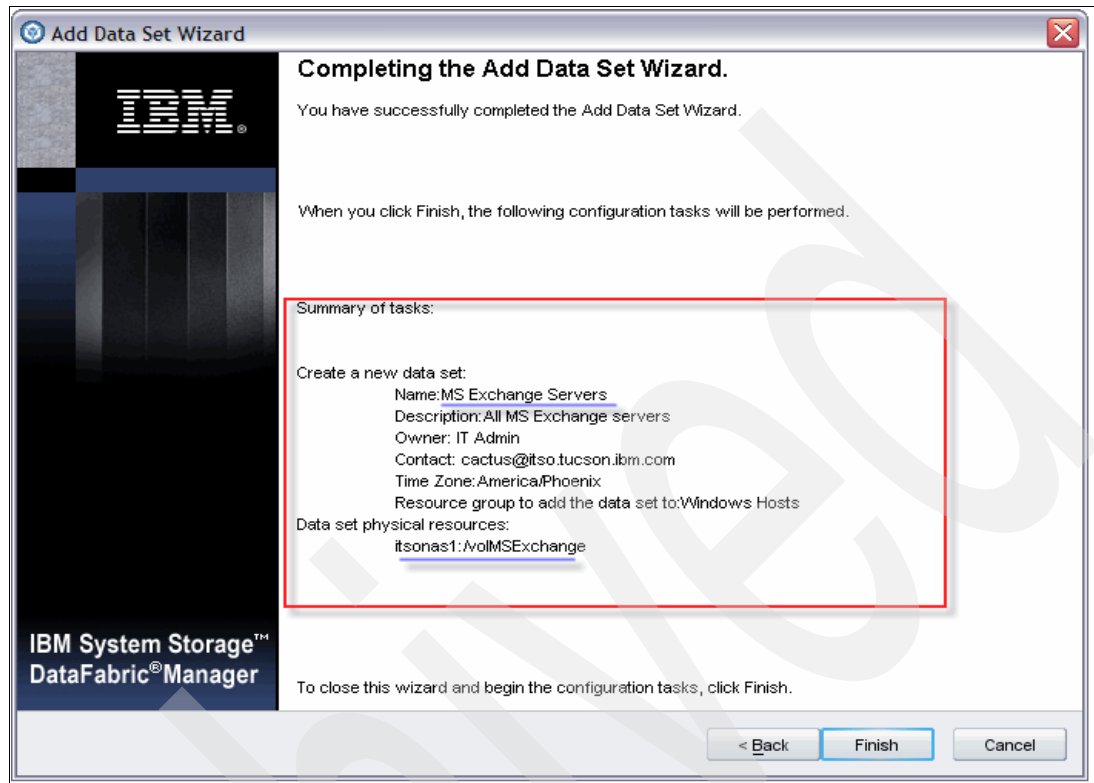


Figure 13-25 Final window for Add Data Set Wizard

Once the data set is created, you will be presented with the Summary window for unprotected data sets. The data set you have just created is ready but is not yet protected, as shown in Figure 13-26. To protect this data set, you need to navigate to the data sets summary window by clicking the **Protect** button.

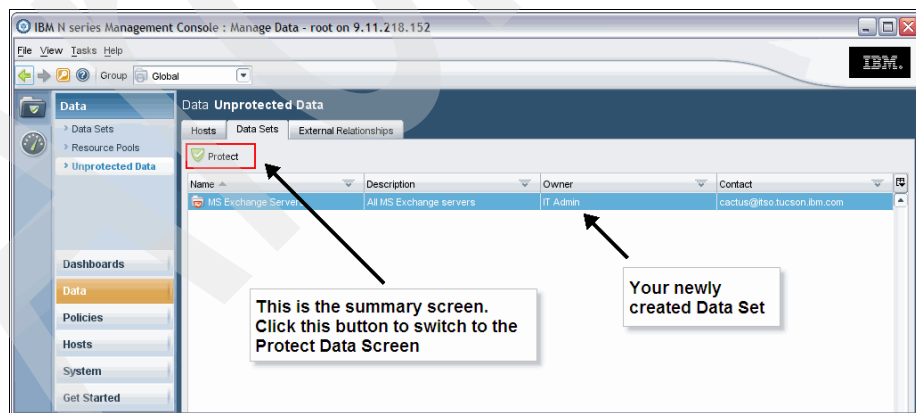


Figure 13-26 Unprotected data sets summary window

13.4.3 Reviewing the resources to be protected

The ownership tab shows all the data sets currently created and a brief summary of the highlighted data set is also shown. From this window, we see that the data set is in the Arizona/Phoenix time zone, there are no issues with the data set (it is conformant), and that it has no protection policy and is currently unprotected.

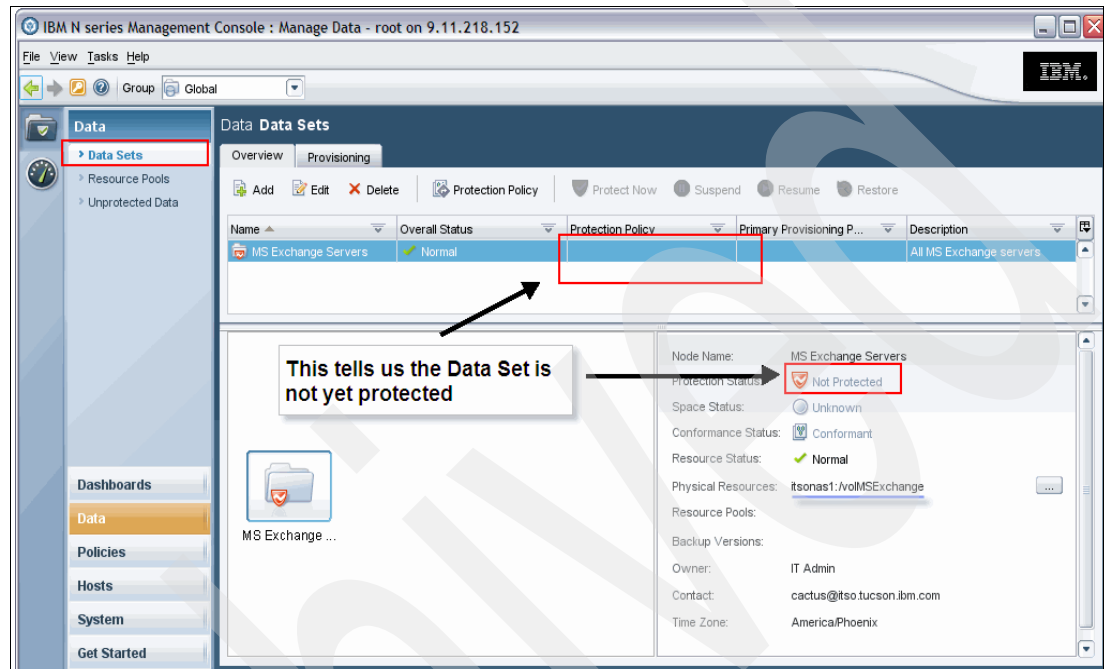


Figure 13-27 Overview of a data set

Before we begin the steps to protect the data set, let us take a look at the resources to be protected. To do this, we click the **Provisioning** tab, as shown in Figure 13-28.

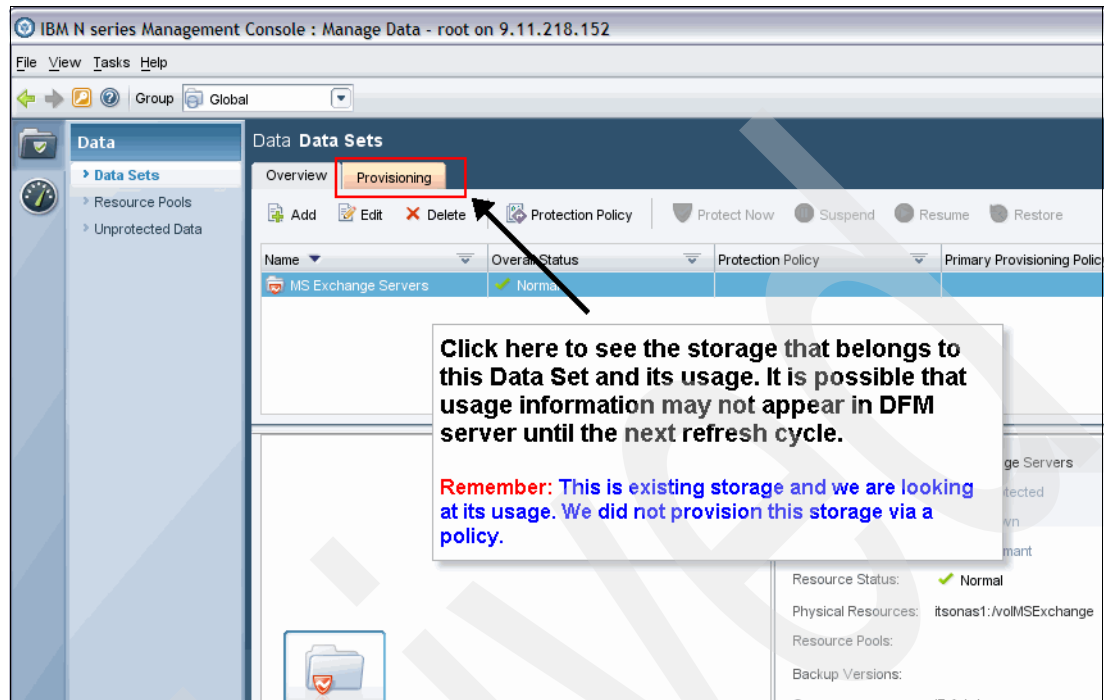


Figure 13-28 Data sets (protected or not) summary window

The provisioning tab for this data set shows the resources that belong to it. Figure 13-29 shows the current space utilization of a resource within the data set. You can look at detailed information about each resource by selecting a different resource in the resources area of the window.

Note: This view may not necessarily show any detail to begin with, as the information about the contents of these resources will only be available at the next DFM Server refresh.

In Figure 13-29, you can see the highlighted volume's details as well as LUNs residing within the volume.

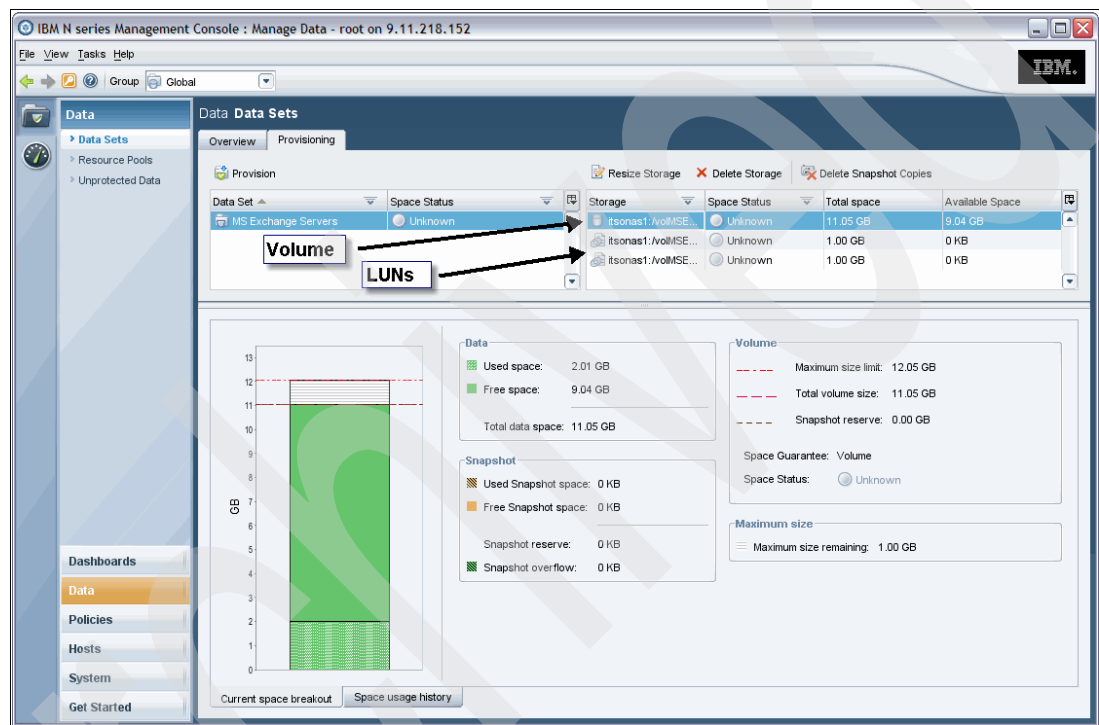


Figure 13-29 Data set provisioning details - current space breakout

Figure 13-30 on page 343 shows the space usage history for each resource within this data set. This information will initially be of little use but will become progressively more useful over days and weeks as trending information is collected by the DFM Server and summarized here by Protection Manager.

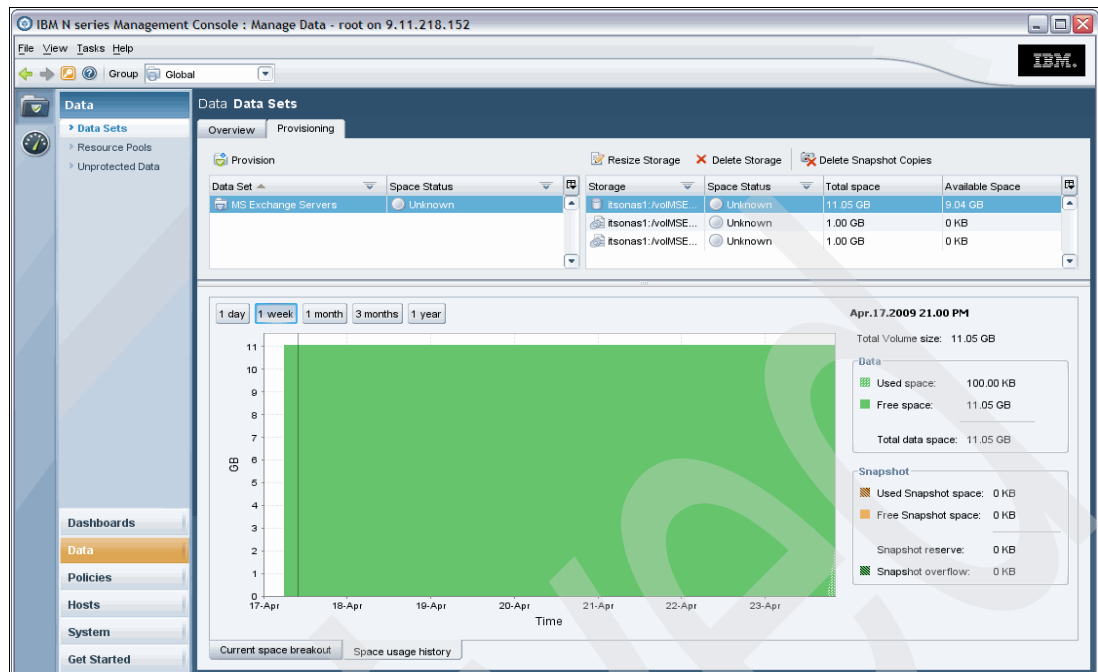


Figure 13-30 Data set provisioning details - space usage history

Note: If you use FilerView to rename the aggregates and volumes for this data set, this action will not cause any problems. The DFM Server is able to track each aggregate, volume, and LUN through name changes. Changes to the names will be reflected in the next refresh cycle.

Figure 13-31 shows the level of protection the data set currently has. In the next section, we assign a protection policy, provision resources to support that policy, and review the job and event logs. We also discuss the use of on demand protection using the **Protect Now** button in 13.4.6, “Protecting the data set manually or on demand” on page 359.

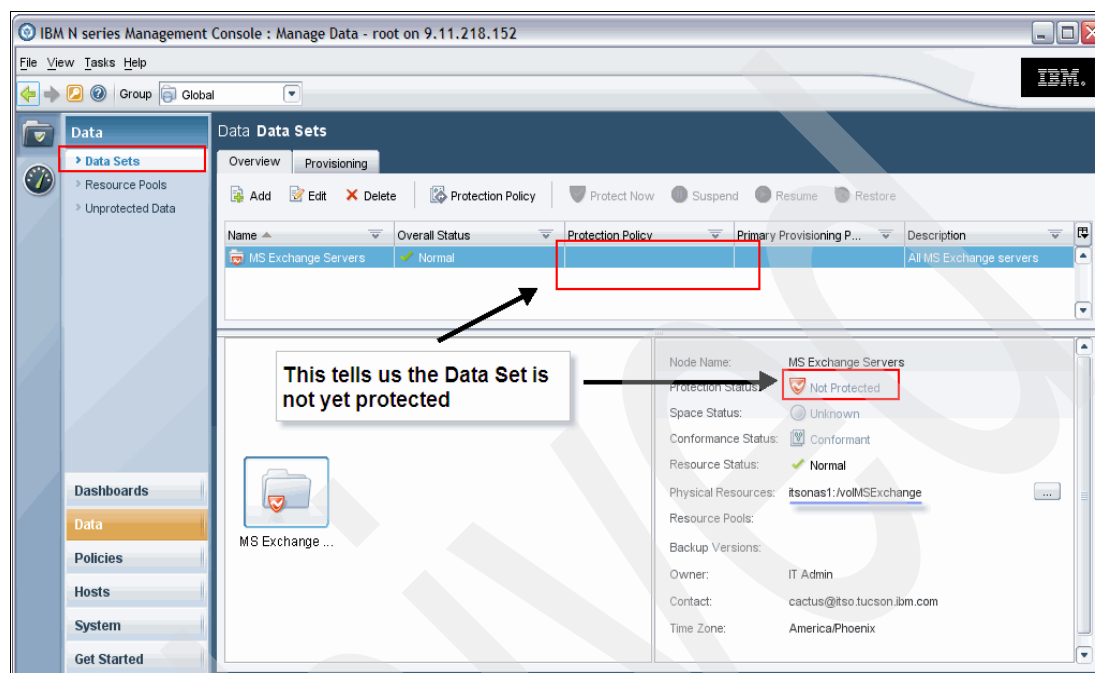


Figure 13-31 Data sets summary window showing the current protection status

13.4.4 Assigning a protection policy to the data set

In this section, we identify an appropriate protection policy and assign it to this data set. To specify a protection policy for this data set, click the **Protection Policy** button, as shown in Figure 13-32 on page 345.

Note: The protection policy you select must match the recovery or Business Continuity outcomes your organization has determined for the application and data you are trying to protect. Protection Manager is not able to predetermine what these outcomes should be based on the type of application or data you have.

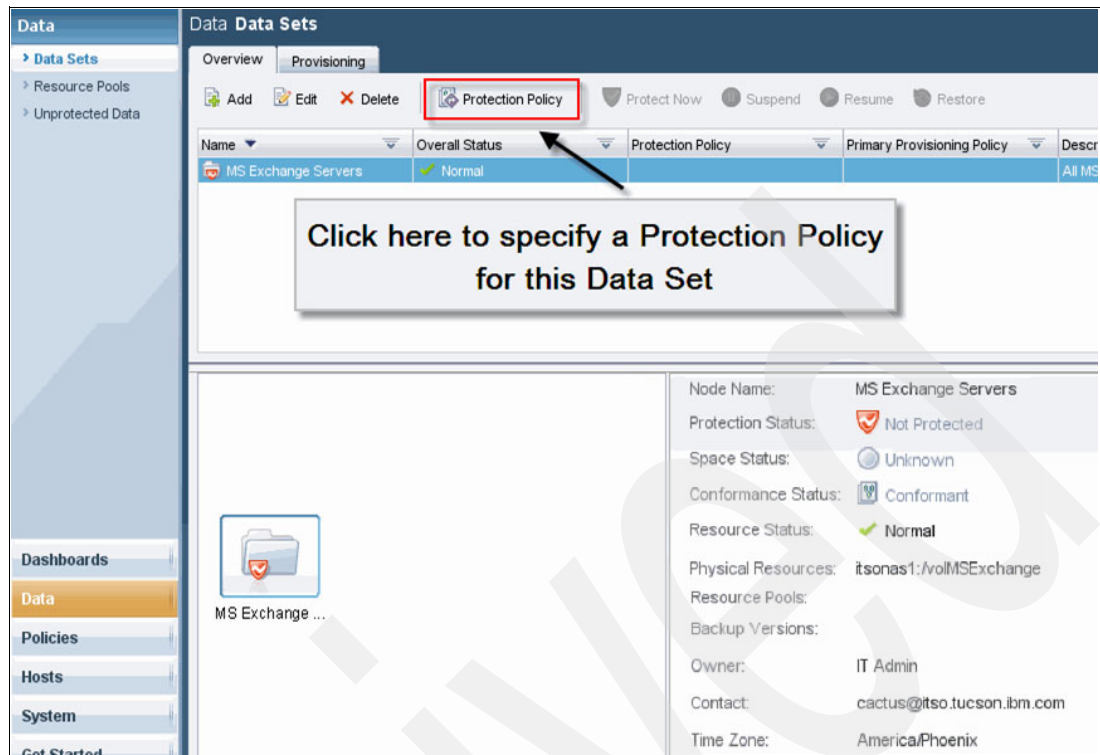


Figure 13-32 Data sets overview window

This action will launch the Data Set Policy Change Wizard, as shown in Figure 13-33, which we use to set the protection policy for this data set.

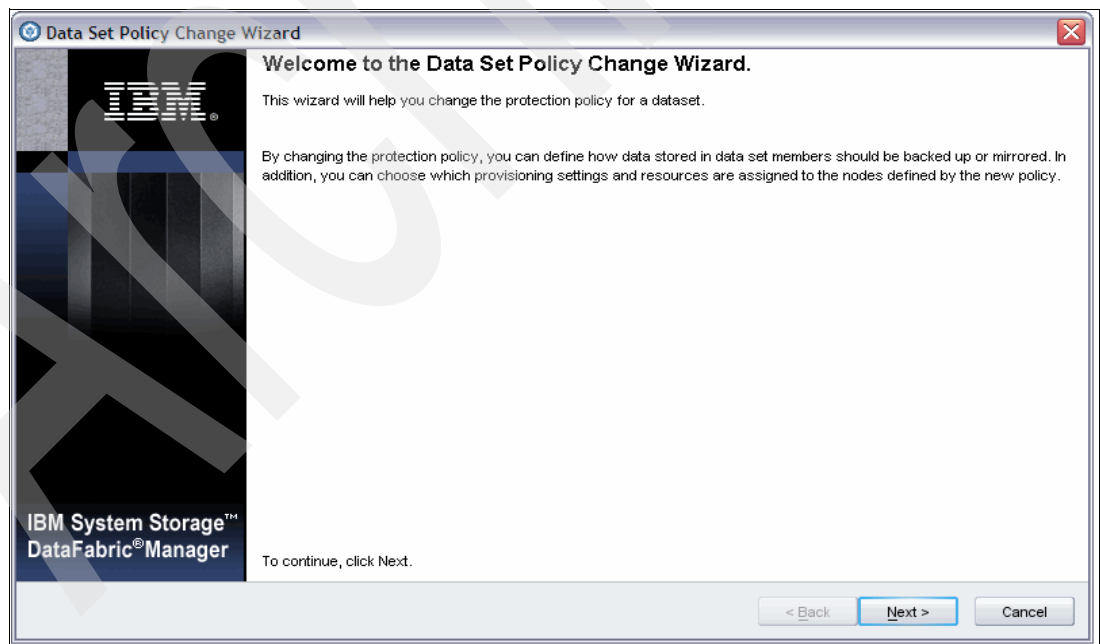


Figure 13-33 Data Set Policy Change Wizard launch window

Clicking **Next** brings us to the Protection Policy window, where we can select an appropriate protection policy. We have a policy called “Mirror MS Exchange LUNs”, as shown in Figure 13-34, which was created by making a copy of the “Chain of Two Mirrors” protection template and then renaming it.

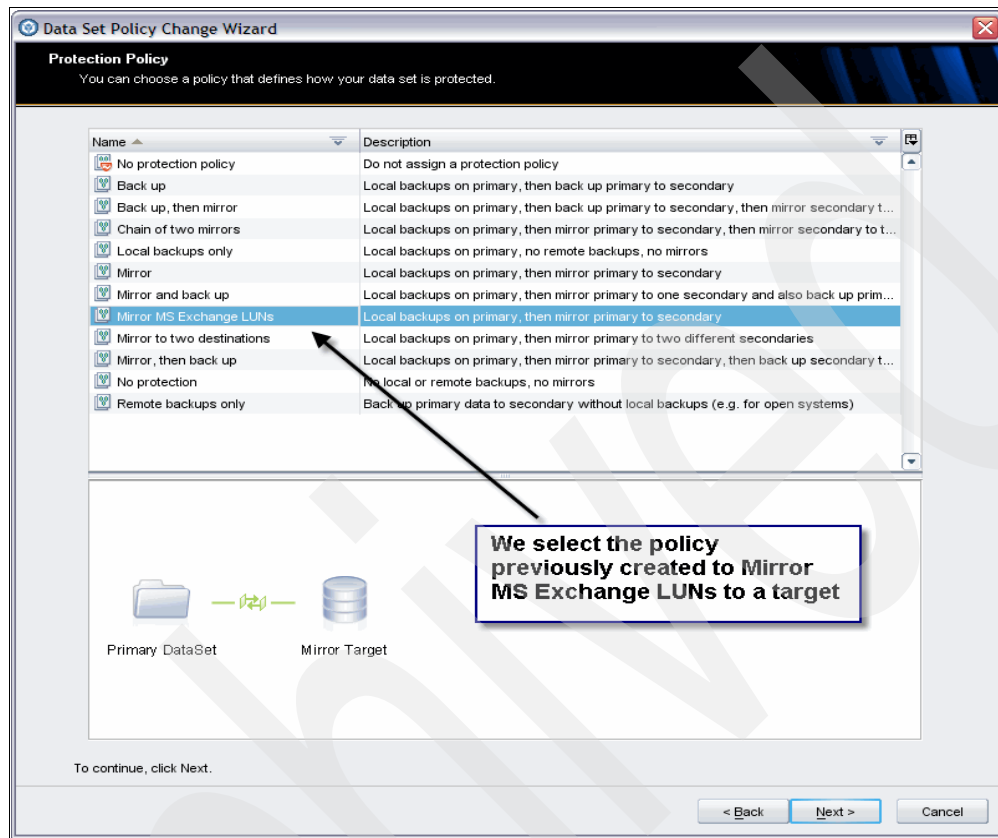


Figure 13-34 Protection policies that can be selected

As the protection schedules were also previously set, as shown in 13.3.1, “Overview of the protection policies templates” on page 326, the wizard is now only interested in determining the target location for the mirror policies, as shown in Figure 13-35 on page 347.

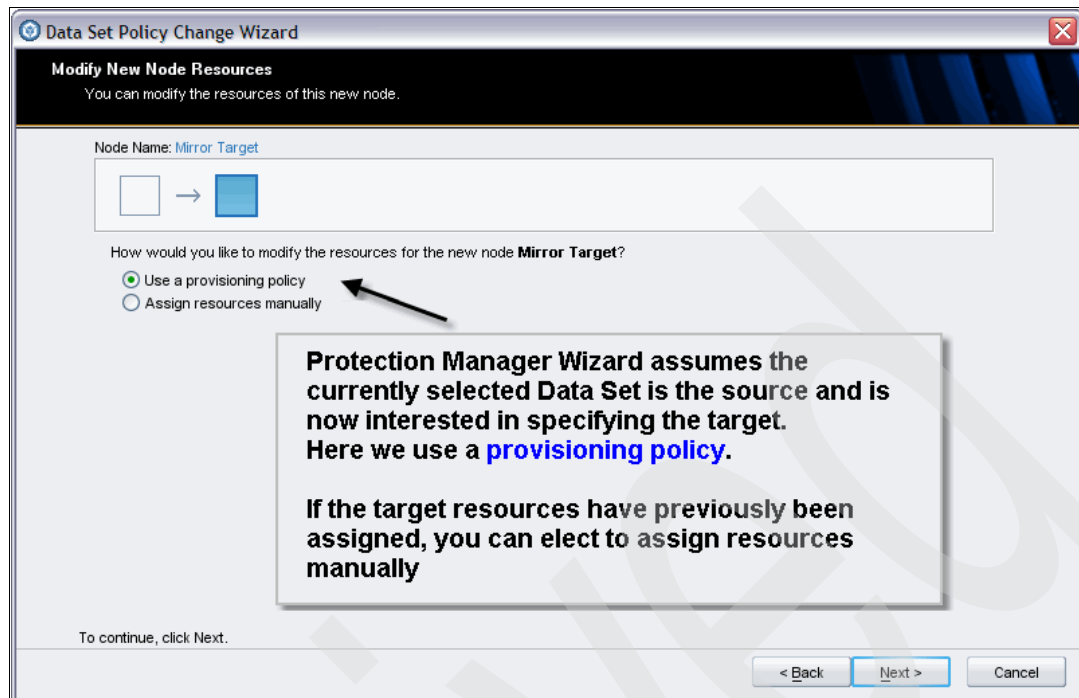


Figure 13-35 Electing to use a provisioning policy to protect the data set

As we plan to provision new targets for this protection policy, we need to select a provisioning policy that correctly matches the desired outcomes. From the drop-down menu shown in Figure 13-36, you can see that the provisioning policies shown are only those of the classification “Secondary” storage type. Depending on the provisioning policy you select, you also need to select a suitable resource pool. The term “Secondary” in this diagram implies that a Resource Pool can be a SnapMirror destination or SnapVault secondary.

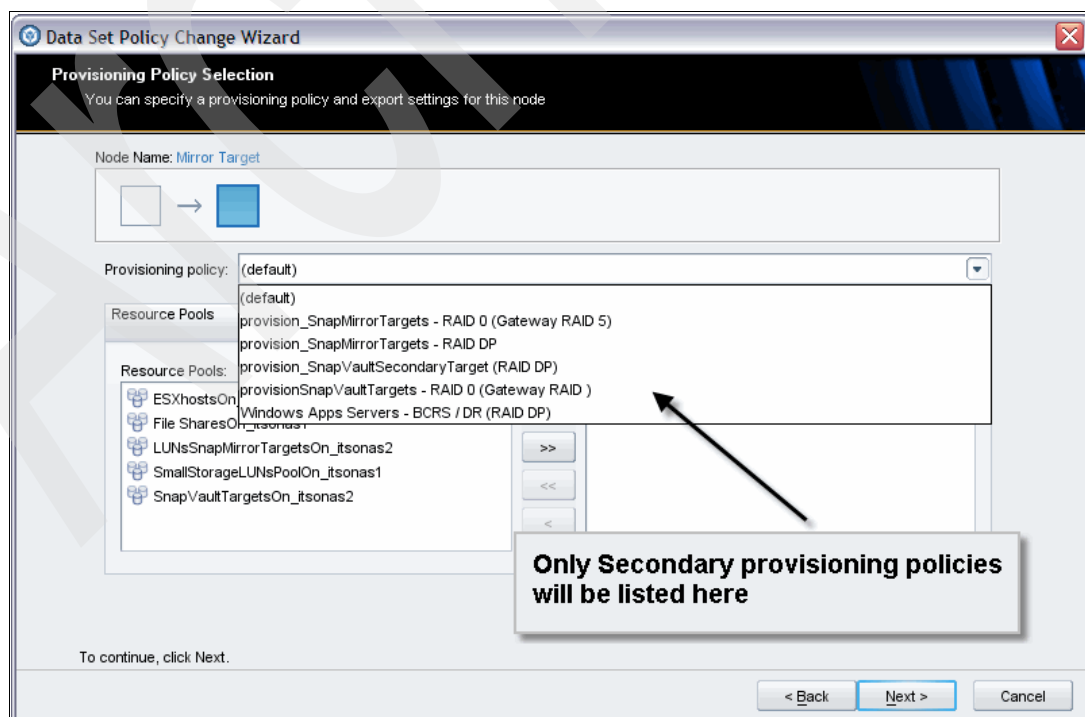


Figure 13-36 List of provisioning policies for Secondary Storage type

When you select a provisioning policy, the wizard will automatically show you which resource pools will match the provisioning requirements as set within that policy. The resource pools that are not compatible will have a “?” next to them, as shown in Figure 13-37.

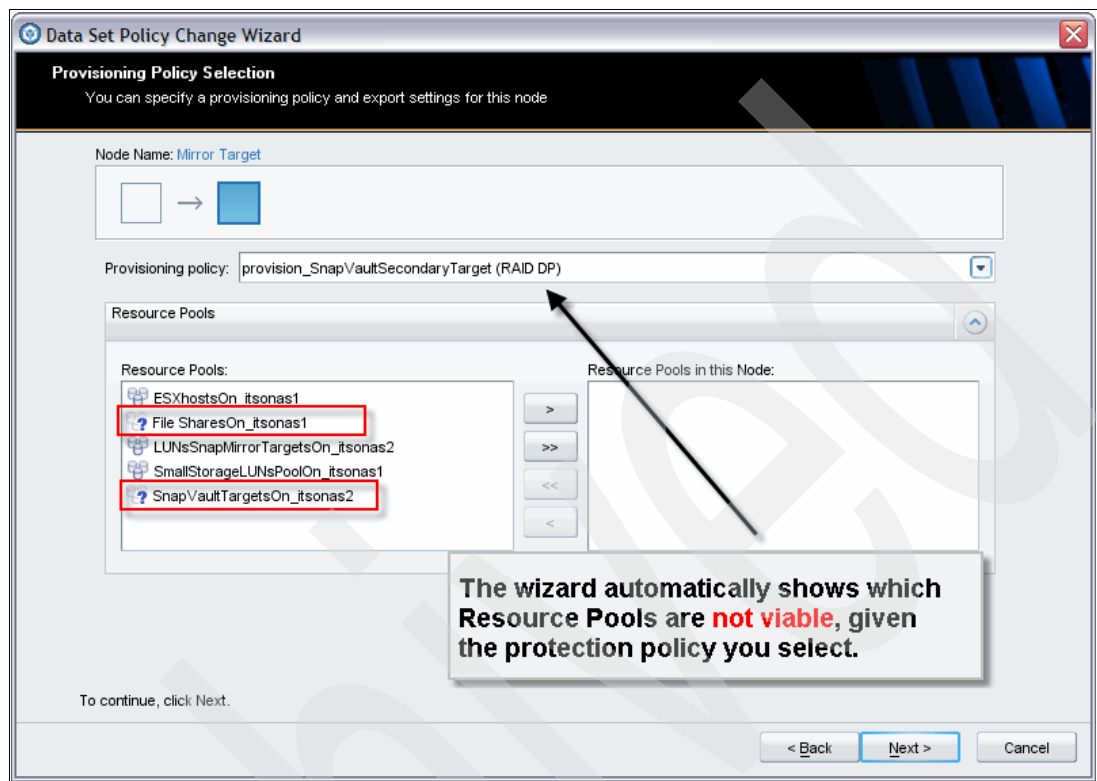


Figure 13-37 List of compatible resources pools for the selected provisioning policy

We select the “SnapMirrorTargets - RAID 0 (Gateway RAID 5)” provisioning policy, as shown in Figure 13-38 on page 349. We then select the resource pool called “LUNsSnapMirrorTargetsOn_itsnas2” as the desired resource pool for our mirror targets to be stored on.

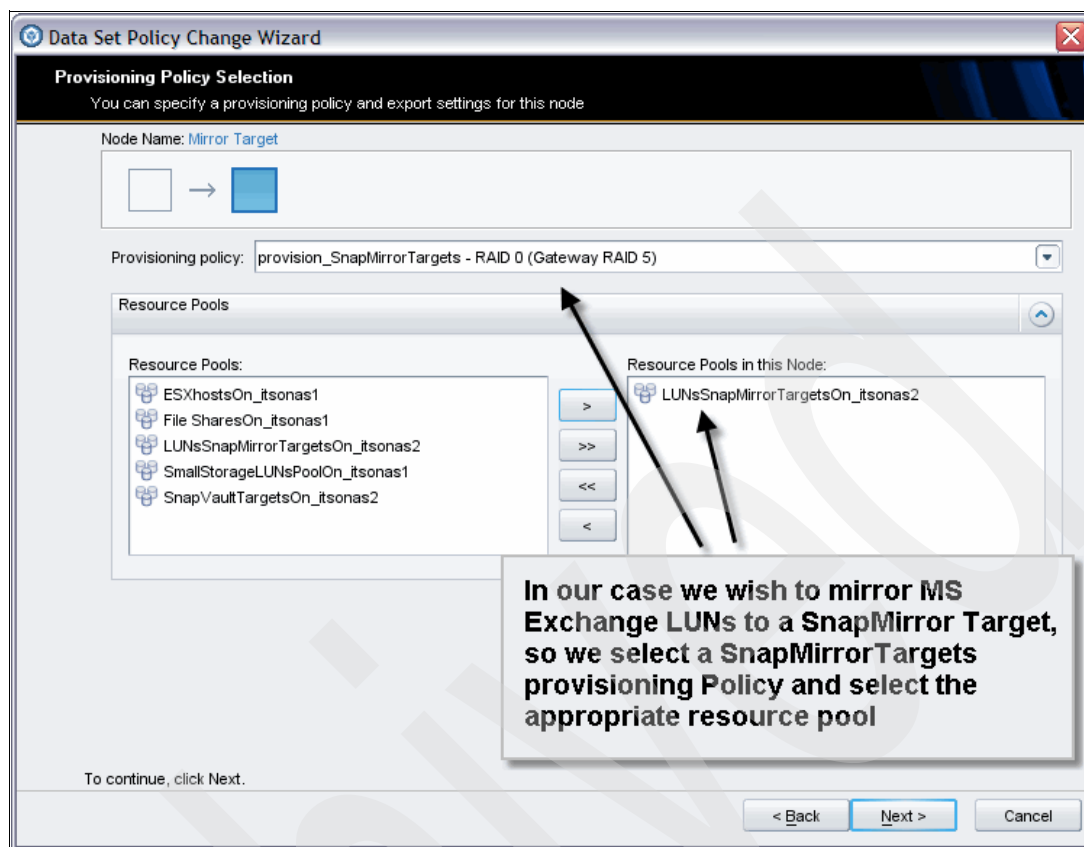


Figure 13-38 Selecting the correct resource pool for the provisioning policy

The wizard gives us the opportunity to specify a vFiler as a target (if there are any), as shown in Figure 13-39.




Figure 13-39 Selecting a vFiler target in the Data Set Policy Change Wizard

The Preview window, shown in Figure 13-40, details all the steps the wizard will take to satisfy the request. It will:

1. Analyze the data set to determine how many aggregates, volumes, and LUNs need to be protected
2. Use the protection policy requirements to determine the scheduling priorities
3. Use the provisioning policies and the resource pools to create the necessary aggregates and volumes to host the mirror targets
4. Use the protection policy to derive the names of the SnapMirror targets that it must create
5. Attempt to set up the relationships between the data sets and the SnapMirror targets
6. Create the scheduling entries to kick off the SnapMirror activities as per the protection schedules

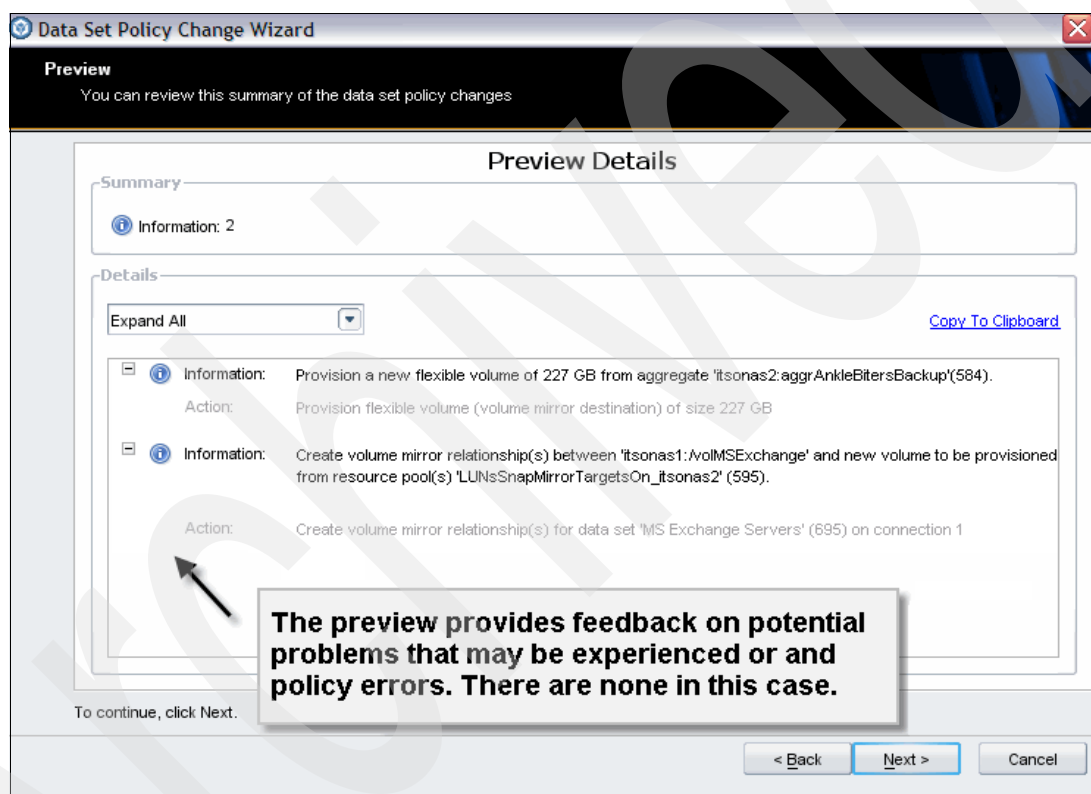


Figure 13-40 Summary window for the Data Set Policy Change Wizard

If, in the previous windows, we have incorrectly selected a resource pool, this window will show us any errors, such as the one shown in Figure 13-41 on page 351.

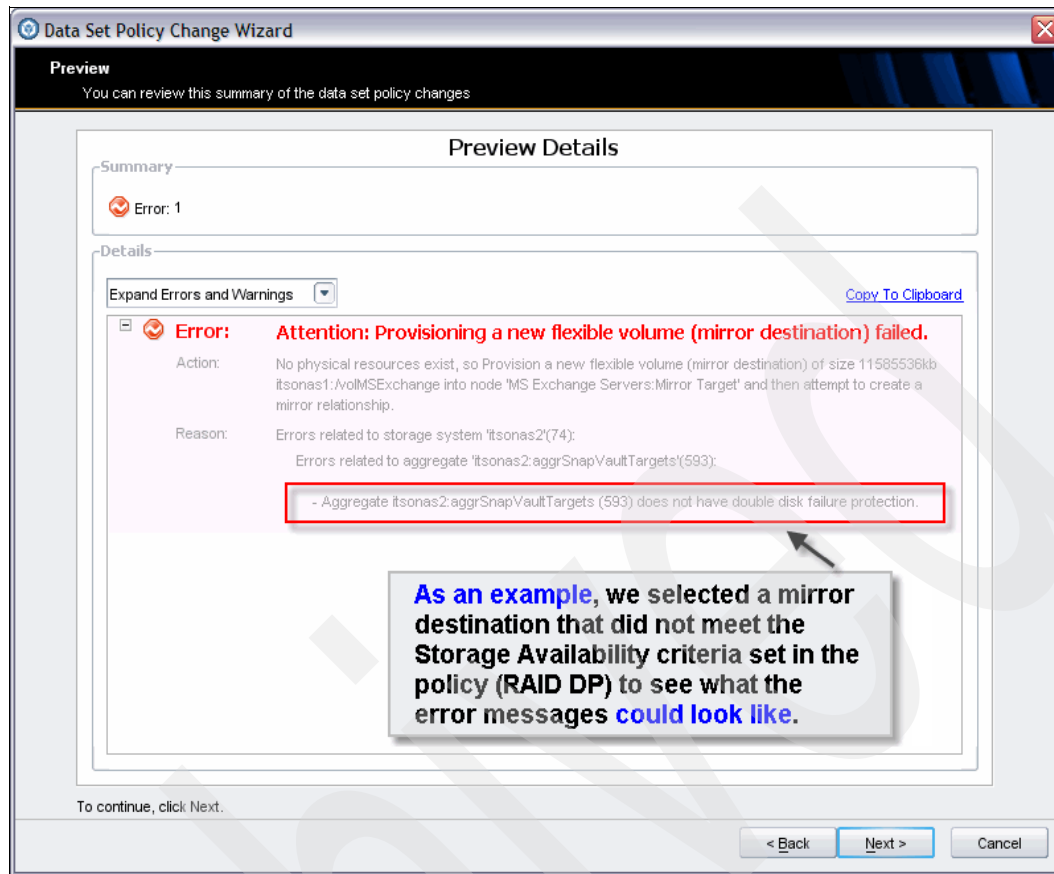


Figure 13-41 Example of errors generated by the Data Set Policy Change Wizard

After any errors are reviewed and rectified, we are presented with the final summary window of actions to be run, as shown in Figure 13-42.

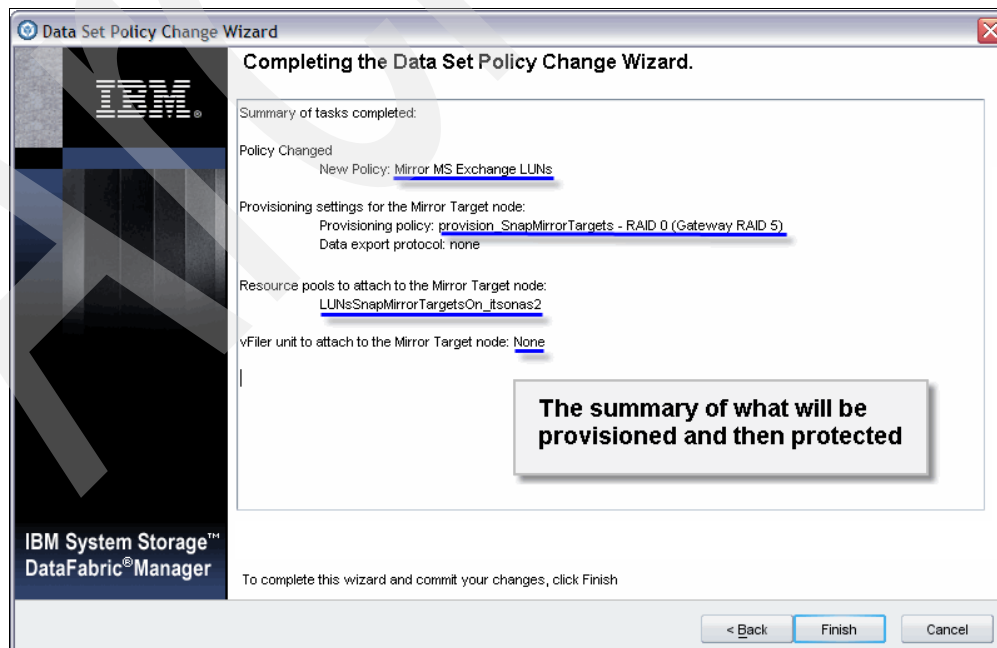


Figure 13-42 Summary window for the Data Set Policy Change Wizard

When you click the **Finish** button, the Data Set Policy Change Wizard will begin executing the scripts it has created to achieve the outcome we specified. We can see that jobs are running and the Conformance status field for the data set has moved from non conformant to “In Progress,” as shown in Figure 13-43. This view is made possible by clicking the Primary DataSet icon in the details pane for this data set.

Before the policy was applied, the data set was in a conformant state in accordance with its then policy requirements. After the protection policy is applied, the data set automatically becomes nonconformant and changes to the “In Progress” state, as shown in Figure 13-43, until it has completed all outstanding steps that are needed for it to be once again in conformant state. In this case, a list of steps are defined to attain that new state and these steps are translated by Protection Manager into jobs that will be executed by the DFM Server and by the storage arrays involved. Then, when all the jobs are completed and verification is satisfied, the data set will once again enter the conformant state.

Review the Jobs and Events windows for details about the jobs that were executed by Protection Manager and the DFM Server as well as the events that were registered. Event messages often refer to a job ID. You can identify the relevant jobs in the jobs window by updating the table preferences to show the Job ID column.

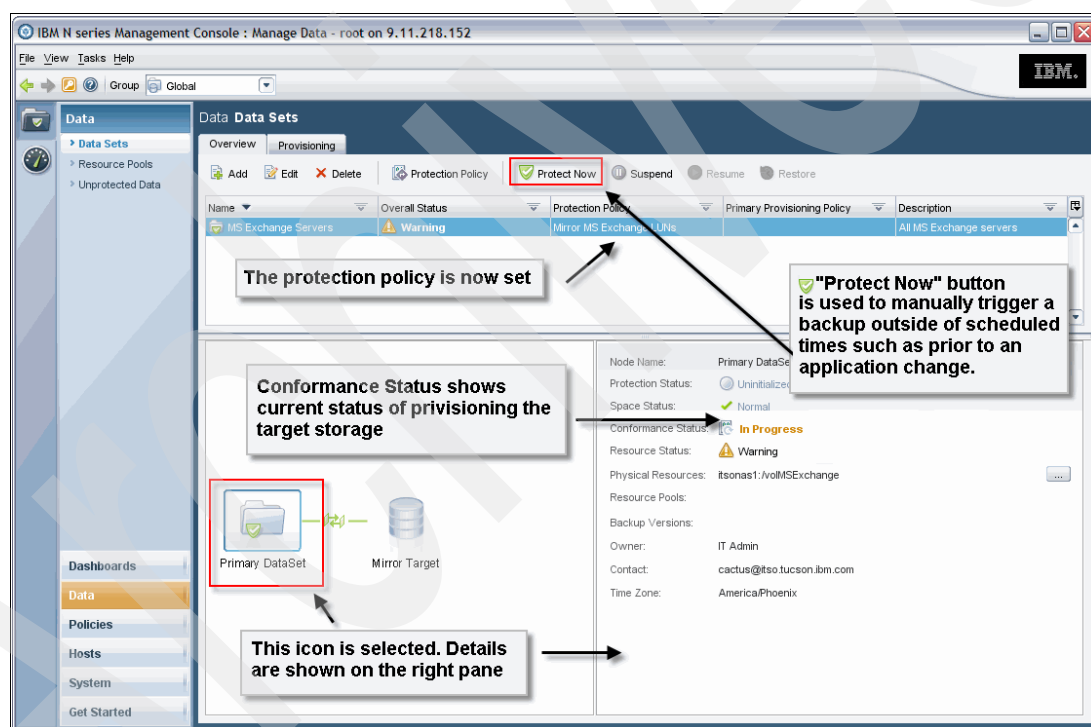


Figure 13-43 Data set protection status overview

The Protect Now button is used to trigger a manual backup of the data set. This is particularly useful when you need a special backup to be run prior to a major change activity, such as a database re-organization, application patch, or code upgrade.

The overview details of the data set shown in Figure 13-43 are visible because the Primary Target icon was selected.

When we click the SnapMirror relationship icon in the same window, we are then presented with details of the mirroring schedule, as shown in Figure 13-44. For completeness, we have also shown the event windows that will be presented when the respective buttons for Relationships and Schedules are clicked. Note that you are able to filter All, Successful, Warning, or Error events while in the Relationship Details window.

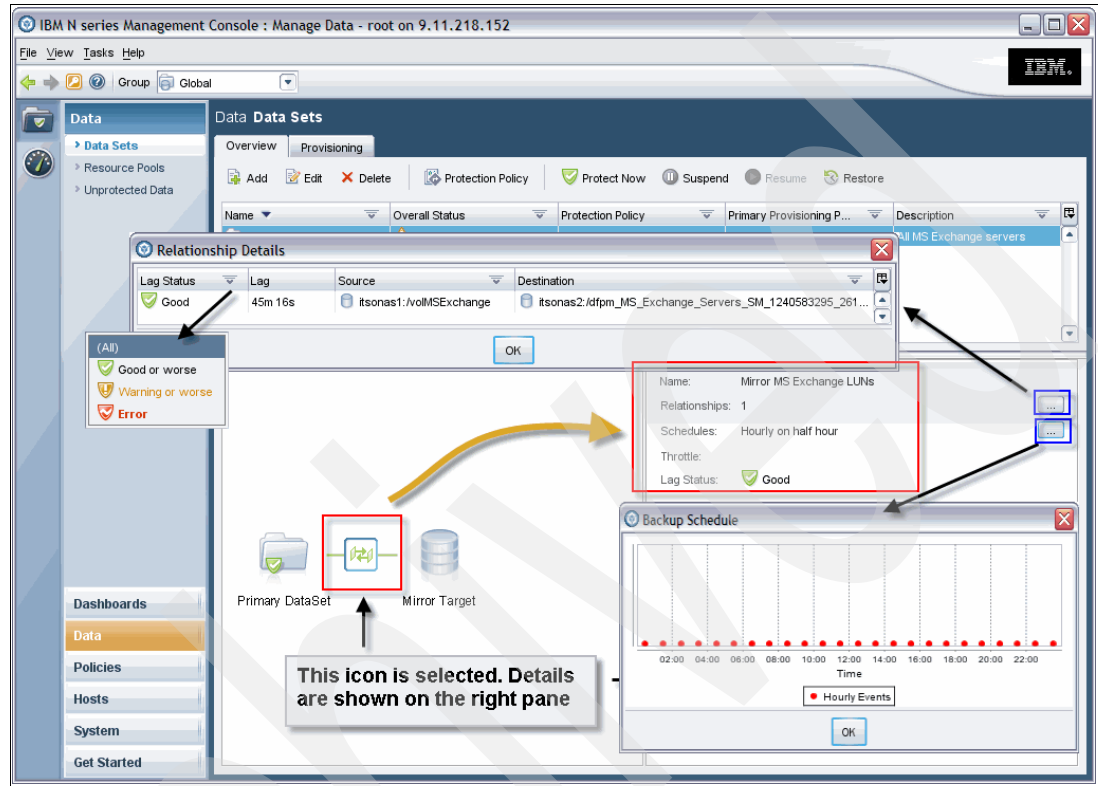


Figure 13-44 Relationship and schedule windows for this data set

When we click the mirror target, as shown in Figure 13-45, we are able to view the physical resources and any Snapshots being used to host the mirror targets as well as the names of the resource pools used. Clicking the respective “details” buttons bring up the Physical Resources details window and the Resource Pools details window.

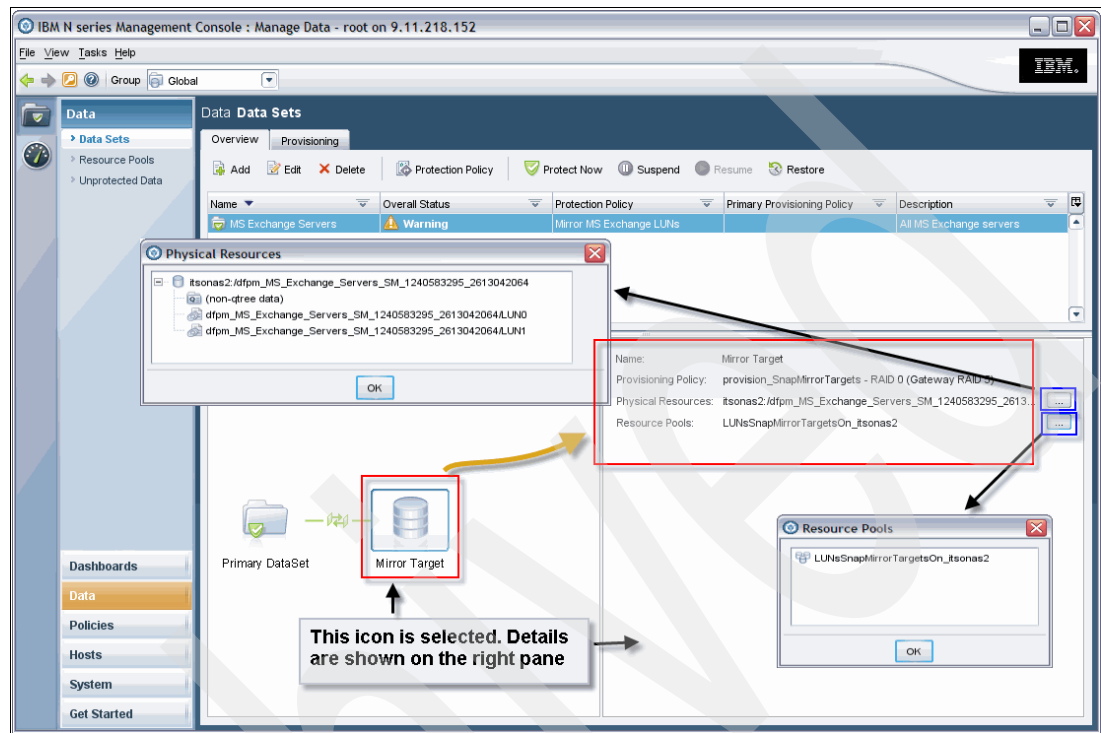


Figure 13-45 Mirror target and details for this data set

When every task is complete, and if there are no errors, we have the Conformance indicator for the data set move to Conformant, as shown in Figure 13-46 on page 355. At this point, all dependencies have now been provisioned and we are finally able to enforce protection for this data set. The data set's Protection Status is currently in the Uninitialized state, reflecting the fact that the background jobs are still in progress.

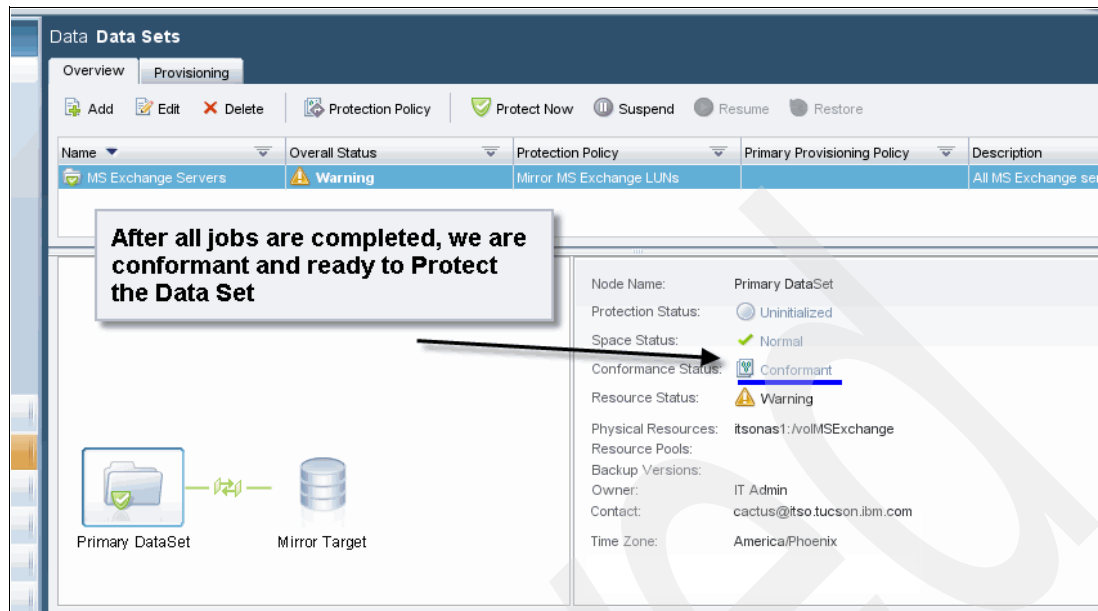


Figure 13-46 Conformant status of the data set

Once the background jobs have completed (including initialization), the status will change to protected, as shown in Figure 13-47.

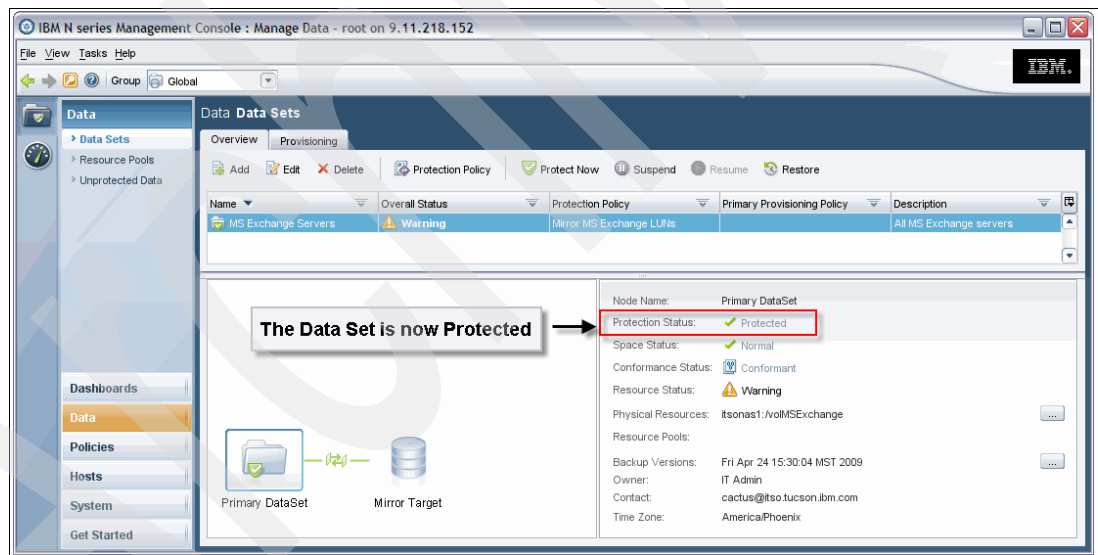


Figure 13-47 Status now shows "Protected"

Finally, if you right-click the data set, you will be presented with a list of possible actions. These actions are also found in the toolbar above the row headings, as shown in Figure 13-48.

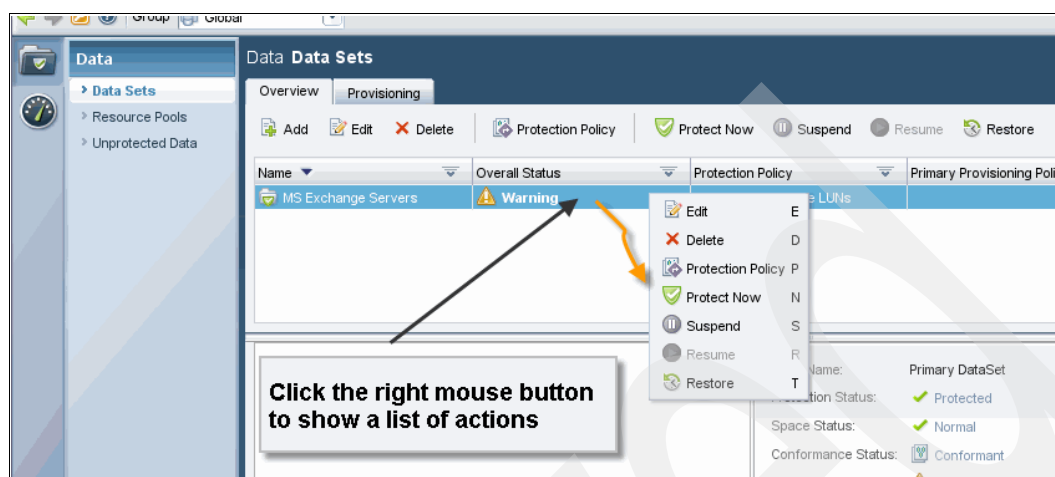


Figure 13-48 Allowable actions for a data set

13.4.5 Reviewing background jobs that ran in order to create the data set

At this point, it is important to note that the Data Set Policy Change Wizard had executed scripts to achieve the outcomes we specified. You can review the jobs that ran and the events generated to achieve this by navigating to the System Jobs window in Protection Manager, as shown in Figure 13-49.

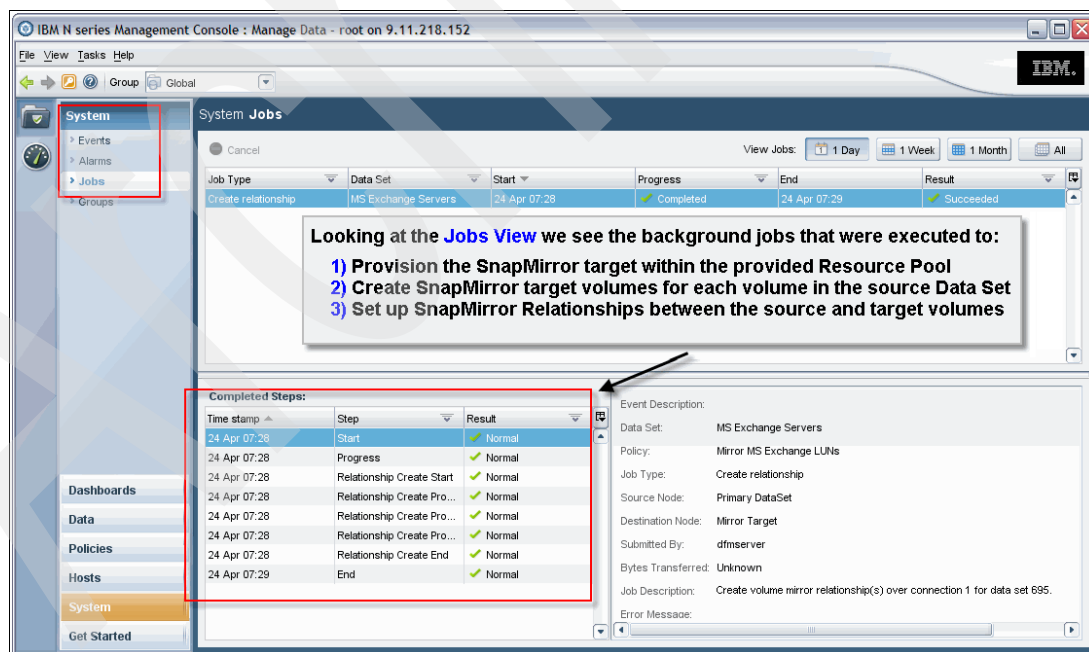


Figure 13-49 Review of jobs and status events in N series Management Console System Jobs window

You can change the display view by clicking the icon at the top right corner of the table, as shown in Figure 13-50 on page 357. This will allow you to include the Job ID as part of the viewable columns. This column is useful when reviewing job status results that refer to specific Job IDs.

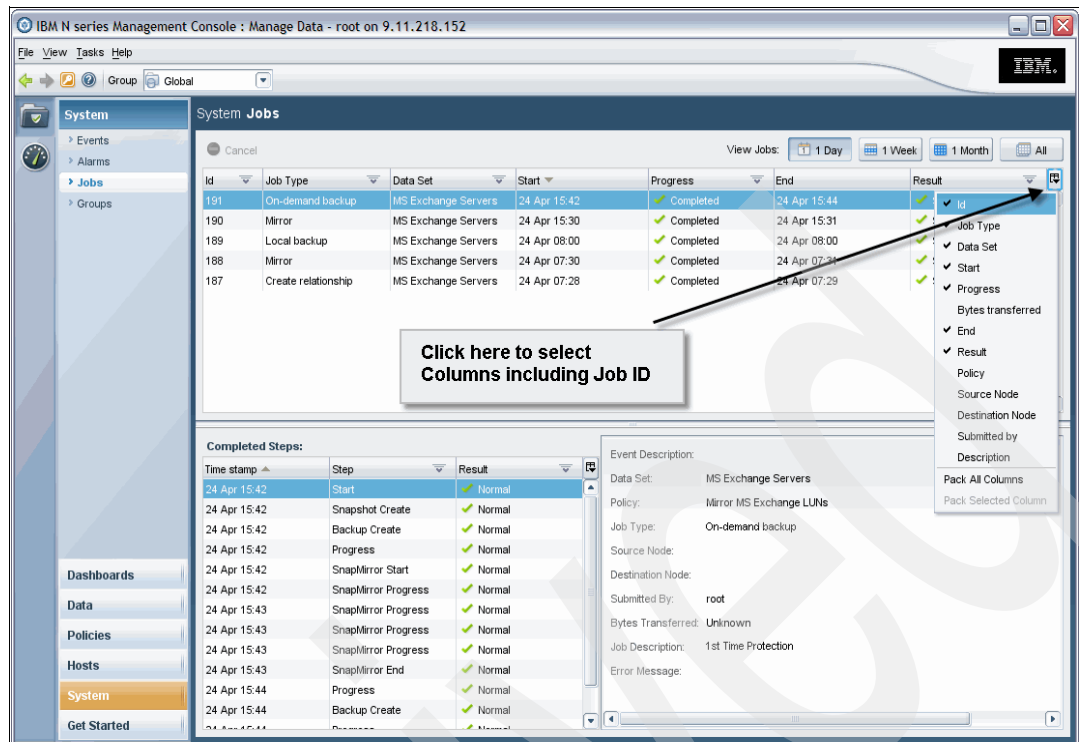


Figure 13-50 Customizing the jobs table view in System Jobs

For each job in this view, you are able to show all the events and messages that were generated. You can see the details of an event by selecting it in the events pane, as shown in Figure 13-51.

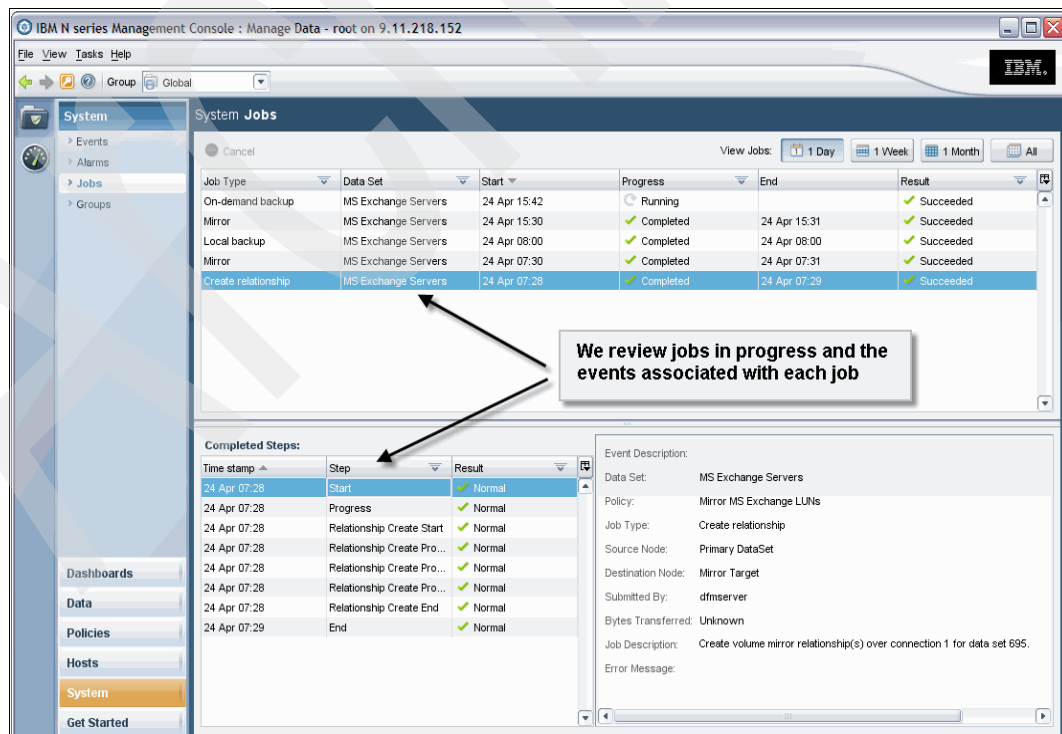


Figure 13-51 Listing event details in the System Jobs view

Figure 13-52 shows you another job selected and the events associated with that job as well as event details for the top listed event.

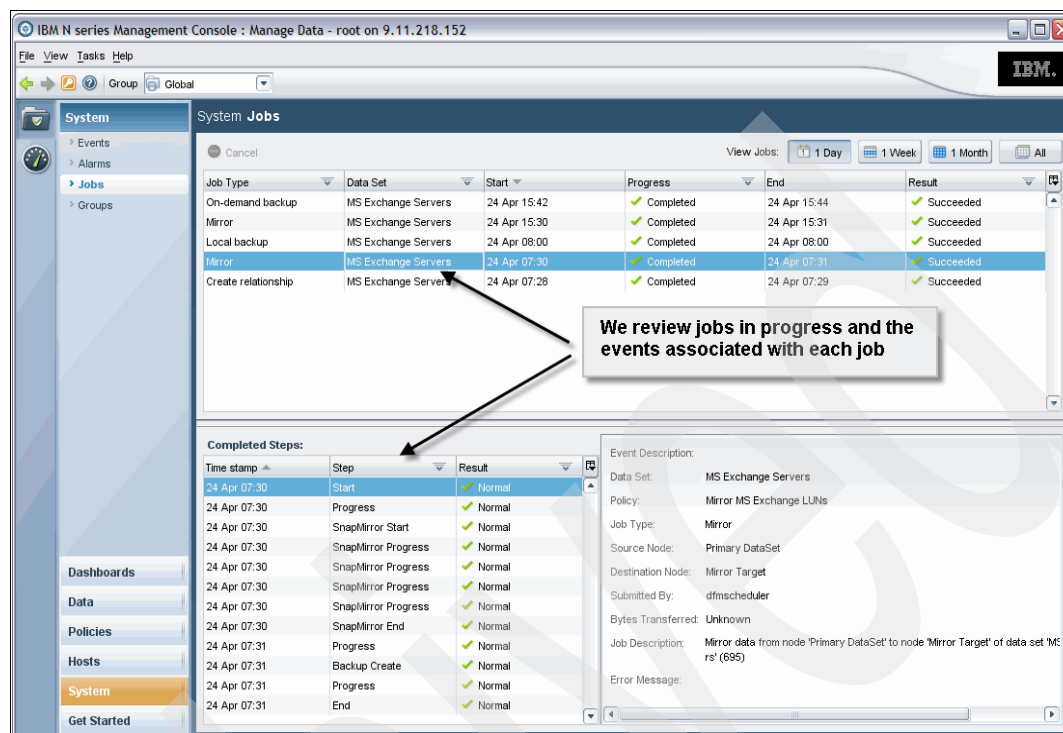


Figure 13-52 Listing more jobs and event details in the System Jobs view

To review a system event, we can navigate to the System Events window and review all events that need to be acknowledged, as shown in Figure 13-53. Here you are able to review the event in detail as well as acknowledge it and take corrective action. You are also able to delete events that are no longer relevant here.

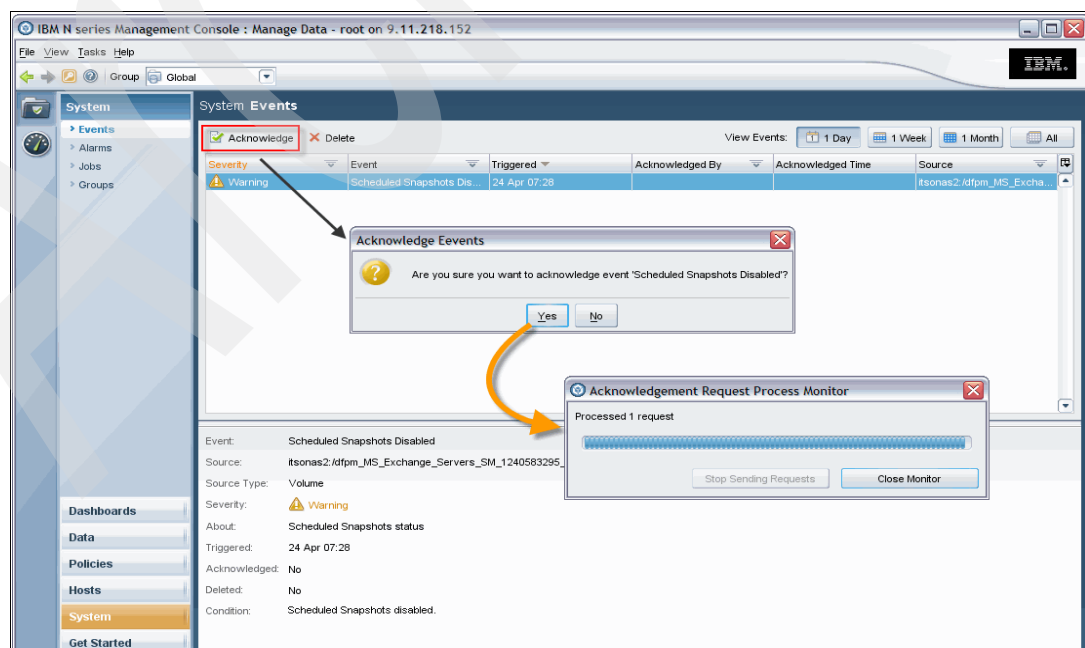


Figure 13-53 Acknowledging an event in the System Events window

13.4.6 Protecting the data set manually or on demand

We are able to manually enforce protection of the data set. We do this by clicking the **Protect Now** button for the selected data set, as shown in Figure 13-54.




Figure 13-54 Selecting Protect Now to enforce protection

This is particularly useful in situations where you need a special backup of the data, perhaps just before the opening of a change window for some maintenance on the application and its data or on the storage system.

The Protect Now button launches a wizard to obtain some basic information for the first time protection event for this data set, as shown in Figure 13-55.

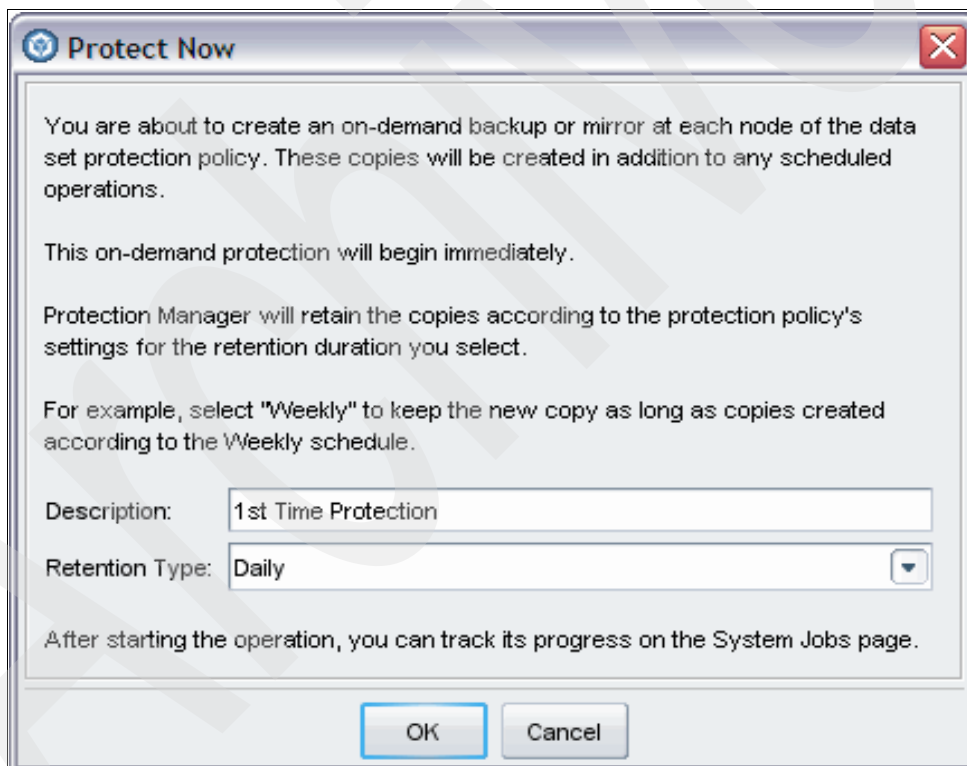


Figure 13-55 Protect Now wizard

You are able to select a retention type for this activity, as shown in Figure 13-56.

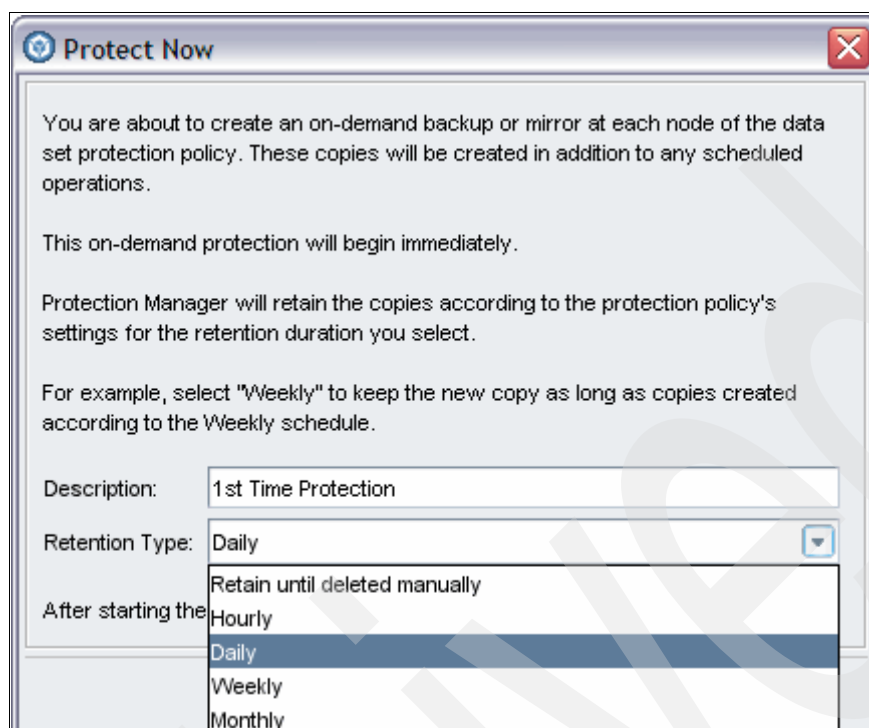


Figure 13-56 Retention periods for manual protection

This special backup will be retained according to the retention policy selected and will be deleted at the end of that retention period. For example, if you select **Weekly**, this backup will be considered equivalent to a weekly backup and will be deleted only at the end of the retention policy for weekly backups for this data set.

13.5 SnapVault, Open Systems SnapVault, and SnapMirror management

In the previous sections, we demonstrated the steps associated with volume or LUN based backups and restores. In this section, we demonstrate how to use SnapVault to back up CIFS and NFS data as well as using Open Systems SnapVault (OSSV) to back up data residing on open systems hosts, such as Windows.

To successfully manage backups of an OSSV host, you first need to register the host with the DFM Server. The following section discusses this step in detail. The essential information you will require is the name or IP address of the OSSV host, the administrator name and password, the Network Data Management Protocol (NDMP) credentials that will be used by the OSSV agent running on the host, and verification that port 1000 on the host will be available for the NDMP agent to run. If this port is not available due to a company security policy or the presence of another service, you can configure the OSSV agent module on the OSSV host to use another IP port. You will then need that information to provide to DFM Server to make sure DFM Server can find the OSSV agent on the OSSV host and manage it.

13.5.1 Adding an OSSV host to a DFM Server

In this section, we demonstrate the steps to add an OSSV host to the DFM Server so that we can then manage data protection from within Protection Manager. An OSSV host is a host on which you have been successful in installing and configuring the OSSV agent.

In Operations Manager V3.7.1 and earlier, the OSSV Agent is currently available for the following operating systems:

- ▶ Microsoft Windows
- ▶ AIX®
- ▶ Linux
- ▶ Solaris

In the Hosts section of N series Management Console, we click the OSSV section to select the OSSV hosts pane. Here we can click **Add** to add a new OSSV host, as shown in Figure 13-57.

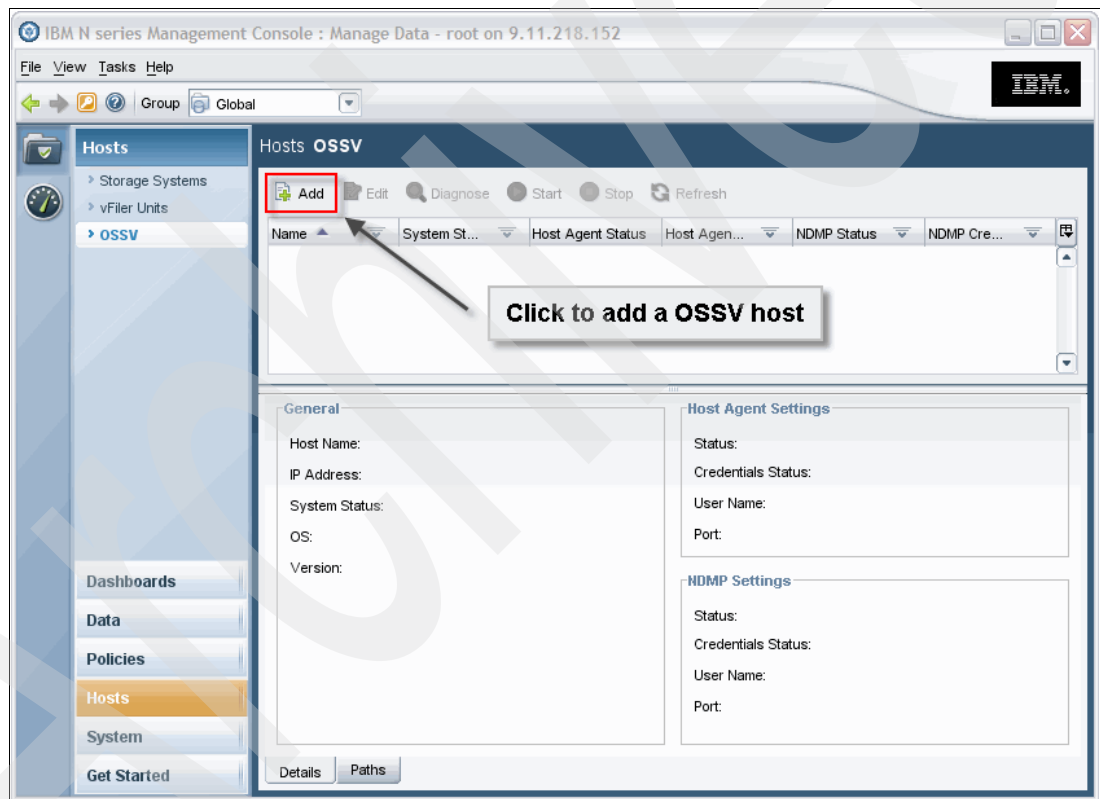


Figure 13-57 OSSV hosts pane

Note: OSSV uses Network Data Management Protocol (NDMP) to copy data from the host to a SnapVault backup target known as a SnapVault Secondary. You will need to have the user name and password that was set on the OSSV host to enable the DFM Server to connect to it and coordinate backups.

The Add button will launch the Add OSSV Host Wizard, as shown in Figure 13-58.

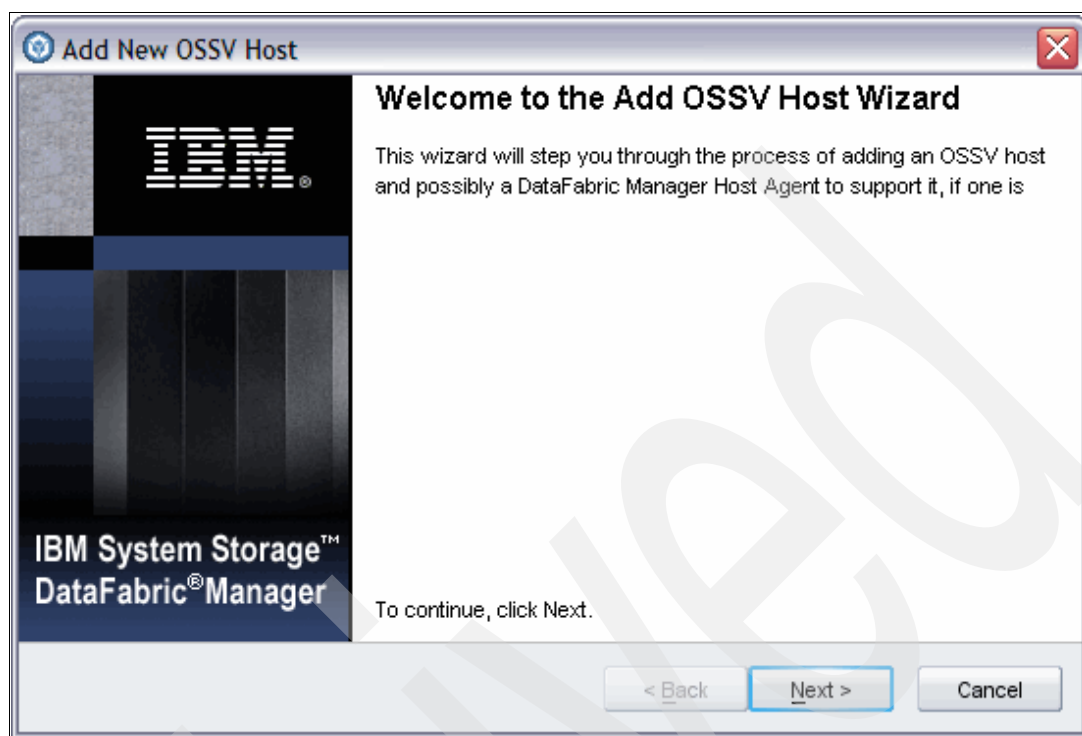


Figure 13-58 Add OSSV Host Wizard

After you click **Next**, you will be presented with a window asking for the name or Fully Qualified DNS Name (FQDN). If the host is not able to be located by name, you can provide its IP address, as shown in Figure 13-59.




Figure 13-59 Providing the IP address of the new OSSV host

Once you have specified the IP address, click **Next**. In Figure 13-60 on page 363, you have the option to ask Protection Manager to manage the credentials of the Operations Manager Host Agent on the host. If you select **Yes, manage the credentials for me**, you will be asked for the operating system admin credentials, as shown in Figure 13-62 on page 364.



Figure 13-60 Letting the DMF Server manage the host agent credentials

In this manner, the DFM Server will log in to the OSSV host and set the admin credentials and store the information in the DFM Server's database.

Note: An OSSV Agent can only be managed and configured by a DFM Server if the host running the OSSV Agent also has Operations Manager Host Agent installed. However, Operations Manager is able to communicate with the host, query its file system, and coordinate backups and restores with the presence of OSSV host, even without using the Operations Manager Host Agent through the NDMP protocol.

If you elect to click **No, I will supply the credentials myself**, as shown in Figure 13-61, you will be prompted to put in the credentials, as shown in Figure 13-62.



Figure 13-61 Supplying credentials

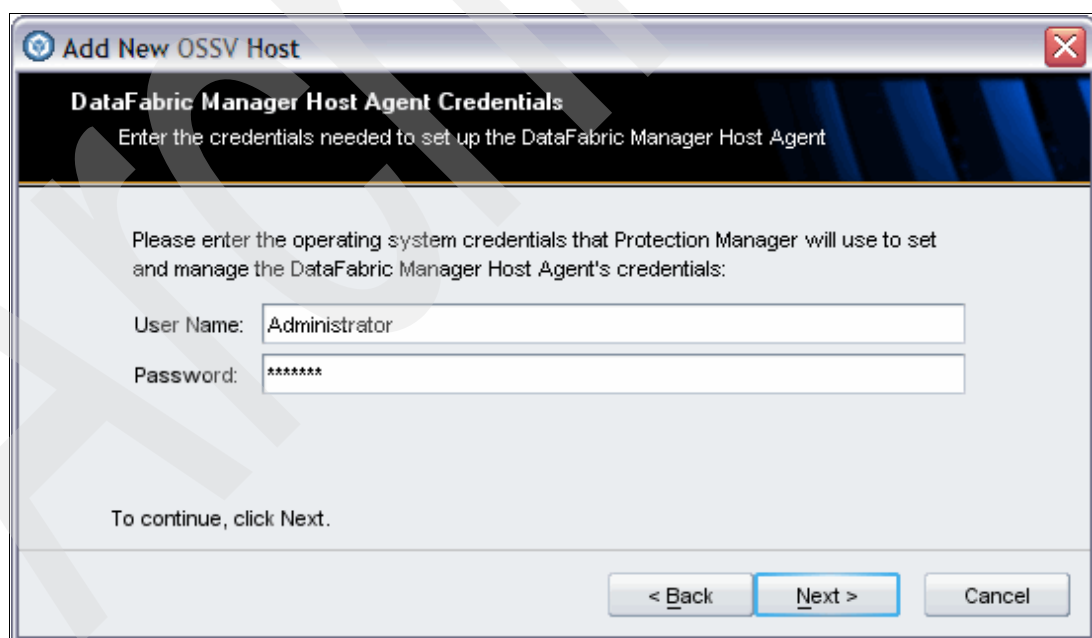


Figure 13-62 Entering Operations Manager Host Agent credentials

Enter the user name and password so that the DFM Server can log on to the host and manage the Operations Manager Host Agent's credentials.

Add New OSSV Host

DataFabric Manager Host Agent Credentials
Enter the credentials needed to set up the DataFabric Manager Host Agent

Please enter the DataFabric Manager Host Agent credentials used for this OSSV's Host Agent:

User Name:

Password:

To continue, click Next.

< Back Next > Cancel

Figure 13-63 Operations Manager Host Agent credentials

Operations Manager will attempt to contact the host using the credentials provided. If the DFM Server was unable to contact the Operations Manager agent on the OSSV host, it could be that it is not running or not installed, as shown in Figure 13-64. You will be prompted to supply the NDMP user name, password, and port ID to use by the DFM Server when contacting the host, as shown in Figure 13-64.

Note: While in this demonstration we use Administrator or root user accounts, this practice is not advisable in production environments. Ensure you have created designated user accounts for this purpose and employ them.

Add New OSSV Host

NDMP Credentials
Enter the credentials needed to set up NDMP

OSSV Status: Unknown

If the host's OSSV status is Unknown, this means that the DataFabric Manager Host Agent was unavailable and can not provide the current status

NDMP Credentials

User Name:

Password:

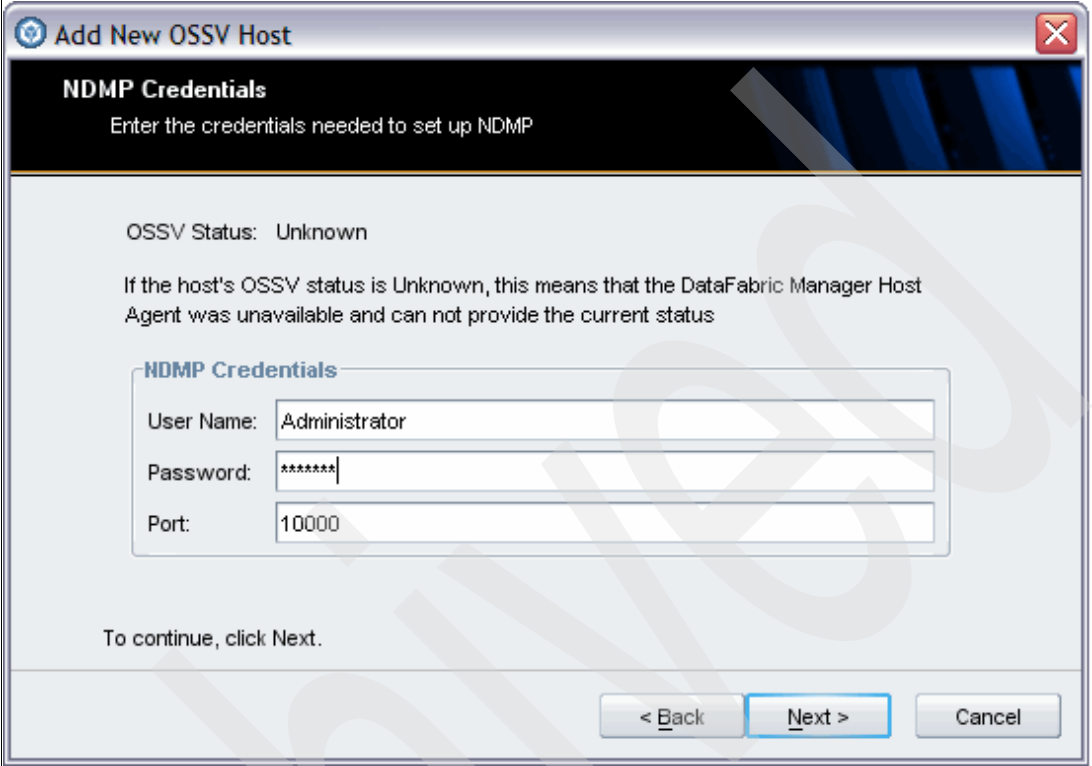
Port:

To continue, click Next.

< Back Next > Cancel

Figure 13-64 Adding a new OSSV host when the OSSV status is unknown

You will need to provide credentials to enable the DFM Server to communicate with the OSSV host agent, as shown in Figure 13-65. The credentials must match the credentials you provided to the OSSV agent when you installed it or last configured it on the host.



The screenshot shows a Windows-style dialog box titled "Add New OSSV Host". Inside, there's a section titled "NDMP Credentials" with the instruction "Enter the credentials needed to set up NDMP". Below this, it says "OSSV Status: Unknown" and provides a note: "If the host's OSSV status is Unknown, this means that the DataFabric Manager Host Agent was unavailable and can not provide the current status". A form titled "NDMP Credentials" contains three fields: "User Name:" with the value "Administrator", "Password:" with masked characters "*****", and "Port:" with the value "10000". At the bottom, it says "To continue, click Next." and has three buttons: "< Back", "Next >" (which is highlighted with a blue border), and "Cancel".

Figure 13-65 NDMP credentials

The summary window shown in Figure 13-66 indicates that the OSSV host was successfully added to the DFM Server.

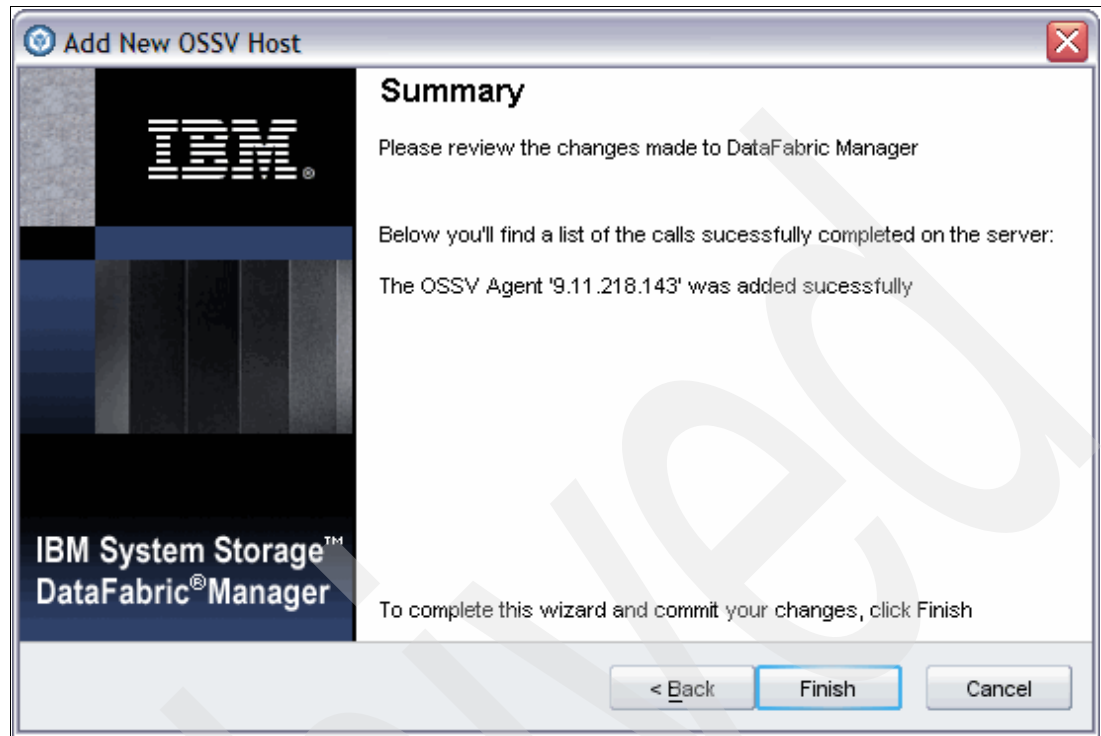


Figure 13-66 Summary window for Add New OSSV Host Wizard

This means that you are now able to create a data set that includes OSSV host folders and assign a protection policy to it. We will do this task in 13.5.3, “Creating a OSSV data set” on page 381.

As you can see in Figure 13-67 on page 369, although we added the OSSV host using an IP address, the DFM Server was able to query the agent and obtain its actual name.

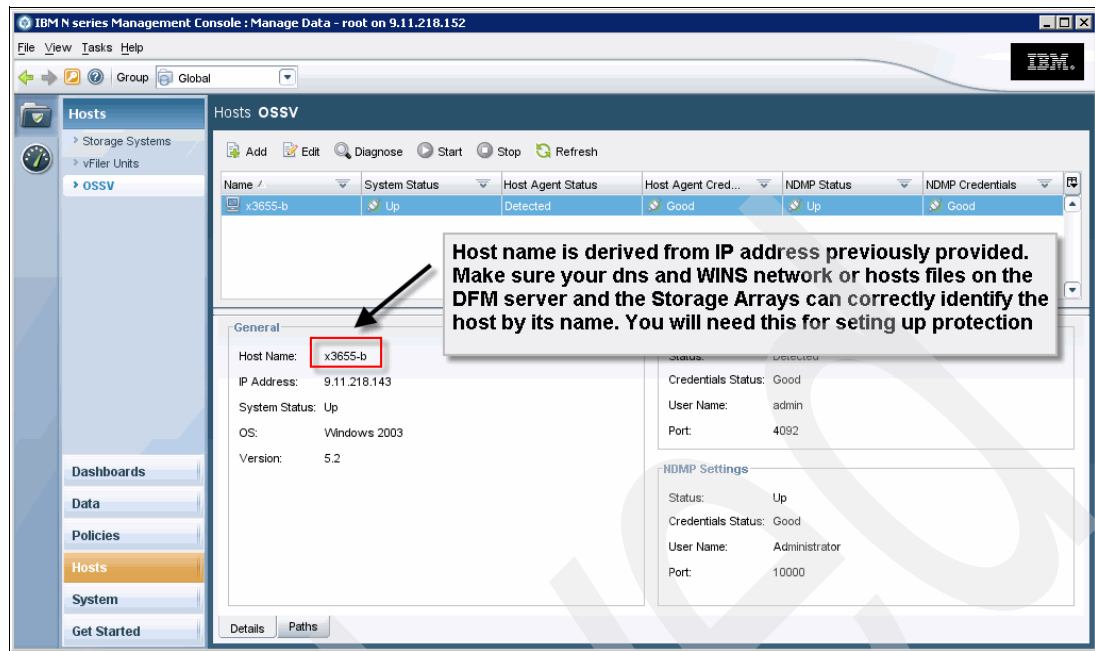


Figure 13-67 The OSSV host name is derived by DFM Server by querying the host

We are now able to interrogate the host, review the file system, and identify data that needs to be protected, as shown in Figure 13-68. The Details tab in Figure 13-68 shows the status of the highlighted host and its NDMP status.

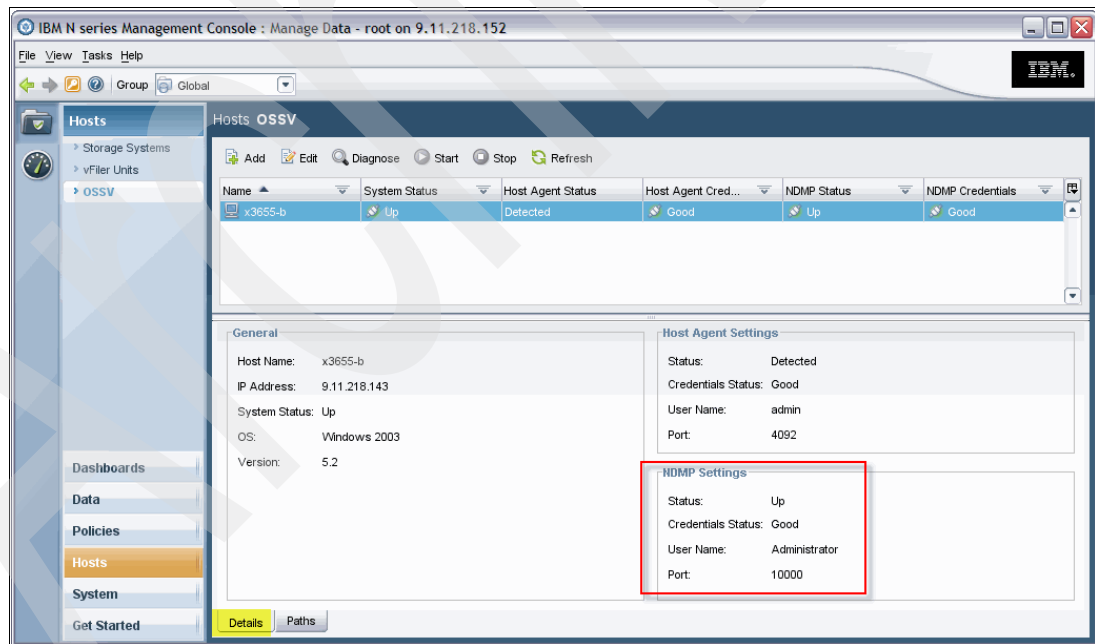


Figure 13-68 Identifying protected data

The Paths tab selected in Figure 13-69 shows the highlighted host's file system and the current protection status of its data.

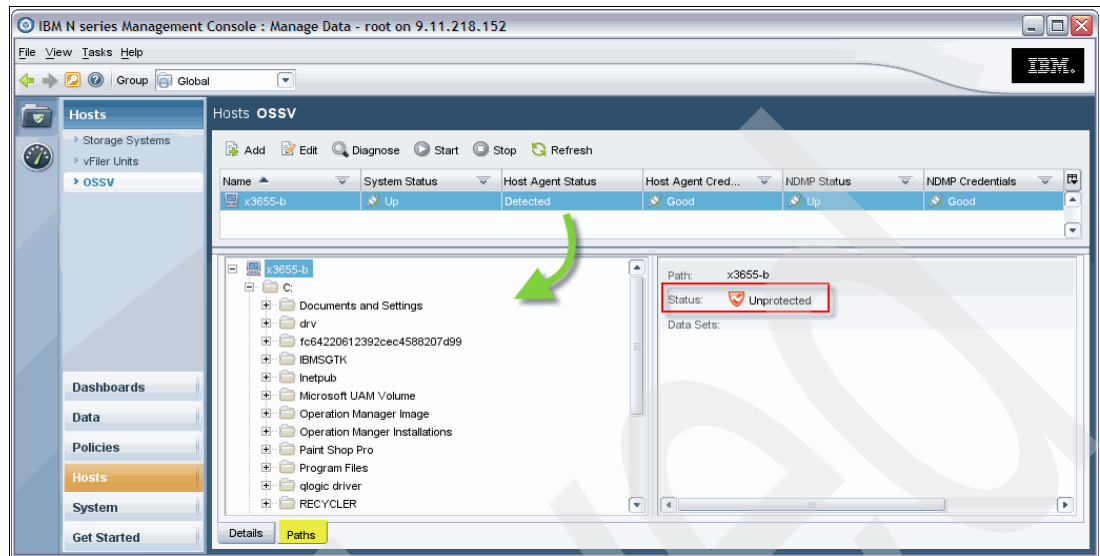


Figure 13-69 Host's file system and current protection status

If you need to change any of the passwords or NDMP details for an OSSV host, you can do so by clicking the **Edit** button, as shown in Figure 13-70.

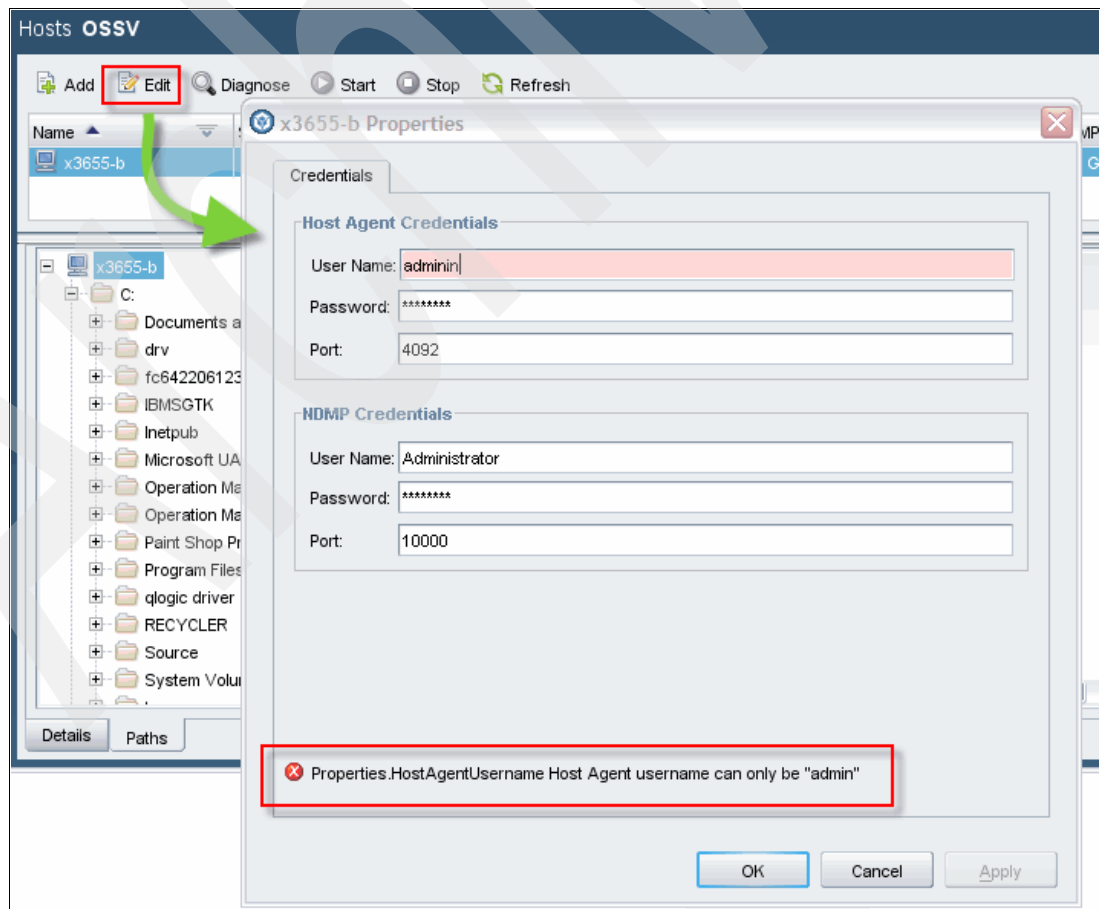


Figure 13-70 Host credentials

Note: The port number should not be changed unless you have configured the NDMP service on the host agent to listen on another port number.

You can start and stop the NDMP service or a host by highlighting the host and clicking the **Start** or **Stop** buttons, as shown in Figure 13-71.

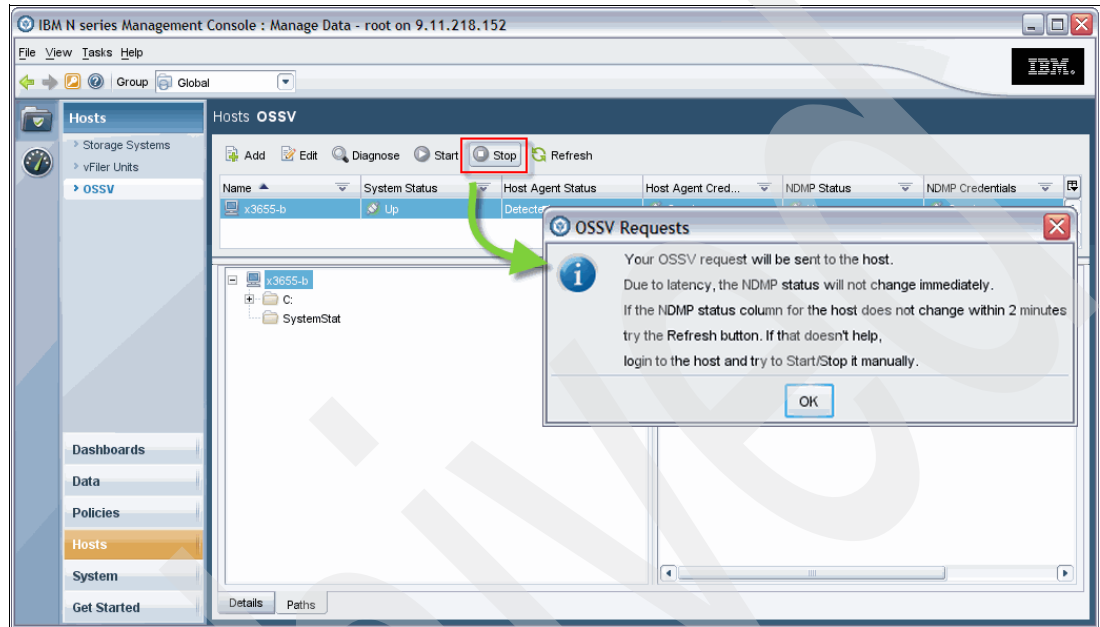


Figure 13-71 Stopping the NDMP services

Stopping the NDMP service, as shown in Figure 13-72, also means that the OSSV host will no longer be protected until the service is restarted. You will not be able to browse the file system of the host as well.




Figure 13-72 Stopping the NDMP service has repercussions

Clicking **Start**, as shown in Figure 13-73, will cause the DFM Server to send a start message to the OSSV host agent on the highlighted host. The outcome of the request will be reflected in the window shown in Figure 13-74.

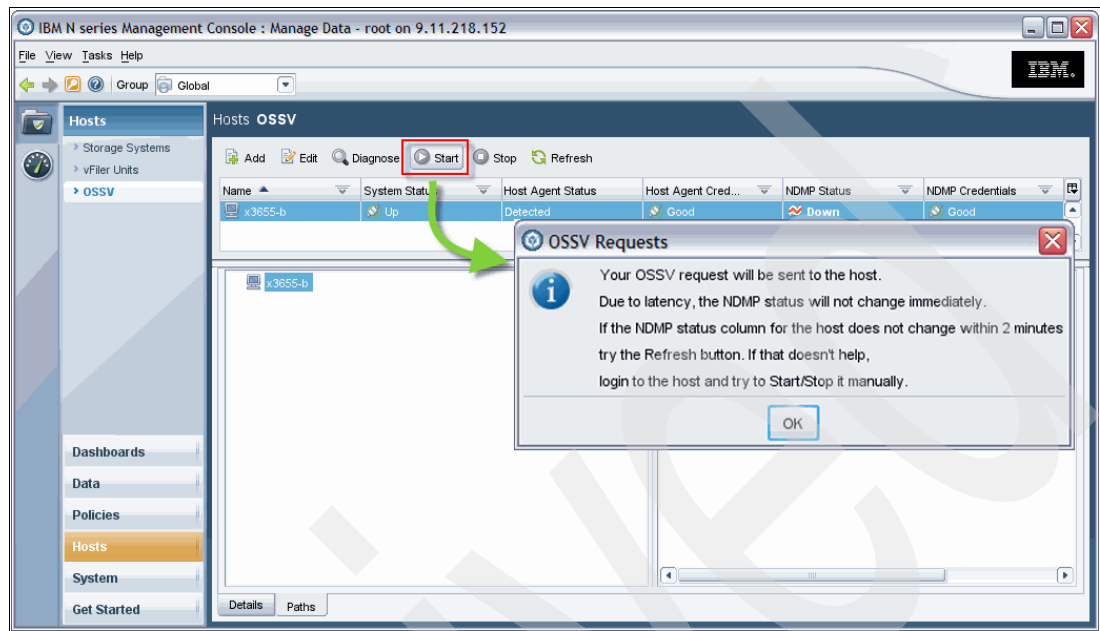


Figure 13-73 OSSV request window




Figure 13-74 Starting OSSV

As shown in Figure 13-74, if the Start request successfully completes, you will then be able to view the file system of the highlighted host and review the status of the NDMP agent.

13.5.2 Adding a suitable storage host as an OSSV secondary

SnapVault is an N series feature that is used to back up file data from one storage location, such as a CIFS or NFS volume, to another, residing typically on another storage system. You can, with SnapVault, have a general repository on the target storage system that is the collection point of all your SnapVault primary sources, whether they reside on N series hosts, on OSSV hosts, or both.

SnapVault makes it possible to back up only what has changed in a file. This is possible for volumes that reside on the N series storage systems because of the Write Anywhere File Layout (WAFL®) file system. For OSSV hosts, this is also possible because OSSV maintains a shadow repository of host files being backed up to enable it to do a before-and-after comparison and then only send the changes to the SnapVault Secondary repository on the N series system.

With SnapVault, you can maintain an update history of your files while having the convenience of storing all this historical information in one or more convenient locations on an N series storage system for easy and rapid restores. SnapVault technology facilitates easy and rapid restoration of individual files or entire directories to a previous location over TCP/IP. OSSV hosts are configured to back up their data to the SnapVault Secondary location. If the OSSV host is integrated with Protection Manager through Operations Manager, you can manage this host remotely and schedule backups to run and monitor the status of the backups.

OSSV backups are scheduled on the N series system with which the OSSV hosts are associated. At the scheduled time, the N series storage system will use the NDMP protocol to initiate a backup from the OSSV host and pull down changed data.

In this section, we demonstrate how to set up a SnapVault Secondary target from within Protection Manager. The prerequisite licenses for a N series storage system to be a SnapVault Secondary are:

- ▶ NearStore® license.
- ▶ SnapVault Secondary license.
- ▶ The appropriate licenses for the OSSV host. For example, for OSSV hosts on Windows, we require an *OSSV SnapVault primary license for Microsoft Windows*.

Selecting the **Storage Systems** menu on the N series Management Console will display the current storage systems that you manage, as shown in Figure 13-75. This window will also provide you with an overview of the highlighted host and the licenses that it currently has.

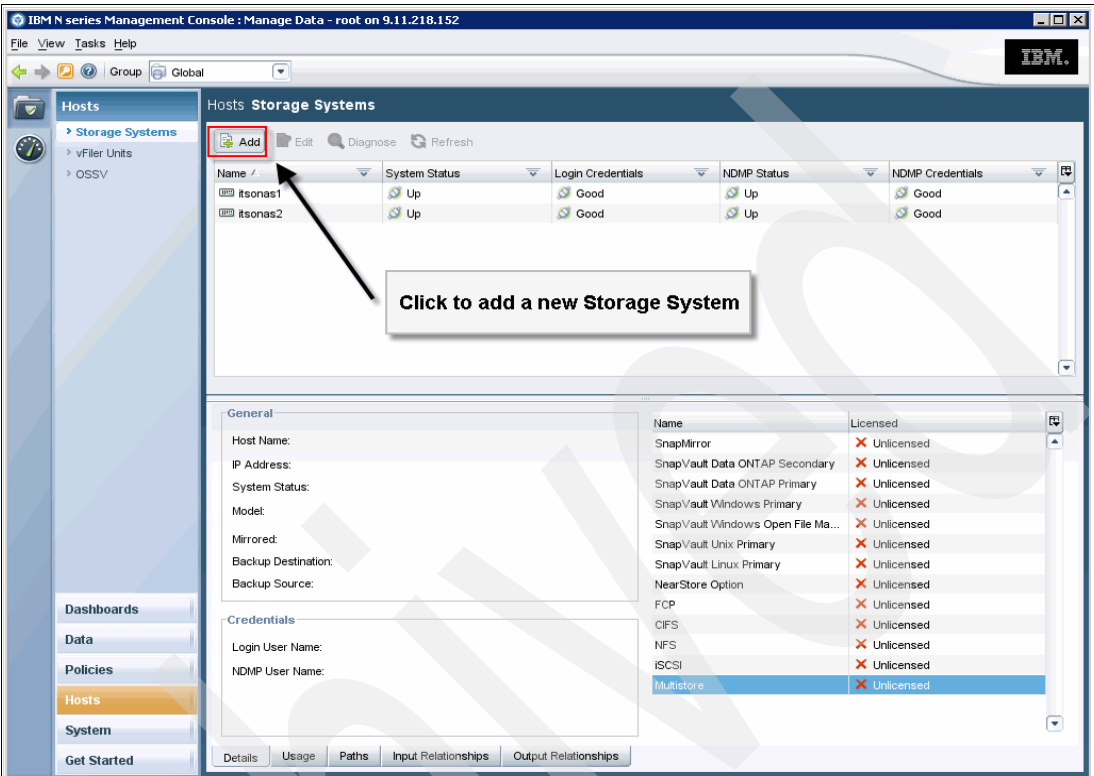


Figure 13-75 Storage systems window on N console

To add a new storage system, click the **Add** button, as shown in Figure 13-76.

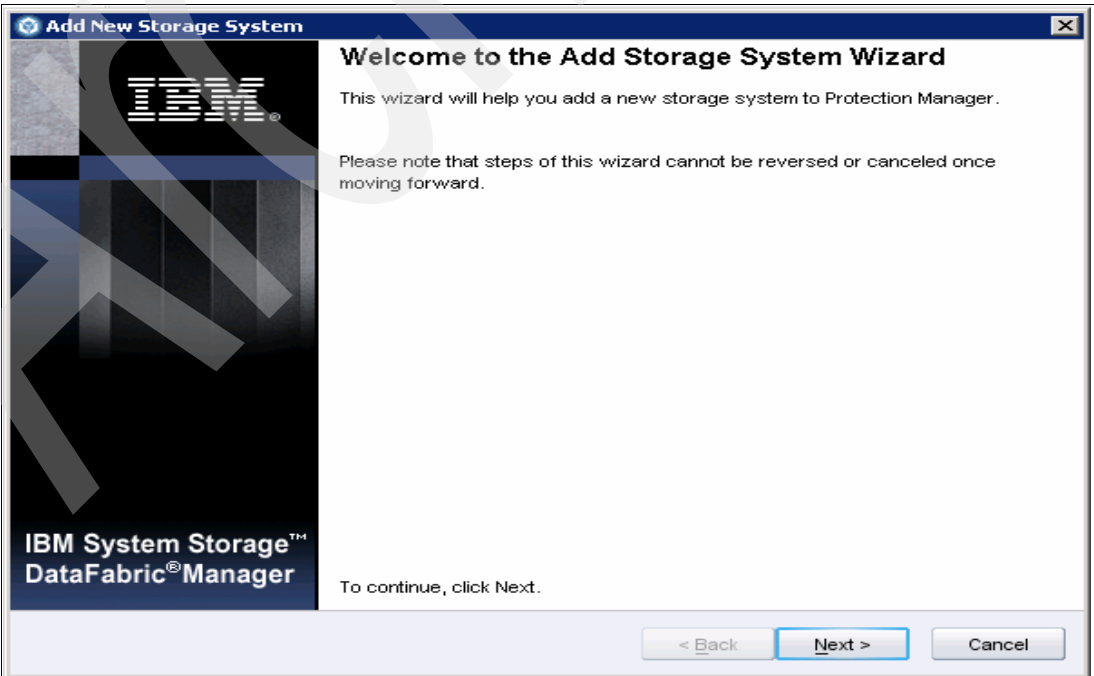
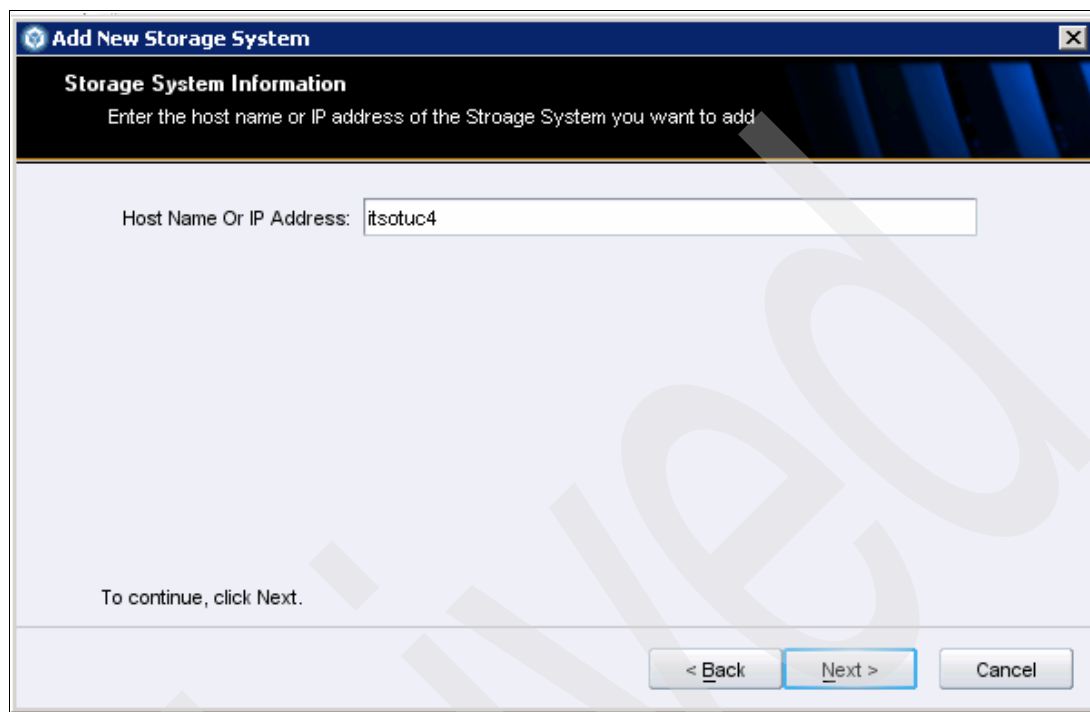


Figure 13-76 Add Storage System Wizard

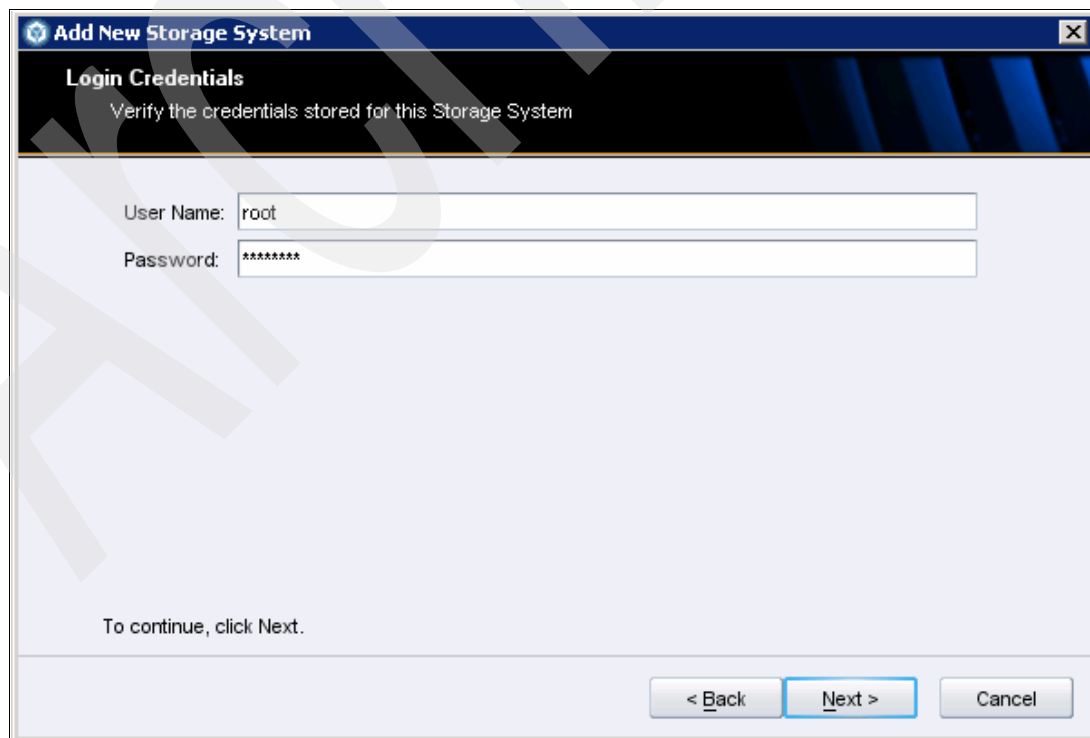
Clicking **Next** takes you to the next window, where you can enter the storage system name or IP address, as shown in Figure 13-77.



The screenshot shows a window titled "Add New Storage System" with a close button (X) in the top right corner. The main heading is "Storage System Information" with a subtitle "Enter the host name or IP address of the Storage System you want to add". Below this, there is a text input field labeled "Host Name Or IP Address:" containing the text "itsotuc4". At the bottom left, it says "To continue, click Next." At the bottom right, there are three buttons: "< Back", "Next >" (which is highlighted with a blue border), and "Cancel".

Figure 13-77 Enter the host name of storage array to be managed

You will be prompted to provide the login credentials of the storage system for the DFM Server to manage, as shown in Figure 13-78.



The screenshot shows a window titled "Add New Storage System" with a close button (X) in the top right corner. The main heading is "Login Credentials" with a subtitle "Verify the credentials stored for this Storage System". Below this, there are two text input fields: "User Name:" containing the text "root" and "Password:" containing the text "*****". At the bottom left, it says "To continue, click Next." At the bottom right, there are three buttons: "< Back", "Next >" (which is highlighted with a blue border), and "Cancel".

Figure 13-78 Host names and credentials of the storage system

Note: You must also review the firmware versions on the storage system you plan to add to the DFM Server. Ensure that you download and install the correct plug-ins for the DFM Server to properly manage the N series storage system.

As you add the storage array, you also have the opportunity to add licenses, as shown in Figure 13-79 and Figure 13-80 on page 377.

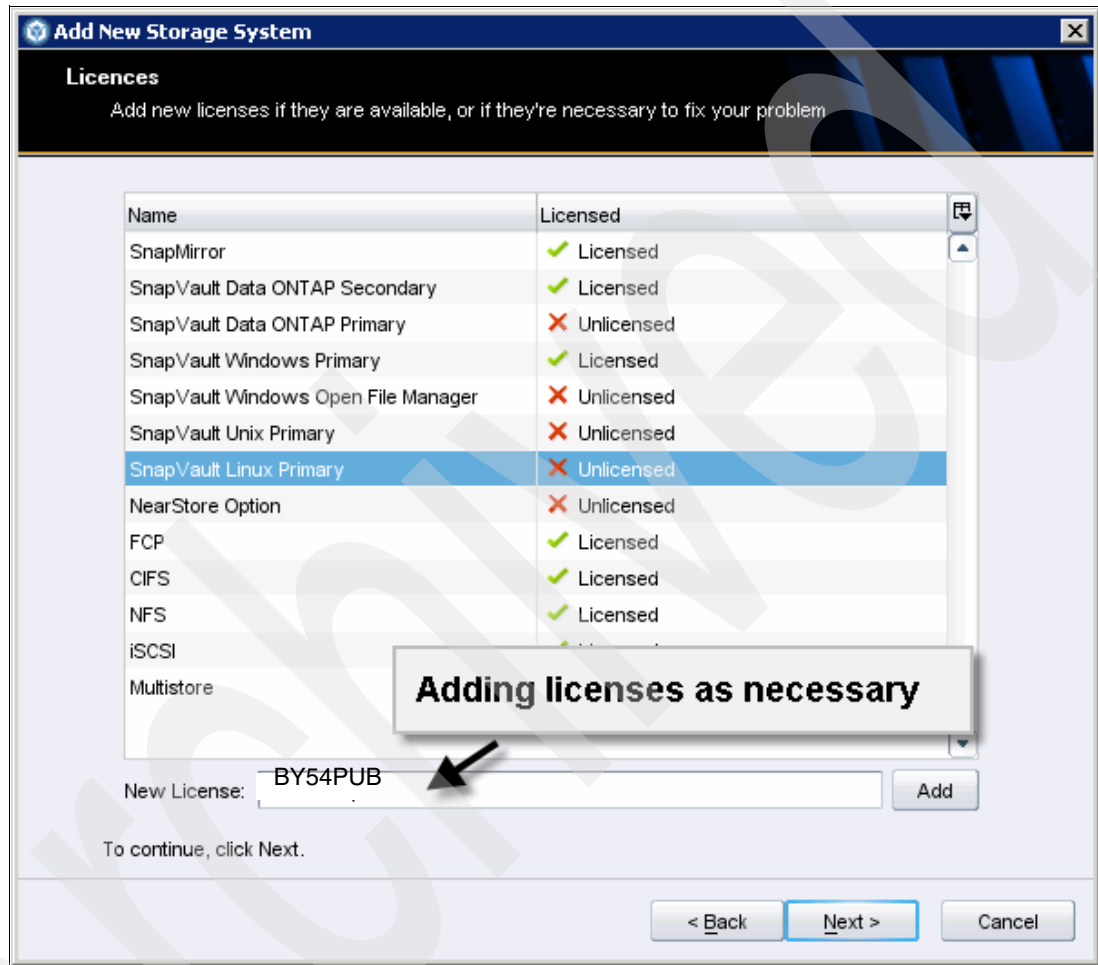


Figure 13-79 Adding licenses as necessary

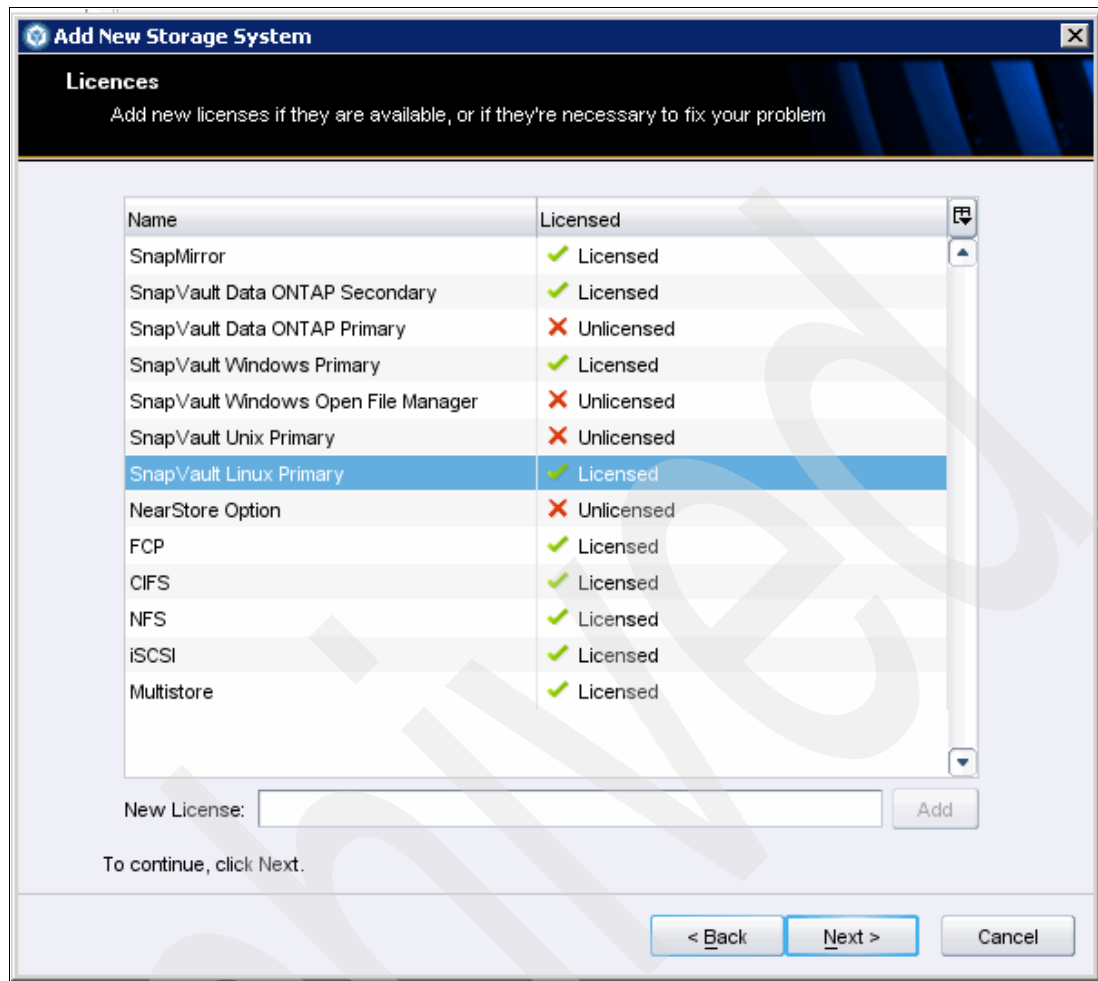


Figure 13-80 Adding licenses to the new storage system

You will also be asked to update the SnapVault Access Control List, as shown in Figure 13-81. Here you can type in the DFM Server that will be used to manage the storage array.

Add New Storage System

SnapVault Settings
Specify the SnapVault settings and access control list

Enable SnapVault: ☒ Yes ☐ No

SnapVault Access Control List:

host=9.11.218.114

Specifying the Management Host (DFM Server)

Enable all access with a *. Disable all access by leaving the field blank.

To continue, click Next.

< Back Next > Cancel

Figure 13-81 Specifying a management host for access control

In order to successfully manage a storage array from Protection Manager, NDMP must be enabled, as shown in Figure 13-82 on page 379, on the storage system, and the appropriate user credentials must be provided, as shown in Figure 13-83 on page 379.

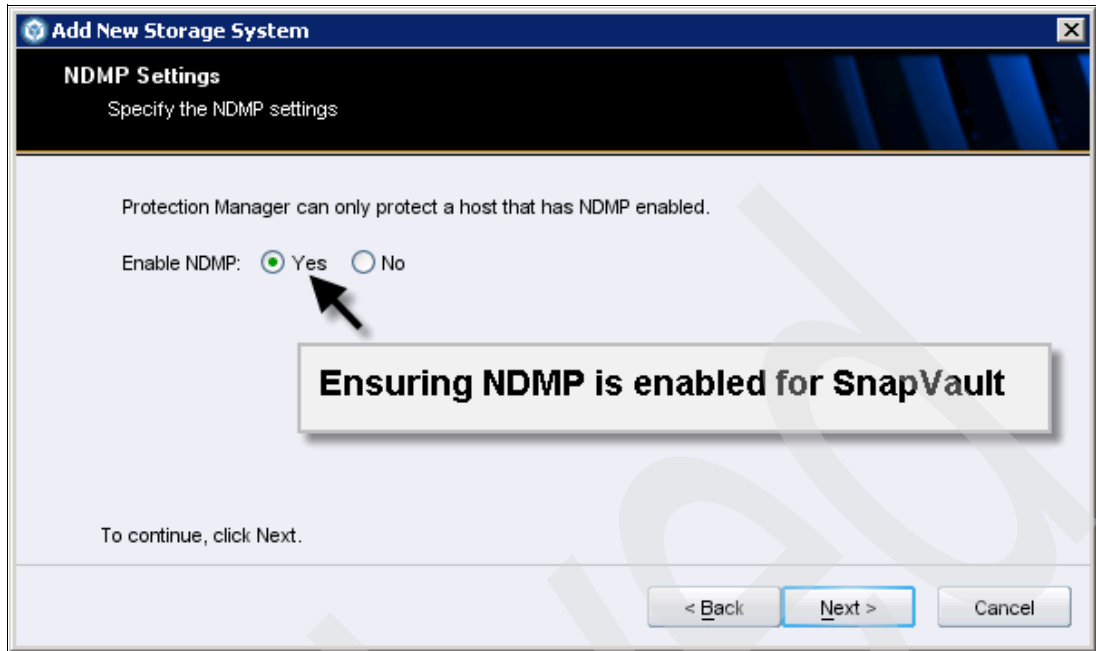


Figure 13-82 Ensuring NDMP is enabled on the new storage system

The NDMP protocol is used to manage data for the purposes of conducting backups and restores. It is the protocol used to communicate with the host (storage system) to instruct it on which volumes to be backed up, obtain the Table of Contents for the volume being backed up, and other management related activities.

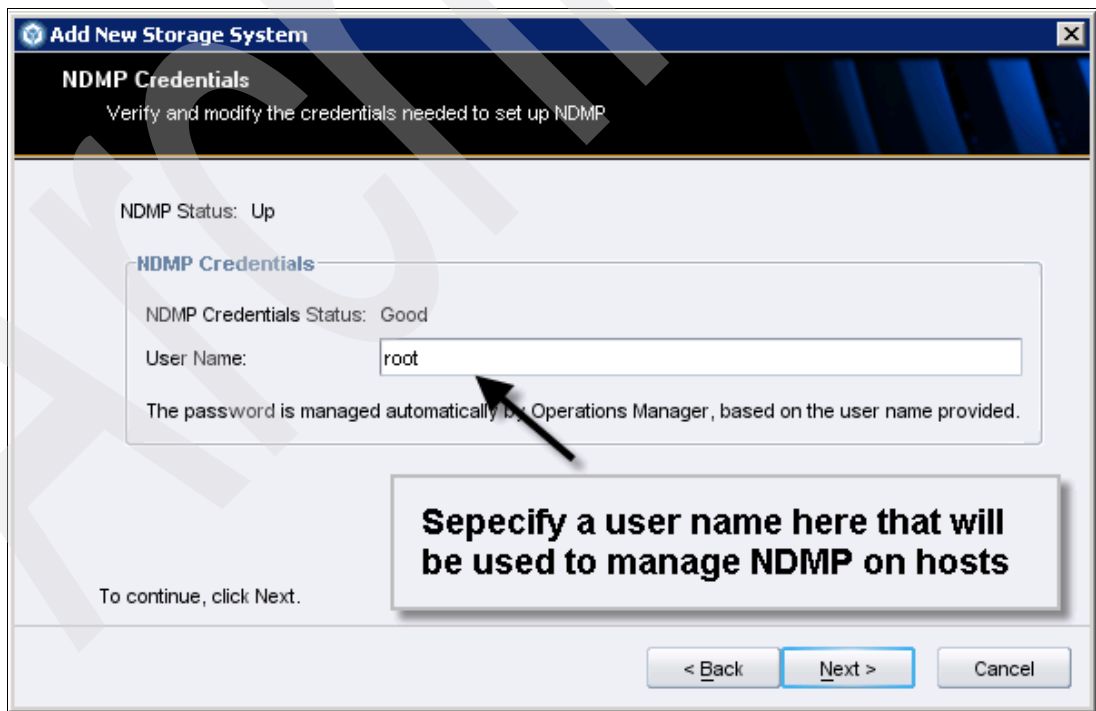


Figure 13-83 Setting the NDMP credentials

After you click **Next**, you will be presented with the summary window shown in Figure 13-84. Click **Finish** to complete the process.

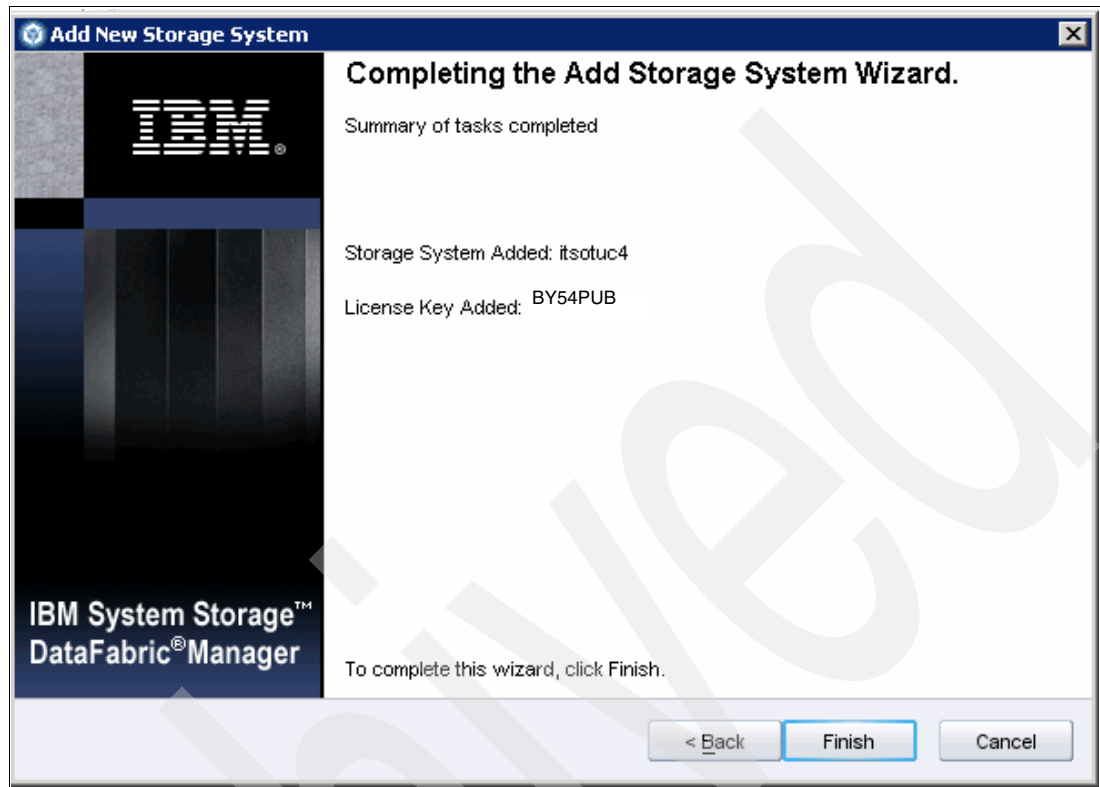


Figure 13-84 Completing the Add Storage System Wizard window

The new storage system will now be added and the results will be visible on the window, as shown in Figure 13-85 on page 381.

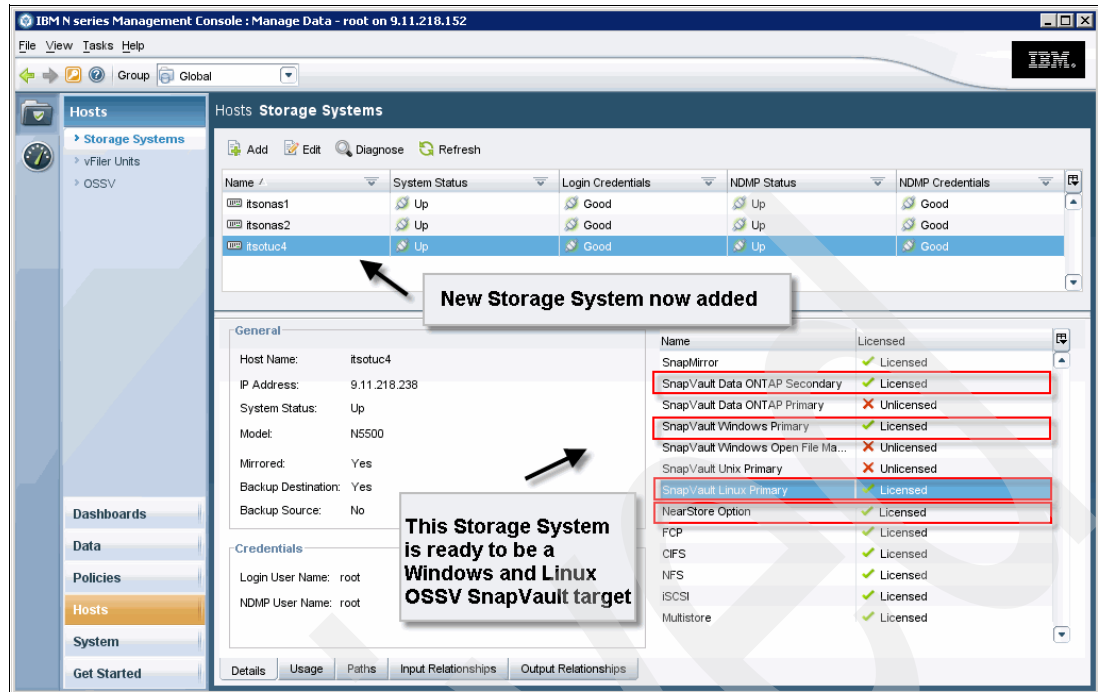


Figure 13-85 Reviewing the new storage system

13.5.3 Creating a OSSV data set

The next step to protect the data on an OSSV host is to create a data set in Protection Manager that defines the OSSV host and details what folders need to be protected.

We begin with the wizard, as shown in Figure 13-86.

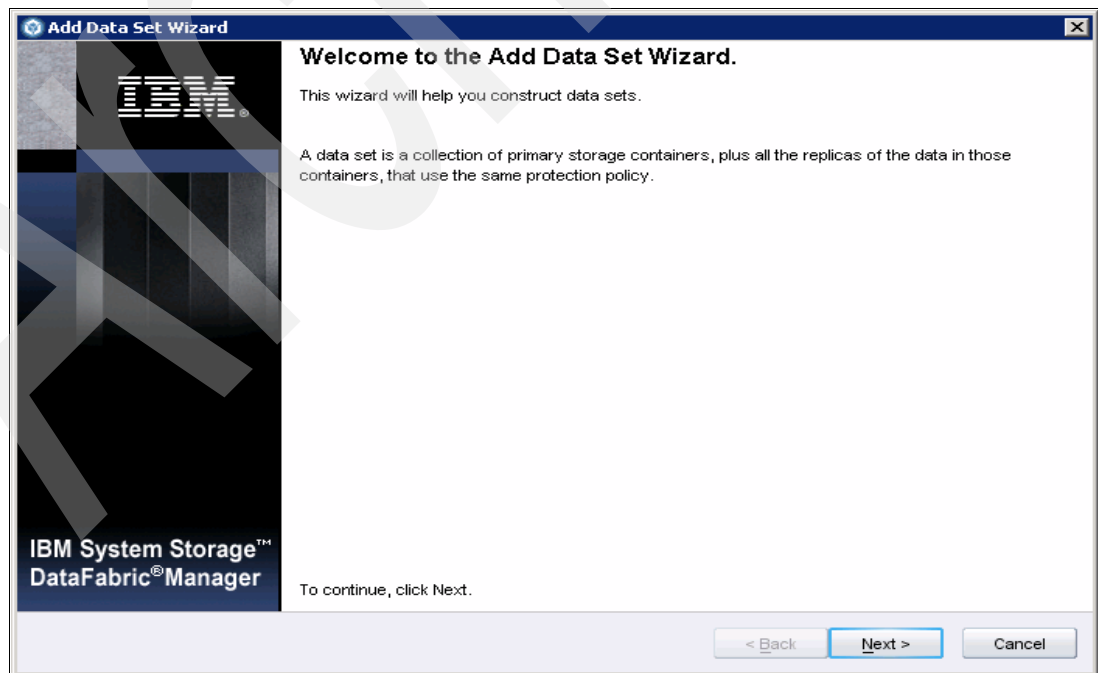


Figure 13-86 Add Data Set Wizard

Using OSSV data sets implies that you want to use the SnapVault to back up host data onto a SnapVault Secondary target.

Ensure that your target storage array has the necessary licenses to facilitate this. They include:

- ▶ SnapVault Secondary licenses
- ▶ NearStore licenses
- ▶ SnapVault Primary licenses for the operating system of your OSSV host

We provide the details necessary to uniquely identify this data set and the time zone it is located in, as shown in Figure 13-87.

Add Data Set Wizard

General Properties
You must provide a name for the new data set. Other properties are optional.

Name: OSSV_Windows_Hosts_Arizona

Description: All Windows OSSV Hosts in Arizona TimeZone

Owner: IT Support

Contact: cactus@itso.tucson.ibm.com

Time Zone: America/Phoenix

To continue, click Next.

< Back Next > Cancel

Figure 13-87 Entering data set details

If your organization spans multiple time zones, you might wish to ensure that your data sets are limited to one time zone. Create a data set for each time zone. You can then add the data sets into the relevant groups for ease of management and reporting, as shown in Figure 13-88 on page 383.

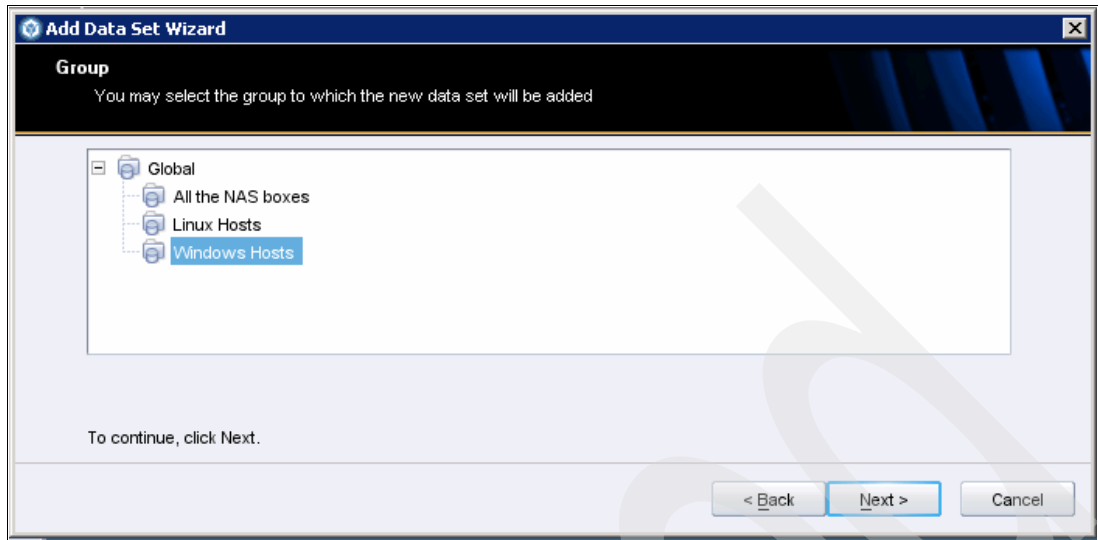


Figure 13-88 Adding the data set to a group

In Figure 13-89, you have the option to create new data resources using a provisioning policy or to assign existing resources manually. As we are planning to backup existing resources on an OSSV host, we select the latter option.

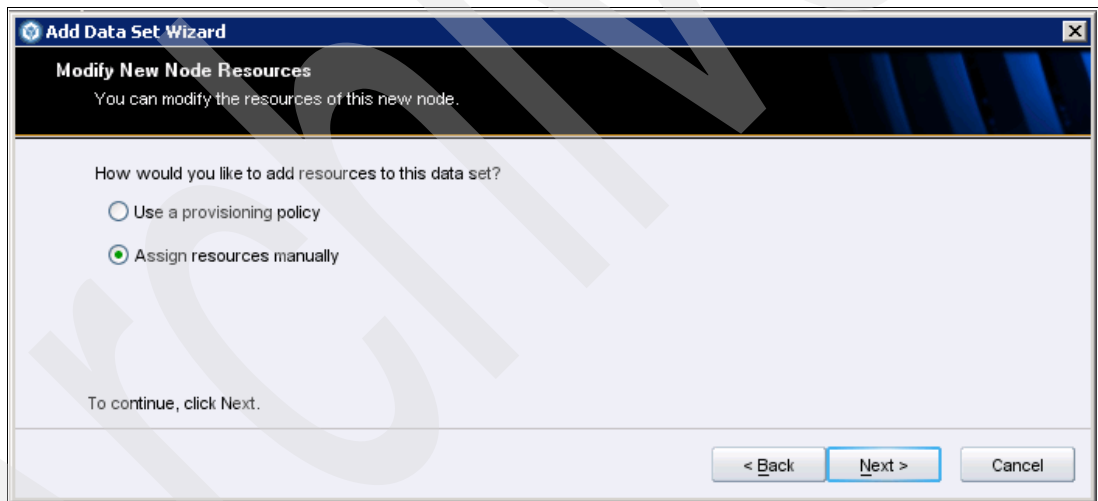


Figure 13-89 Assigning resources to the Data Set manually

Clicking **Next** will take you to a resource selection window, as shown in Figure 13-90.

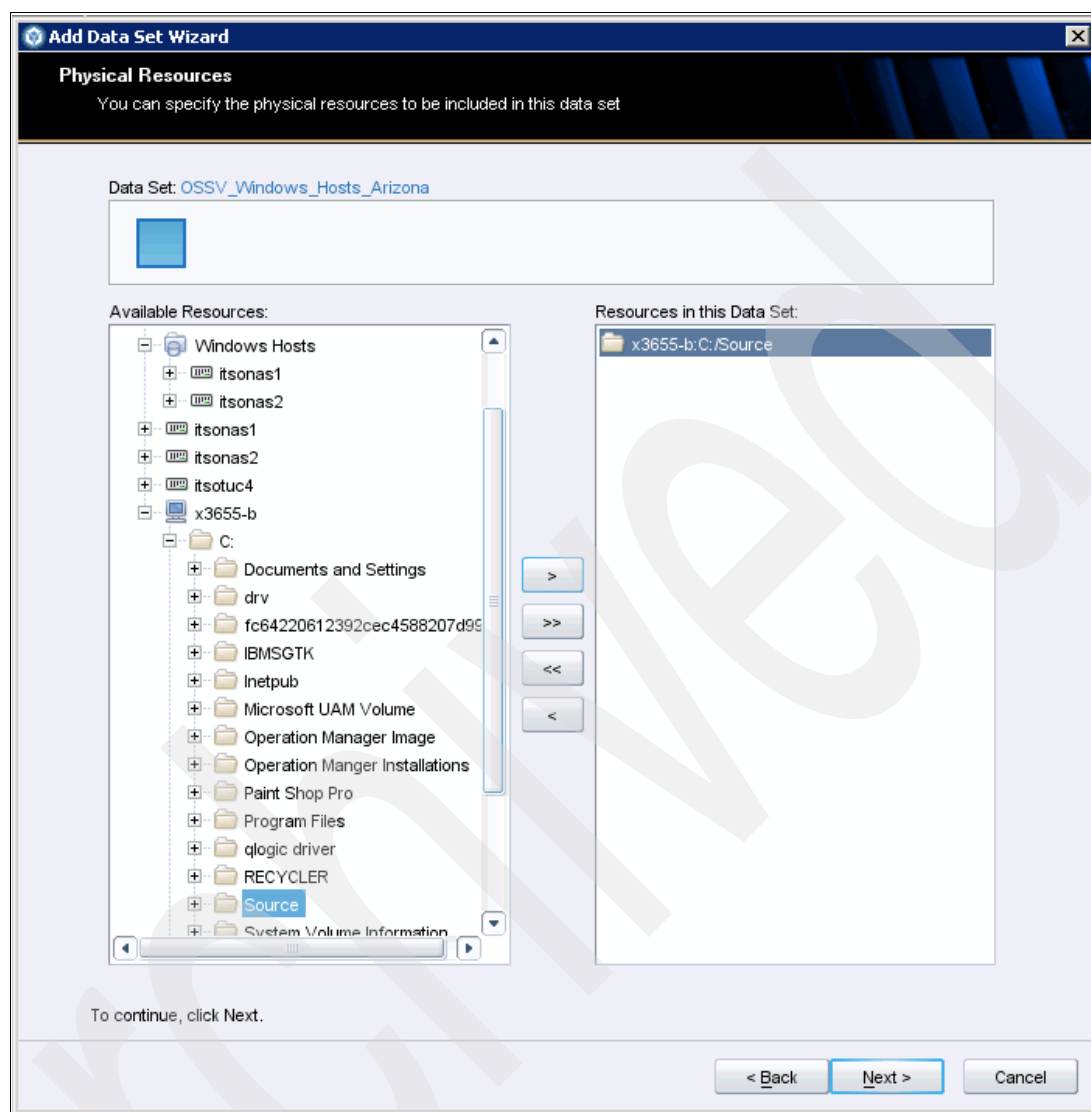


Figure 13-90 Selecting the resources on an OSSV host

Select the folders that you want to be included in this data set and click **Next**. You then have the opportunity to review the details, as shown in Figure 13-91. The Preview window also shown the results of tests run by Protection Manager to validate your request. Any errors found will be shown here for you to remediate.

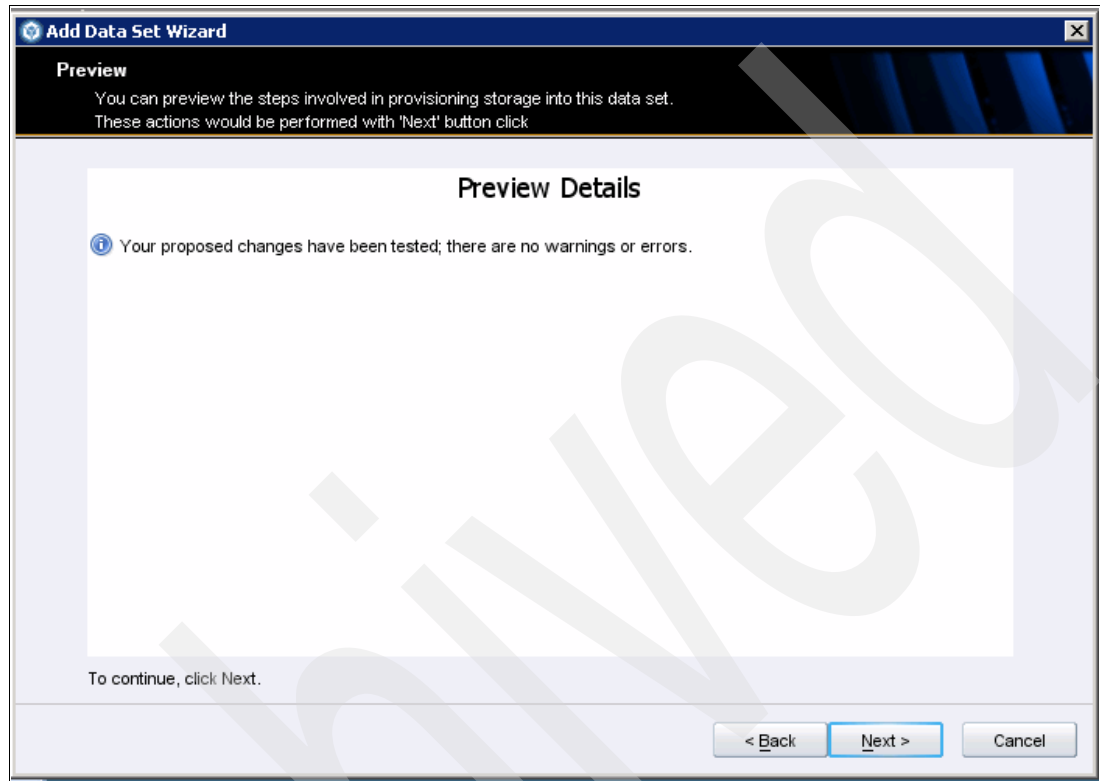


Figure 13-91 Previewing the execution steps

Clicking **Next** will bring you to the completion window, where Protection Manager will attempt to set up the data set as per your instructions, as shown in Figure 13-92.

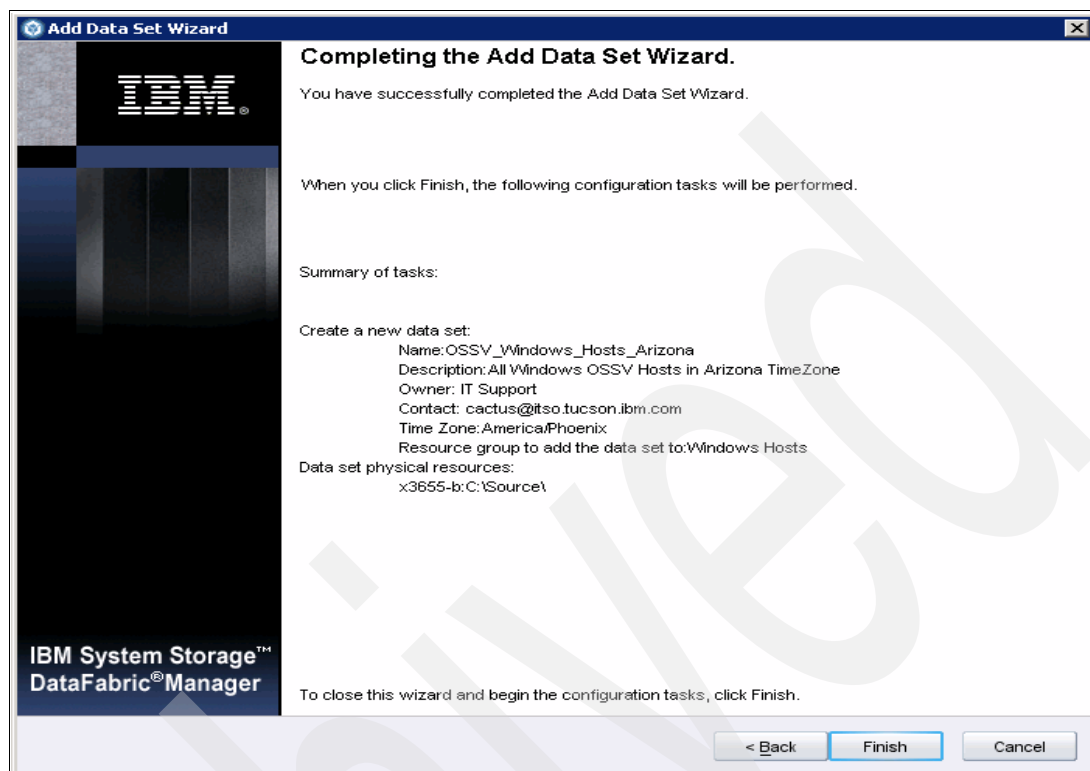


Figure 13-92 Final step in creating a new OSSV data set

13.5.4 Protecting the OSSV data set

Once you have defined your data set, it is time to assign a protection policy, as shown in Figure 13-93 on page 387, for the express purpose of ensuring that the data defined in your data set will be protected. If a suitable backup target has not yet been identified, you will have, in this wizard, to use a suitable provisioning policy to create the necessary resource pools.

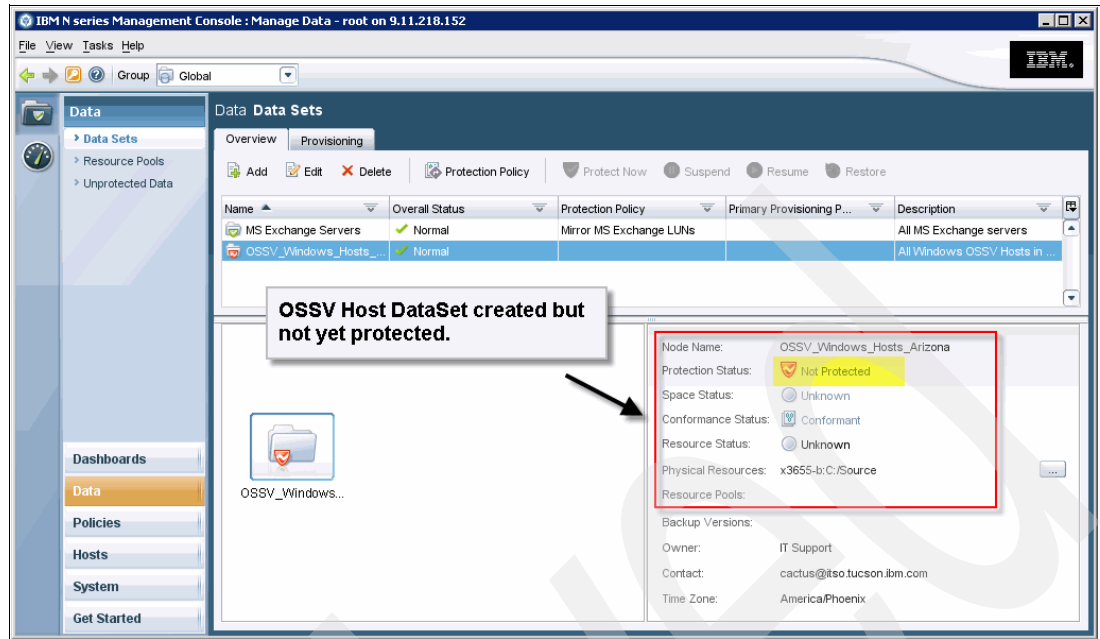


Figure 13-93 Reviewing the data set's protection status

Click **Protection Policy** to launch the Data Set Policy Change Wizard, as shown in Figure 13-94.

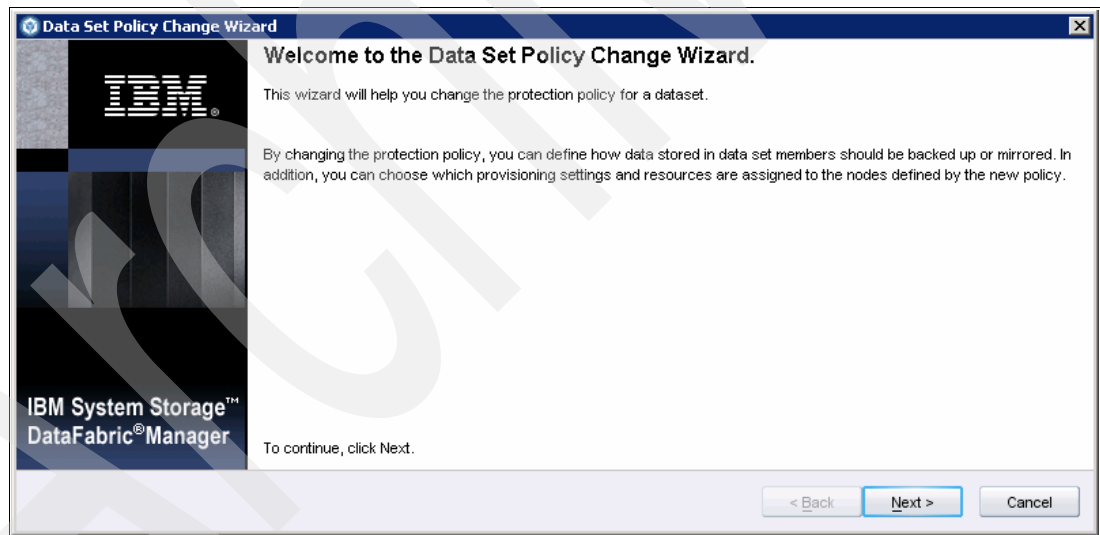


Figure 13-94 Launching the Data Set Policy Change Wizard

The wizard will present you with a list of templates or created policies for you to choose from, as shown in Figure 13-95. You will find that there is only one suitable policy (called “Remote backups only”) that is set up to facilitate backing up OSSV hosts.

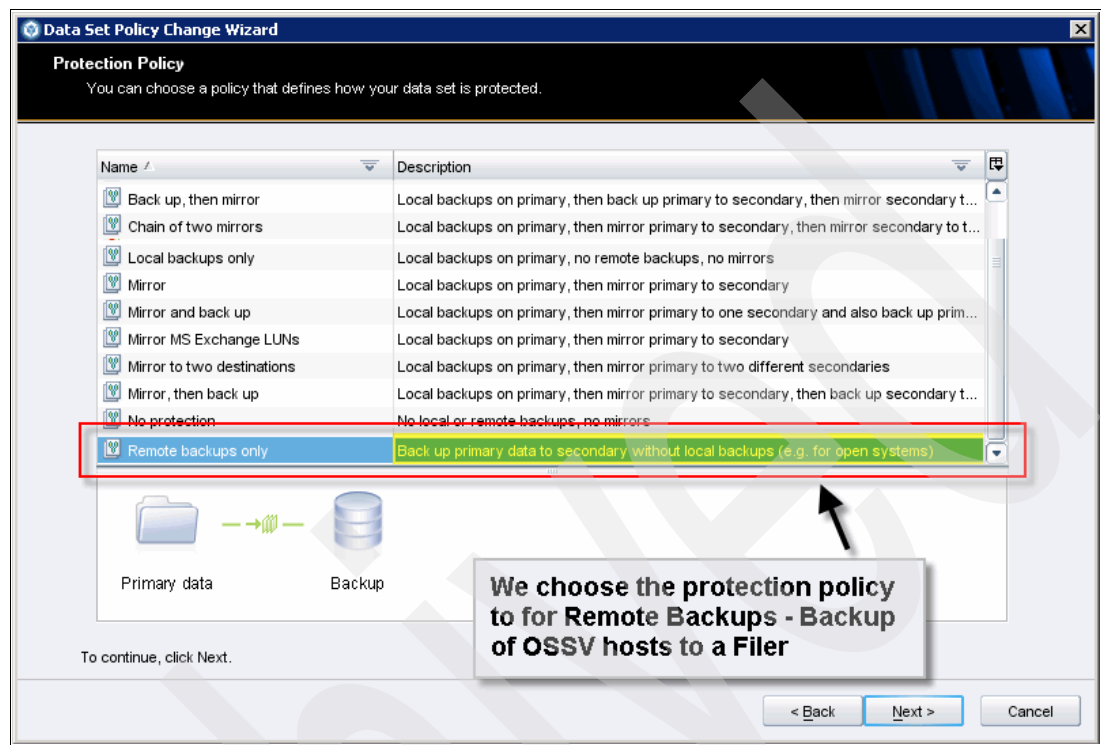


Figure 13-95 Selecting the (only) suitable backup policy for OSSV hosts

Select the “Remote backups only” policy and click **Next**. You will then be prompted to assign resources using a provisioning policy or manually, as shown in in Figure 13-96 on page 389.

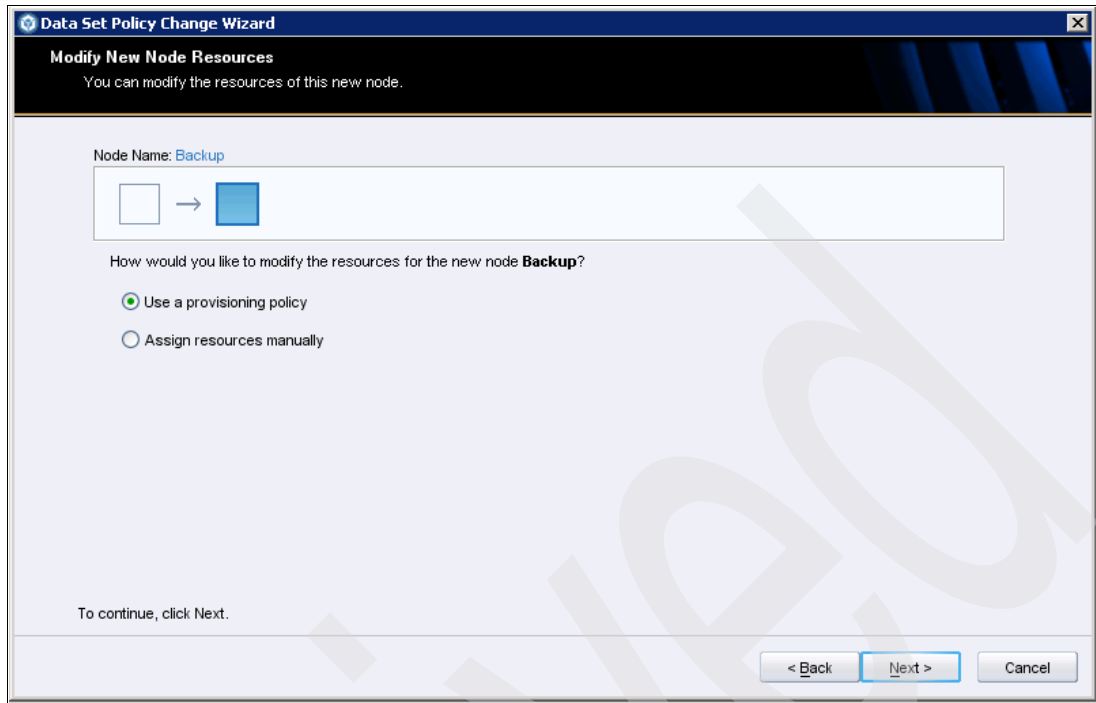


Figure 13-96 Provisioning backup targets using a provisioning policy

In this case, we elect to provision resources using a provisioning policy. Clicking **Next** will take us to the window shown in Figure 13-97.

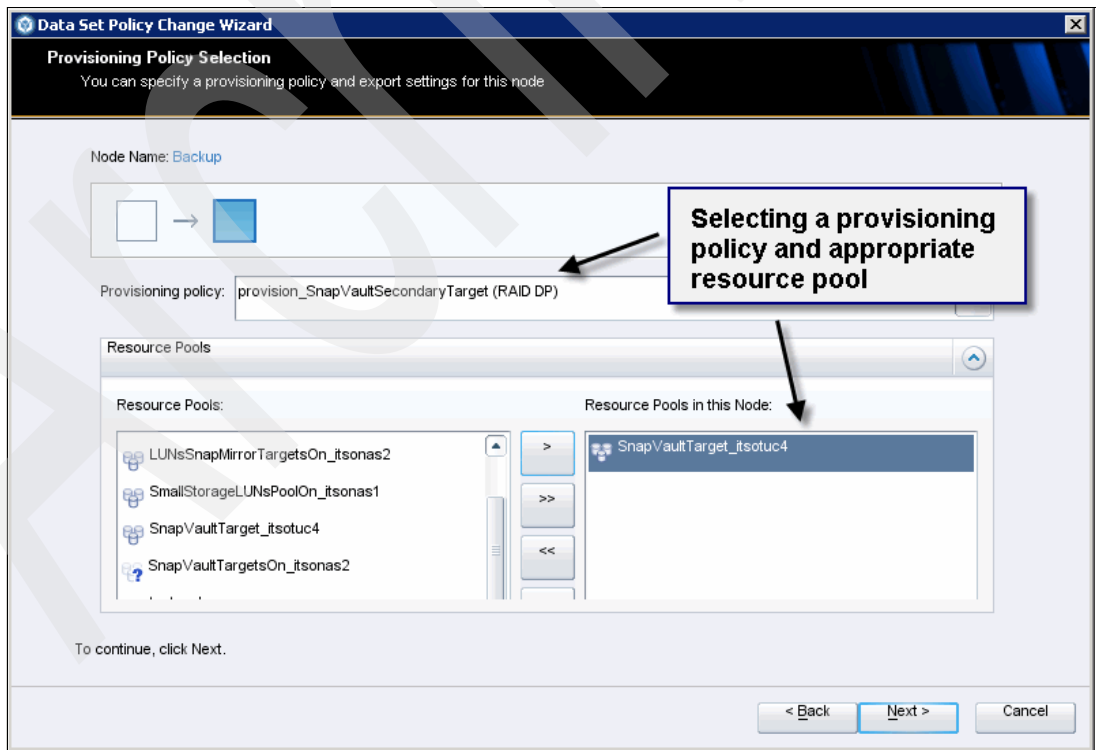


Figure 13-97 Selecting the appropriate protection policy

In Figure 13-97 on page 389, we select “SnapVaultSecondaryTargetsOn_itsotuc4” as a suitable target for the backup. We will not select a vFiler unit, as shown in Figure 13-98.

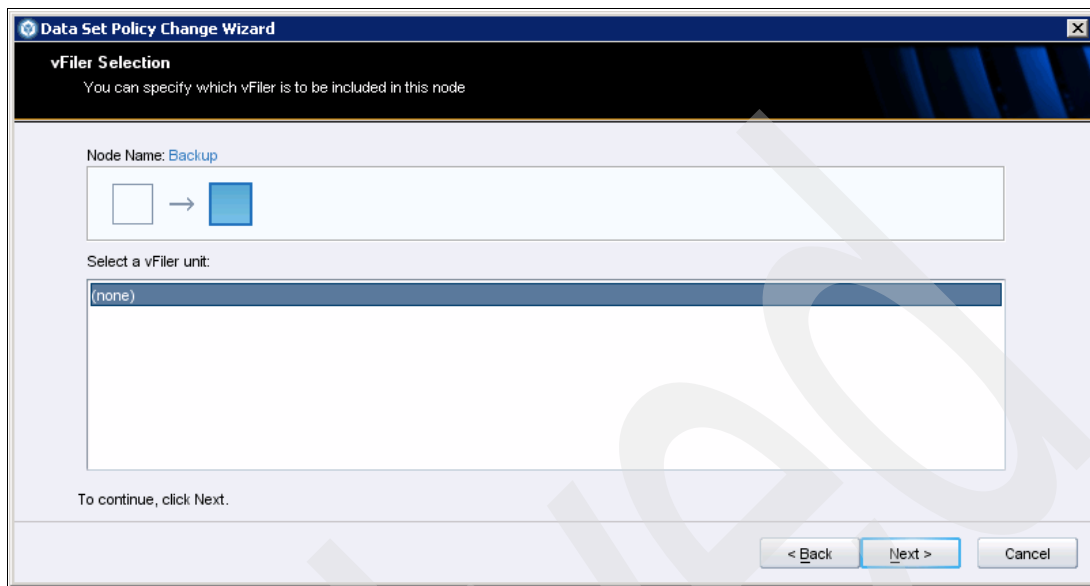


Figure 13-98 vFiler Selection window

In Figure 13-99, we review the changes we would apply. The wizard reports that there is an issue with the available bandwidth and it is unable to start up the OSSV mirror.

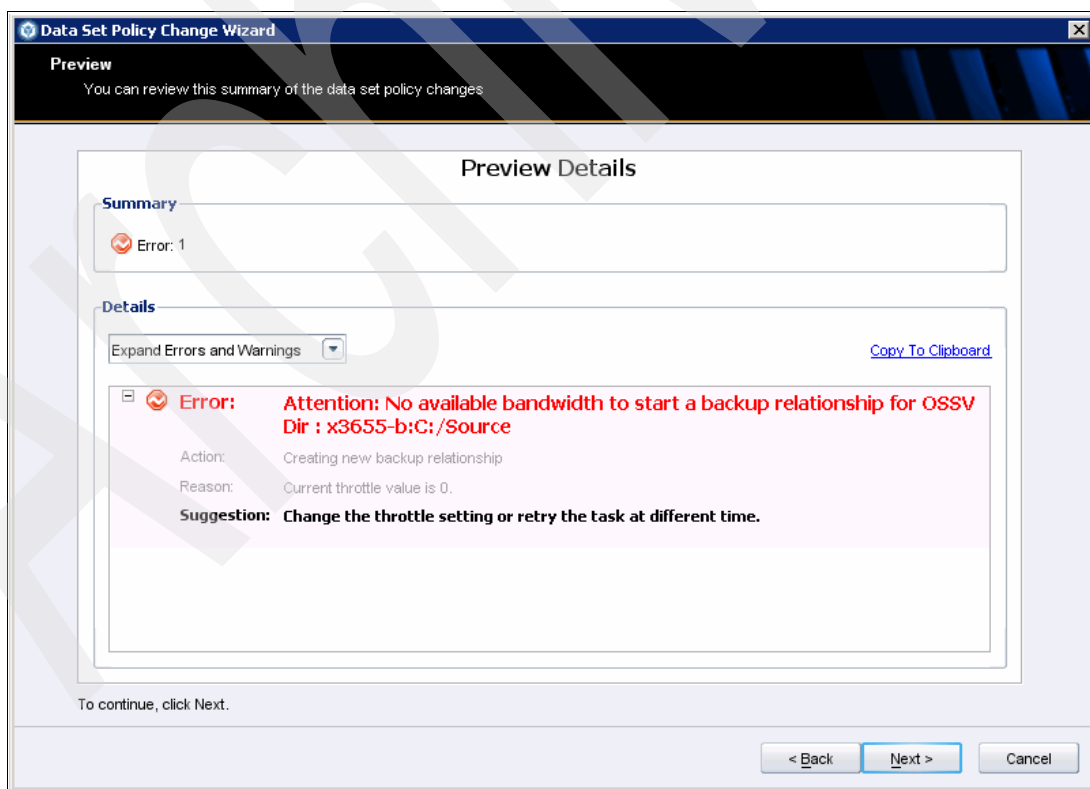


Figure 13-99 Failed preview prior to applying the change

Upon investigation, we discover that the issue had to do with an incorrectly entered license. We rectified the issue and retraced the data set creation process to arrive at a successful preview, as shown in Figure 13-100.

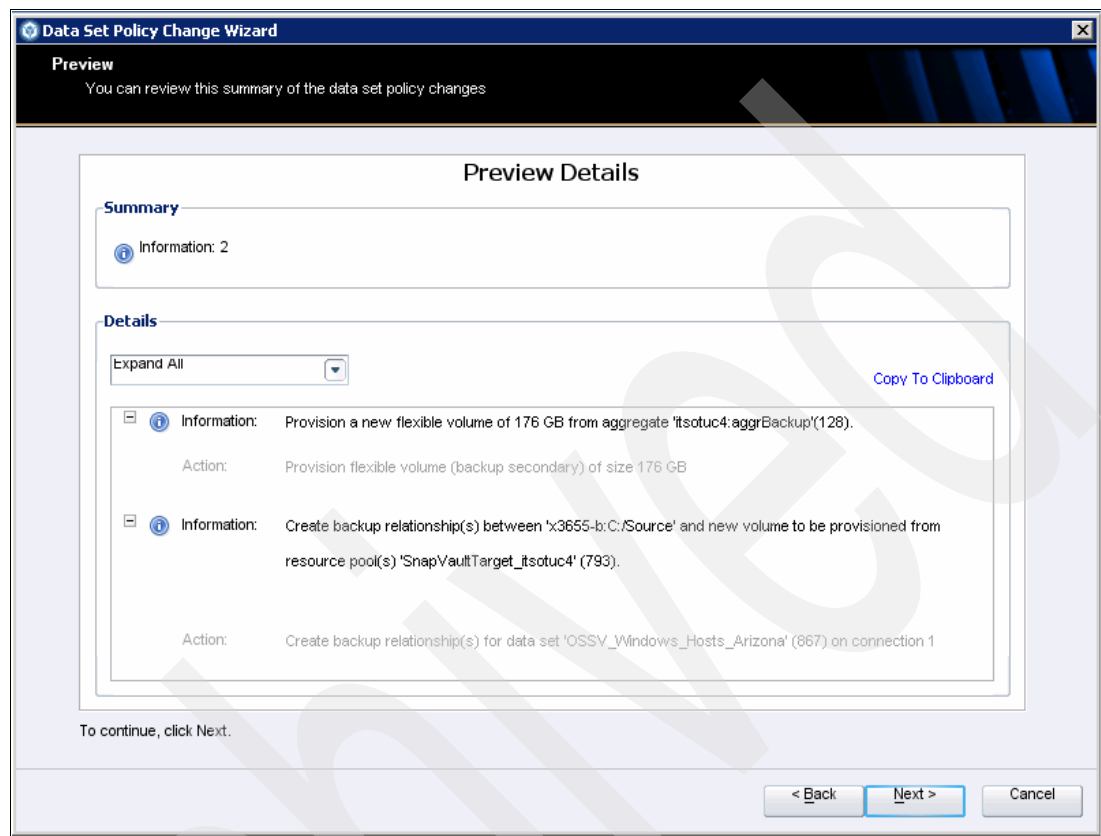


Figure 13-100 Successful preview prior to applying the change

We proceed to the next step and establish the policy, as shown in Figure 13-102 on page 393.

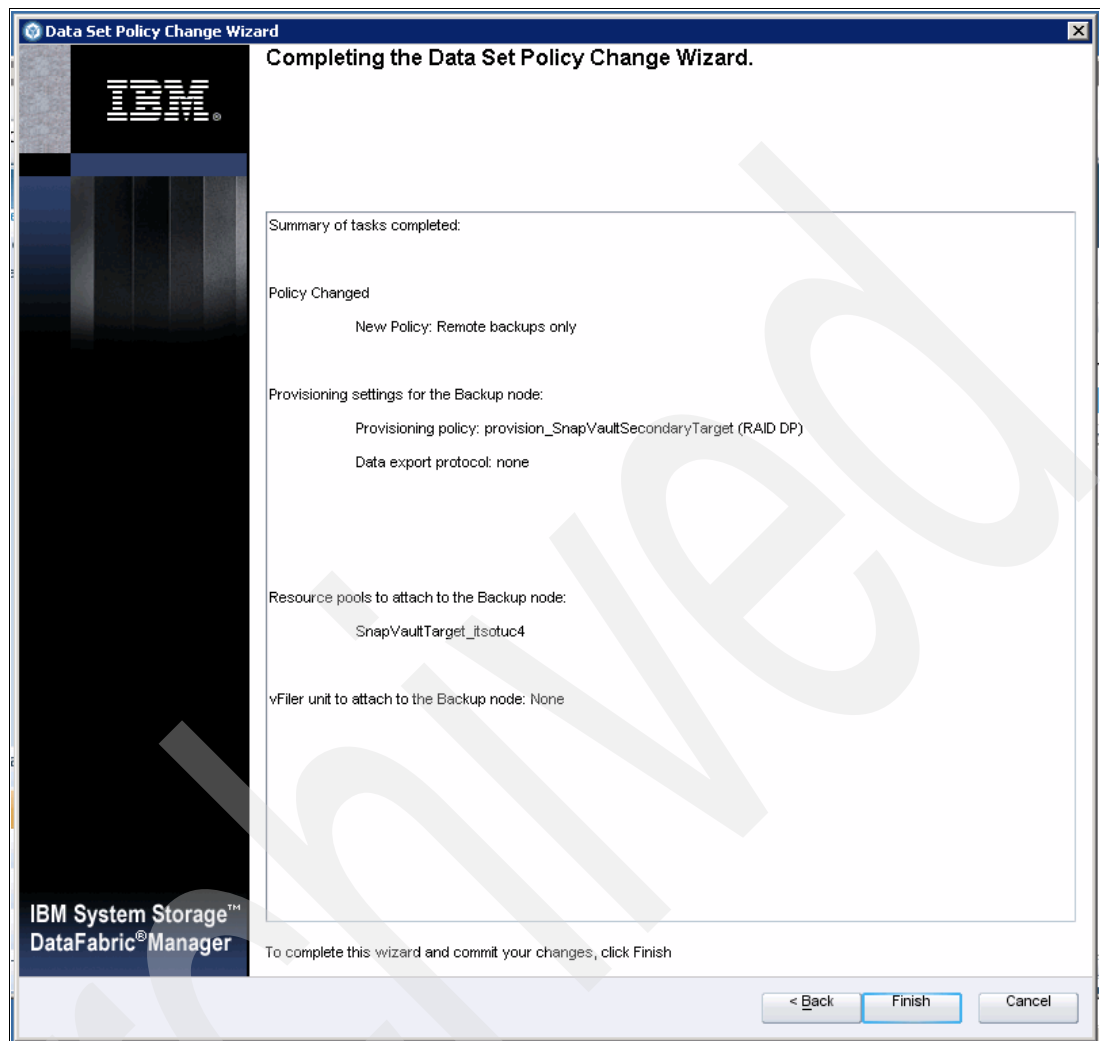


Figure 13-101 Completing the Data Set Policy Change Wizard

In Figure 13-101 on page 392, we review all the steps that will be taken and click **Finish** to confirm. We are then presented with the “Data Sets” window shown in Figure 13-102, where the policy is being enforced through background jobs.

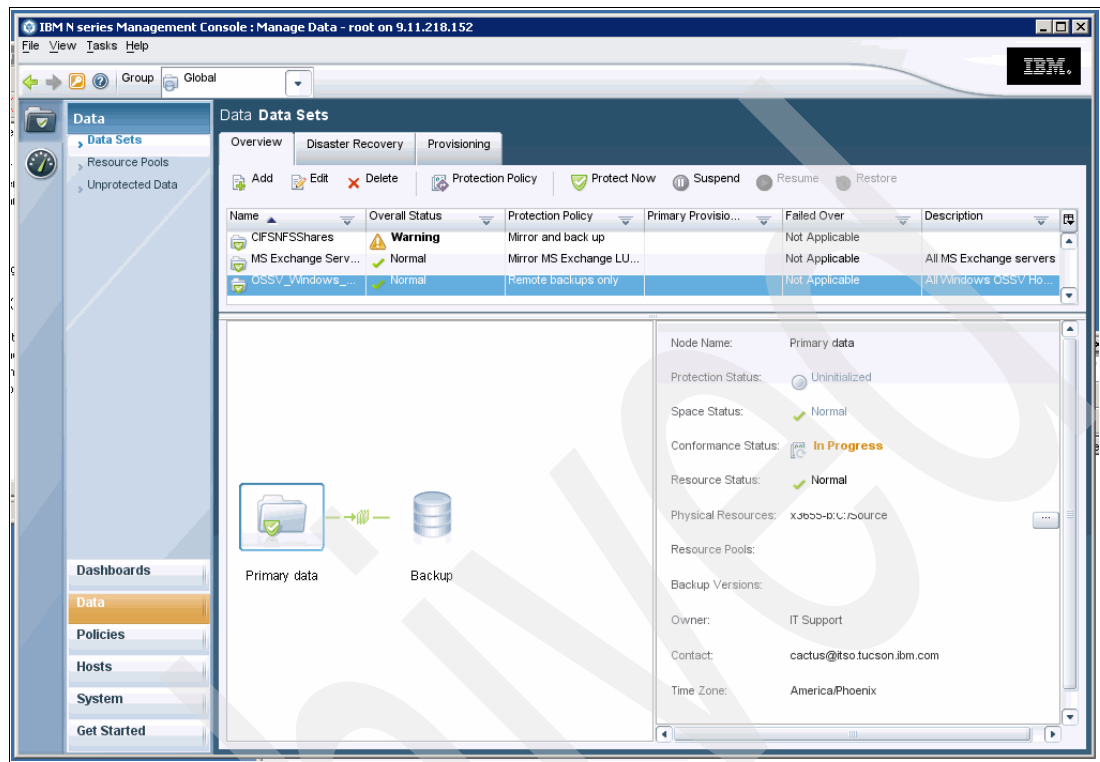


Figure 13-102 Data set policy enforcement in progress

We can review the jobs that are running by navigating to the Jobs pane in the N console System Jobs window, as shown in Figure 13-103.

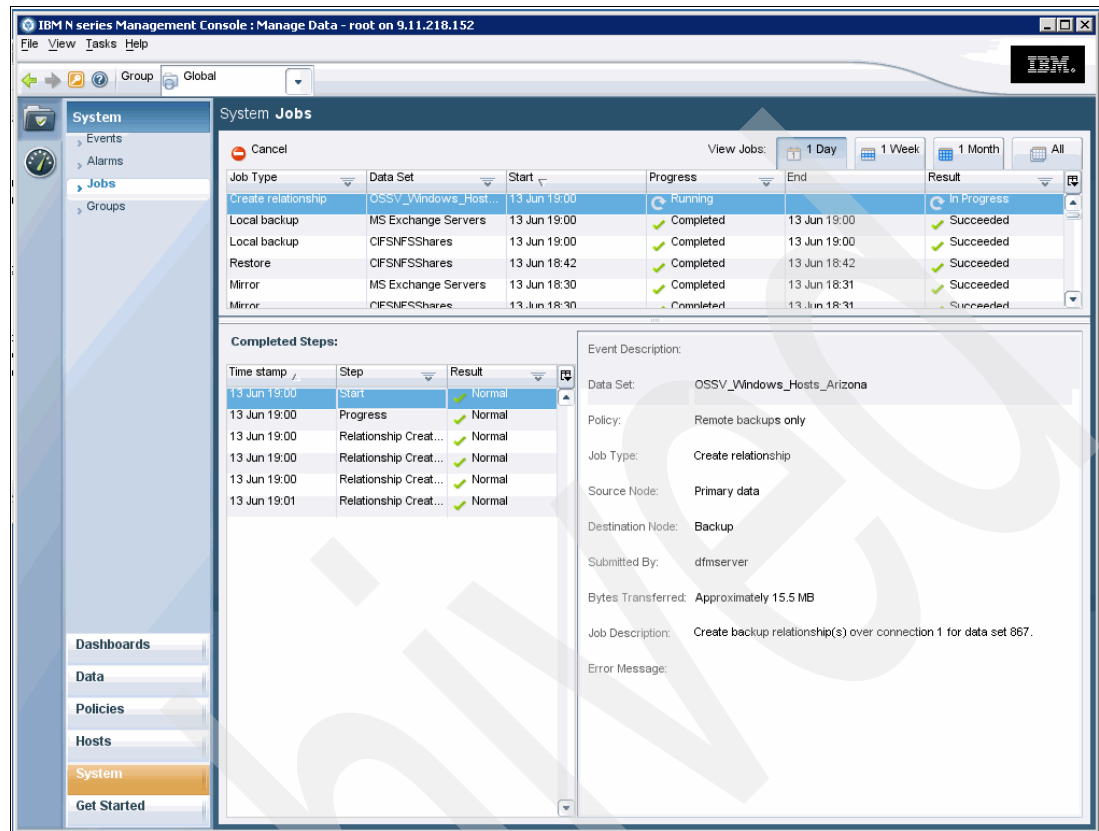


Figure 13-103 Reviewing the System Jobs window while the OSSV protection policy is being applied

As each stage of the job runs, Operations Manager will post event messages providing information about the status of the activity. You can click any one of the events to review the message details, as shown in Figure 13-103.

When the job is finally finished, you have effectively an audit trail of every action taken to provision the data set, its outcomes, and the subsequent steps it took to protect the data set using the policy and resource pools we specified, as shown in Figure 13-104. The event messages will reflect any issues encountered by the DFM Server in this window as well.

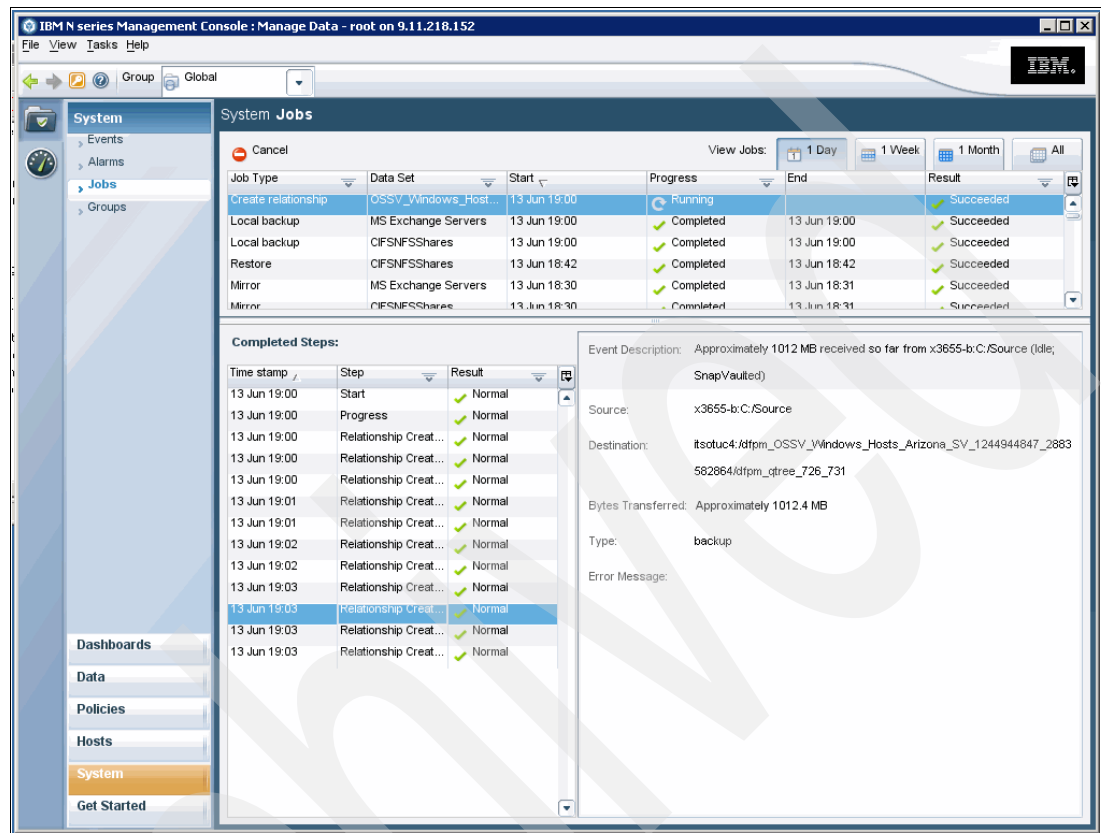


Figure 13-104 Reviewing event messages for the completed job

In Figure 13-105, we can see the details of the OSSV data set we created and its protection status. It will continue to show as Uninitialized until all the replication jobs are completed. Figure 13-106 shows the replication plan and schedule.

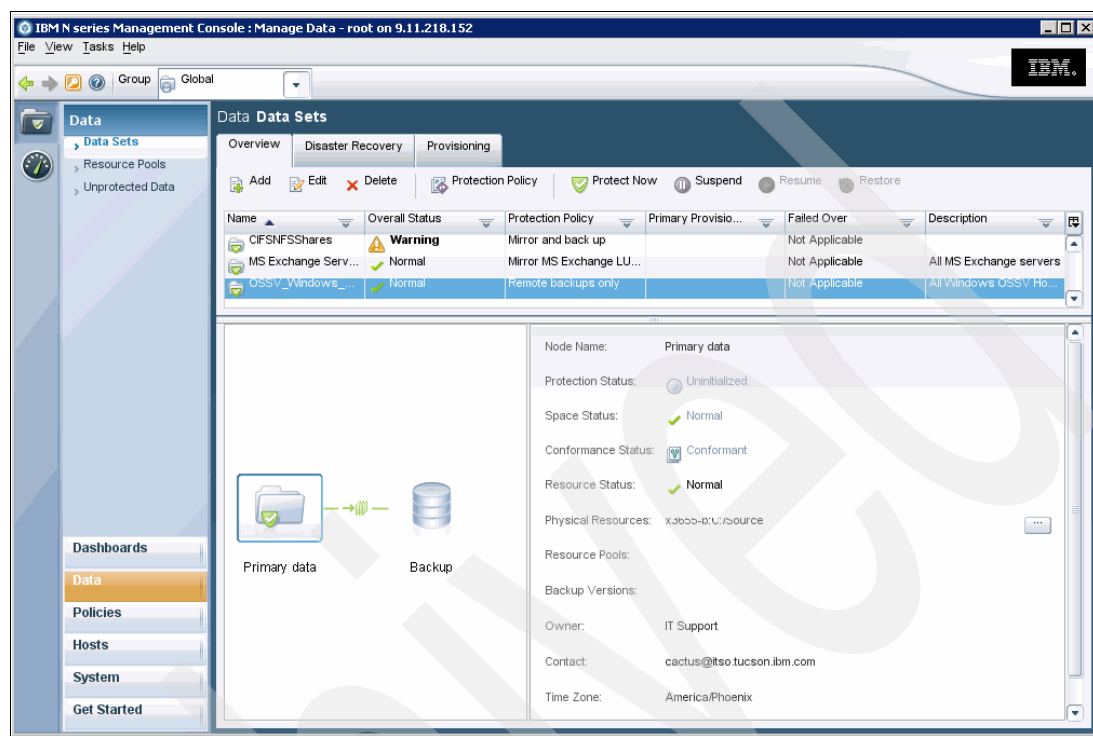


Figure 13-105 Reviewing the OSSV data set - primary data

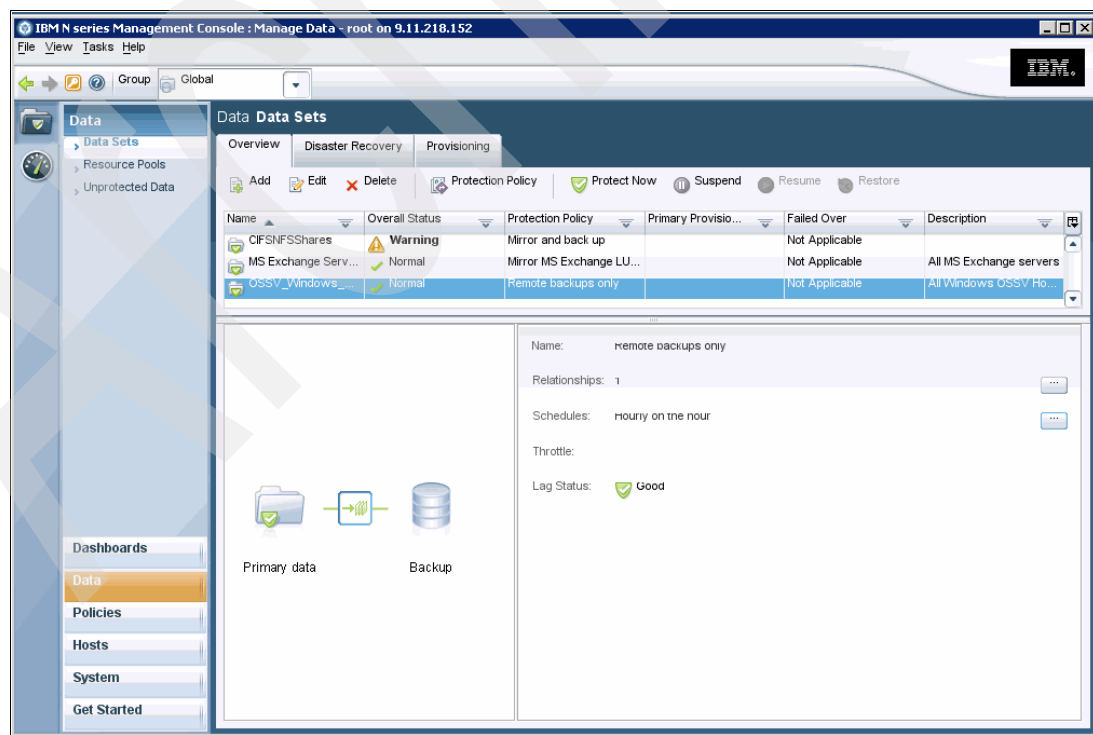


Figure 13-106 Reviewing the OSSV data set - replication method and schedule

Figure 13-107 shows the details of the backup target.

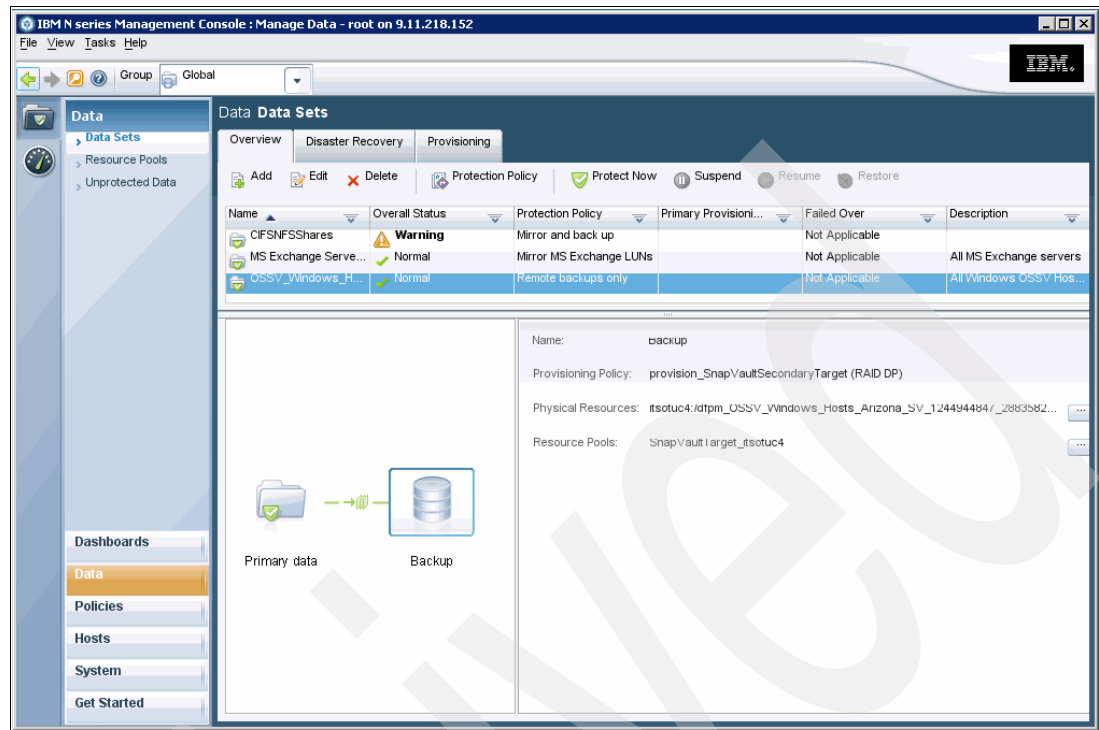


Figure 13-107 Reviewing the OSSV data set - backup target

13.6 Demonstration of restoring data

Protection Manager is able to restore data that was previously successfully backed up with a protection policy and has not yet been deleted.

Two types of data restoration exists:

- ▶ LUN / Volume based restore
- ▶ Qtree based restore

LUN and Volume based restores are used to completely overwrite the current contents of a volume or LUN. Qtree based restores are ideal for restoring individual files or entire directories. Users accessing files stored on N series volumes will do so using CIFS or NFS file access protocols.

In this section, we demonstrate how easy it is to restore data using the Restore command on Protection Manager. The Restore command is available for any data set that was successfully backed using Protection Manager policies.

In our scenario, we wish to restore files that were accidentally deleted from a data set called “CIFS NFS Shares”. This data set represents the files backed up on the CIFS / NFS volume on itsonas1. Figure 13-108 shows the files in the folder. We will use this as a reference to compare the folder to after the restore activity.

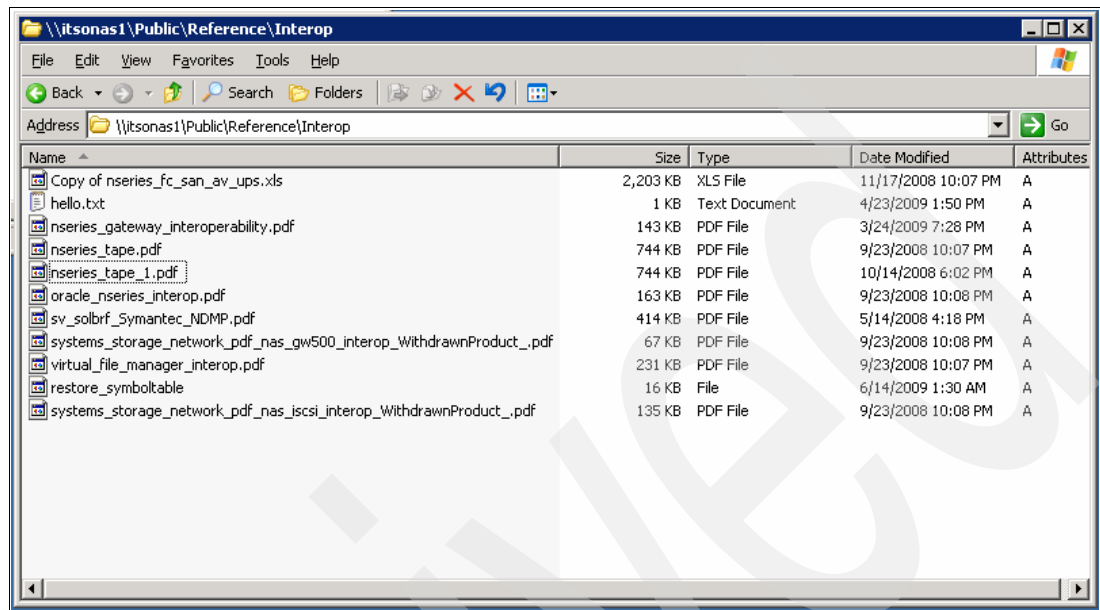


Figure 13-108 Data folder on itsonas1 that has missing files

When you highlight a data set, as shown in Figure 13-109, Protection Manager will indicate if there are successful backups to restore from by enabling or disabling the Restore button on the ToolBar. Click the **Restore** button to launch the Restore wizard.




Figure 13-109 Selecting a data set that requires a restore action

Figure 13-110 shows the beginning of the restore wizard. This wizard will help you select the backup set and files within that set that you wish to restore. Click **Next** to proceed to the next window.




Figure 13-110 Launching the Restore Wizard

Figure 13-111 on page 401 shows the list of backups that have been taken for this data set. You will recall from Figure 13-109 on page 399 that the data set employs two types of backups: a mirror and a SnapVault backup. In this case, we use the SnapVault backup set to locate the files we want to restore.

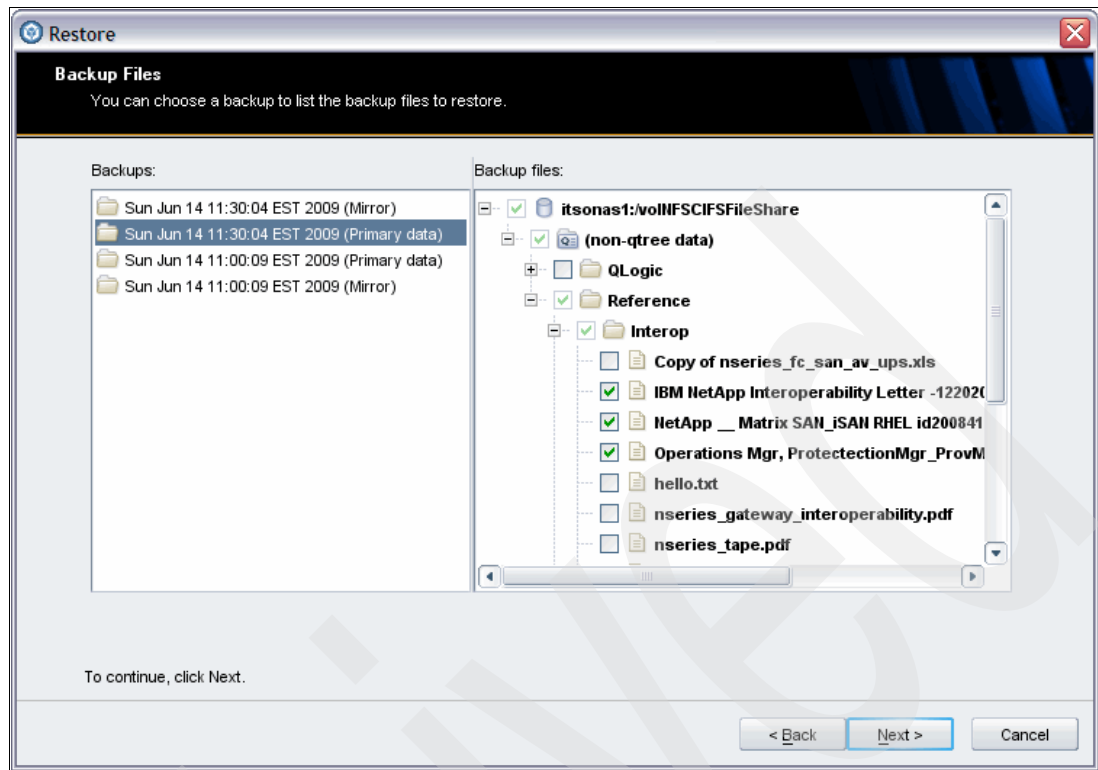


Figure 13-111 Reviewing backup history to select files for restore

Once you have identified the files you wish to restore, click **Next** to proceed to the next window.

In Figure 13-112, you are prompted to indicate if you want the files restored to their original location or to an alternative location. In our case, we elect to restore to the original location.

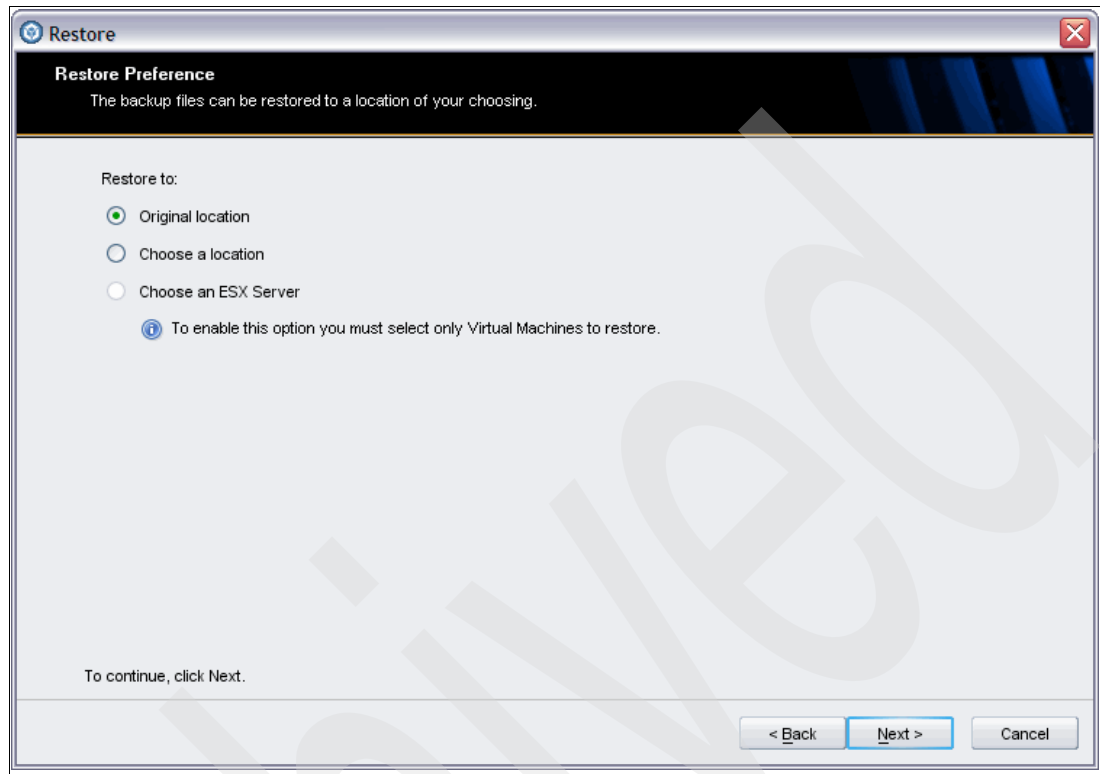


Figure 13-112 Restoring to the original location

Once you have identified the files and restore location, you are taken to the summary window shown in Figure 13-113. Click **Finish** to begin the restore.

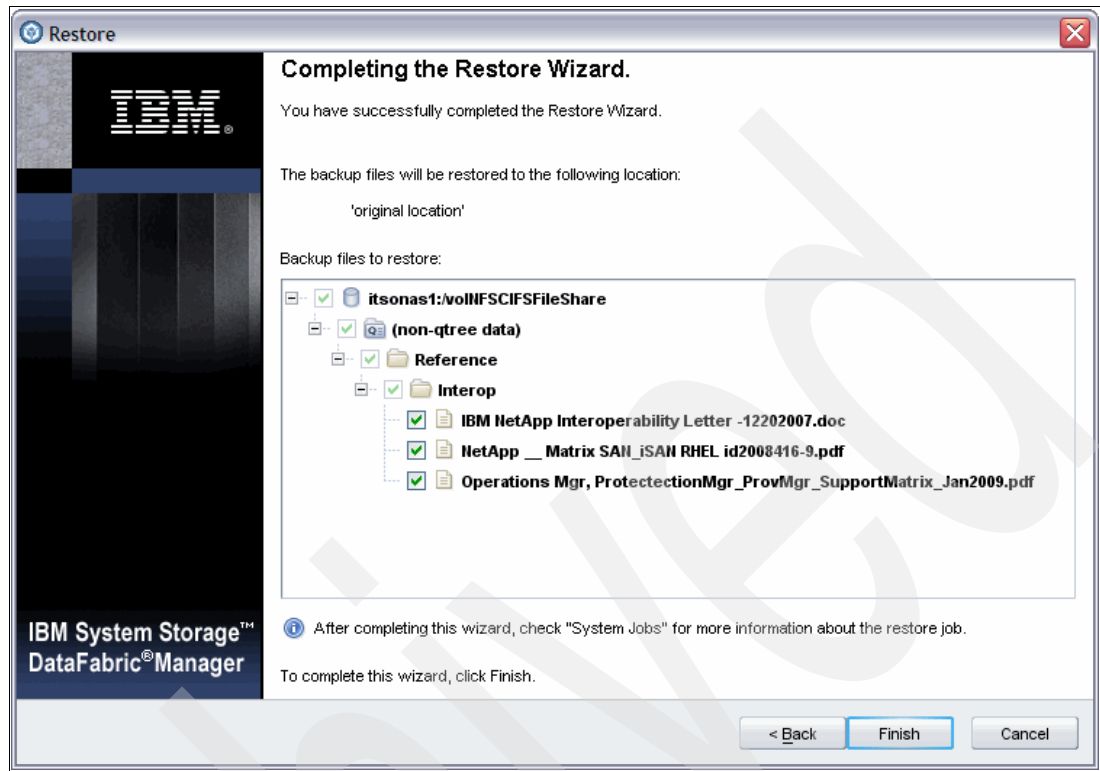


Figure 13-113 Finalizing the restore request with the wizard

[illegible]

At the end of the restore action, we review the CIFS / NFS share on istonas1 and can see that the files we selected to be restored are now visible, as shown in Figure 13-115 on page 405.

Name	Size	Type	Date Modified	Attributes
Copy of nseries_fc_san_av_ups.xls	2,203 KB	XLS File	11/17/2008 10:07 PM	A
hello.txt	1 KB	Text Document	4/23/2009 1:50 PM	A
nseries_gateway_interoperability.pdf	143 KB	PDF File	3/24/2009 7:28 PM	A
nseries_tape.pdf	744 KB	PDF File	9/23/2008 10:07 PM	A
nseries_tape_1.pdf	744 KB	PDF File	10/14/2008 6:02 PM	A
oracle_nseries_interop.pdf	163 KB	PDF File	9/23/2008 10:08 PM	A
sv_solbrf_Symantec_NDMP.pdf	414 KB	PDF File	5/14/2008 4:18 PM	A
systems_storage_network_pdf_nas_gw500_interop_WithdrawnProduct_.pdf	67 KB	PDF File	9/23/2008 10:08 PM	A
virtual_file_manager_interop.pdf	231 KB	PDF File	9/23/2008 10:07 PM	A
restore_symboltable	16 KB	File	6/14/2009 1:43 AM	A
systems_storage_network_pdf_nas_iscsi_interop_WithdrawnProduct_.pdf	135 KB	PDF File	9/23/2008 10:08 PM	A
IBM NetApp Interoperability Letter -12202007.doc	46 KB	Wordpad Document	12/2/2008 12:23 AM	A
NetApp_Matrix SAN_iSAN RHEL id2008416-9.pdf	103 KB	PDF File	4/6/2009 2:13 PM	A
Operations Mgr, ProtectionMgr_ProvMgr_SupportMatrix_Jan2009.pdf	100 KB	PDF File	3/16/2009 3:57 PM	A

Figure 13-115 Missing files are restored to the itsonas1 folder

In this section, we reviewed the steps associated with restoring individual files. The process for restoring entire LUNs is effectively the same, except that any host server that is currently accessing the LUN will need to unmount the LUN or be shut down prior to the restoration. This helps ensure that the file system information in cache memory on the host that is no longer relevant to the LUN will not be accidentally overwritten or flushed back to the restored LUN by the host.

13.7 Creating resource pools, provisioning policies, and data sets

In this section, we review the steps taken to create resource pools, provisioning policies and data sets. We also review the steps taken to provision a data set's storage using one of the provisioning policies. We review these steps here because they are the foundation for successful protection strategies.

13.7.1 Creating a resource pool

For this example, we create a resource pool called *SmallStorageLUNsPoolOn_itsonas1*. These steps are the same to create a resource pool for backup purposes as it is for production data purposes.

The following figures show you how to create a resource pool. We used this technique to create the resource pools listed in this chapter.

We begin by navigating to the “Resource Pools” subsection within the “Data” section of the N series Management Console. To do this task, we follow the navigation steps shown in Figure 13-116.

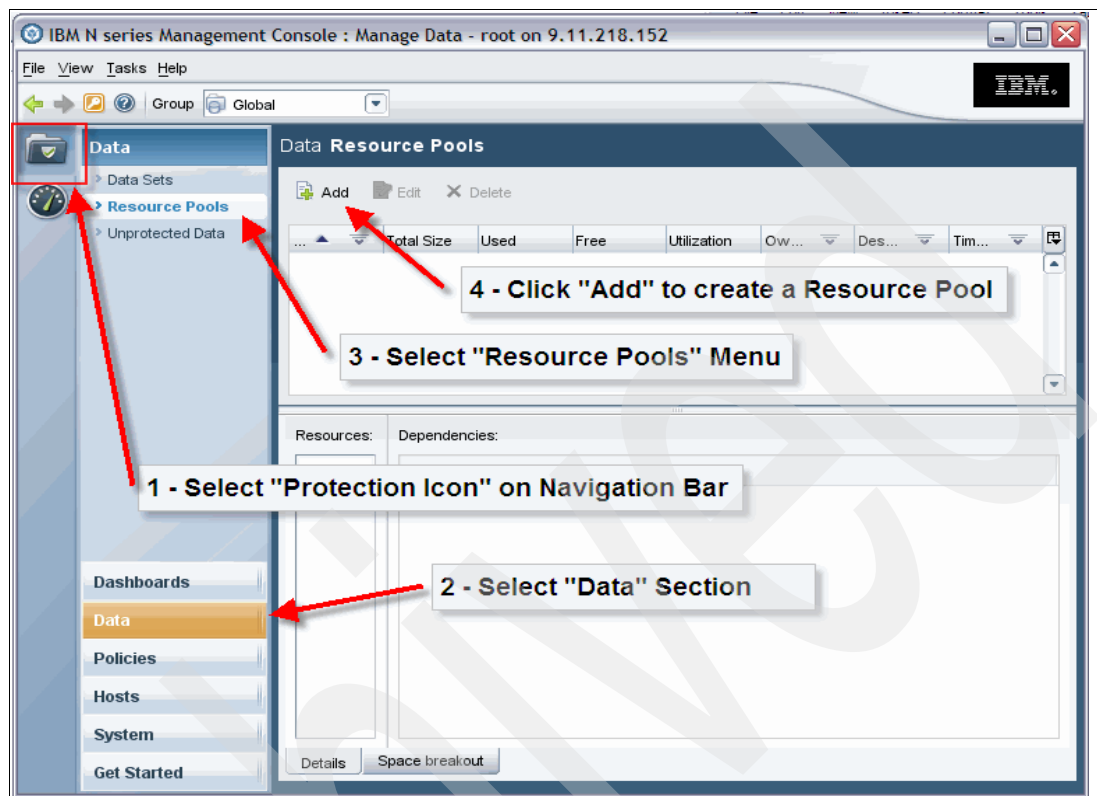


Figure 13-116 Creating a resource pool overview

The Add Resource Pool wizard, shown in Figure 13-117 on page 407, simplifies the process of creating a resource pool. The information required, shown in Figure 13-118 on page 408, is necessary to ensure the resource pool can be uniquely identified.

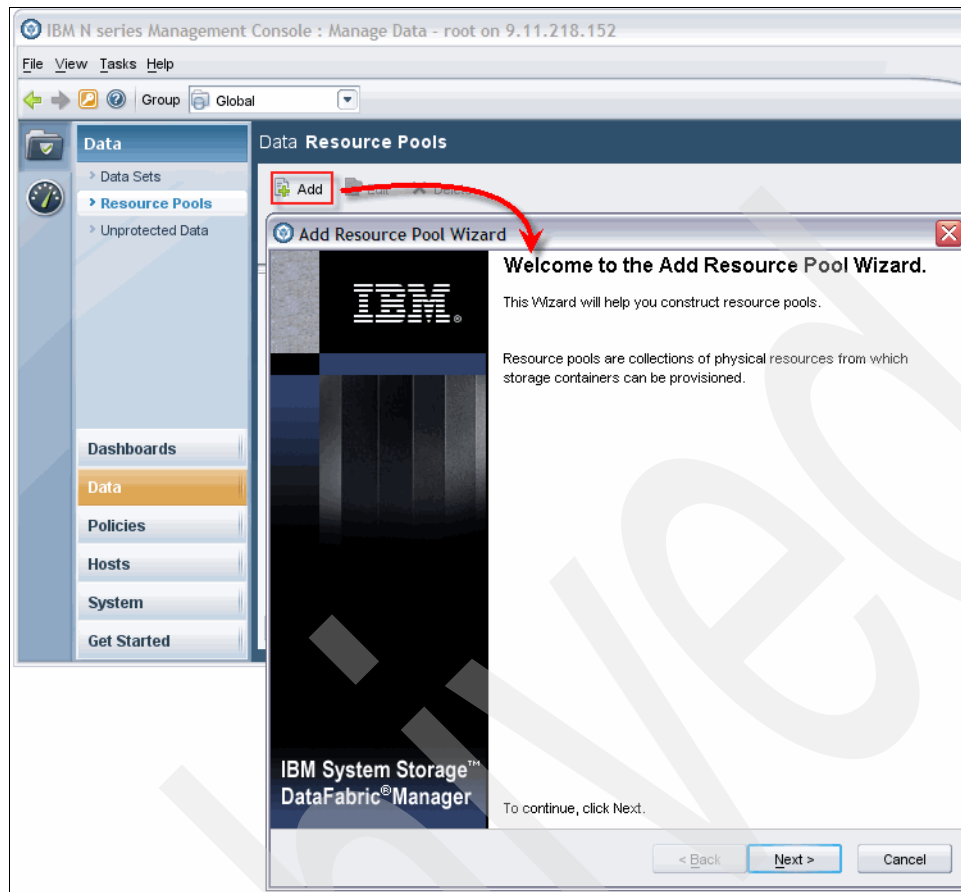


Figure 13-117 Launching the Add Resource Pool wizard

The time zone where the physical resource is located must be correctly specified, as it plays an important role when setting protection schedules later on. The time zone information makes it possible to schedule backups from a central location even though the physical resources may be distributed globally.

Add Resource Pool Wizard

General Properties
You should name your new resource pool for easier identification.

Name:

Description:

Owner:

Contact:

Time Zone:
 America/Panama
 America/Pangnirtung
 America/Paramaribo
America/Phoenix
 America/Port-au-Prince
 America/Porto_Acre
 America/Port_of_Spain
 America/Porto_Velho
 America/Puerto_Rico

To continue, click Next.

< Back Next > Cancel

Figure 13-118 Uniquely identifying the resource pool

As shown in Figure 13-119 on page 409, a resource pool represents physical storage and can contain one or more aggregates from one or more N series storage systems. We specify aggregates when we wish to restrict Protection Manager and Provisioning Manager to automatically creating volumes within the designated aggregates on shared storage systems.

It is also possible to designate entire storage systems as a resource pool. In this scenario, Protection Manager and Provisioning Manager use machine generated names to manage the actual aggregates and volumes.

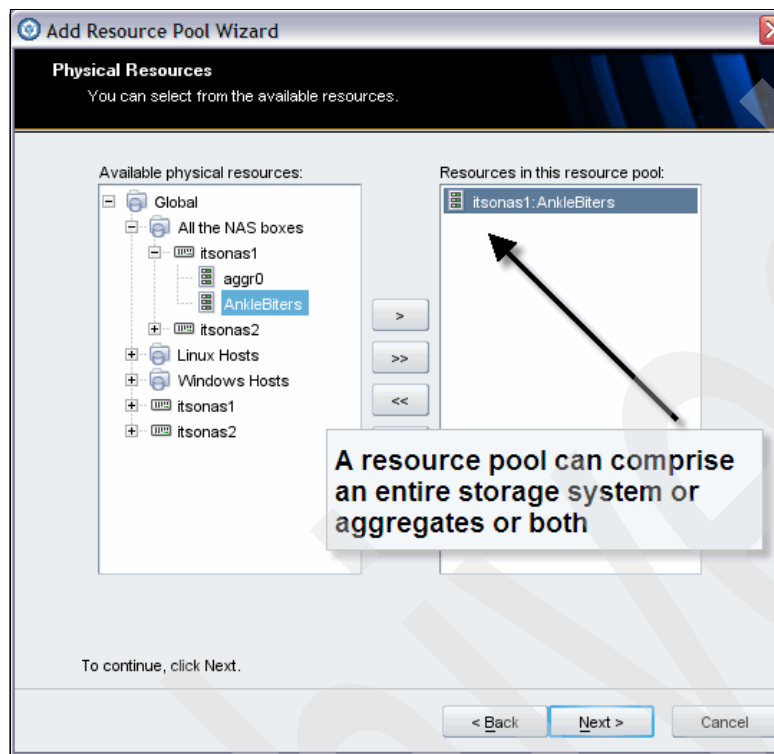


Figure 13-119 Adding members to the resource pool

The wizard window shown in Figure 13-120 provides you with a very powerful method of limiting the choices Provision Manager or Protection Manager will have when provisioning storage for a particular use.

For example, you may have a N series storage system with the NearStore option or a low order N series storage array you wish to use as a SnapVault target. You could designate specific resources (aggregates or the entire storage system) on that storage system with the label “SnapVault”. Similarly, you may have other aggregates and storage systems labeled “SnapMirror”, all defined within the same resource pool. Then, when provisioning data sets for SnapVault usage, you specify “SnapVault” in the data set or policy label. When provisioning, you specify this resource pool, and Provisioning Manager will restrict its selection of candidates to those with the resource label “SnapVault”.

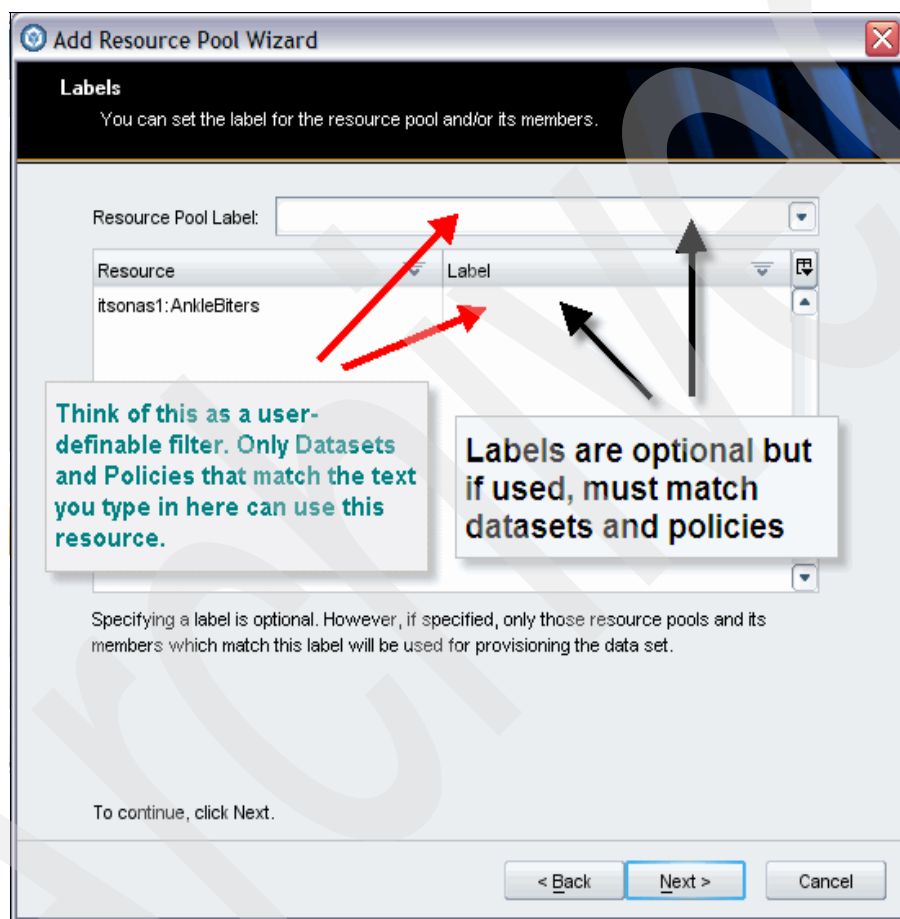


Figure 13-120 Assigning labels to resources

Note: Each resource within the resource pool can have a different label. This can get very complicated and will be challenging to troubleshoot. Therefore, care must be taken if you plan on using labels to further differentiate resource pools.

Figure 13-121 is an example of what happens when the resource label specified does not match in the provisioning policy; Provisioning Manager is unable to find a resource that satisfies the criteria during a process to create a data set and fails to provision the resources for that data set.

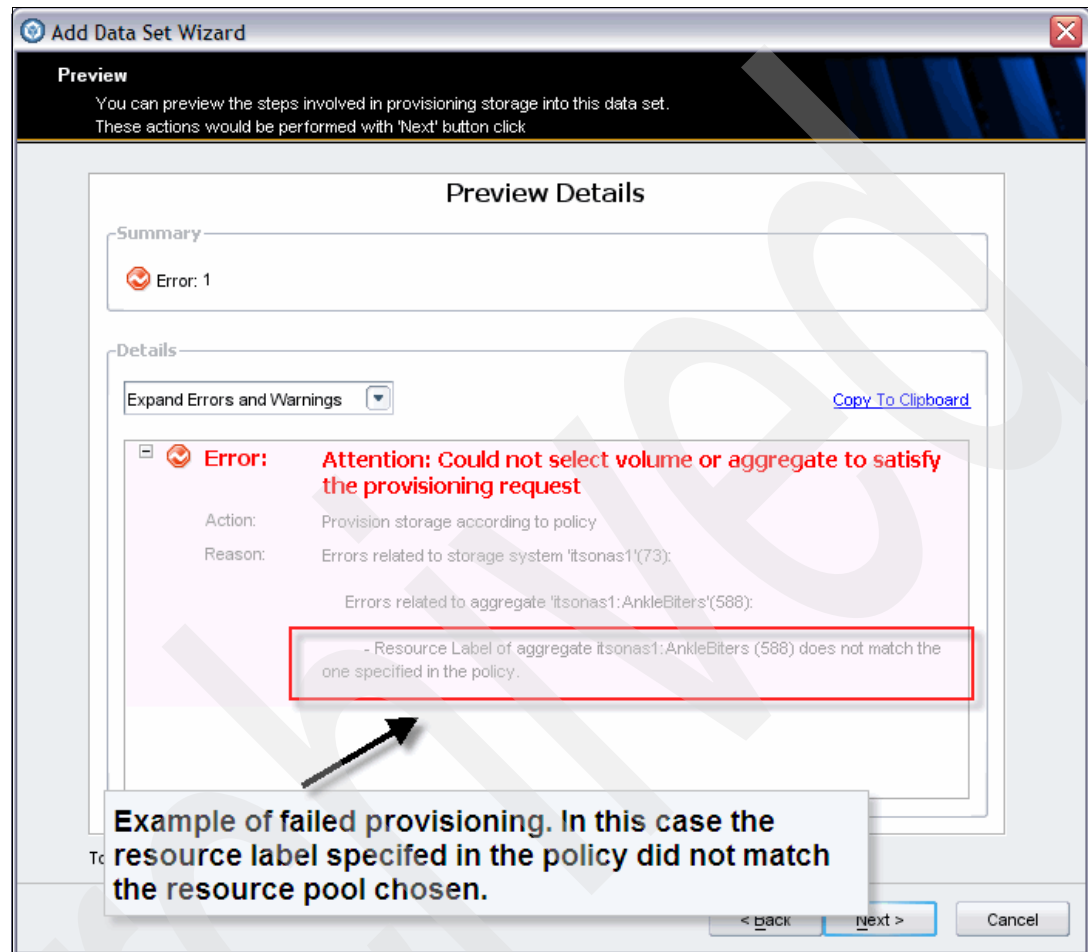


Figure 13-121 Example error in locating suitably labeled resources

In our resource pool provisioning example, we omit the use of labels, as we plan to create a resource pool for SnapVault targets and another resource pool for SnapMirror targets. We continue with the Resource Pool Wizard, as shown in Figure 13-122.

Add Resource Pool Wizard

Space Thresholds
You can configure thresholds on the space usage for resource pool.

Space thresholds

- ☒ Enable event generation
- Nearly Full threshold (%): 80 (1st alert threshold)
- Full threshold (%): 90 (2nd alert threshold)

Aggregate overcommitted thresholds

- ☒ Enable aggregate overcommitted thresholds
- Nearly overcommitted threshold (%): 300 (Thin provisioning overcommit alert thresholds for this resource pool)
- Overcommitted threshold (%): 400 (Thin provisioning overcommit alert thresholds for this resource pool)

To continue, click Next.

< Back Next > Cancel

Figure 13-122 Setting the space threshold for the resource pool

Space thresholds are very important tools to help ensure that timely and adequate warnings are generated and sent to the administrator. Armed with these warnings, the administrator is able to take corrective action, including procuring additional storage if necessary. It is important you set these values to match the budgetary and procurement lead times your organization requires to help ensure that your SLAs for protection will not be compromised.

As we can see in Figure 13-122, you have two alert thresholds as well as special thresholds to warn you if you have over committed an aggregate as a result of aggressive thin provisioning. Aggressive over commitment will increase the likelihood of space thresholds being exceeded quickly.

Figure 13-123 show the summary window showing what was created. Click **Finish** to exit this wizard.



Figure 13-123 Completing the Add Resource Pool Wizard

After the resource pool has been created, you will once again be presented with the Data Resources Pool pane, as shown in Figure 13-124.

You can see, for the highlighted resources pool, details about the pool, such as the physical resources that make up the pool, the total capacity of the pool (across all its resources), the capacity currently in use, its utilization, owner, and the time zone it is in.

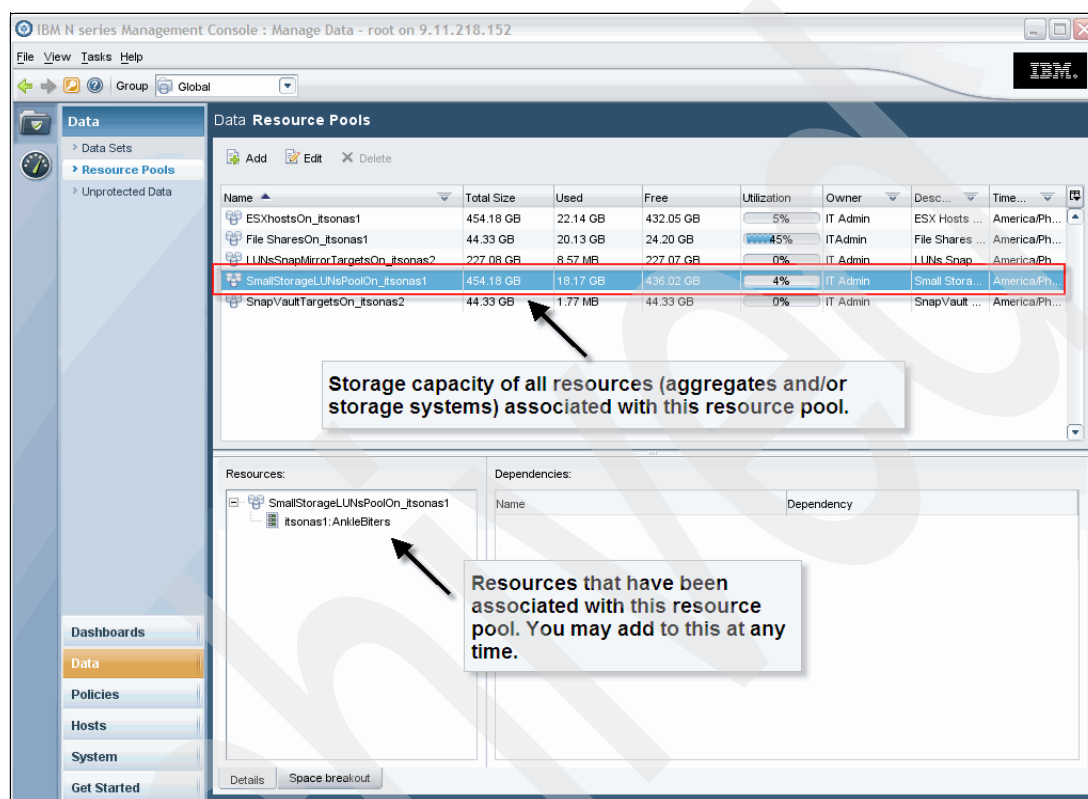


Figure 13-124 Resources Pool pane within data section of N series Management Console

We use this information to keep track of the usage of the pool. The time zone information is useful for us to ensure that protection policies that run on the resource pool will run at off peak times in relation to the location of that pool.

Accordingly, we generally do not place resources from different time zones in the same resource pool unless the protection policy that needs to be run is being applied simultaneously against resources across geographies.

13.7.2 Creating the provisioning policies

This section will demonstrate how we created provisioning policies. The method used here was used to create the provisioning policy used in this chapter.

In Figure 13-125 on page 415, we can see the provisioning pane of the Policies section for the N series Management Console. We labeled policies we created to reflect the type of data protection the policy will require to make it easier for us to track. The actual data protection requirement of the policy is set in the policy details of each policy.

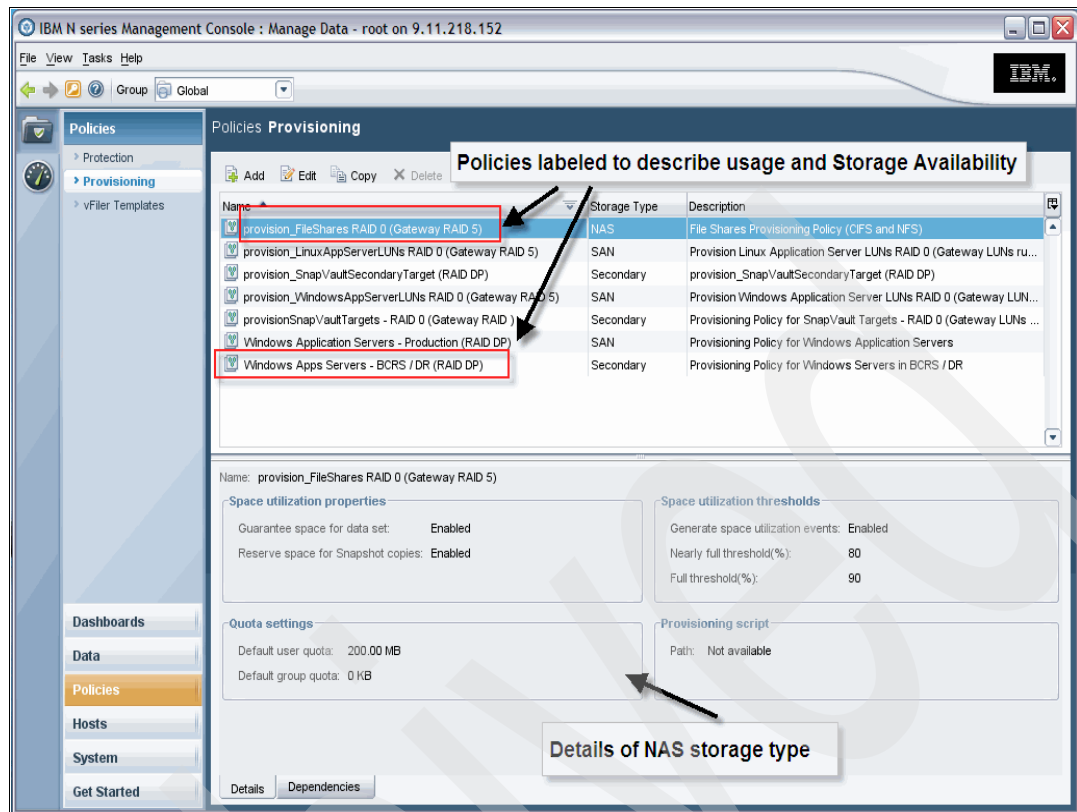


Figure 13-125 Overview of provisioning policies created

In our scenario, we elect to create provisioning policies that reflect the various classes of servers (Windows, Linux, and so on), usage types (CIFS, NFS, and so on), and their use (production, development, testing, backup, and so on). With this method, we are able to specify the data protection requirements of the different types of servers.

The scenario assumes a single location for all the resources. If we had resources and servers distributed across different geographic locations, we would be able to use our provisioning policies against resource pools in those geographic locations, as the policies do not have a means to be limited by time zone or geographic location.

The data protection tab of a policy detail, shown in Figure 13-126, is a little misleading, as it implies that you do not get any data failure protection if you use LUNs presented from an external storage array. This is obviously incorrect. What this diagram is trying to say is that the gateway will not be using its *own* failure protection routines (RAID DP or RAID 4) to protect the contents of this storage resource. The responsibility of protecting the data integrity of this resource will reside with the storage system delivering the storage, such as an IBM System Storage DS4000® or IBM System Storage DS5000 system, IBM System Storage DS8000®, or SAN Volume Controller, which would employ RAID 1, RAID 10, RAID 5, or RAID 6 to protect any LUNs being presented to the N series gateway.

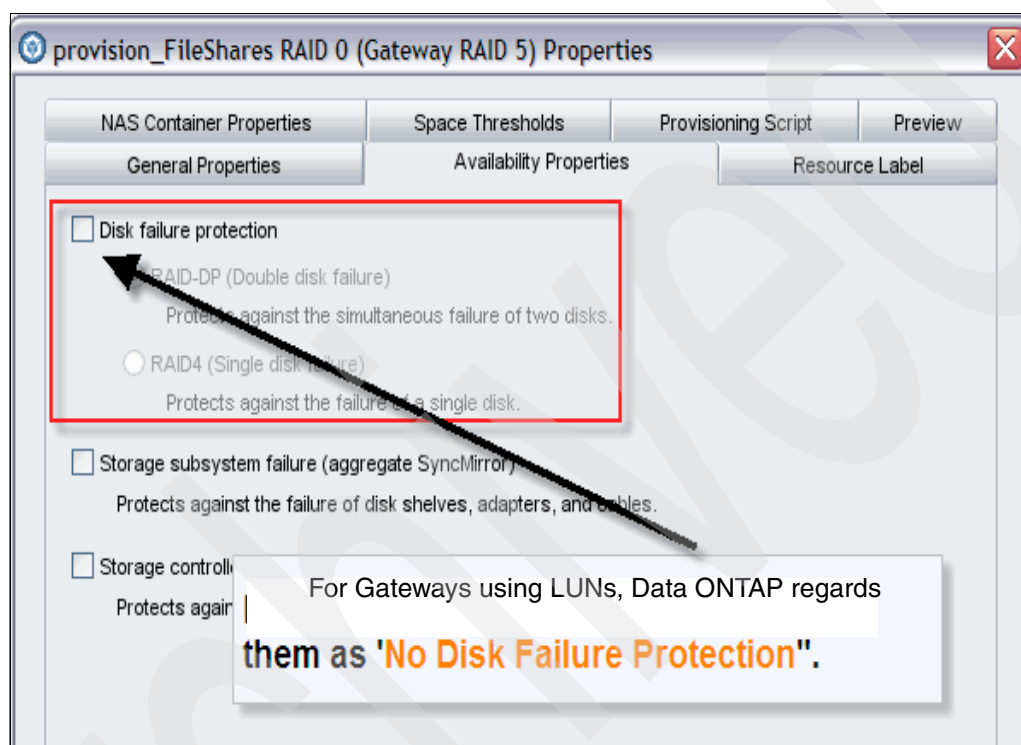


Figure 13-126 Unchecked selection implies that a gateway with LUNs provided by external storage array, such as DS4000 or DS5000 with RAID 1 or RAID 5 LUNs

When Protection Manager needs to provision storage and has been instructed to use a provisioning policy, it will check to ensure that only resources in resource pools that comply with the protection or availability requirements of the policy will be used.

In Figure 13-127 on page 417, we continue to demonstrate the practice of including the protection type in the name of the policy. In the details section, you can see the SAN container properties, provisioning script (if there is one), and the space utilization thresholds set for the highlighted policy.

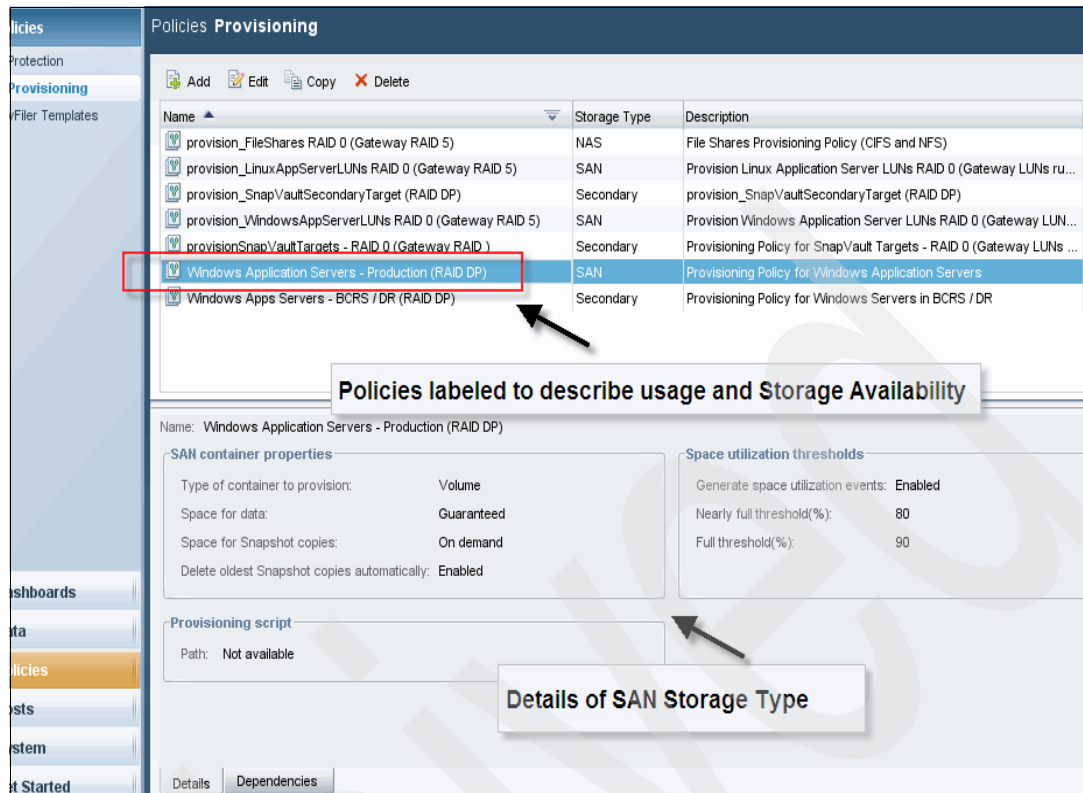


Figure 13-127 Include the storage availability type in the name for easy of use

The properties of this policy reveal the data protection or availability requirements, as shown in Figure 13-128.

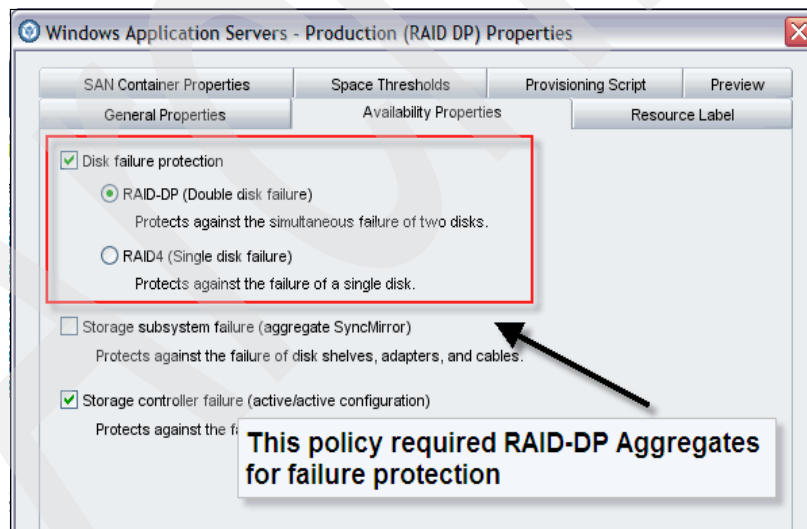


Figure 13-128 RAID DP selection implies vFiler with local storage enclosures (no gateway LUNs)

Note: The data protection or availability requirements shown in Figure 13-128 indicate the minimum data protection level required to satisfy the requirements of this policy. If we cleared the check box, the minimum protection requirements will be RAID 0. In this case, resource pools with RAID DP protection would still qualify.

As stated before, care must be taken when using resource labels, as they form an additional filter, as shown in Figure 13-129. Only resources that have a matching resource label can be used with such policies.

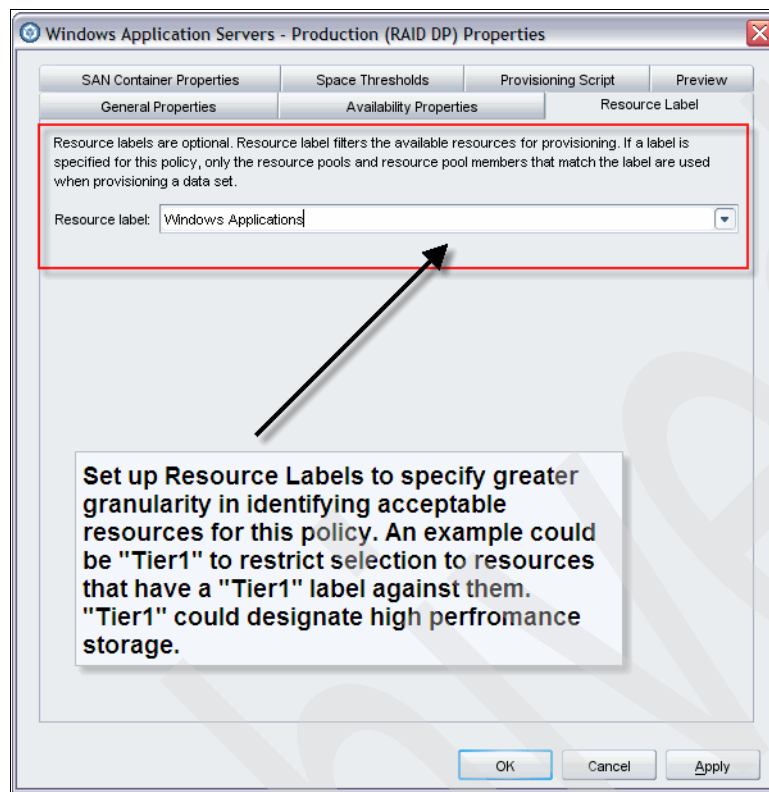


Figure 13-129 Using resource labels to filter resource choices

13.7.3 Provisioning the data sets using provisioning policies

This section discusses how we provisioned a data set for use by our Linux host. Operations Manager introduced the capability to create data sets by applying provisioning policies against resource pools designated for production use. We used this capability to allocate storage for our Linux host.

Note: Additional steps will still be required on the host after the data set is provisioned to the Linux host, such as adding the provisioned LUN to a LVM, formatting and mounting, and so on.

In Figure 13-130, we begin with the process of creating a data set by clicking the **Add** button to launch the Add Data Set Wizard. Click Next to get to the next window, as shown in Figure 13-131 on page 420.

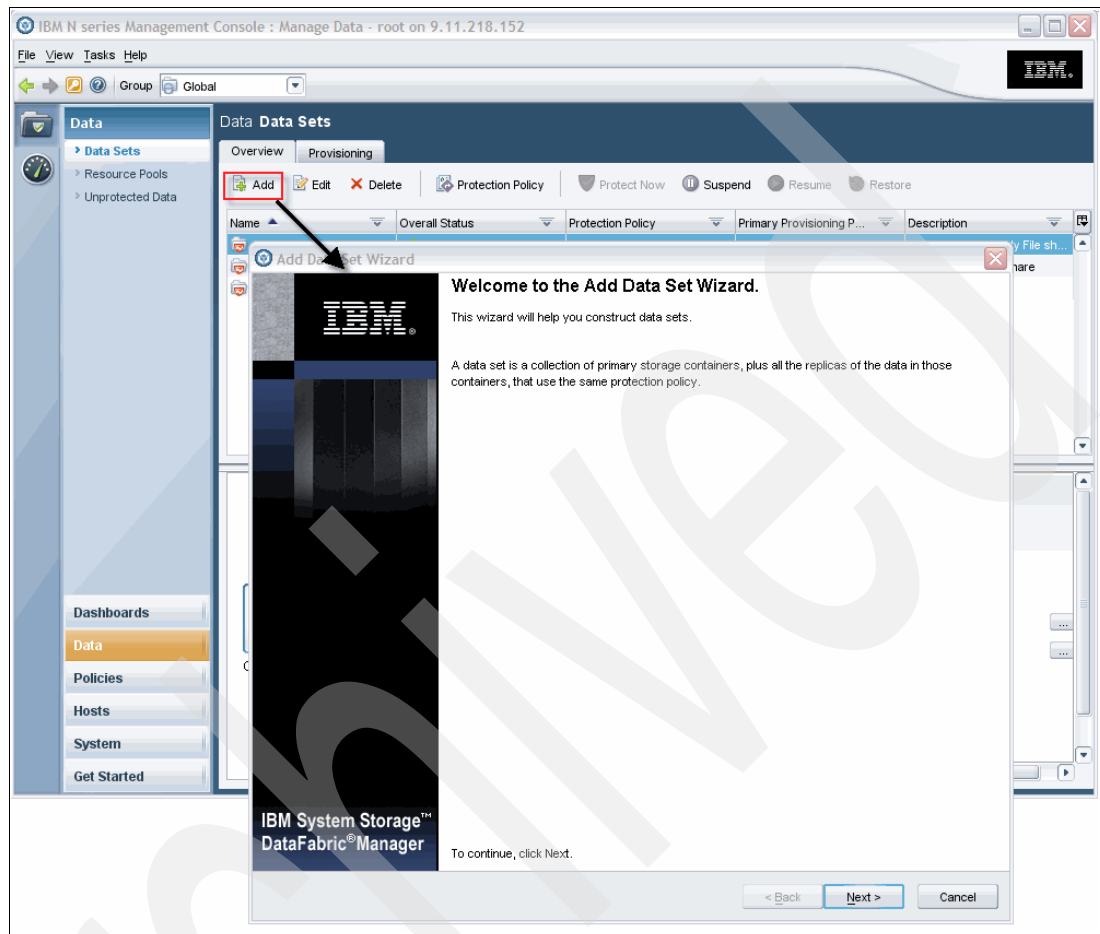


Figure 13-130 Launching the Create / Add Data Set Wizard

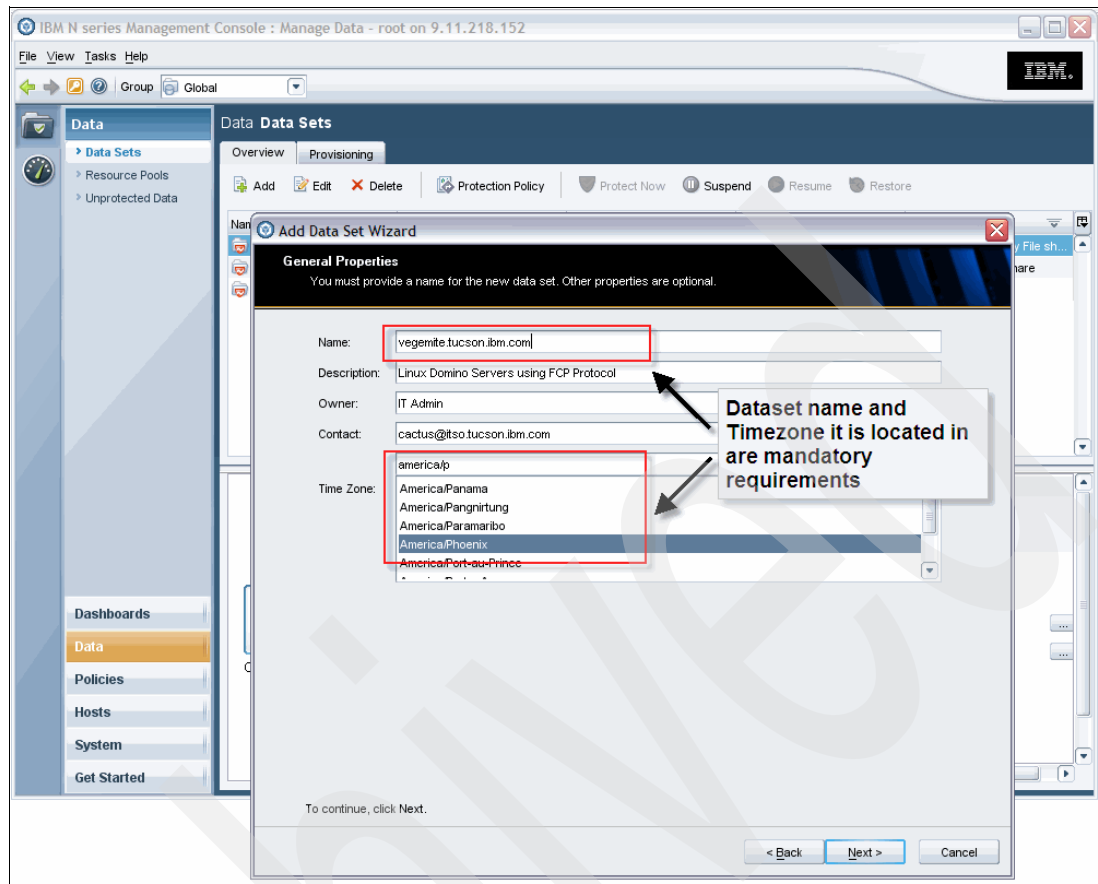


Figure 13-131 Specifying pertinent information for the creation of a data set

We decided to use the host name as the data set name in this example and specified the location of the host and its data set by selecting the America/Phoenix time zone. This naming convention might be applicable for small server installations; however, for more complex server installations, such as an SAP® implementation that runs across several hosts, a different naming convention may be more appropriate.

In more complex implementations, the data set may need to accommodate all the servers for a particular application if there are transactional or data interdependencies between the servers at an application level. This is because only one protection policy can be assigned to a data set. If you have transactional and data interdependencies across several servers within your application, you may need to collectively back up or Snapshot the volumes in the data set to ensure data consistency at an application level.

The most appropriate way to determine this situation is to consider what recovery methods would be successful for your application and select the one that is least complex to work with but will deliver the recovery outcomes your business requires. Then, with this information, you will be able to determine if one data set or more should be used to represent your complex application's components and assign the protection policies accordingly.

Note: Scripting is also available for more targeted and complex coordination activities to help ensure that you are able to back up and restore the application's data in a consistent state.

A data set must also be associated with a group to simplify management, as shown in Figure 13-132. In our example, we selected the Linux Host group and clicked **Next**.

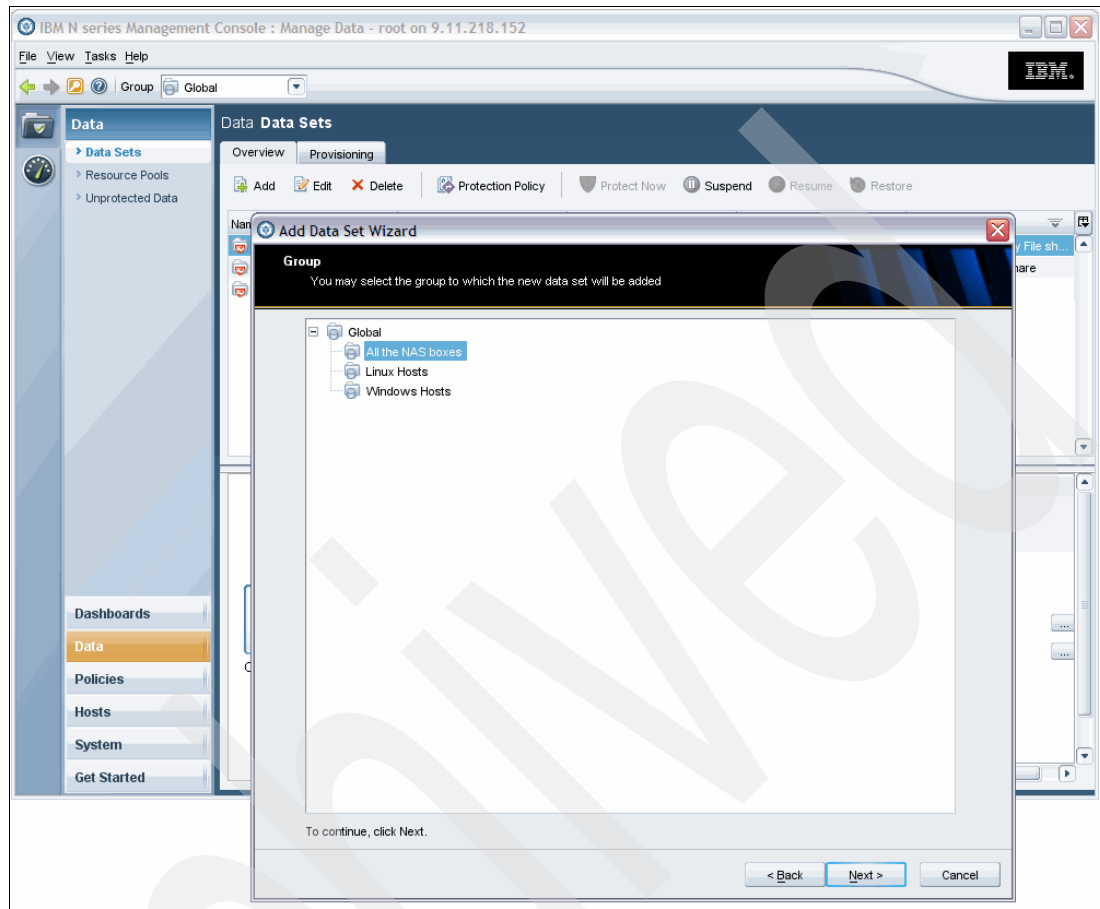


Figure 13-132 Identifying the group within which to create the data set

In Figure 13-133, we elect to provision the resources using a provisioning policy. We selected the provisioning policy called “*provision_LinuxAppServerLUNs RAID 0 (Gateway RAID 5)*”, as shown in Figure 13-134 on page 423.

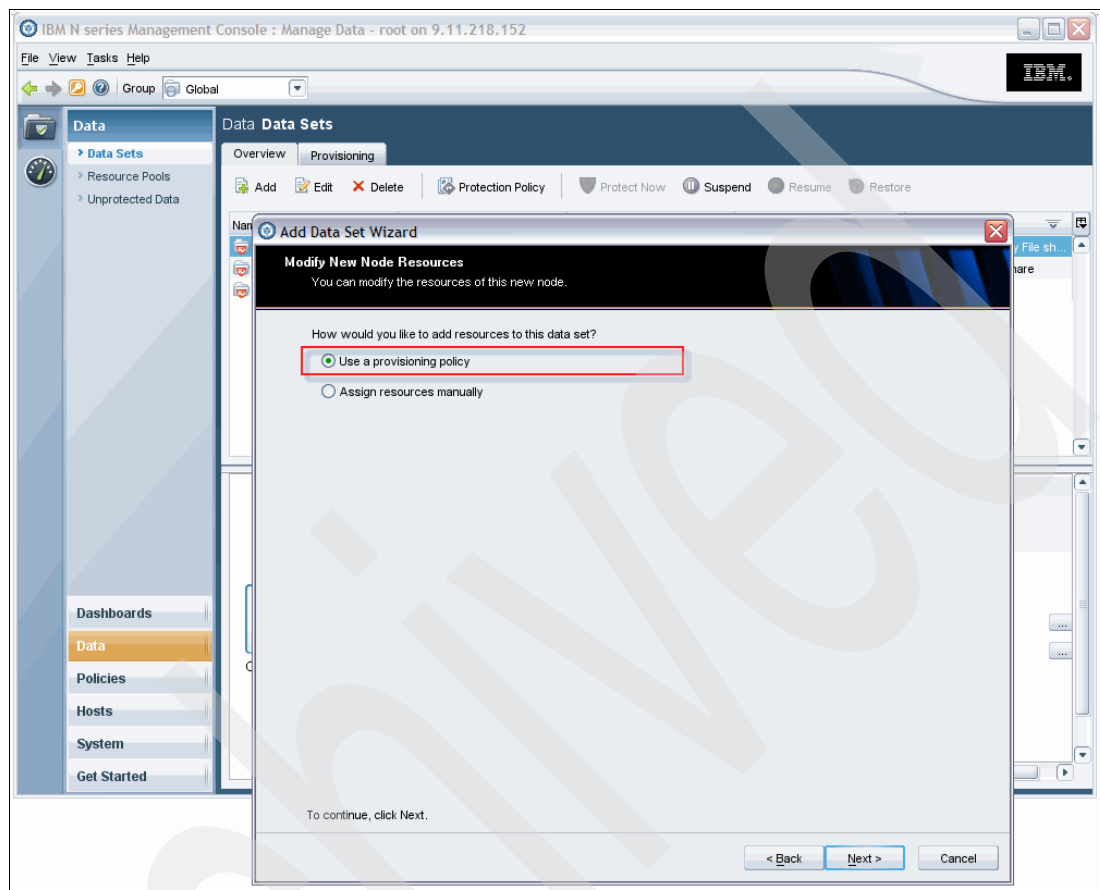


Figure 13-133 Using a provisioning policy to provision storage

Of the available compatible resource pools, we selected the resource pool called “SmallStorageLUNsPoolOn_itsonas1”, as shown in Figure 13-134, and clicked **Next**.

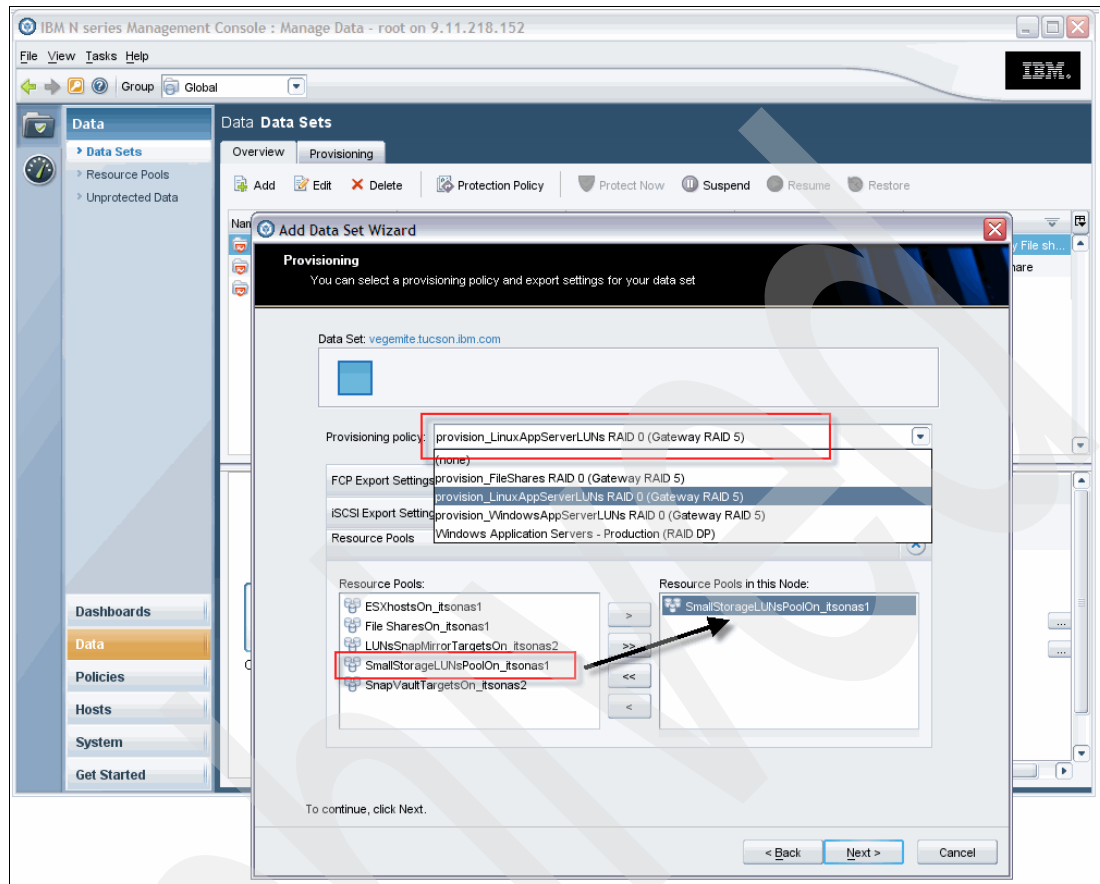


Figure 13-134 Identifying the resource pool from which to provision storage given the provisioning policy

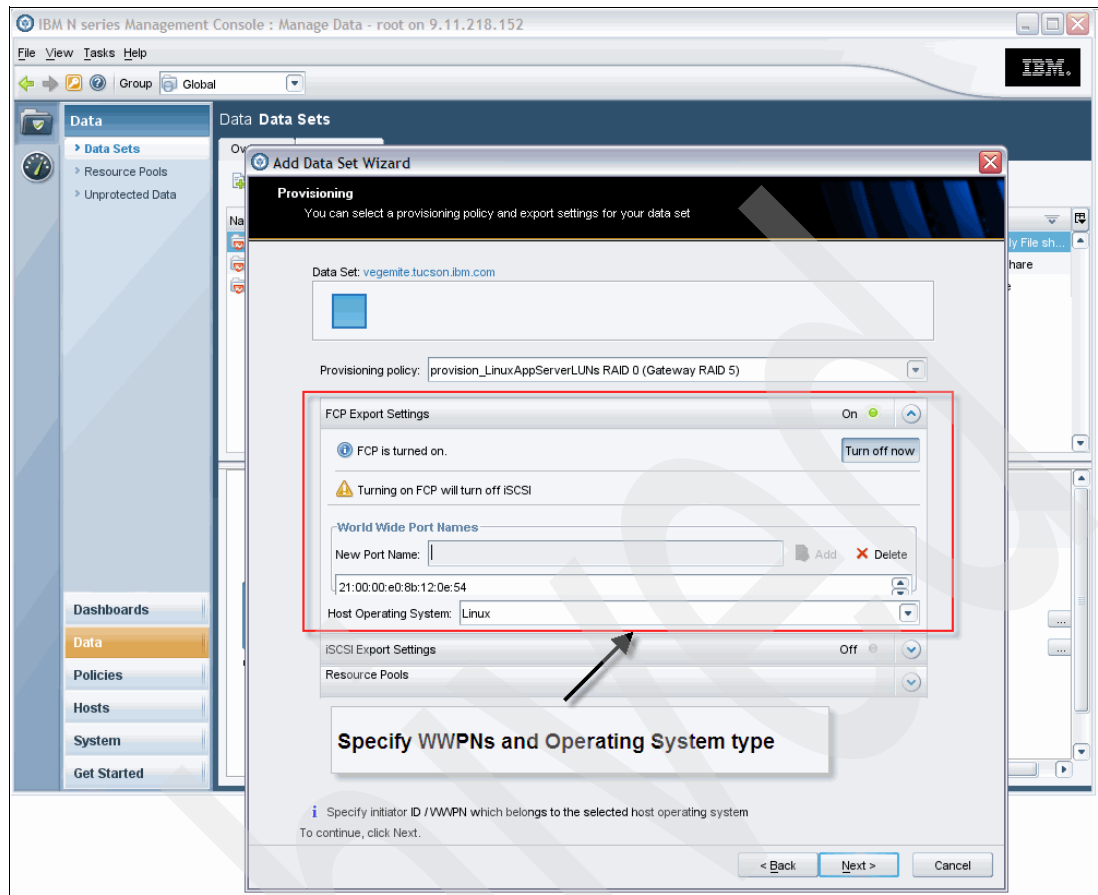


Figure 13-135 Defining WWPNS to be assigned to this data set for FCP access

In Figure 13-135, you can see that we elect to map the LUNs to be created to the Linux host using the FCP protocol only.

In this window, we can type in the World Wide Port Name (WWPN) of each Fibre Channel HBA port the Linux host will use to access the LUNs. The WWPN may be obtained through a number of methods including querying the FCP utils package on the host, running SANSurfer (we were using Qlogic HBAs), interrogating the SAN Switches, querying the storage system to get a list of logged in WWPNS, or consulting the HBA labels if you are able to get physical access to the HBAs.

Note: You will notice that in Operations Manager we cannot have both FCP and iSCSI configured for the same client at the same time to access the resource pool.

The format of the WWPN is specific, as shown in the diagram. It is also important to correctly specify the host type ("Linux" in our case), as this will instruct Data ONTAP on the storage system where the resource is to be created to set the correct byte offset for partition alignment when creating LUNs. This will help ensure the maximum performance of the LUN.

Note: Ensure you specify the correct WWPNs at this point. The provisioning wizard will use the provided WWPNs to locate and remove the WWPNs from and igroups that currently contain these WWPNs and then create a new igroup containing these WWPNs only. The wizard will not warn you about what it is about to do and if you specify the wrong WWPNs, you run the risk of the unintended outage on a server that has resources assigned to it on this storage array, as it loses access to LUNs it previously had.

Note: Extra care must also be taken if you plan to use a provisioning policy for an existing host that already has LUNs assigned to it on the same storage system the new resources will be located because of what the wizard will do to the igroups. You will need to ensure that the data set will encompass those resources as well and that they will be manually assigned instead of being provisioned.

In Figure 13-136, we elect to provision the resource at a later time so that we can observe the process from the main data sets pane. We click **Next** to go to the Preview Details window, as shown in Figure 13-137 on page 426. The wizard will test the provisioning steps and provide an indication of the success or failure of the operation.

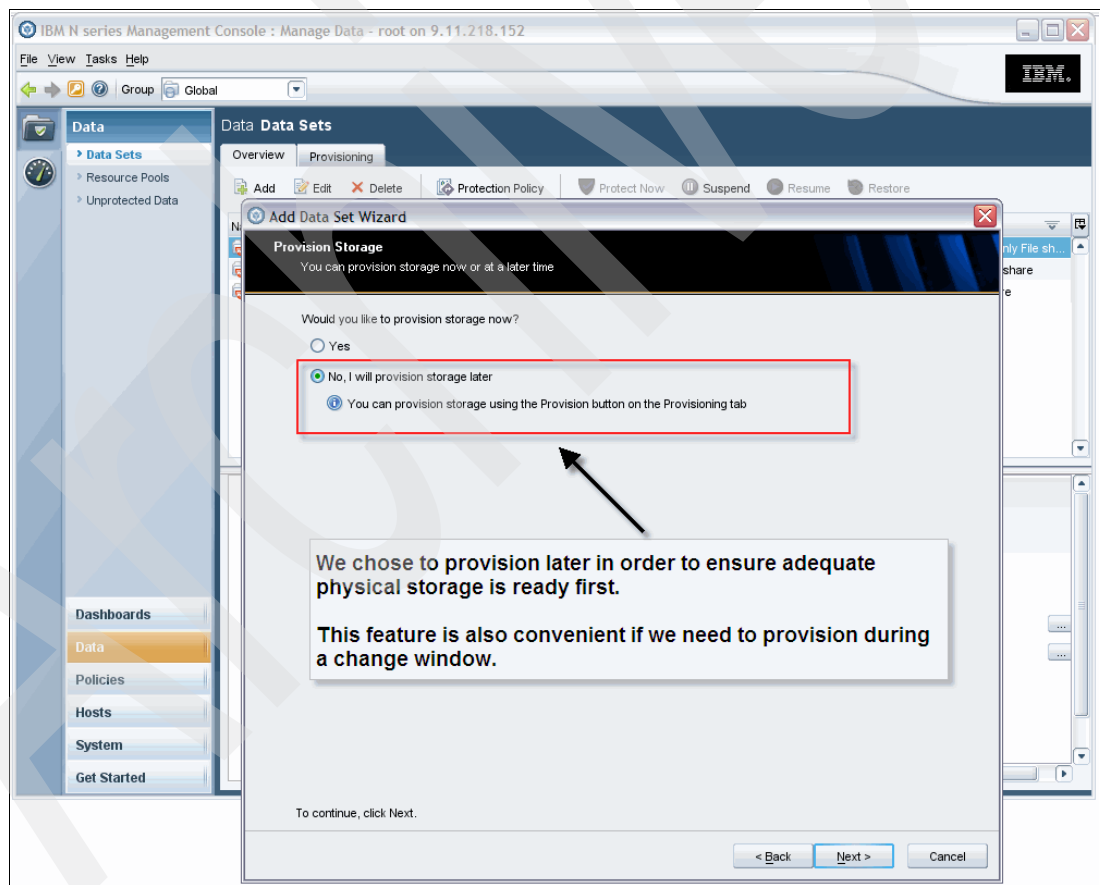


Figure 13-136 Provisioning physical storage at a later time

As shown in Figure 13-137, the wizard completed the testing and has found no errors.

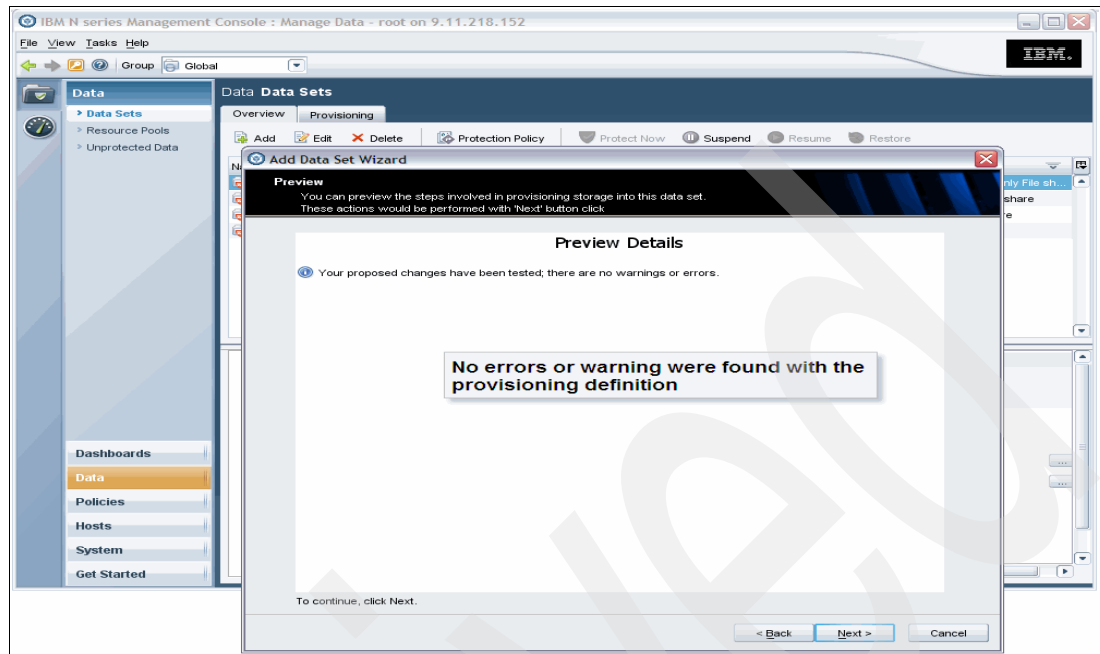


Figure 13-137 Preview reveals no errors

Figure 13-138 shows the summary of what the wizard will do when you click Finish. The results of this activity will be visible in the data sets pane, as shown in Figure 13-139 on page 427.

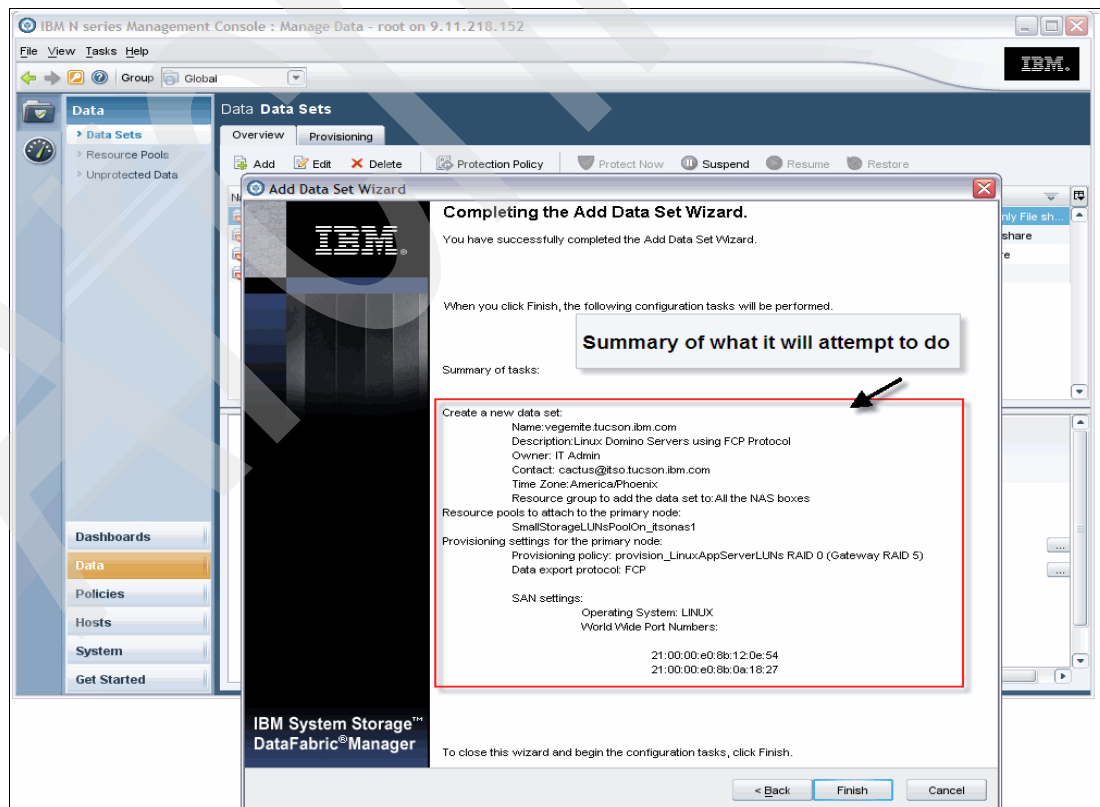


Figure 13-138 Summary of tasks to be executed to define the data set

You can review the jobs that were executed and the corresponding events for this activity in the respective jobs and events panes in the System section of the N series Management Console.

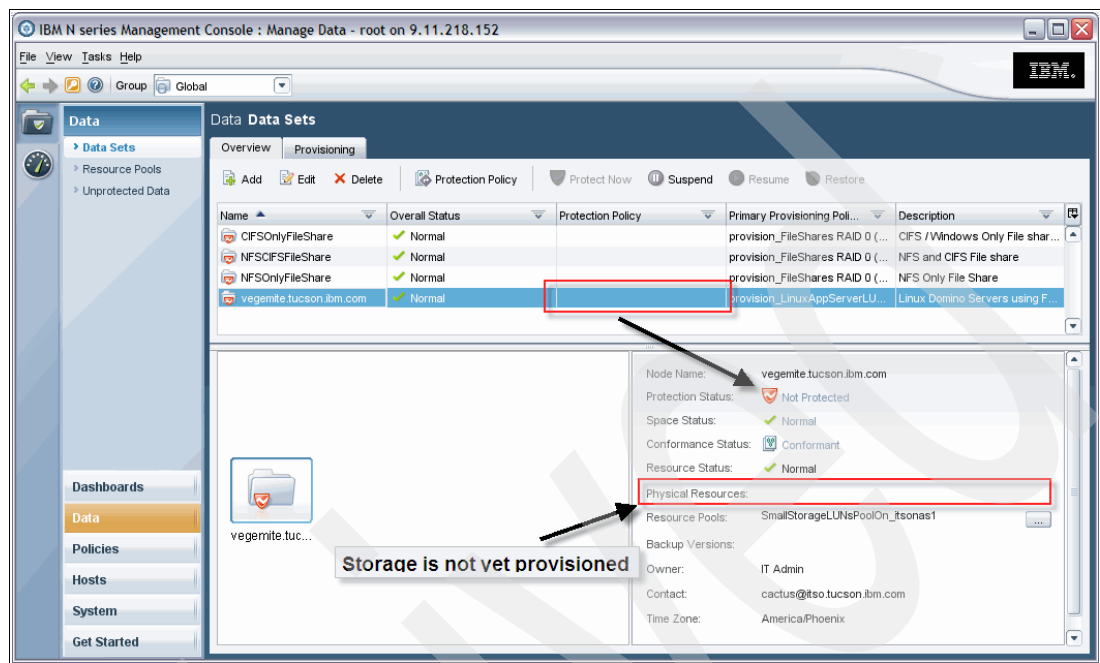


Figure 13-139 List of data sets defined but no physical resources have yet been provisioned

As shown in Figure 13-139, the Overview tab displays the data sets. The newly created data set does not have any resources assigned to it, and there is no protection policy assigned to it, because we elected to not have the wizard provision the resources automatically.

We will now review the Provisioning tab and then provision some storage and apply a protection policy.

As you can see in Figure 13-140, the Provisioning tab displays details about provisioned or allocated resources to each data set. Here we can click the Provision button to provision storage for the highlighted data set.




Figure 13-140 Navigate to the data sets Provisioning tab to launch the Provision Wizard

Clicking the **Provision** button will launch the Provisioning Wizard, as shown in Figure 13-141.

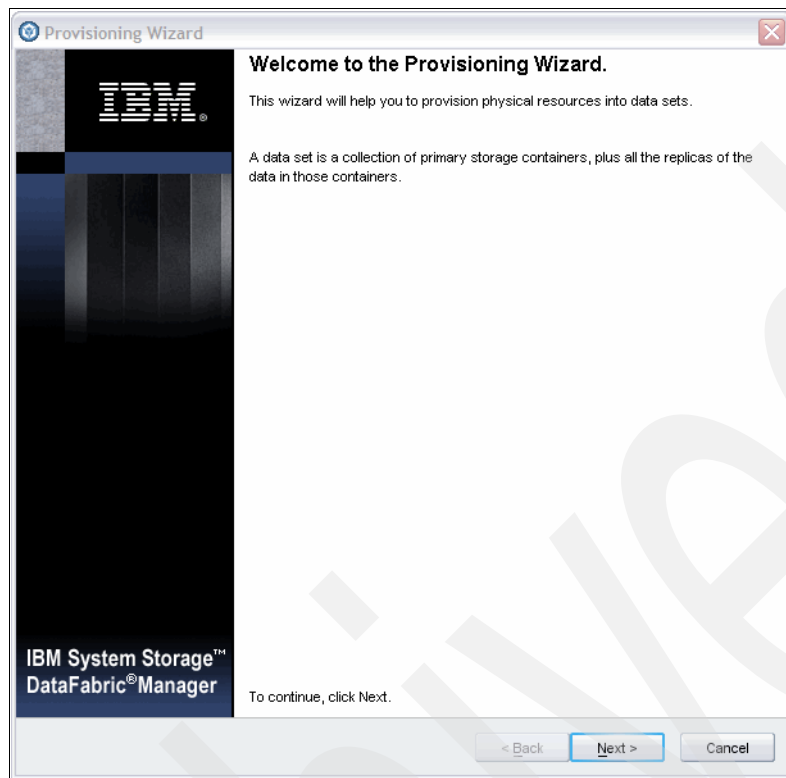


Figure 13-141 Data set Provisioning Wizard

We click **Next** to go to the window shown in Figure 13-142.

Provisioning Wizard
Container Name and Size
You should specify the name and size of the container to be provisioned

Volume name:
Description:
Data size: GB
Maximum snapshot size: GB

Space break out

Component	Size (GB)
Data size	3.0
Overwrite space	3.0
Space for Snapshot copies	1.0
Total size	7.0

Space usage:
Guaranteed space (Solid Green)
Space allocated on demand (Hatched Green)
Space components:
Overwrite space (Hatched Green)
Space for Snapshot copies (Hatched Green)

This information here is based on the resource pool previously specified

< Back Next > Cancel

Figure 13-142 Specifying pertinent information to create the physical resources

In this window, we can provide a name for the volume, a description, the (usable) data size, and the size of the maximum space you wish to allocate for Snapshots.

The wizard uses this information to calculate the overwrite space and space for Snapshots copies and tells you what the actual size you need to allocate is. The calculations are also influenced by the type of resource pool you specified earlier.

Click **Next** to go to the Preview Details window shown in Figure 13-143 on page 431. The Preview Details window shows the steps that will be executed and the expected outcomes for these steps.

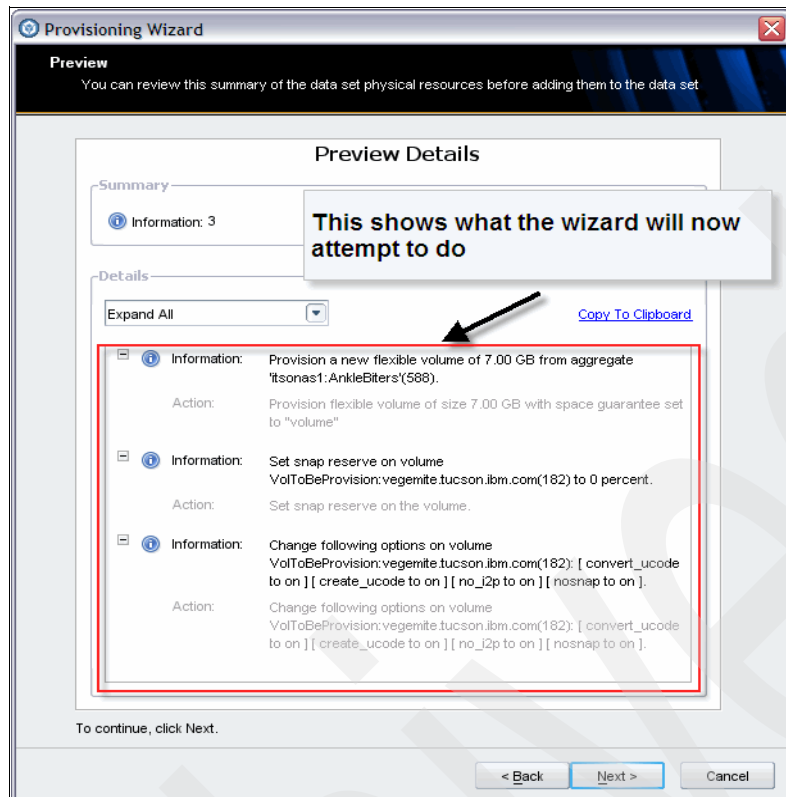


Figure 13-143 Summary of the steps the wizard will carry out to physically allocate storage

When you click **Next**, the wizard will create and execute the necessary jobs to achieve the outcomes specified here, as shown in Figure 13-144.

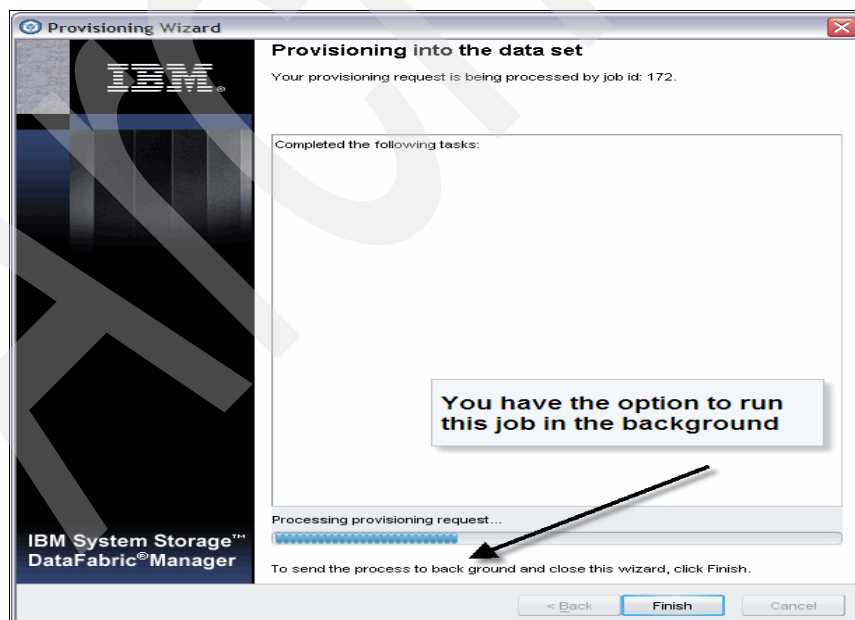


Figure 13-144 The wizard can run in the background while physically allocating storage

As these jobs may take an extended amount of time to run, you have the option to click the **Finish** button and let the jobs continue to run in the background.

You will be able to review the progress, status, and outcomes of the jobs in the Jobs pane in the System section of the N series Management Console.

As each task is completed, the results of that task will be displayed, as shown in Figure 13-145.

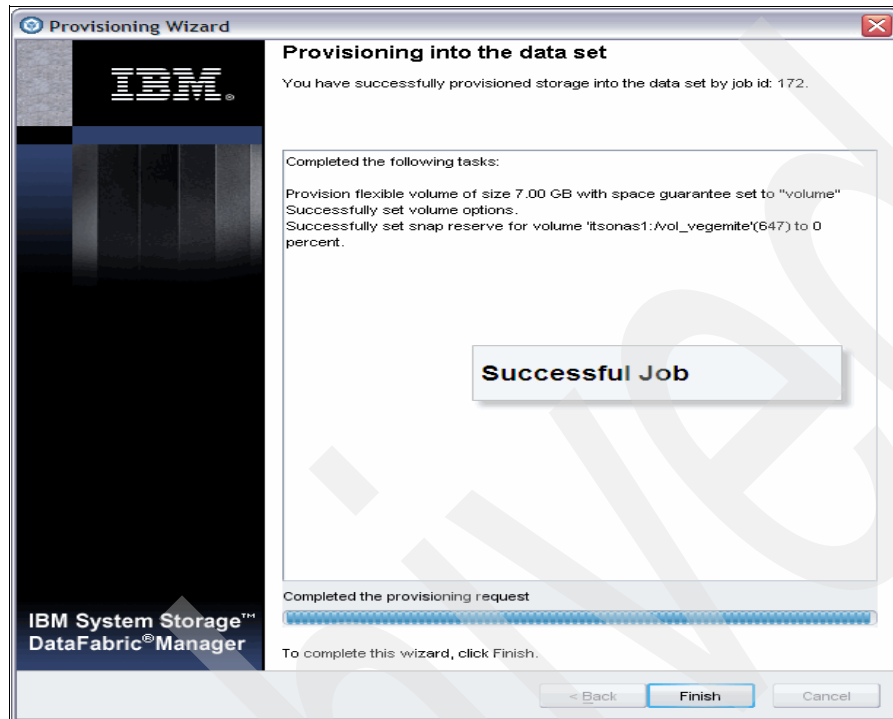


Figure 13-145 Job completion window

The details of the storage allocated to this host is shown in Figure 13-146 on page 433. As you can see, a volume on itsonas1 called /vol_vegemite.tucson.ibm.com has been created and space has been allocated according to the specified provisioning policy.

Figure 13-146 is an example of provisioning storage for SAN type data sets. In this case, we have a volume created where a host can allocate storage through SnapDrive.

Note: Coordinating and allocating LUNs through SnapDrive from Provisioning Manager only works for Windows host running SnapDrive Version 6 or later. This capability currently does not work for hosts running Linux, Solaris, or AIX. For Linux hosts, you need to run the SnapDrive command manually from within the host.

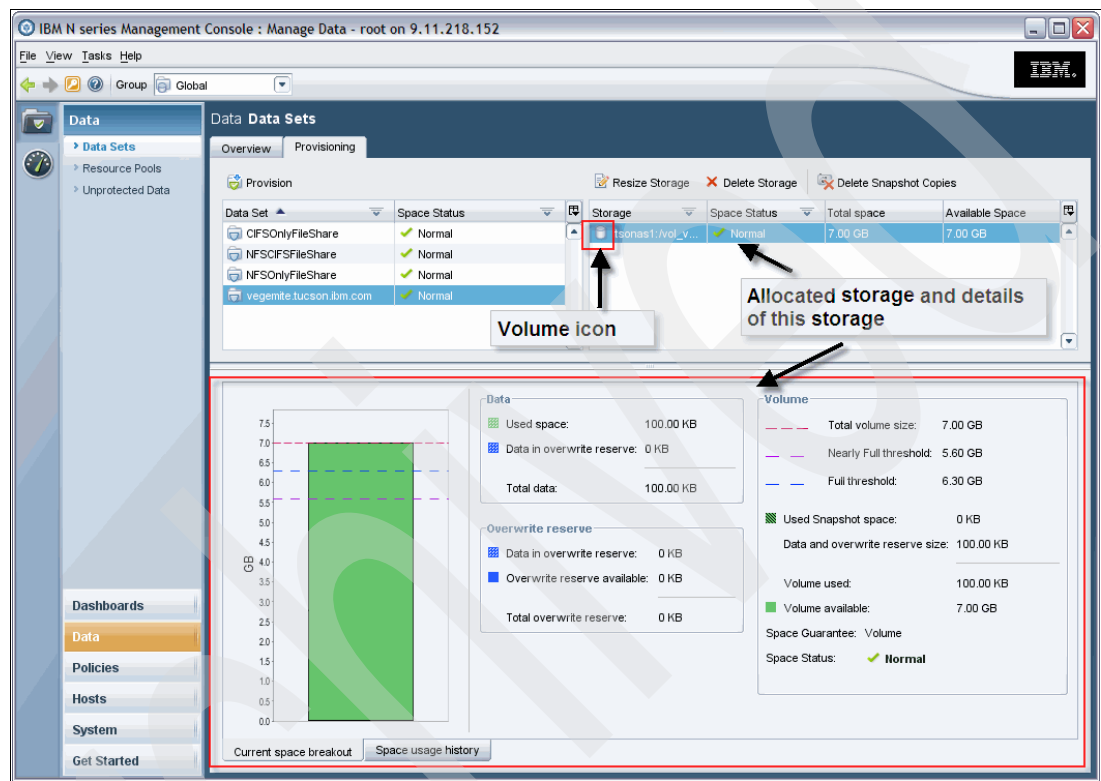


Figure 13-146 Example of provisioning for SAN type data sets

We employ a similar methodology to provision CIFS and NFS volumes, as shown in Figure 13-147.

Details about creating NFS or CIFS volumes are shown in Chapter 14, “Provisioning Manager setup” on page 437. It is important to note that while these volumes are created using a wizard, they still need to be protected in accordance with your business continuity requirements. You will also need to manually share or export these volumes, to enable hosts to connect to and mount these volumes.

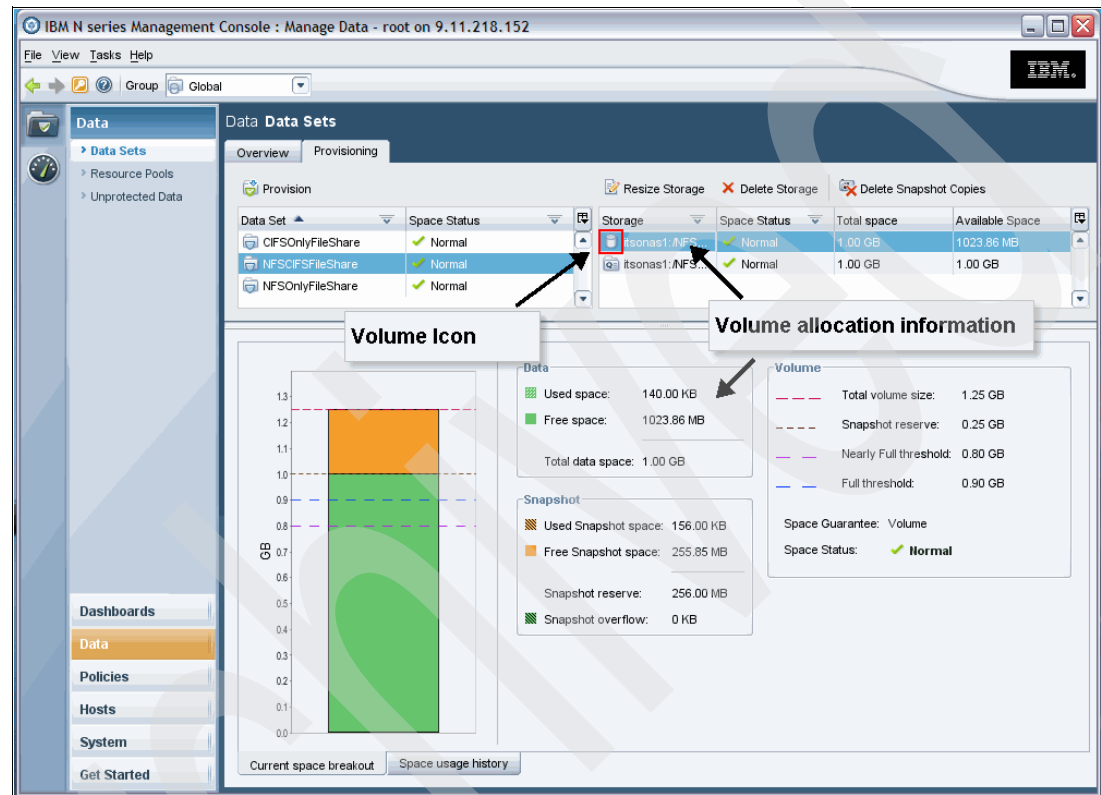


Figure 13-147 Example of provisioning for NAS (CIFS or NFS) type data sets - volume information

As you can see in Figure 13-147, the details pane shows all the pertinent information regarding the volume, such as capacity allocated, Snapshots details, and so on.

Figure 13-148 on page 435 shows an example of CIFS and NFS volumes being provisioned and their qtree allocation information. This information will become more useful over time as usage statistics continue to be collected by the DFM Server, as will the Space Usage History tab.

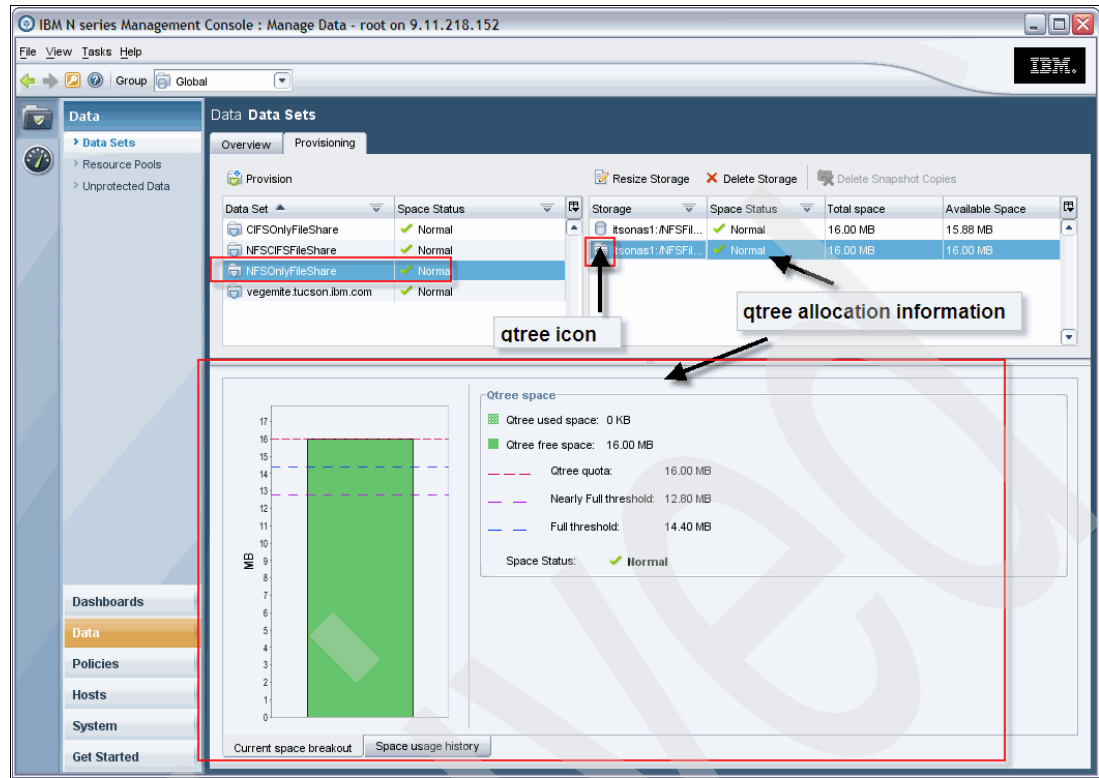


Figure 13-148 Example of provisioning for NAS (CIFS or NFS) type data sets - qtrees information

Protection Manager is a very sophisticated application and is an integral part of the data protection arsenal an organization should use. IBM System Storage N series technology has many data protection protocols to help facilitate proper protection of data to meet your business continuity and recovery requirements. Operations Manager and Protection Manager come together to simplify and automate repetitive data protection tasks, making it easier for IT support staff to manage and quickly assess the recover ability of their organization's critical data.

Archived



Provisioning Manager setup

This chapter will cover the preparation, setup, and installation of Provisioning Manager.

14.1 Host prerequisites

Provisioning Manager is a part of the N series Management Console and the Integrated component of Operations Manager. In order to enable this feature, you need to purchase the Provisioning manager license key and add it to Operations Manager.

The host requirements for Provisioning Manager is same as the host requirements for Operations Manager.

To view and configure the Provisioning Manager, you need to have the N series Management Console installed, which is installed on another server or on the same server where Operations Manager is installed

14.2 N series prerequisites


The N series prerequisites for Provisioning Manager is same as the N series prerequisites for Operations Manager.

14.3 Setup

In this chapter, we cover the Provisioning Manager license setup in Operations Manager, the Provisioning Manager information, and a few configurations.

14.3.1 Adding the Provisioning Manager license key for Operations Manager

In this section, we will add the Provisioning Manager license key for Operations Manager.



IBM System Storage™ N series

Events: [Emergency: 0](#) [Critical: 0](#) [Error: 26](#) [Warning: 18](#)

Control Center

Home
Setup
Reports
Management
Help

Group
Options
Discovery
Network Credentials
Groups
Administrative Users
Roles
Alarms
Database Backup
Download Management Console

Group Summary

Group Status
Member Details

Primary
File SRM
Streaming
Events

Global

Status
Error

Group Members

Hosts 7 Volumes 42
Qtrees 4 Lun Paths 25
SRM Paths 2 Aggregates 12
Disks 93

Events
Emergency 0 Warning 18
Critical 0 Information 38
Error 26

Monitored Devices
Storage Systems 4 Active/Active Controllers 2
Host Agents 3

Storage Capacity
Aggregates 12 3.33 TB of 6.01 TB (55.5% of total)
Volumes 42 726 GB of 3.44 TB (20.6% of total)
Qtrees 4 0 bytes of 80.0 GB (0.0% of limit)
LUNs 25 290 GB

Expand All
Collapse All

Current Group:

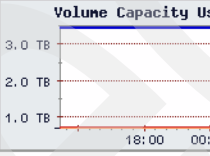
Group Status Reports
Member Details Reports

Edit Settings
Edit Membership
Edit Storage System Configuration
Edit vFiler Configuration
Edit NetCache Configuration
Edit NetCache Software
Edit NetCache Files

Run a Command

Report

Volume Capacity Used vs Total



Graph:
1d | 1w | 1m | 3m | 1y

Physical Space

Data Disks 36 (7.30 TB)
Parity Disks 18 (3.35 TB)
Spare Disks 19 (4.20 TB)
Total Disks 73 (14.8 TB)

Array LUNs

Data LUNs 10 (87.5 GB)
Spare LUNs 10 (87.5 GB)
Total LUNs 20 (175 GB)

Figure 14-2 shows the Operations Manager Options window. Click **Setup** → **Options** to view the current license or add the new license.

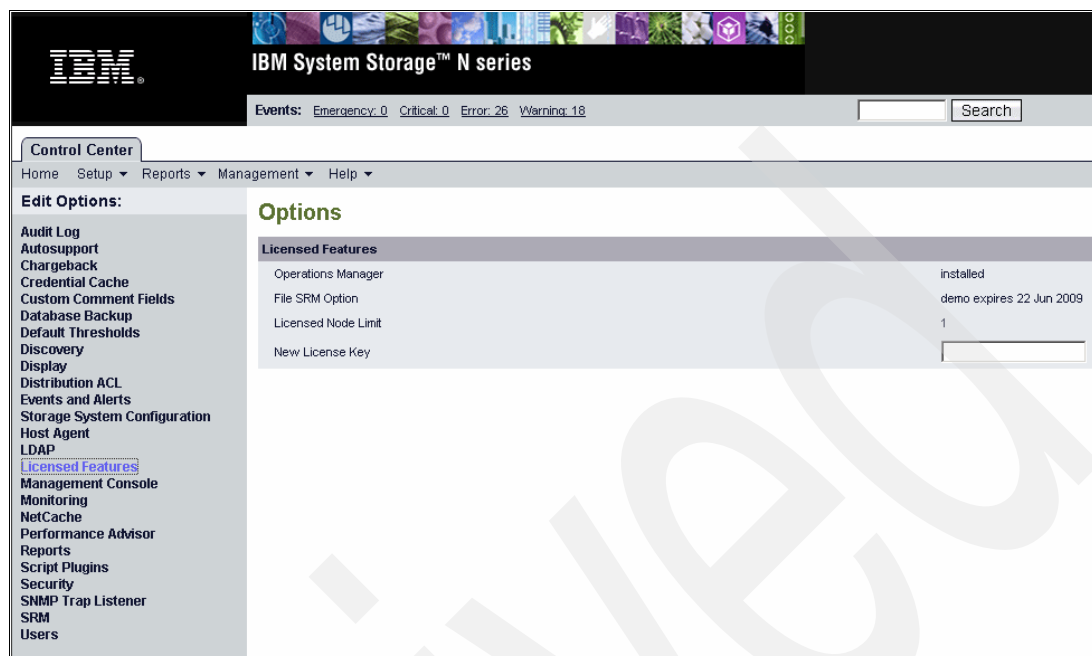


Figure 14-2 Operations Manager add license options

Figure 14-3 shows the Operations Manager Licensed Features option, which allows you to view the current license and add a new license. You can reach this window by selecting **Control Center** → **Setup** → **Options** → **Licensed Features**.



Figure 14-3 Operations Manager license add and update option window

The Provisioning Manager option provides:

- ▶ Policy-based provisioning of storage in data sets
- ▶ Conformance of storage to the provisioning policy attached to the data set
- ▶ Space management

If you have both Protection Manager and Provisioning Manager licensed, then the following features are enabled:

- ▶ Assigning provisioning policies to nonprimary nodes of a data set
- ▶ Policy-based provisioning of nonprimary nodes

Figure 14-4 shows the Operations Manager Licensed Features Add Option and Add New License window, where you can add the license for the Provisioning Manager. This is also where you add the Provisioning Manager's 14 digit license key. After you have entered the key, click **Update**.

Licensed Features	
Operations Manager	installed
File SRM Option	demo expires 22 Jun 2009
Licensed Node Limit	1
New License Key	<input type="text" value="HBMERVWZSRAGGD"/>

Back Update

Figure 14-4 Operations Manager Provisioning Manager license add and update option

Figure 14-5 shows the Operations Manager License Add Option and Add New License window, where the Provisioning Manager license is updated and displayed.

Licensed Features	
Operations Manager	installed
File SRM Option	demo expires 22 Jun 2009
Provisioning Manager	demo expires 01 Jun 2009
Licensed Node Limit	1
New License Key	<input type="text" value="HBMERVWZSRAGGD"/>

Back Update

Figure 14-5 Operations Manager Provisioning Manager license added and updated

After the Provisioning Manager license is updated, you have to restart your Operations Manager. Open the Provisioning Manager setup by opening the N series Management Console.

14.4 Using the N series Management Console to view Provisioning Manager

In this section, we show how to access the Provisioning Manager through the N series Management Console.

To open the N series Management Console, select **Start → Programs → IBM → N series Management Console → N series Management Console**, as shown in Figure 14-6.

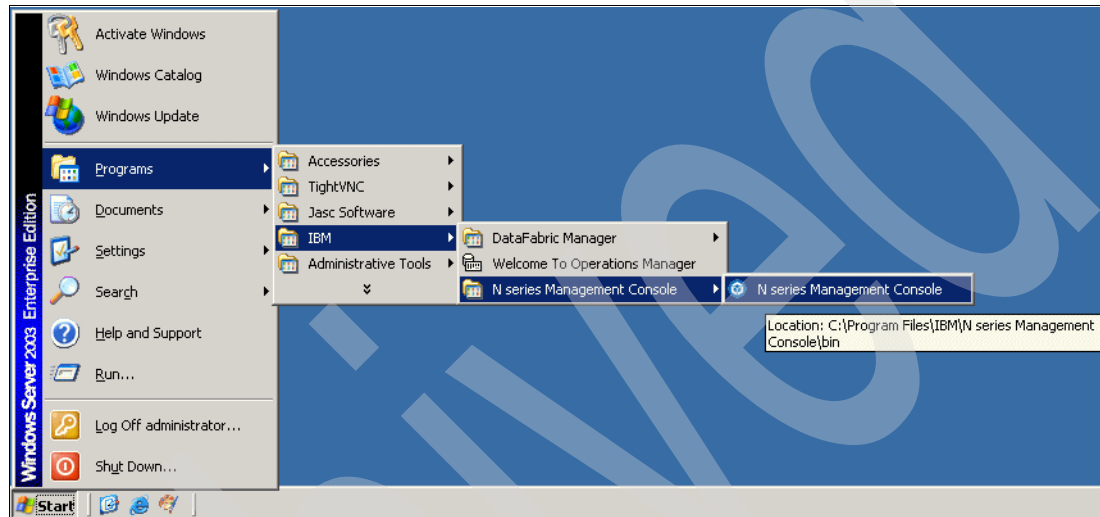


Figure 14-6 Accessing the N series Management Console

Provisioning Manager simplifies and automates the tasks of provisioning and managing storage for NAS and SAN access, and it improves efficiency in storage utilization.

Provisioning Manager is a licensed application for use with Operations Manager. The provisioning application can be accessed through the N series Management Console, which is the client platform for IBM management software applications. The provisioning application provides the following capabilities:

- ▶ User-defined policies to automate storage provisioning and configure the default settings for exporting storage
- ▶ Periodic conformance checking to ensure the provisioned storage conforms to the provisioning policy
- ▶ Manual controls for resizing space and capacity of existing storage
- ▶ Manual controls for provisioning new and existing storage

Anyone using Provisioning Manager should be familiar with general storage provisioning concepts.

Note: The procedure remains same for Windows 2003 32-/64-bit and Windows XP.

14.5 Dashboards provisioning information

In this section, we show the Provisioning Manager Dashboard options and some configuration tips.

Provisioning Manager and Protection Manager are integrated components of Operations Manager. Some of the options on the Dashboard are connected to Protection Manager, so here we demonstrate only the Dashboard options for Provisioning Manager; the options for Protection Manager and Performance Advisor and discussed in Chapter 13, “Protecting your data with Protection Manager” on page 313 and Chapter 10, “Performance Advisor operation and configuration” on page 199.

Figure 14-7 shows the N series Management Console Provisioning Manager Dashboard before we make any configurations. The Provisioning options are indicated with arrows.

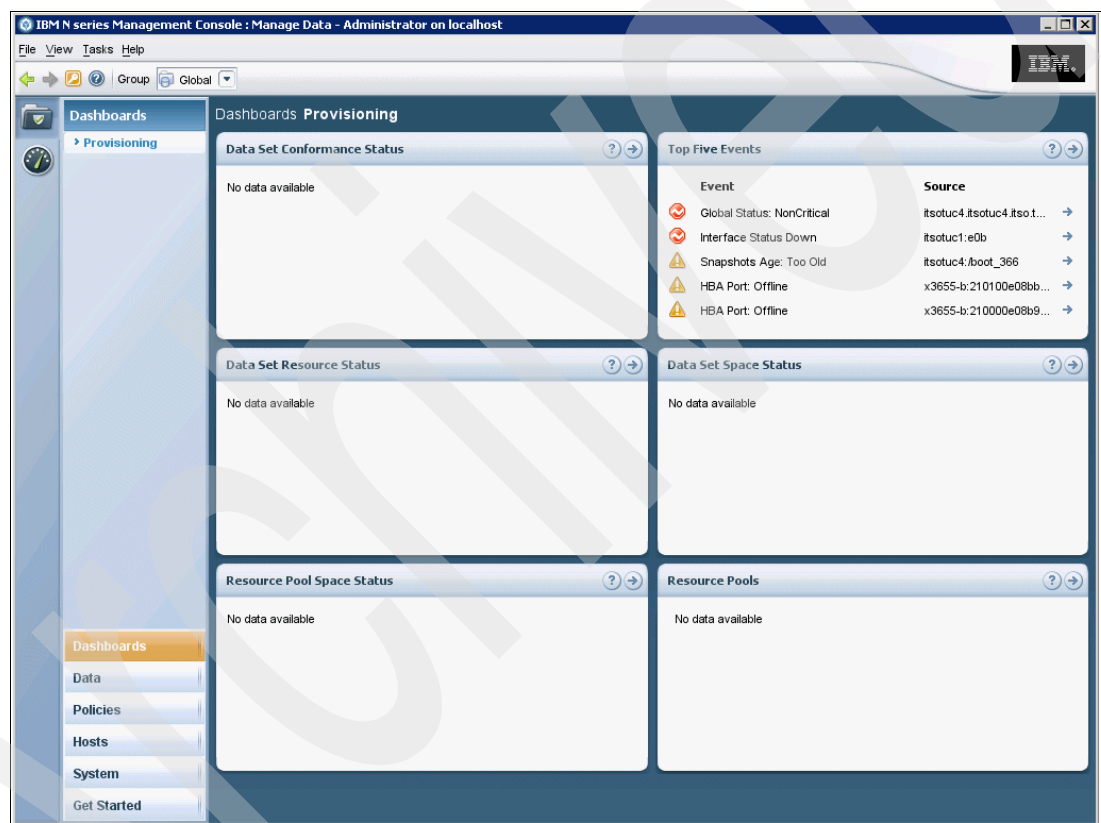


Figure 14-7 N series Management Console Provisioning Manager Dashboard window

To provision the Dashboards, you must first open them by doing these steps:

1. Click the folder icon in the upper left corner of the window to open the Provisioning Manager Dashboard window.
2. Click **Dashboards** to open the provisioning option.
3. Click **Provisioning** to open the Top Five Events pane, for example.

The Provisioning Dashboards window provides cumulative at-a-glance status information for space management issues related to data sets and resource pools. Some of the Provisioning Manager window options are:

- ▶ **Top Five Events:** Shows the five events with the highest severity levels, ordered by time.
- ▶ **Data Set Conformance Status:** Displays the total number of data set members that are conforming to associated policies. Values can be Conformant or Nonconformant.
- ▶ **Data Set Resource Status:** Displays the number of data sets at different levels of resource status severity. The status represents the worst event severity of all current events on all direct and indirect members of the data set nodes. Values can be Emergency, Critical, Error, Warning, or Normal.
- ▶ **Data Set Space Status:** Displays the total number of data sets being managed by N series Management Console, grouped according to their current space status value. The status represents the worst space status of all members in all nodes of the data set. Events are generated at the data set level when the space status of a data set changes. Values can be Error, Warning, Normal, or Unknown.
- ▶ **Resource Pool Space Status:** Displays the total number of resource pools that currently meet or exceed the space thresholds. Values can be Full, Nearly Full, Overcommitted, and Nearly Overcommitted.
- ▶ **Resource Pools:** Displays, by resource pool name, the total space allocated to and the space utilization for each resource pool. Total Size values are indicated in numbers, such as gigabytes or terabytes. Utilization is indicated in percentage of the Total Size. Items are sorted in increasing order of available space.

For more details about any of these options, you can click the arrow button on any of the panes. Here we click the arrow button of the Top Five Events pane, as shown in Figure 14-8.

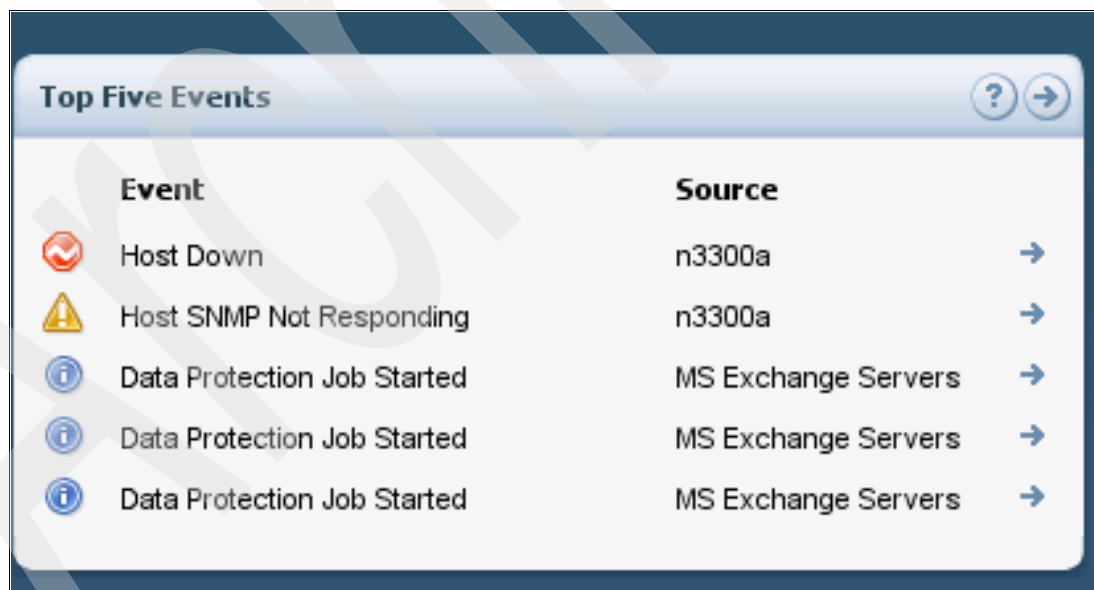


Figure 14-8 Top Five Events pane

The resulting window is shown in Figure 14-9 on page 445.




Figure 14-9 Top Five Events details window

14.5.1 Data sets information

A data set is a collection of physical resources, such as entire storage systems, flexible volumes, qtrees, and copies of backed up data that are used to group data and resource pools to organize storage for simplifying monitoring, provisioning, reporting, and access control of SnapVault and SnapMirror relationships. You can associate a data protection policy with a data set to automate tasks such as:

- ▶ Applying consistent policies to primary data
- ▶ Propagating policy changes
- ▶ Provisioning new volumes

Figure 14-10 shows the images and options for a data set configuration. The Data Sets window's Overview tab provides a single location to monitor the status of all data sets, create and modify data sets, assign protection policies to unprotected data, and assign physical resources manually or by using provisioning policies.

For more information about data set configuration, refer to “Data sets configuration” on page 448 and Chapter 12, “Protection Manager setup” on page 307.

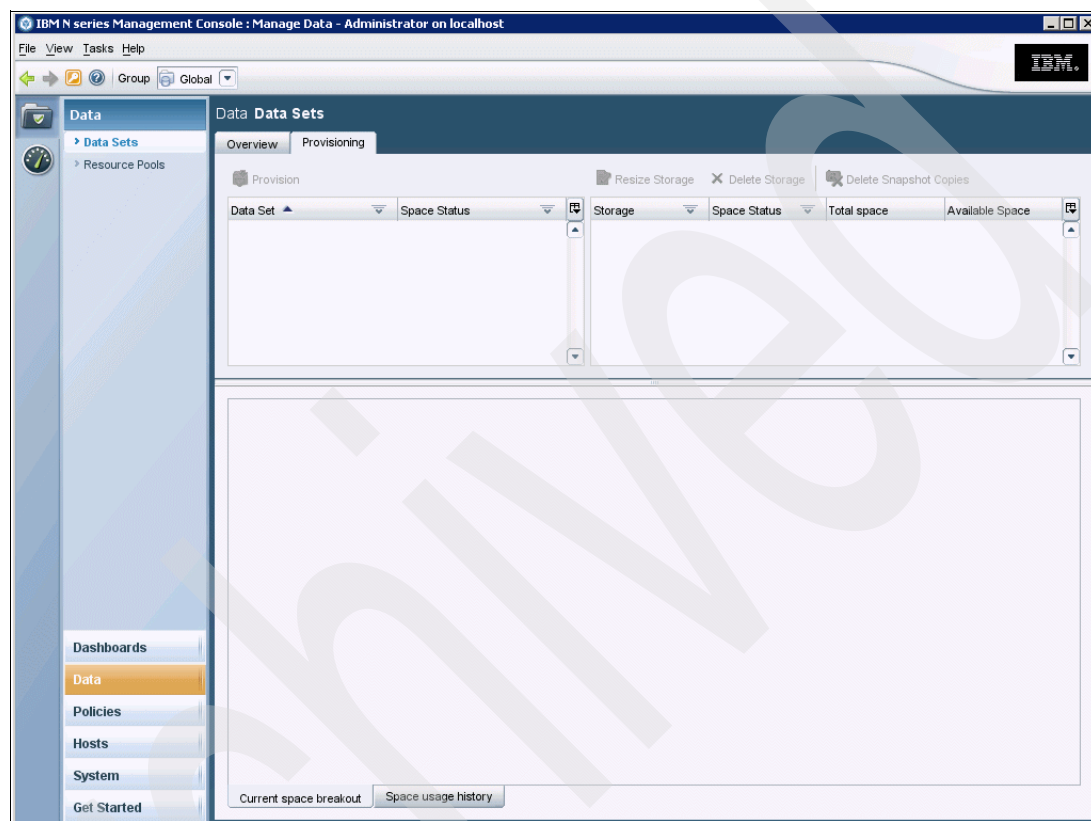


Figure 14-10 Provisioning Manager data sets configuration

To view an overview of the data sets, select **Data**, then **Data Sets**, and then the **Overview** tab.

In the simplest terms, a data set is a collection of user data you manage as a single unit, plus all the replicas of that data. The data is identified by the volume, qtree, or directory in which it is located. Because you manage a data set as a single unit, its members should have common management requirements. In Provisioning Manager, each node in a data set might not share the same provisioning requirements, but all members of each node should share the same provisioning requirements. For example, different types of data supporting the same application would probably share the protection requirements of the application. You would want to collect that data in the same data set, even if the data were stored in different volumes or qtrees.

For provisioning, if a data set had a protection policy that created primary, backup, and mirror nodes, your provisioning requirements for each node might be different. You might want to provision the primary node on high-availability storage but provision the mirror node on less expensive, low-availability storage. Figure 14-11 shows the steps to access the data set provisioning window; select **Data**, then **Data Sets**, and then the Provisioning tab.




Figure 14-11 N series Management Console Provisioning Manager data sets provisioning

Figure 14-12 shows the Space Usage tab. You can access this tab by selecting **Data**, then **Data Set**, then **Provisioning**, and then the **Space breakout** tab. This tab graphically displays the history of space allocation and use on the selected volume or qtree over the selected period of time. This applies only to qtrees that have quotas settings configured. To display space allocation data for a specific date and time, place your cursor at that date and time along the graph's time axis. The information for the cursor-selected date and time is displayed to the right of the graph.

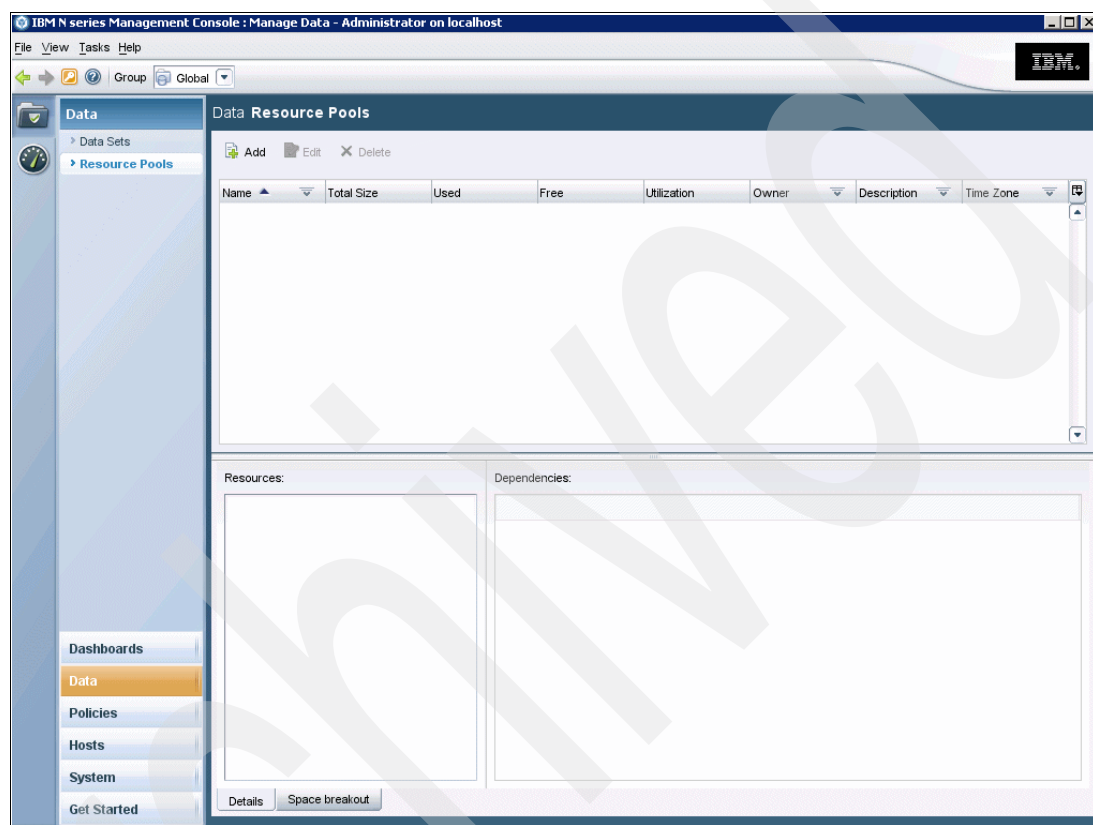


Figure 14-12 Provisioning Manager data sets provisioning Space breakout tab

Data sets configuration

In this section, we cover the step by step configuration of data sets. Data set configuration involves the following steps:

1. Provisioning policy: Associate the provisioning policy.
2. Export settings: Associate the export settings for FCP/iSCSI. This is a list of host initiator IDs to which the LUNs in the data set are mapped after provisioning.
3. Resource pools: Associate resource pools.
4. vFiler units: Optionally, associate a vFiler unit if provisioned storage is to be exported through a vFiler unit.

Do these steps:

1. We start the data set addition process by clicking **Data** → **Datasets**, click the **Overview** tab, as shown in Figure 14-13 on page 449, and click the **Add** button, as shown in Figure 14-14 on page 449. Click **Next** to proceed.

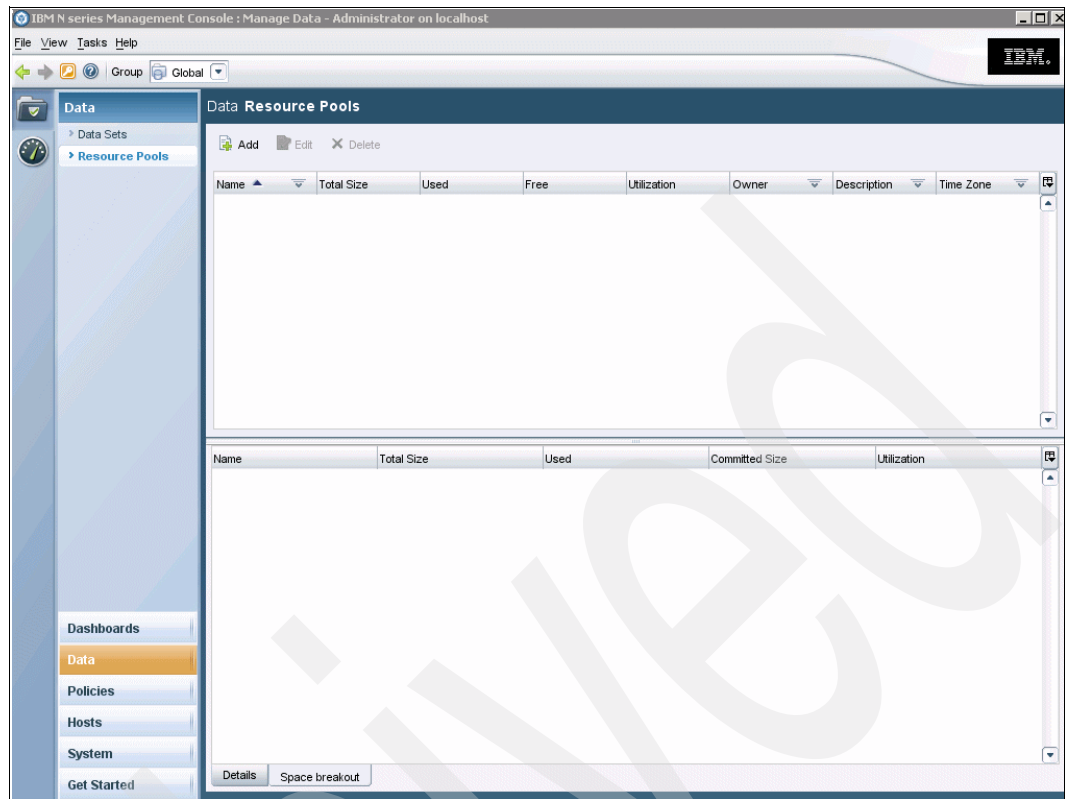


Figure 14-13 Provisioning Manager data sets overview

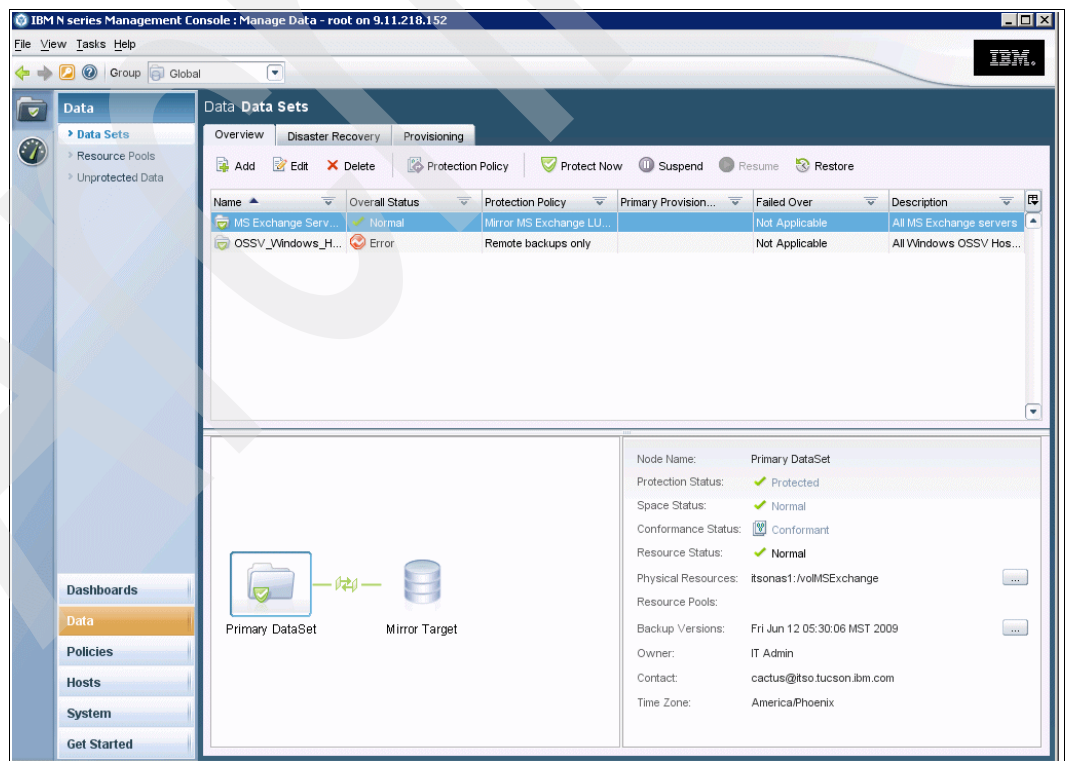


Figure 14-14 Provisioning Manager Add New Data Set window

2. Figure 14-15 shows the Add Data Set Wizard welcome window. Click **Next** to continue.

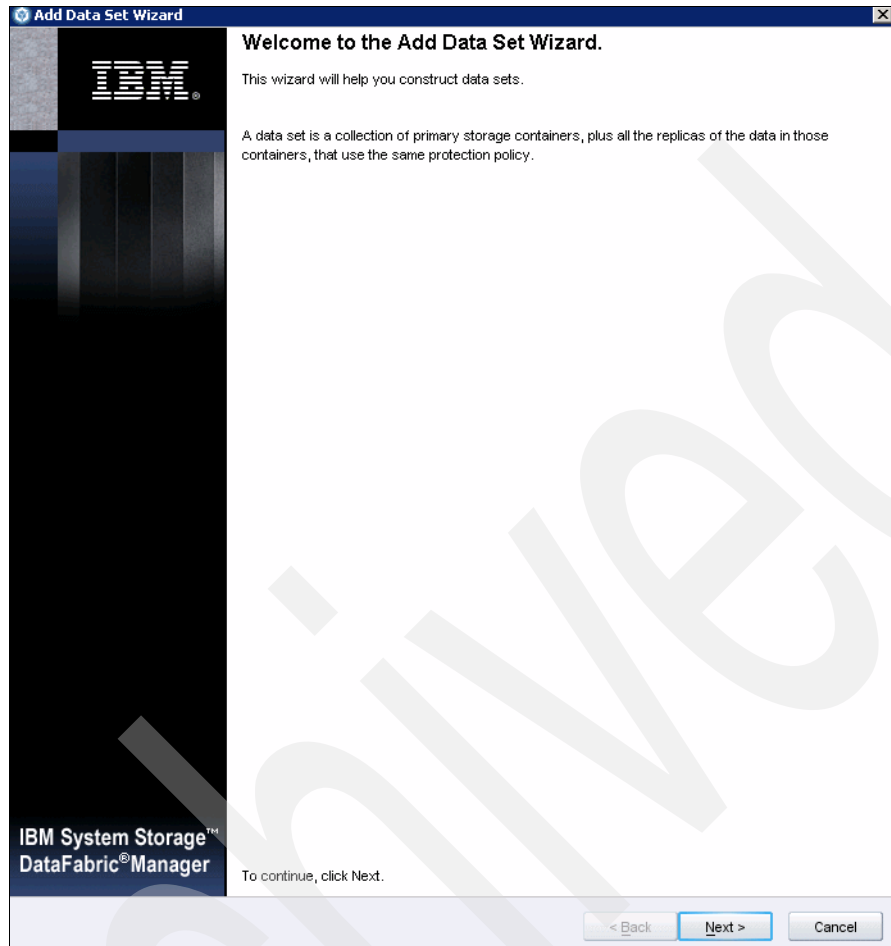


Figure 14-15 Add Data Set Wizard welcome window

3. In the General Properties window, enter the name, description owner of the data set, contact details, and the time zone. Click **Next** to proceed. Click **Cancel** to stop the configuration, if necessary. We have provided the General Properties information for our environment, but you have the option to use the information for your environment in your own configuration.

The screenshot shows the IBM N series Management console with the 'Add Data Set Wizard' window open. The 'General Properties' tab is selected, and the following information is entered:

Field	Value
Name	Test Data1
Description	Test Data1 for ITSO Datacenter1
Owner	ibmer
Contact	ibmer@ibm.com
Time Zone	US/Indiana-Starke

The 'Time Zone' dropdown menu is open, showing options: US/Indiana-Starke, US/Michigan, and US/Mountain. The 'Next >' button is highlighted in blue. The console sidebar on the left shows 'Data' > 'Data Sets' selected.

Figure 14-16 Provisioning Manager add new data - data set General Properties window

4. In the Provisioning window, select the provisioning policy, as shown in Figure 14-17.

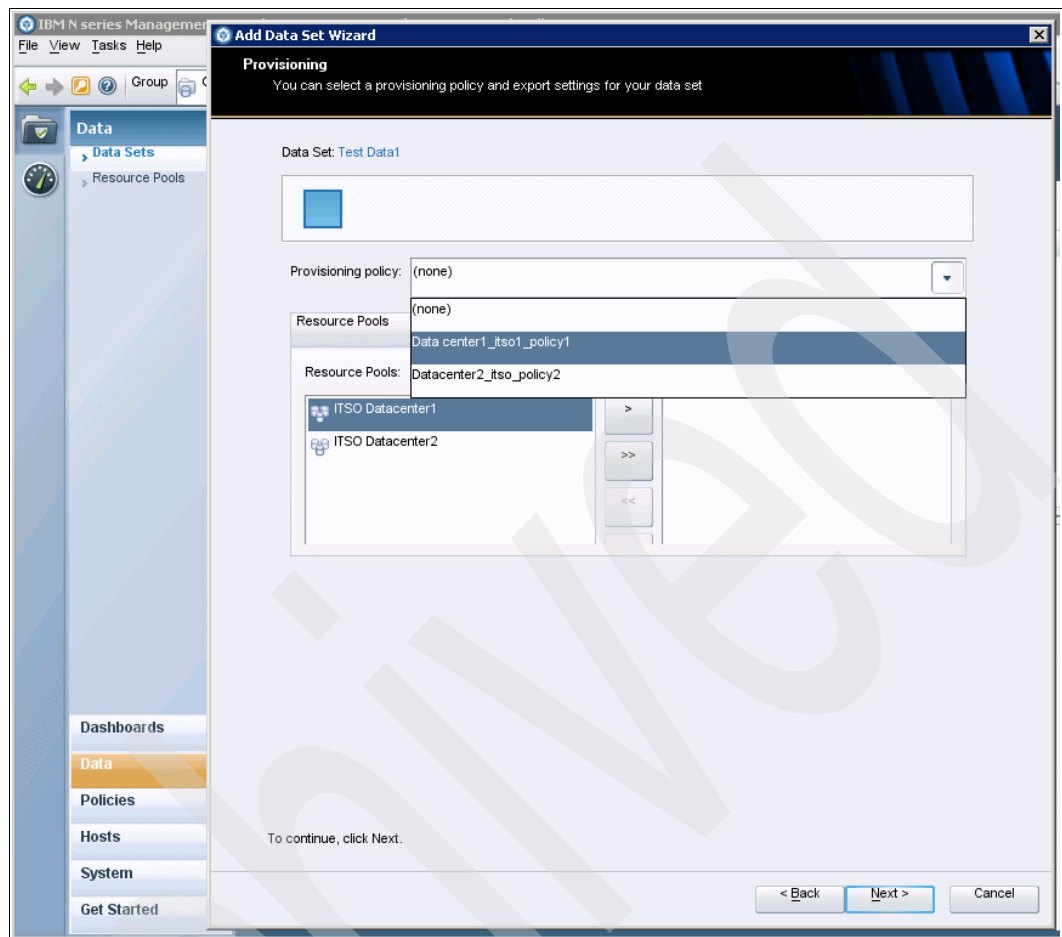


Figure 14-17 Provisioning Manager add policy new data set window

5. Change the Provisioning policy for NFS and CIFS to **Off**, as shown in Figure 14-18.

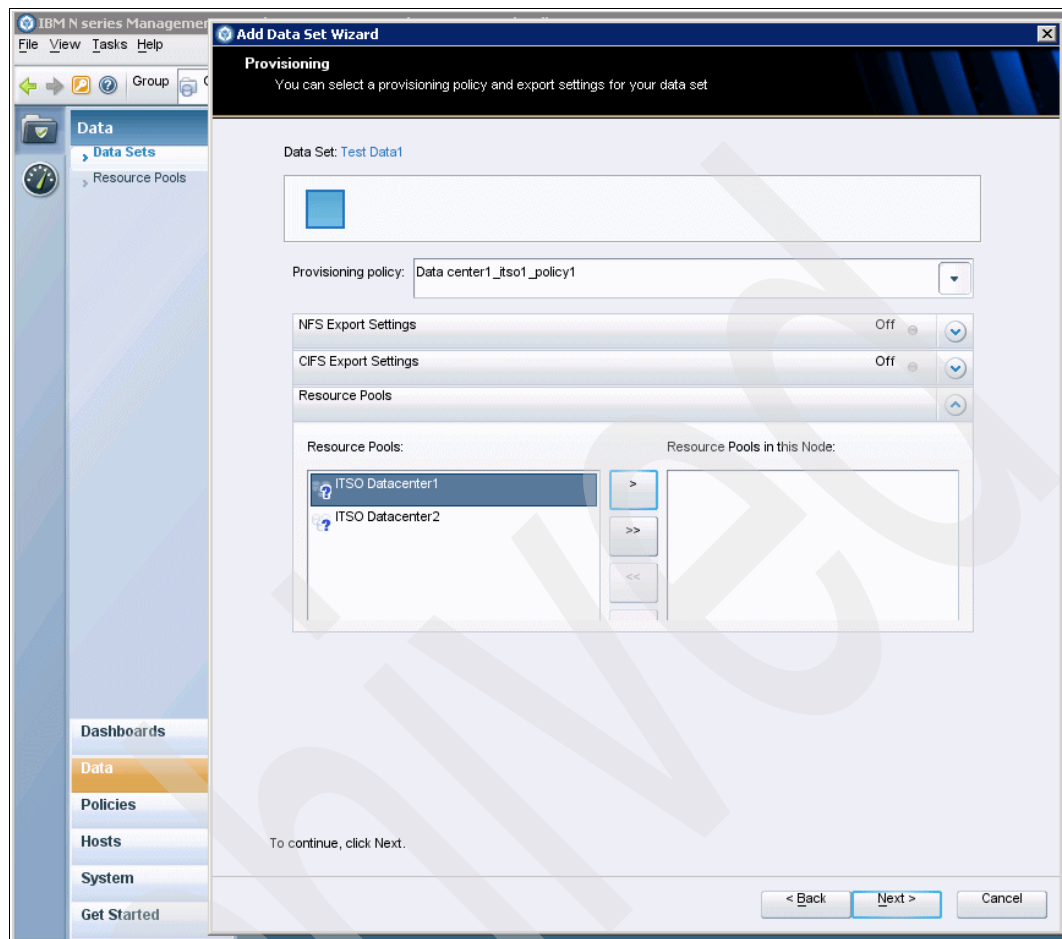


Figure 14-18 Provisioning Manager add new data set for NFS Export Settings, CIFS Export Settings, and Resource Pools

6. Change the Provisioning policy for NFS Export to On and set the permission for the host or add the individual host to set the permission, as shown in Figure 14-19.

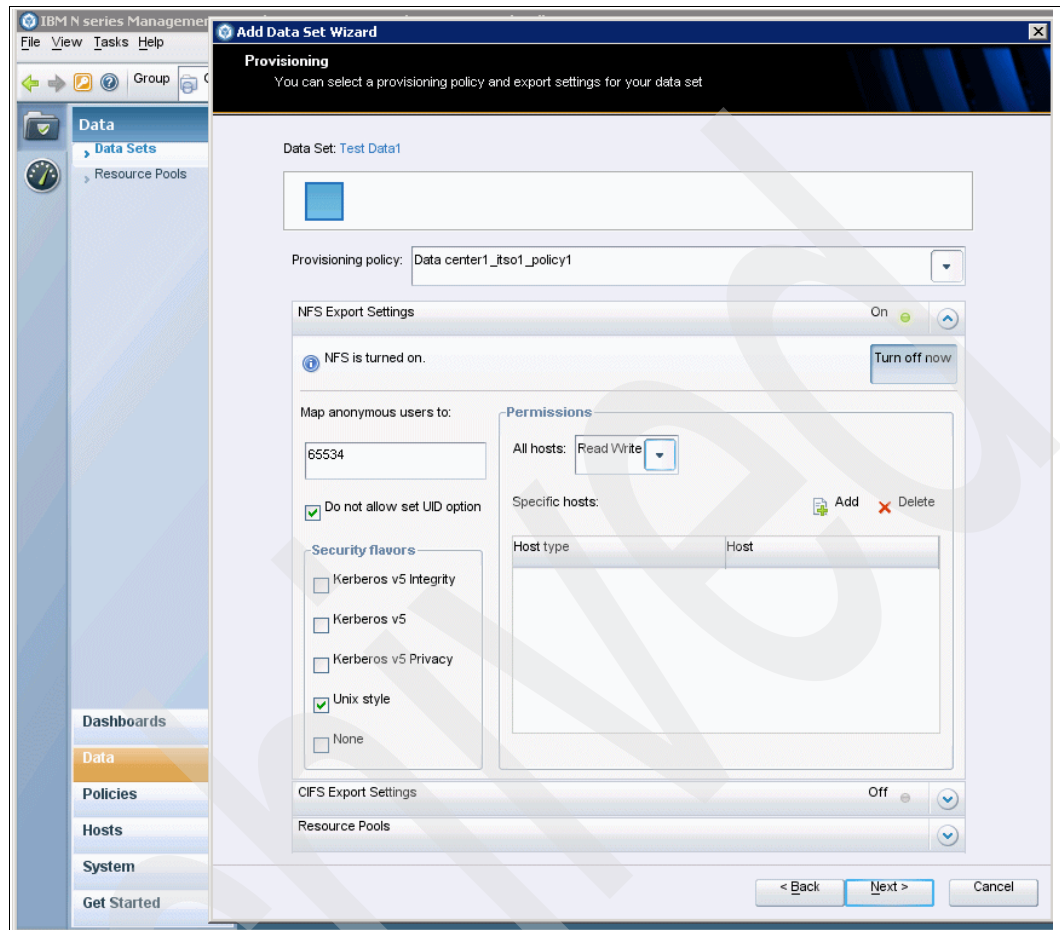


Figure 14-19 Provisioning Manager add new data set provisioning with provisioning policy1 NFS Export setting

The settings in this pane are:

- NFS: Provides file system access for UNIX and Linux environments.
- Permission:
 - All Host
 - Read Only: The user cannot make changes.
 - Read Write: The user can make changes.
 - Specific Host
 - Add: Adds a host from the list and assigns read only or read write permission. After you add the host, it will display under Host type and Host.
 - Delete: Deletes the host from the Host type and Host list after changing the settings.
- Security Flavors: We selected the **Unix style** security level.

After setting the NFS, go to CIFS export setting by clicking **Next**.

7. Change the Provisioning policy for CIFS Export to **On**. In the Windows Domain pane, add the user and permission, as shown in Figure 14-20.

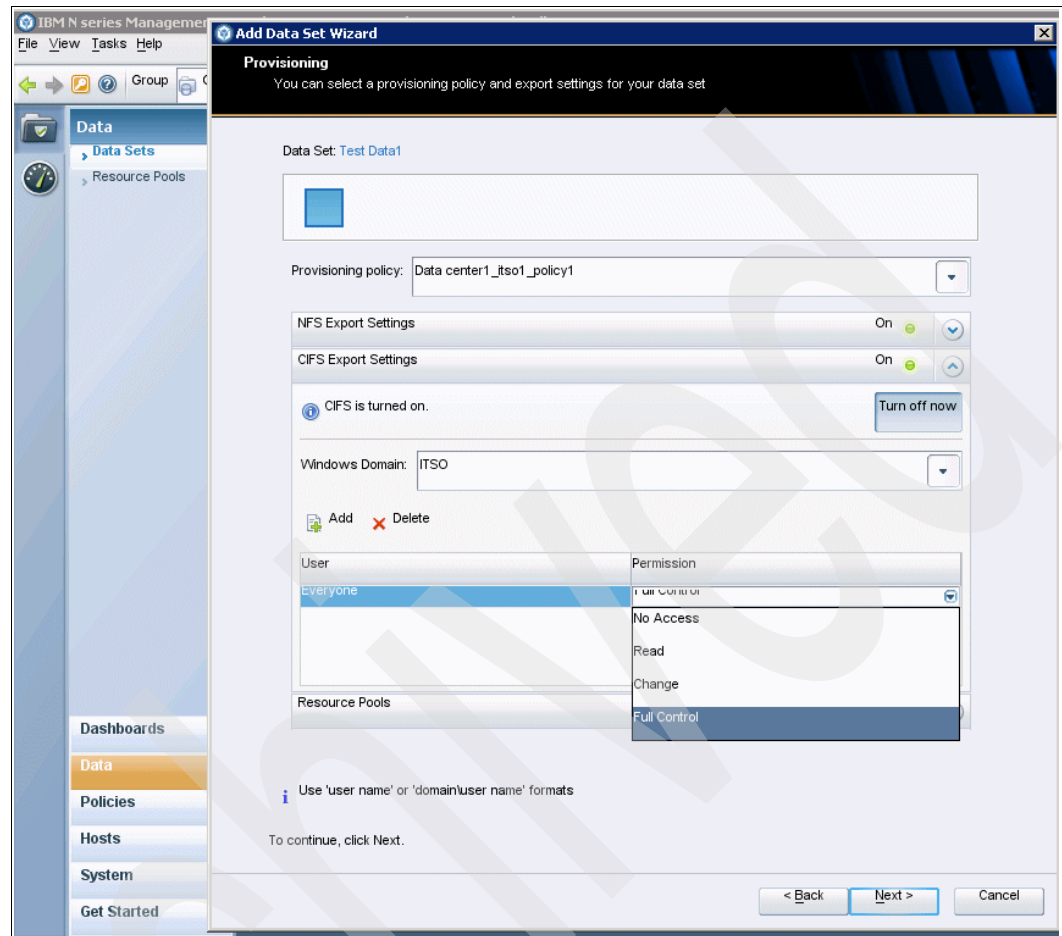


Figure 14-20 Provisioning Manager add new data set provisioning with CIFS export settings

This pane has the following settings:

- CIFS: Provides file system access for Microsoft Windows environments.
- Windows Domain: You can select the domain from the list.

You can add the new user and assign the permissions as No Access, Read, Change, or Full Control.

If you want to delete the user assigned with permission from the User list in Figure 14-20, delete it and, after setting the CIFS, go to Resource Pools.

8. Select the Resource ITSO Data center1 in the Resource Pools pane and move it to the Resource pools in this Node field. Click **Next**.

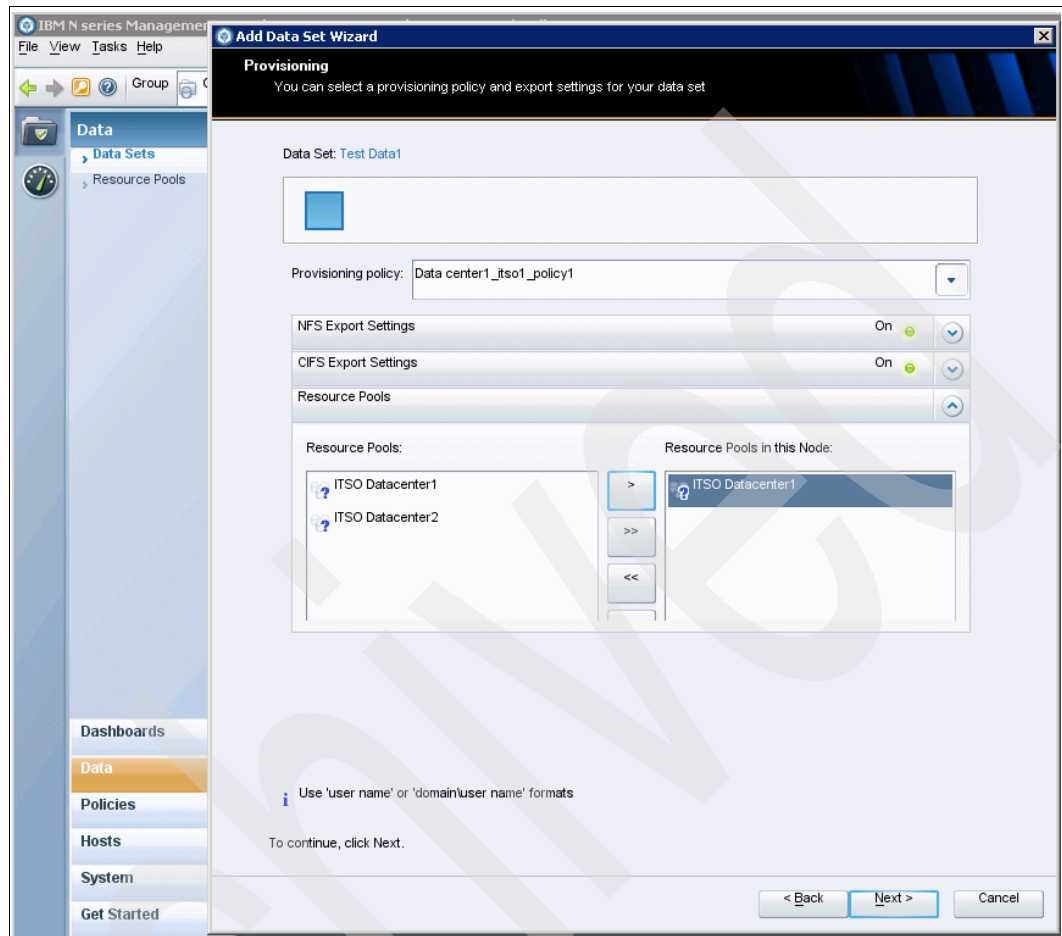


Figure 14-21 Provisioning Manager add new data set provisioning from the resource pool

9. Figure 14-22 shows the vFiler unit properties window. We have not specified vFiler, as it is optional in this configuration. More information about vFiler configuration can be found in 14.5.6, “Host vFiler Units information” on page 536. Click **Next**.

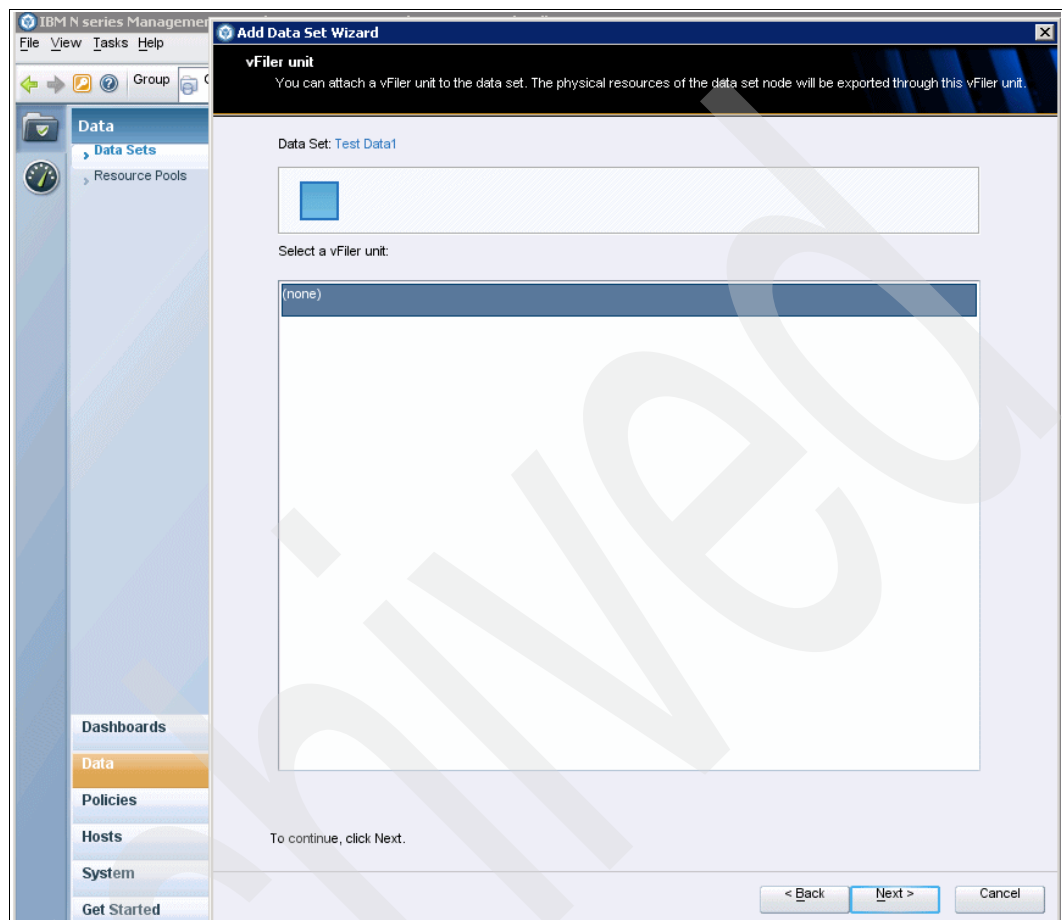


Figure 14-22 Provisioning Manager add new data set provisioning through vFiler

10. Select **No, I will provision storage later**, as shown in Figure 14-23. For more information about Provision Storage, refer to Chapter 13, “Protecting your data with Protection Manager” on page 313.

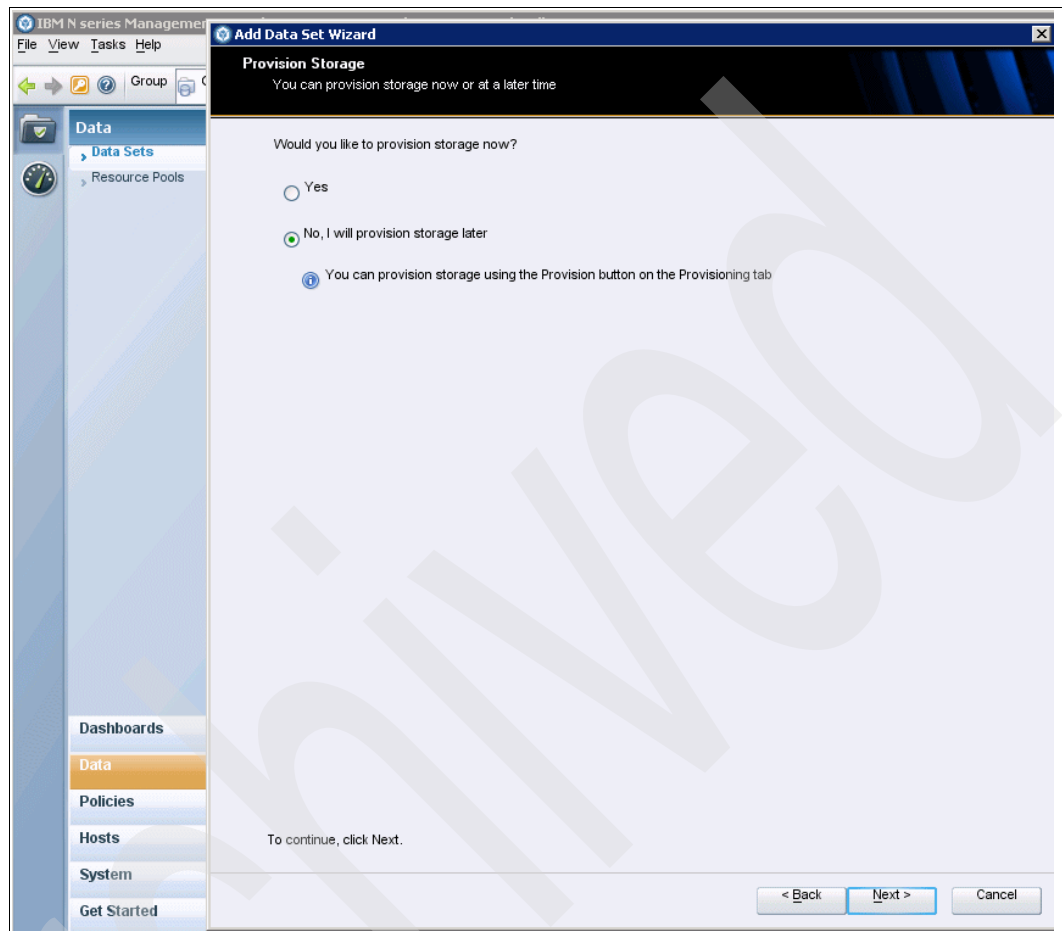


Figure 14-23 Provisioning Manager add new data sets provision storage

11. The Preview window shows a preview of the status, as shown in Figure 14-24. Click **Next**.

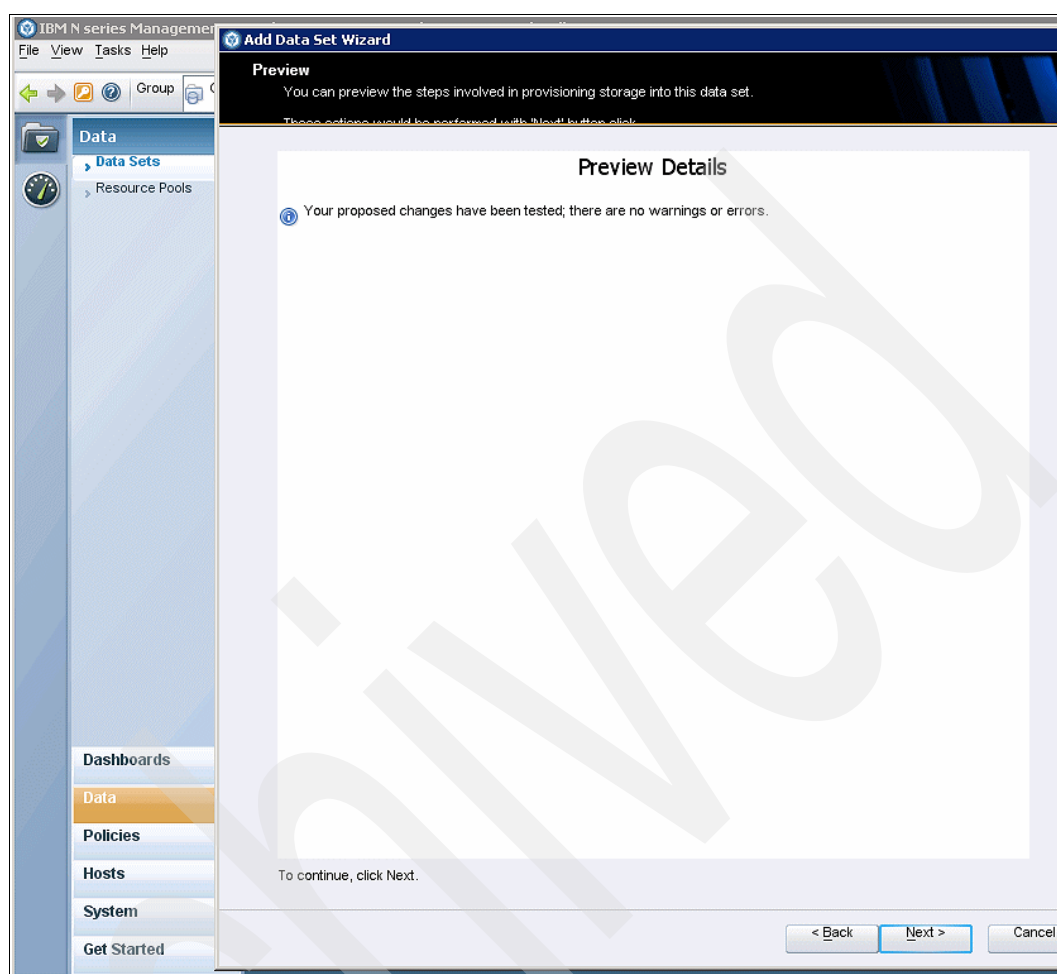


Figure 14-24 Provisioning Manager add new data set preview of the configuration

12. The Completing the Add Data Set Wizard window appears, which displays a summary of the tasks, as shown in Figure 14-25. Click **Finish**.

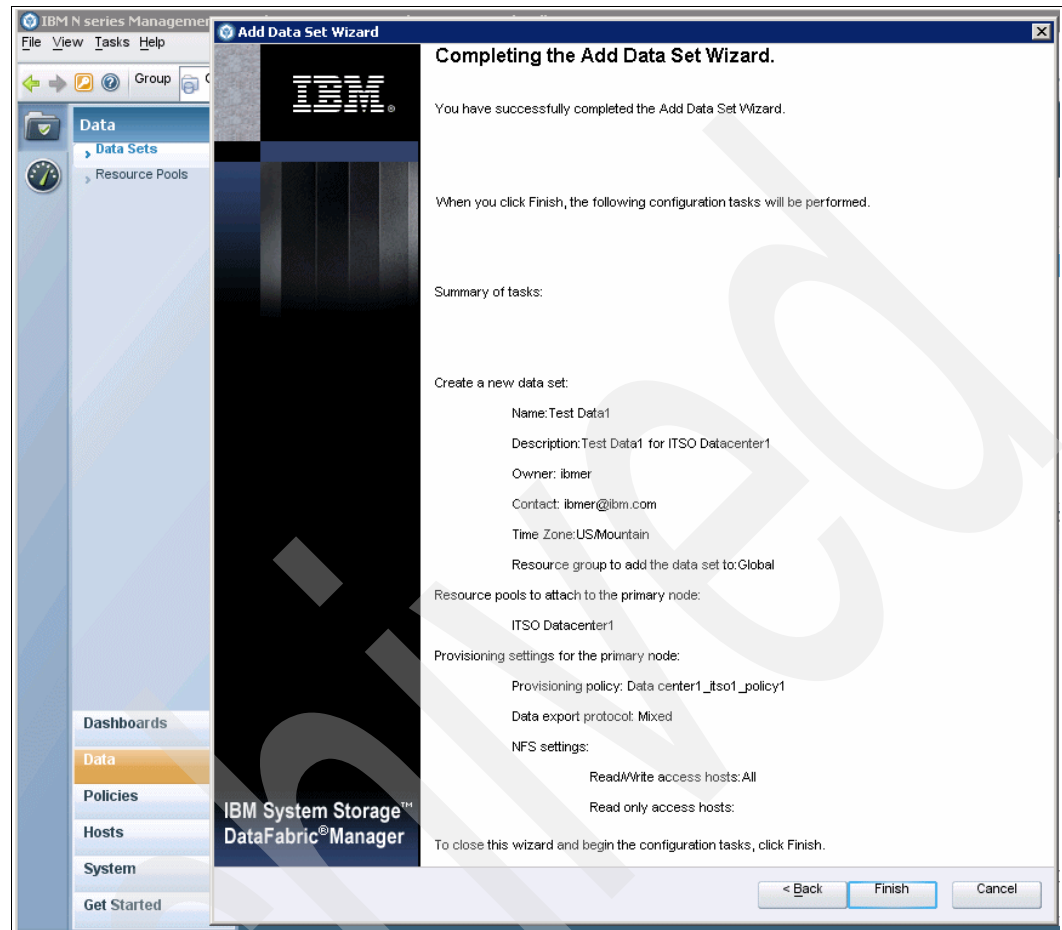


Figure 14-25 Completing the new data set primary node storage configuration

13. Select **Data** → **Datasets**, click the **Overview** tab, and click **View**, which should show that Test Data1 was configured successfully, as shown in Figure 14-26. To create the second data set, follow the same procedure used for creating the first data set, with some modifications.

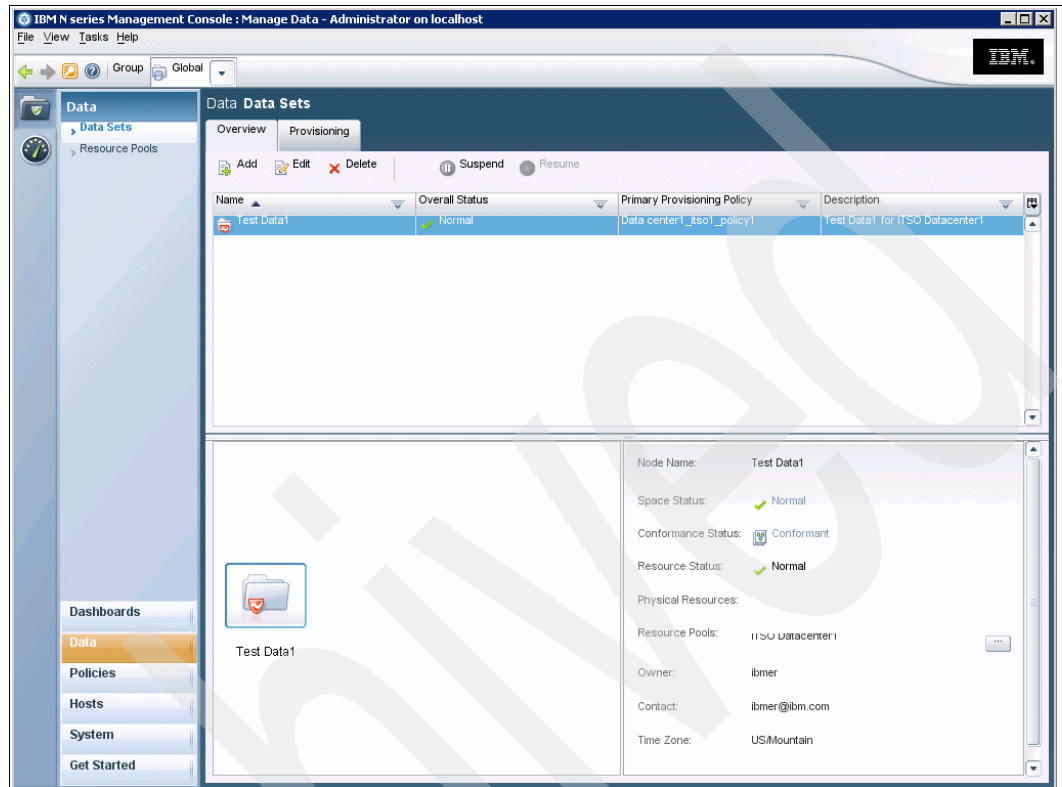


Figure 14-26 Provisioning Manager add new data set NAS Storage view

The Data members list displays information about data sets and the volumes, qtrees, or LUNs in each data set. You can select a data set to see detailed information about that data set. The status of the available space for the data set can be Normal, Warning, Error, or Unknown.

14. Select **Data** → **Data Sets**, click the **Overview** tab, click the **Add** button, and select **General properties** to create the second data set. Follow the same procedure used for creating the first data set, with some modifications.

IBM N series Management

File View Tasks Help

Group

Data

Data Sets

Resource Pools

Dashboards

Data

Policies

Hosts

System

Get Started

Add Data Set Wizard

General Properties

You must provide a name for the new data set. Other properties are optional.

Name: Test Data2

Description: Test Data2 for Datacenter2

Owner: ibmer

Contact: ibmer@ibm.com

us

Time Zone: US/Indiana-Starke
US/Michigan
US/Mountain

To continue, click Next.

< Back Next > Cancel

Figure 14-27 Provisioning Manager add new data set General Properties

15. In the provision policy, change the FCP Export setting to On. Click **Next**.

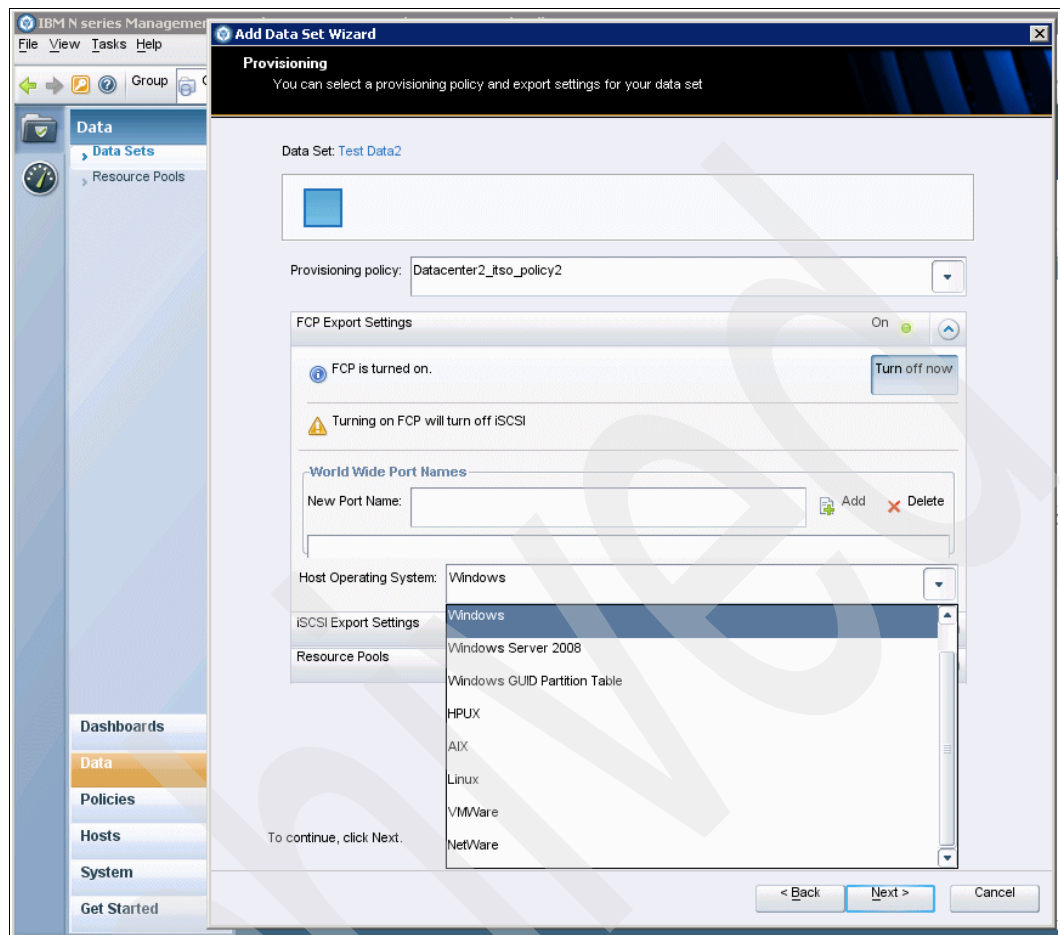


Figure 14-28 Provisioning Manager add new data set FCP Export settings

Fibre Channel Protocol (FCP) is the specification of how the semantics of the SCSI architecture can be communicated across a Fibre Channel network.

Attention: Turning FCP on will turn off the iSCSI, that is, you cannot have both of them on at the same time.

16. In the provision policy, change the iSCSI Export setting to On. Click **Next**.

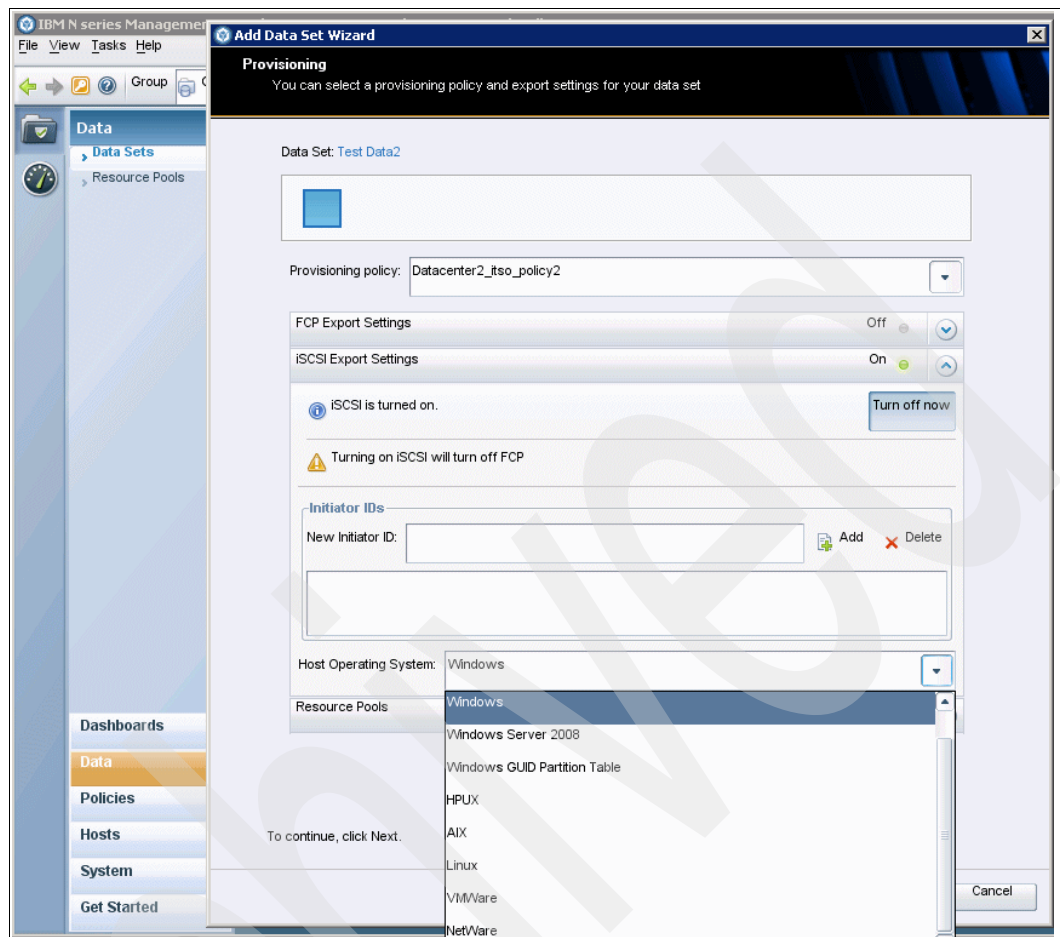


Figure 14-29 Provisioning Manager add new data sets from SAN Storage iSCSI Export settings

Attention: Turning iSCSI on will turn off FCP, that is, you cannot have both of them at the same time.

17. Select the resource ITSO Datacenter2 from the Resource Pools, move it to the Resource Pools in this Node field, and click **Next**.

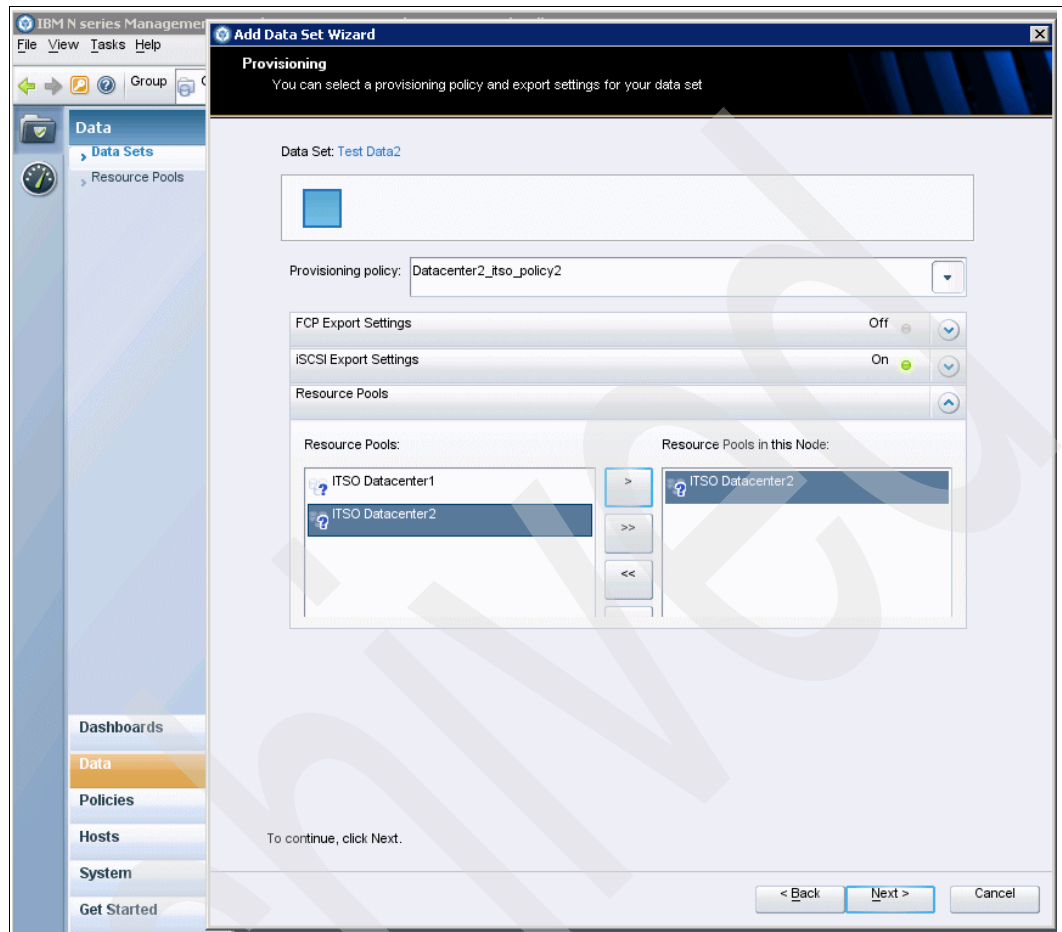


Figure 14-30 Provisioning Manager add new data set from SAN Storage Resource Pool settings

18. Select **No, I will provision storage later**, and click **Next**. For more information about provisioning storage, refer to Chapter 13, “Protecting your data with Protection Manager” on page 313.

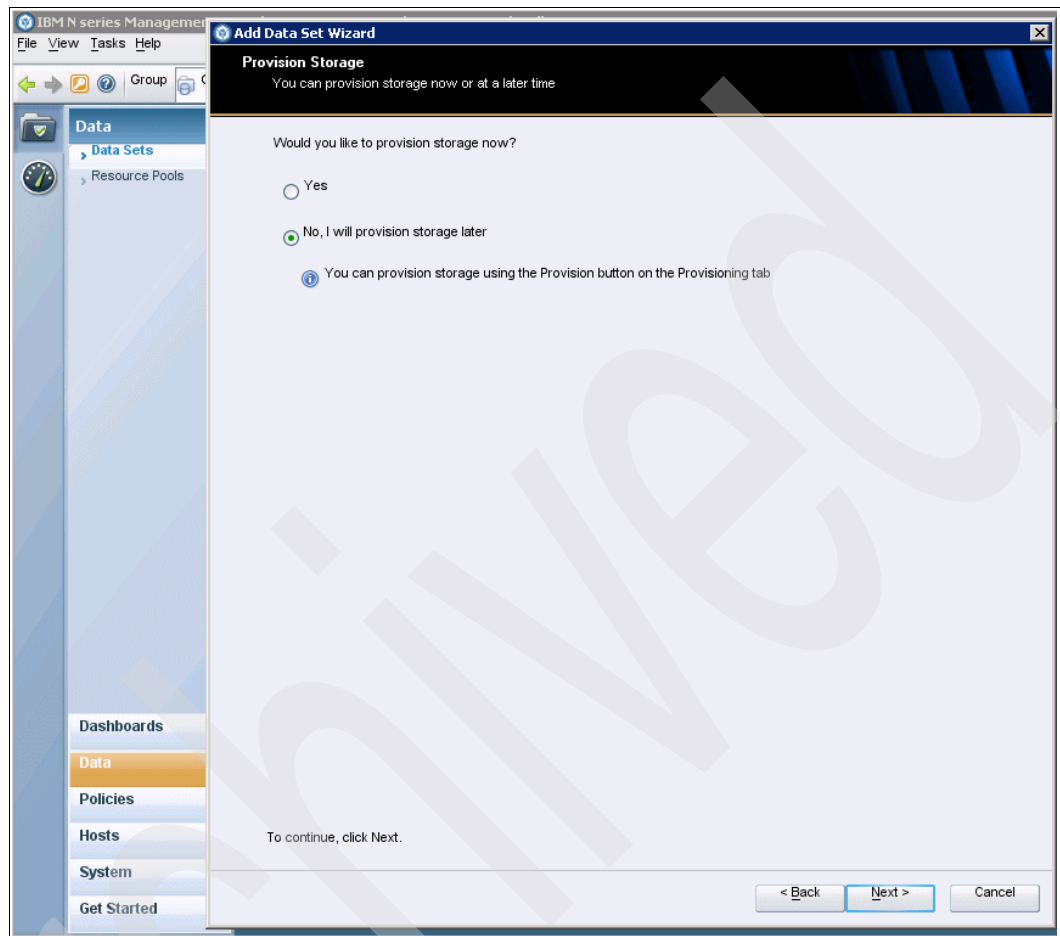


Figure 14-31 Provisioning Manager add new data set storage

You can resize or reallocate data and Snapshot copy space in selected qtrees and volumes.

You can use the Space Management feature to delete selected Snapshot copies from individual volumes in order to make more volume space available for data. You can also delete selected volumes, LUNs, or qtrees to return the space they use to their containing aggregates or volumes.

19. The Completing the Add Data Set Wizard window appears, which displays the summary of the tasks completed, as shown in Figure 14-32. Click **Finish**.

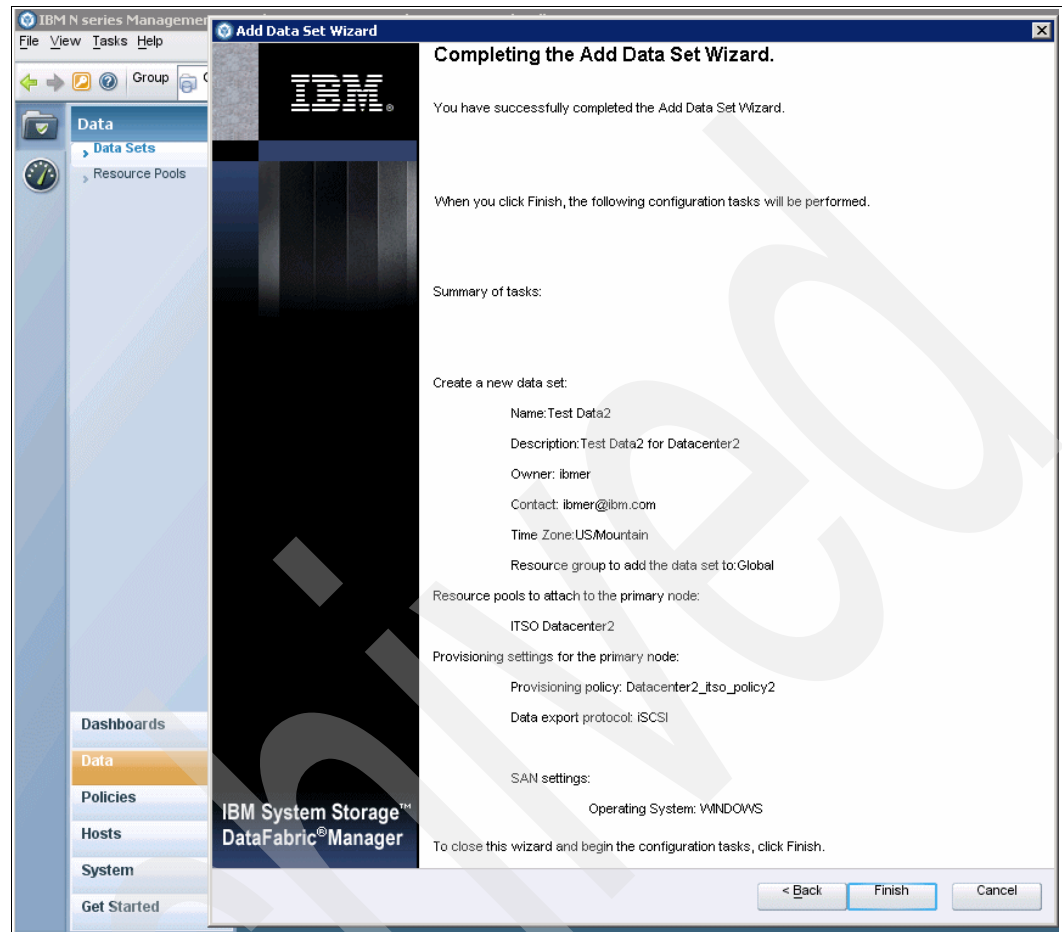


Figure 14-32 Provisioning Manager add new data set settings complete

20. Select **Data** → **Data Sets**, click the **Overview** tab, and click Test Data2. You should see that configuration was successful, as shown in Figure 14-33.

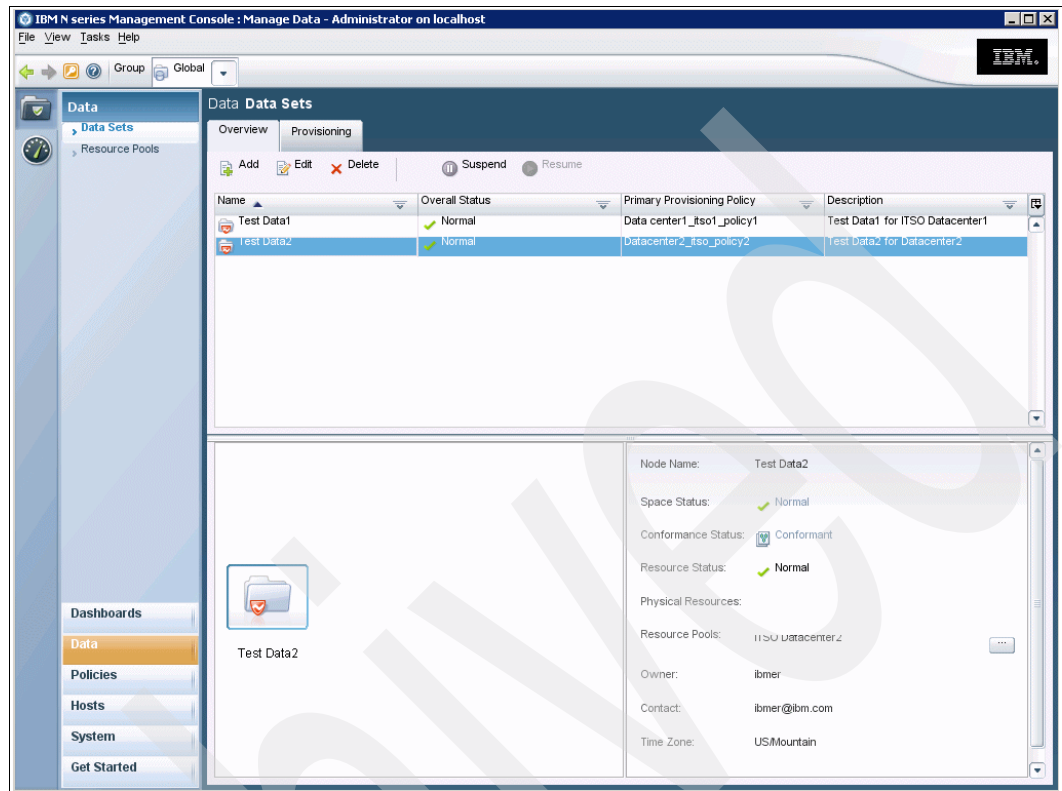


Figure 14-33 Provisioning Manager add new dataset Test Data2

21. Select **Data** → **Data Sets**, click the **Provisioning** tab, and click Test Data1 and Test Data2 to see if they are configured successfully, as shown in Figure 14-34.

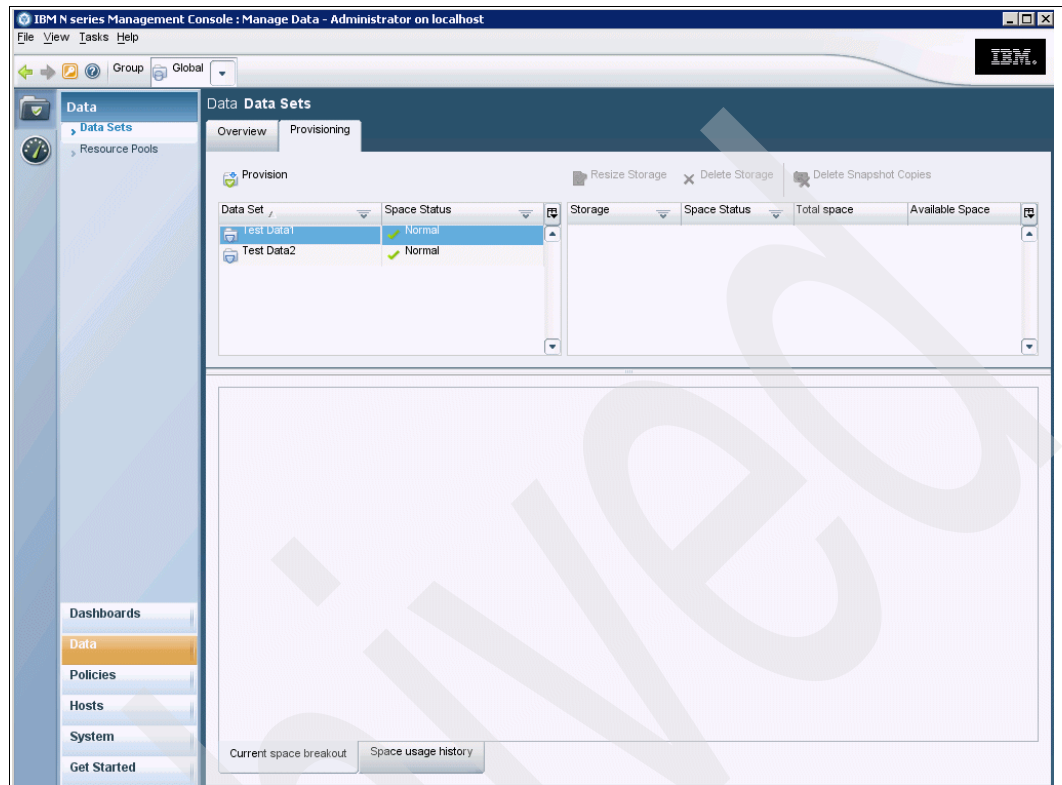


Figure 14-34 Provisioning Manager add new data set Test Data2 Provisioning view

22. The provisioning application user interface allows you to view and manage space usage in storage containers. Figure 14-35 shows the first step in resizing volumes.

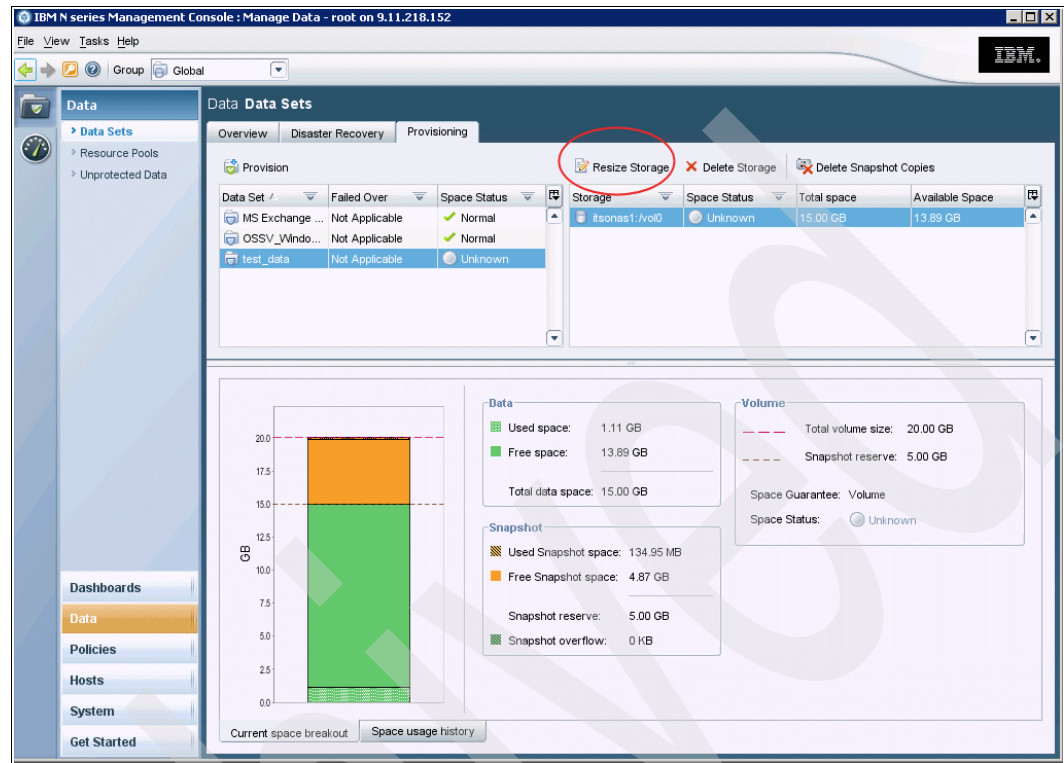


Figure 14-35 Resizing storage

In Figure 14-35, we highlight the data set we want to resize and click **Resize**.

In Figure 14-36, we drag the marker lines to simulate the resizing of the storage. The affected values are shown in red circles.

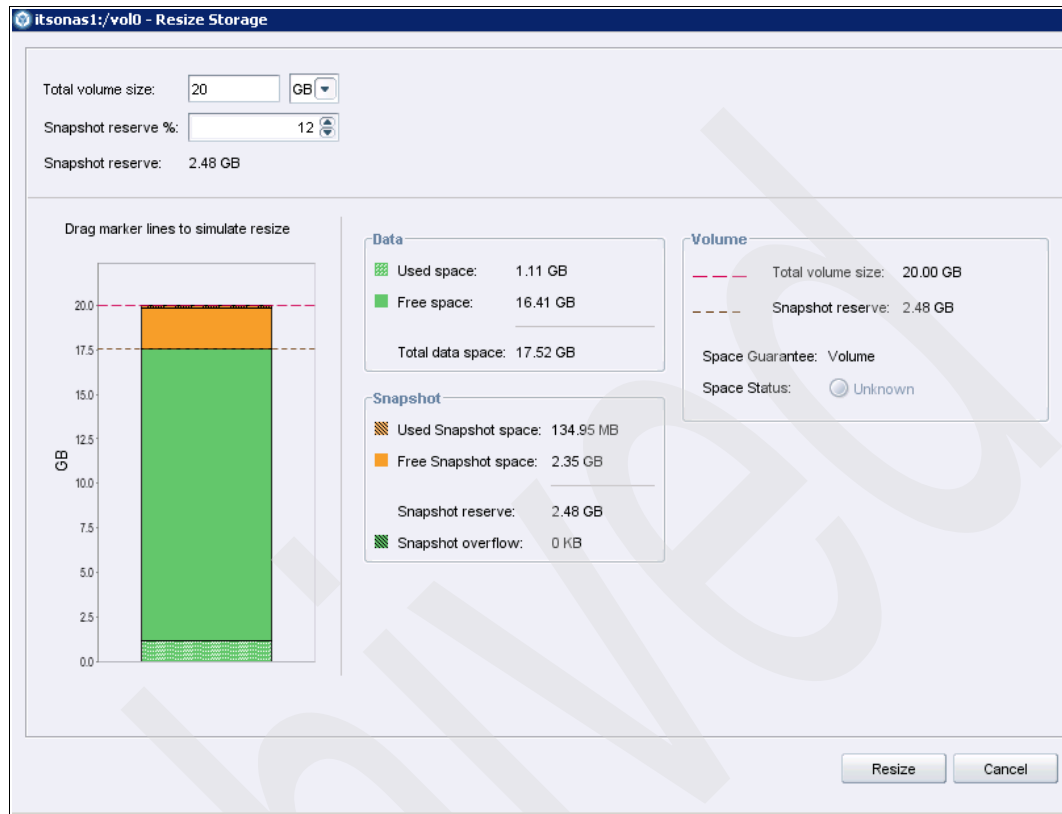


Figure 14-36 Simulating resize

Figure 14-37 shows the results of resizing the data set.

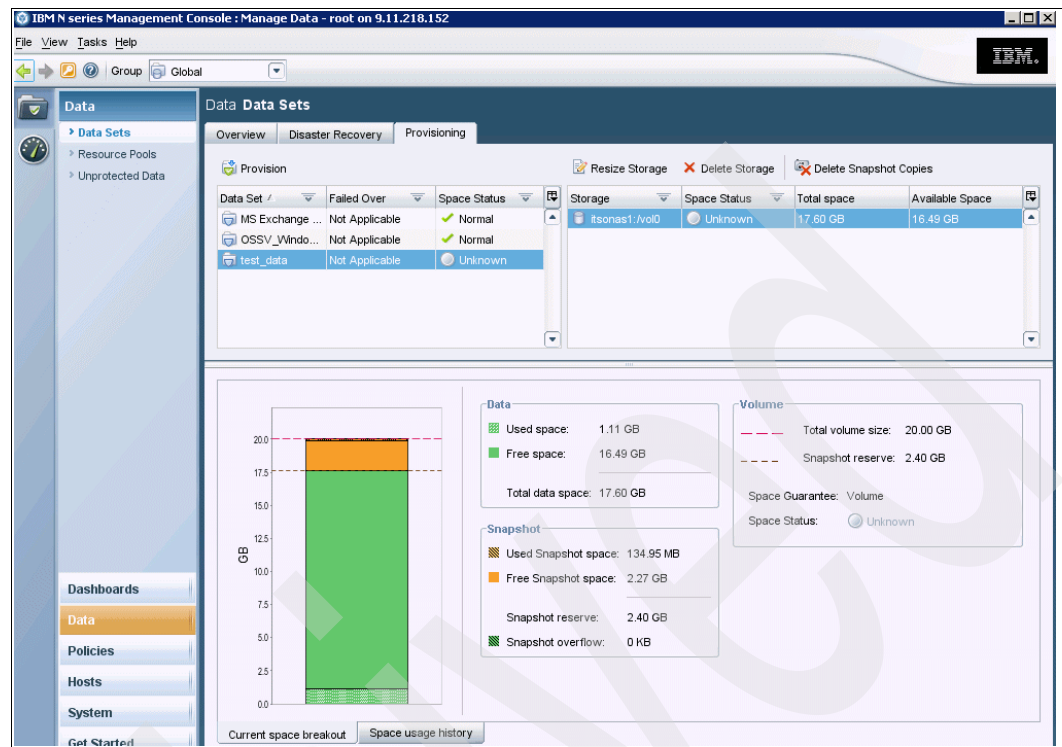


Figure 14-37 Resizing data

14.5.2 Data resource pools information

A *resource pool* is a collection of unused physical storage (such as storage systems or aggregates) from which new volumes or LUNs can be provisioned to contain data, as shown in Figure 14-38. If you assign a storage system to a resource pool, all aggregates on that storage system become available for provisioning. Any unused physical resource in a resource pool is potentially eligible for provisioning. You can organize physical resources into resource pools by location, performance, or other important factors.

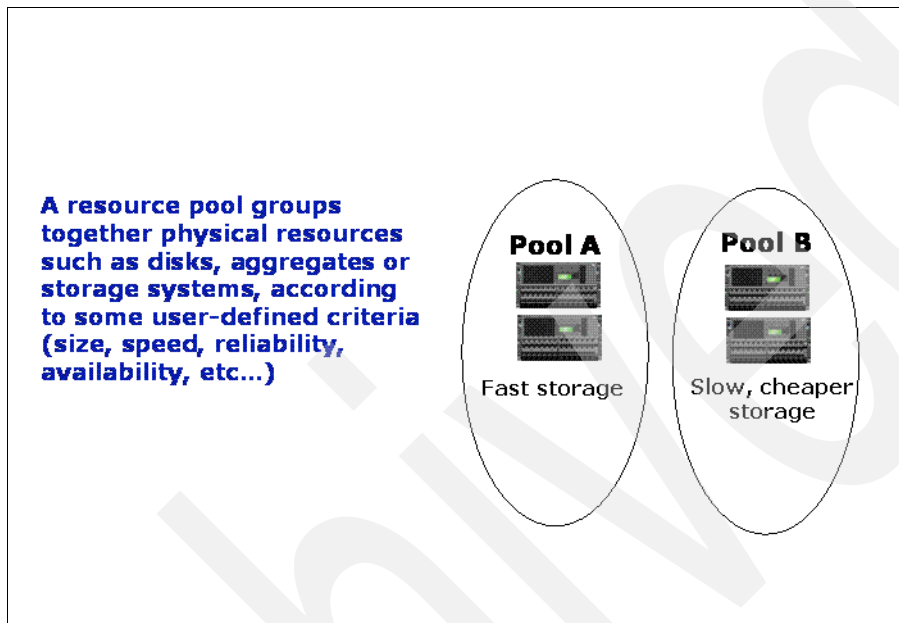


Figure 14-38 Resource pool

With the licensed provisioning application, you can use resource pools to fulfill requests for storage space for the primary or secondary data of a data set. By applying a provisioning policy to a data set node, the provisioning application applies the resiliency characteristics and space settings in the policy to automatically select the resources needed to fulfill a provisioning request.

You can use the Resource Pools window to create, view, and modify collections of physical storage resources, called *resource pools*. When a resource pool is assigned to a data set, data management applications use the resources in the resource pool to provision storage containers as needed by the data set and according to the settings defined in policies assigned to the data set.

For information about resource pool, configuration, refer to “Data resource pools configuration” on page 475.

To view details about a resource pool, select **Data** → **Resource Pools** and then the **Details** tab, as shown in Figure 14-39.

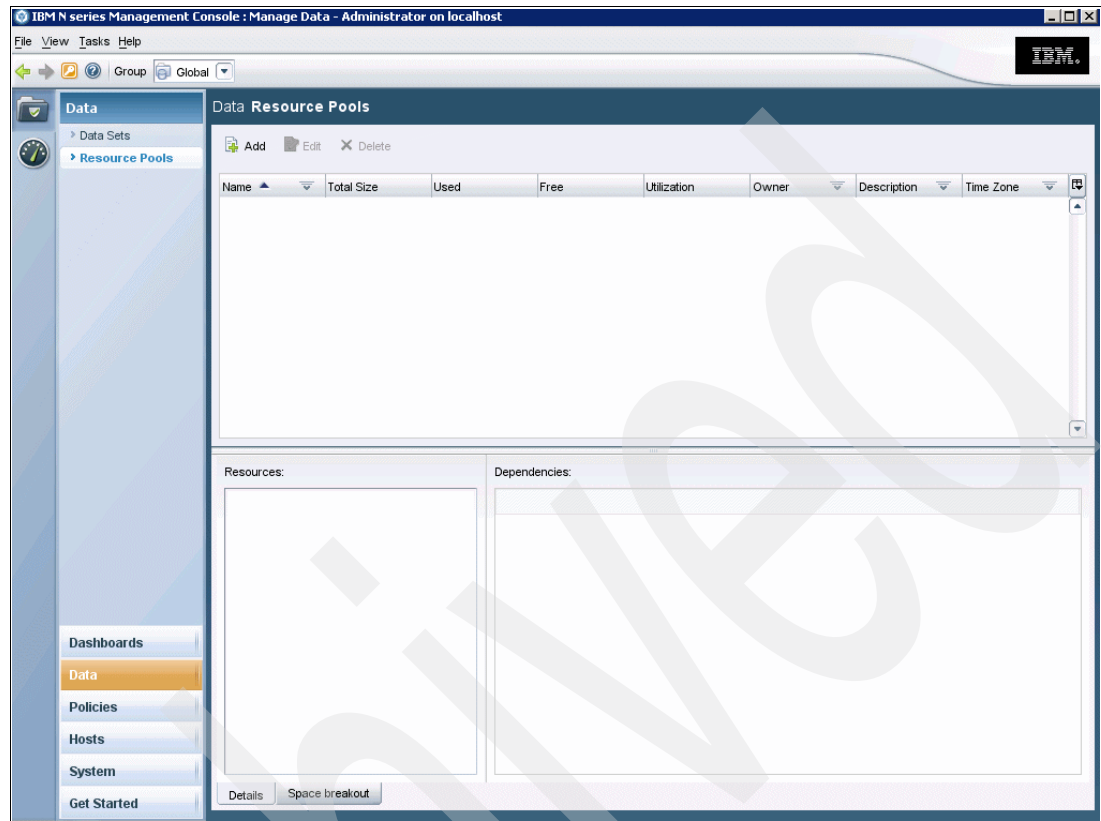


Figure 14-39 Provisioning Manager Data-Resource Pools-Details

The following fields appear:

- ▶ Available resource pools list: This list displays key property settings and space management information for each of the resource pools.

Note: If a resource pool is assigned to a data set using a Protection Manager policy, the time zone you select can impact the protection schedule. Make sure you understand the potential impact to any of your schedules before changing the Time Zone property

- ▶ Resources Details tab: This area displays a tree view of the physical resources assigned to a selected resource pool. The physical resources associated with the selected resource pool are displayed in the expandable Resources list.
- ▶ Dependencies: This area displays the data sets, if any, currently using the resource pool shown in the Resources area.
- ▶ Name: The name of a selected resource pool currently assigned to a data set.
- ▶ Dependency: The name of the data set with which the selected resource pool is associated. A resource pool can be associated with more than one data set.

Data resource pools configuration

This section provides the steps needed to configure data resource pools.

Resource pool creation involves the following steps:

1. Resource pool members: You add storage systems or aggregates from the storage systems. Multiple storage systems and aggregates can be assigned to a resource pool. When a storage system is added to a resource pool, it implies that all its aggregates and disks are part of the resource pool.
2. Resource label: Storage administrators can specify resource pools and their members with labels. For example, they can specify tier 1 (gold), tier 2 (silver), or tier 3 (bronze) for their storage based on performance and cost. These labels can then be used in the provisioning policy to specify the tier of storage. During resource selection, Provisioning Manager will try to find a resource in a resource pool that matches the label.
3. Space thresholds: These sets of thresholds can be used to track the space utilization of a source pool. Administrators get alerts when these thresholds are breached.
4. Aggregate overcommit thresholds: In thin provisioning environments, administrators overcommit more storage than is physically available. These set of thresholds govern the amount of overcommitment. Once these thresholds are breached, an overcommitment alert is generated.

Do the following steps to configure the data resource pools:

1. Select **Data** → **Resource Pools** and click the **Add** button to add a new resource to the resource pool, as shown in Figure 14-40.

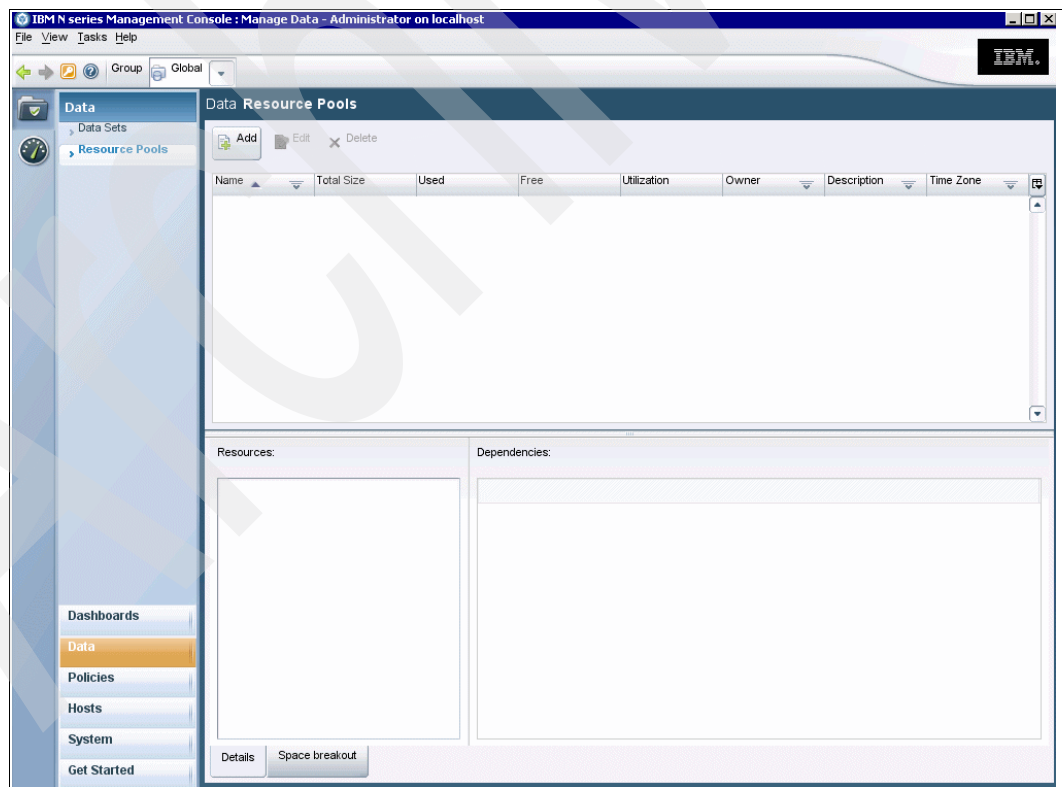


Figure 14-40 Data resource pools add new resource

2. The Add Resource Pool Wizard welcome window appears, as shown in Figure 14-41. Click **Next**.

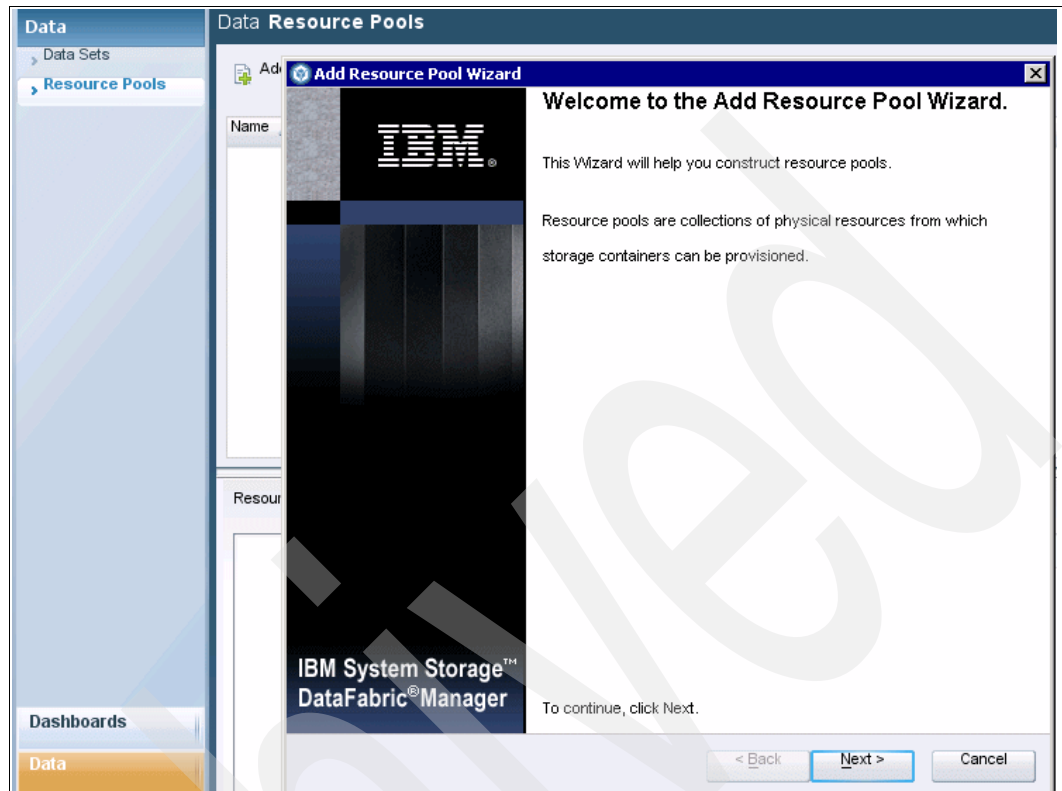


Figure 14-41 Provisioning Manager Add Resource Pool Wizard welcome window

3. In the General Properties window, specify the name, description, owner, and contact details with time zone, as shown in Figure 14-42.

The General Properties options are as follows:

- Name: Type in a name. We use ITSO Datacenter1 here.
- Description: Type in a description. We use Data1 for ITSO Datacenter.
- Owner: Type in the name of the person who is responsible for this data resource pool. We use ibmer.
- Contact: Type in the e-mail of the owner for the data resource. We use ibmer@ibm.com.
- Time Zone: Choose the region where this resource pool is located from the drop-down menu.

Click **Next**.

The screenshot displays the 'Add Resource Pool Wizard' window, specifically the 'General Properties' tab. The window is titled 'Data Resource Pools' and 'Add Resource Pool Wizard'. It contains the following fields and values:

- Name:** ITSO Datacenter1
- Description:** Data1 for ITSO Datacenter
- Owner:** ibmer
- Contact:** ibmer@ibm.com
- Time Zone:** US/Mountain (selected from a dropdown menu showing options: US/Eastern, US/East-Indiana, US/Hawaii, US/Indiana-Starke, US/Michigan, US/Mountain, US/Pacific)

At the bottom of the wizard, there are three buttons: '< Back', 'Next >', and 'Cancel'. The 'Next >' button is highlighted. The background shows the Provisioning Manager interface with a sidebar containing 'Data Sets', 'Resource Pools', 'Dashboards', and 'Data'.

Figure 14-42 Provisioning Manager add resource pool General Properties

- The Physical Resources window shows the list of the physical resources available under the Global group, as shown in Figure 14-43. We select the itsonas1 storage system with all of its aggregates (you can choose individual aggregates instead of the full storage system). Click **Next**.

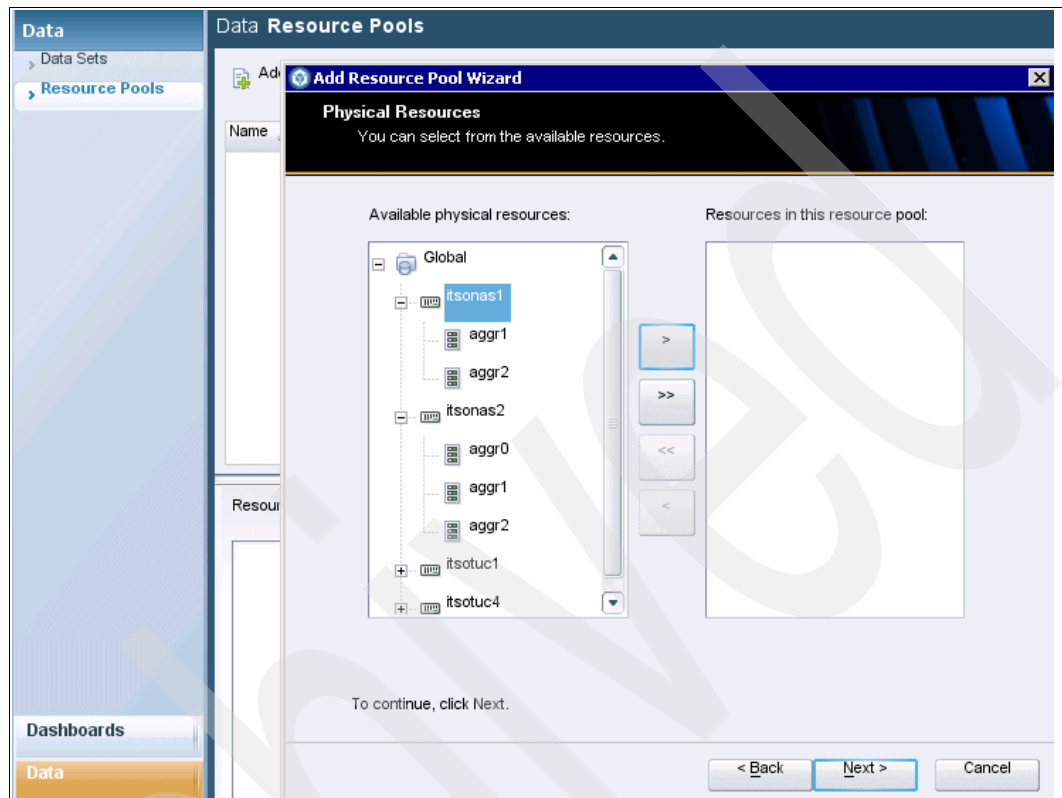


Figure 14-43 Provisioning Manager add resource pools Physical Resources window

5. In the Physical Resources window, select the storage system itsonas1 and move it to the Resource in this resource pool field, as shown in Figure 14-44.

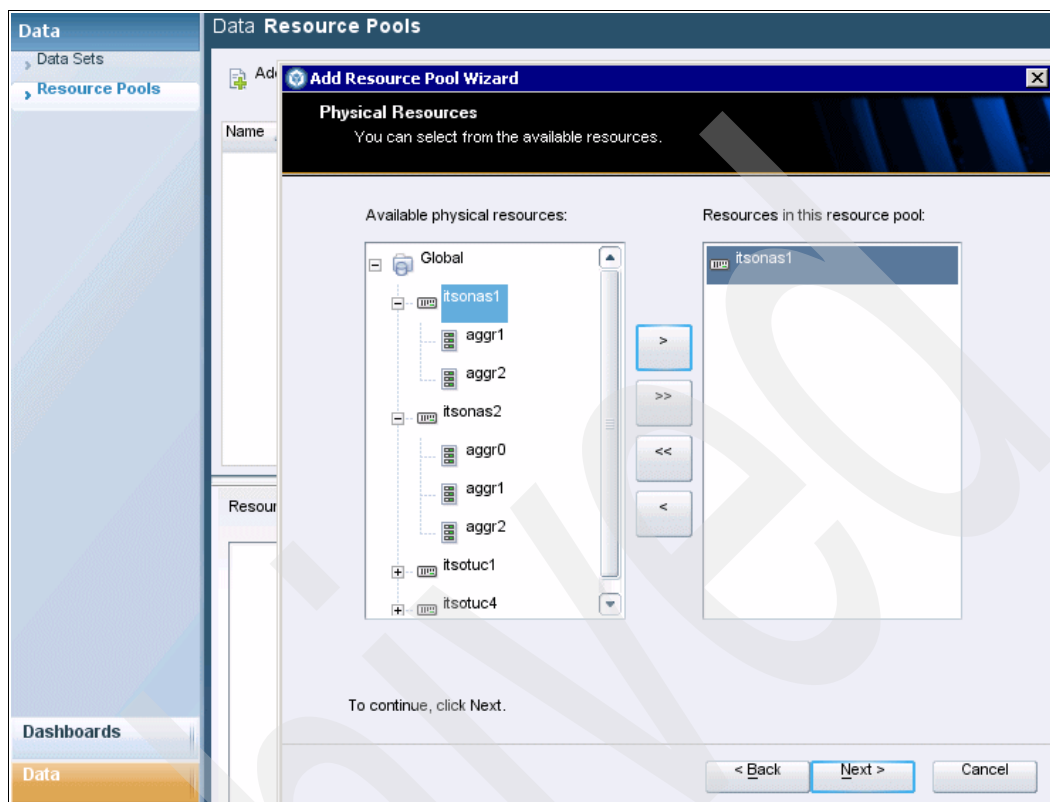


Figure 14-44 Provisioning Manager data physical resource in the resource pool

You can move whole storage systems to the Resources in the resource pool field by clicking the >> button, and you can move them from the Resources in the resource pool field to the Available physical resources field either one by one (by clicking the > button) or all of them (by clicking the >> button).

We have list of the storage systems and respective aggregates, as shown in Figure 14-44, out of which we select the itsonas1 storage system to move to the Resources in this resource pool field. Click **Next**.

6. In the Labels window, assign a resource pool label to our resource, as shown in Figure 14-45. A resource pool label is a user-assigned label associated with a resource pool or the objects in a resource pool. It essentially functions as a filter, allowing you to identify specific resources to be considered when fulfilling a provisioning request. A resource pool label is a text string of any length. You can provide a value for this property in the Add Resource Pool Wizard. A resource pool label is an optional property. You can view this property in the Resource Pools window and change it in the Edit Properties window. We use the label `itso_itsonas1`, where `itsonas1` is our physical resource.

Click **Next**.

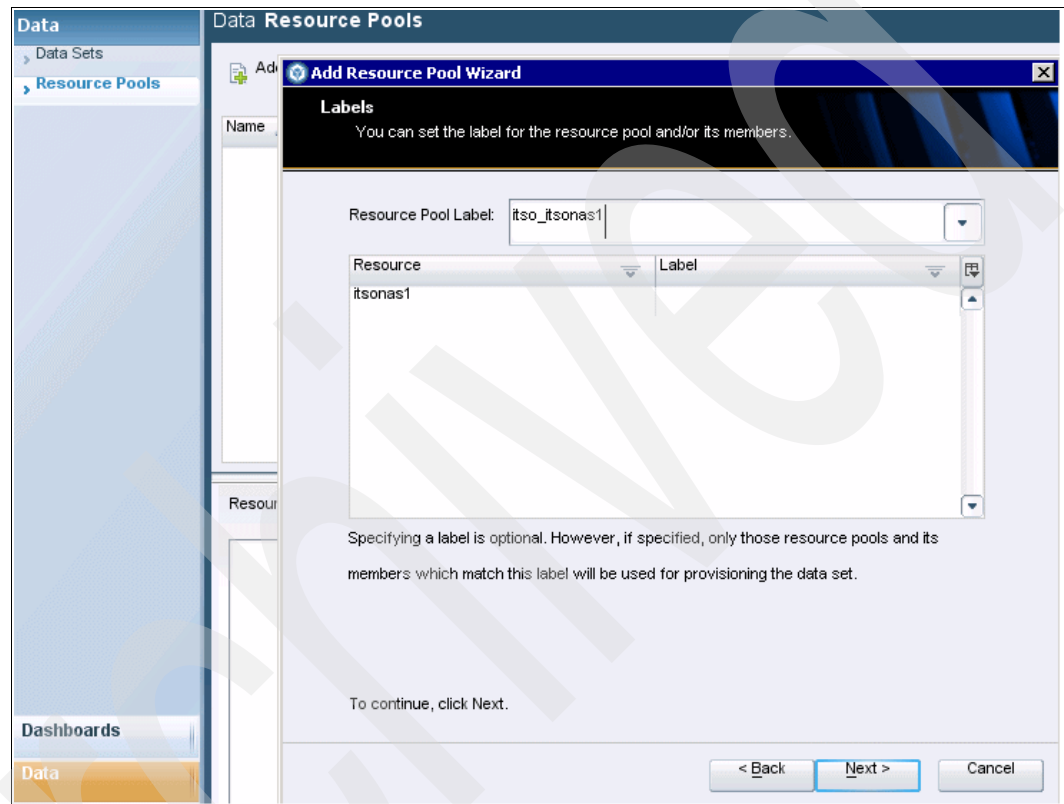


Figure 14-45 Provisioning Manager data resource pool label

7. In the Space Thresholds field, the Space thresholds and Aggregate overcommitted thresholds should be enabled by default; if they are not, enable them both and fill in their fields as follows:

- Space thresholds (use defaults)
 - Nearly Full threshold (%): 80
 - Full threshold (%): 90
- Aggregate overcommitted thresholds (use defaults)
 - Nearly overcommitted threshold (%): 300
 - Overcommitted threshold (%): 400

Click **Next**.

The screenshot shows the 'Add Resource Pool Wizard' dialog box. The 'Space Thresholds' section is active, showing the following configuration:

Section	Option	Value
Space thresholds	Enable event generation	<input checked="" type="checkbox"/>
	Nearly Full threshold (%)	80
	Full threshold (%)	90
Aggregate overcommitted thresholds	Enable aggregate overcommitted thresholds	<input checked="" type="checkbox"/>
	Nearly overcommitted threshold (%)	300
	Overcommitted threshold (%)	400

At the bottom of the dialog, the text 'To continue, click Next.' is displayed. The navigation buttons are '< Back', 'Next >', and 'Cancel'.

Figure 14-46 Provisioning Manager data resource pool space threshold

8. The Completing the Add Resource Pool Wizard window appears, as shown in Figure 14-47. Click **Finish**.

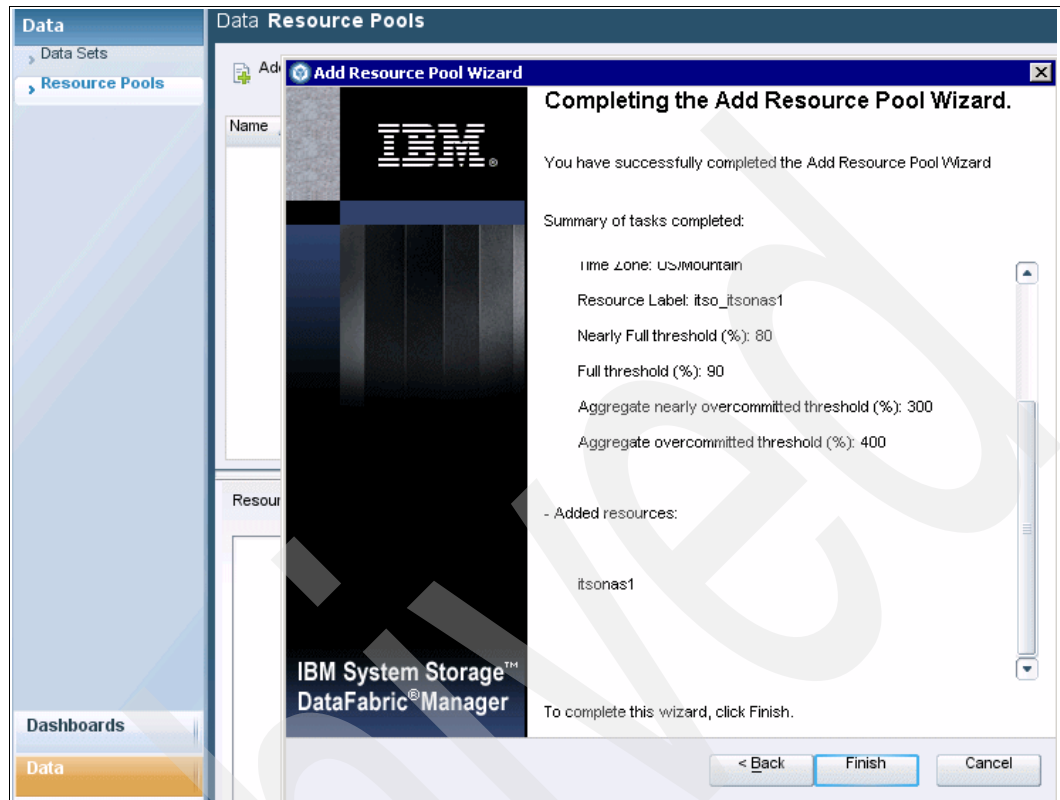


Figure 14-47 Provisioning Manager data resource added and completing the add resource pools

9. Select **Data** → **Resource Pools** and click the **Details** tab, which will show the details about the new data resource we created, as shown in Figure 14-48). The window displays the name of the data, total size of the data, used space, utilization of the space, owner of the data resource, and the region.

You can take the following actions in this window:

- Add: Starts a wizard that helps you create a new resource pool.
- Edit: Opens the Properties window for the selected resource pool. From the Properties window, you can modify the settings (Name, Description, Contact, and Owner), assigned resources, labels, or space thresholds of an existing resource pool.
- Delete: Displays a window that asks you to confirm that you want to delete the selected resource pool. You can either proceed with deleting the selected resource pool or cancel the activity. Deleting a resource pool currently assigned to a data set also deletes any relationships created in accordance with a policy assigned to that data set. Deleting a resource pool does not delete the physical resources that were in the resource pool.

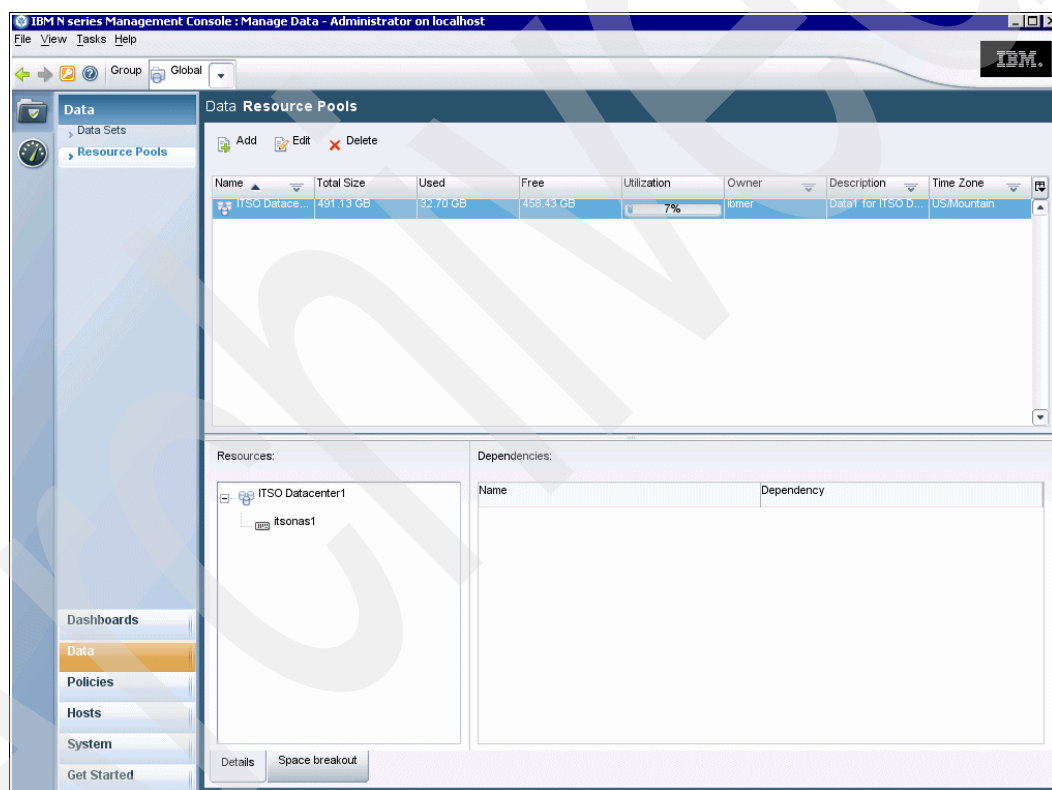


Figure 14-48 Provisioning Manager resource pool data details

10. Select **Data** → **Resource Pools** and then the **Space breakout** tab to show information about the newly created data resource and the aggregates of the storage system, as shown in Figure 14-49 on page 484. There are a number of items that you can view in this window:

- Space breakout tab: The Space breakout tab lists the aggregates assigned to each selected resource pool. If you select one or more items in the list of available resource pools, the aggregates associated with each selected resource pool are displayed in the Space breakout list. If you added a storage system to a selected resource pool, each of its aggregates is listed separately.
- Name: The name of an aggregate in the resource pool.

- **Total Size:** The total amount of storage space provided by the aggregate. This value is expressed in KB, MB, GB, or TB.
- **Used:** The amount of space currently in use on the aggregate. This value is expressed in KB, MB, GB, or TB.
- **Committed Size:** The amount of space guaranteed to the volumes contained in the aggregate. This value can be higher than the total size of the aggregate if you are using the aggregate overcommitment strategy. This value is expressed in KB, MB, GB, or TB.
- **Utilization:** The amount of currently used space out of the total size assigned to the aggregate, expressed in percentages.

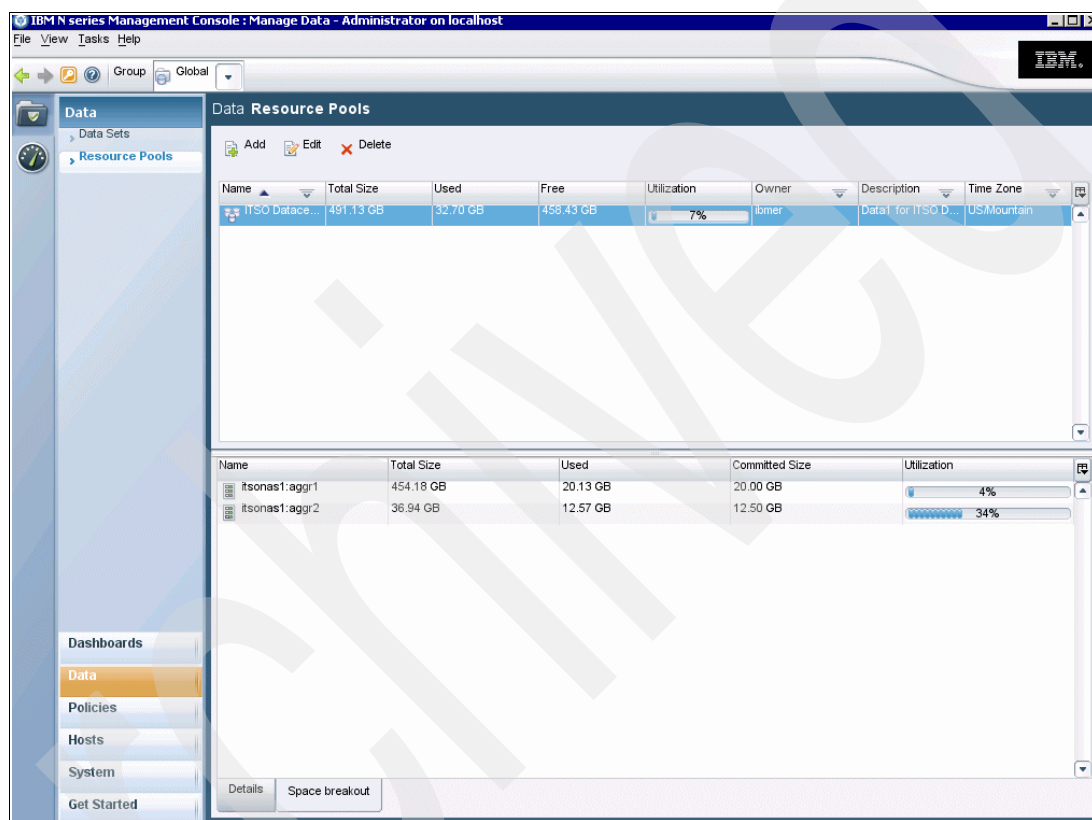


Figure 14-49 Provisioning Manager data resource pools Space breakout tab

14.5.3 Policies provisioning information

This section shows the options for configuring policies provisioning. For more information about this topic, refer to “Policy provisioning configuration” on page 486.

To view the policies provisioning details, select **Policies** → **Provisioning** and click the **Details** tab, as shown in Figure 14-50 on page 485. You can use the Provisioning Policies window to view, add, edit, copy, or delete provisioning policies.

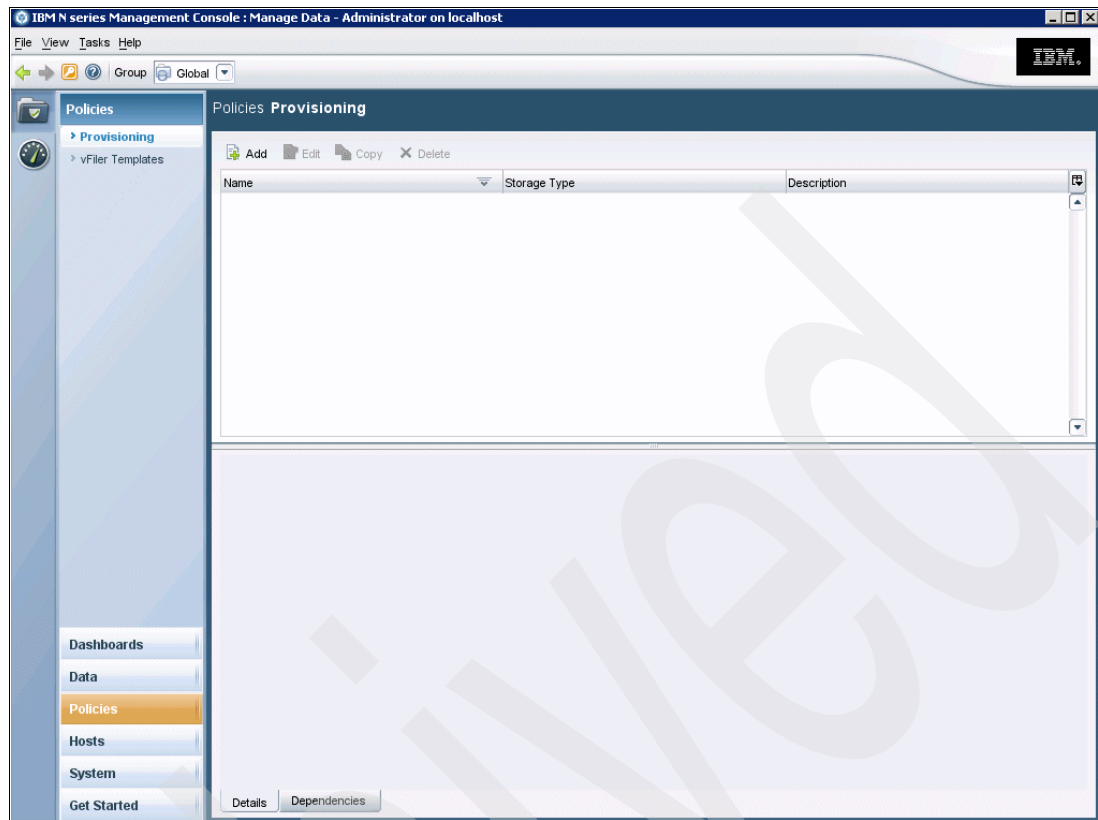


Figure 14-50 Provisioning Manager Policies Provisioning details

The following items are available in this window:

- ▶ **Details tab:** Displays detailed configuration information about the selected policy.
- ▶ **Policies list:** Displays the currently configured provisioning policies and their general property settings. The list is updated dynamically when the status changes. You can customize the display as follows:
 - You can select a policy to see the detailed property settings for that policy.
 - You can use the sort arrows in the column header to specify the sort order of the entries (you click the column header to display the sort arrows).
 - You can click in the upper-right corner of the list to select which columns you want displayed.
 - You can drag the bottom of the policies list area up or down to resize that area.

Provisioning policies define the desired provisioning features of NAS or SAN storage for data sets. For data sets that also have a protection policy assigned, you can create a separate provisioning policy to define the desired provisioning features of a secondary data set node.

Provisioning policies define the storage reliability requirement, space management settings, and appropriate actions when a storage container needs more space. The policy settings specify how you want to have storage provisioned, exported, and managed for the data sets to which you apply the policy.

To view the dependencies of the policies, select **Policies** → **Provisioning** and then click the **Dependencies** tab, as shown in Figure 14-51.

The following items are available in this window:

- ▶ Dependencies tab: Displays all the data sets that are provisioned using the selected policy.
- ▶ Data Set Name: The name of a data set or data set member to which the selected provisioning policy is assigned.
- ▶ Data Set Node: The type of storage, that is, primary storage node or secondary storage node.

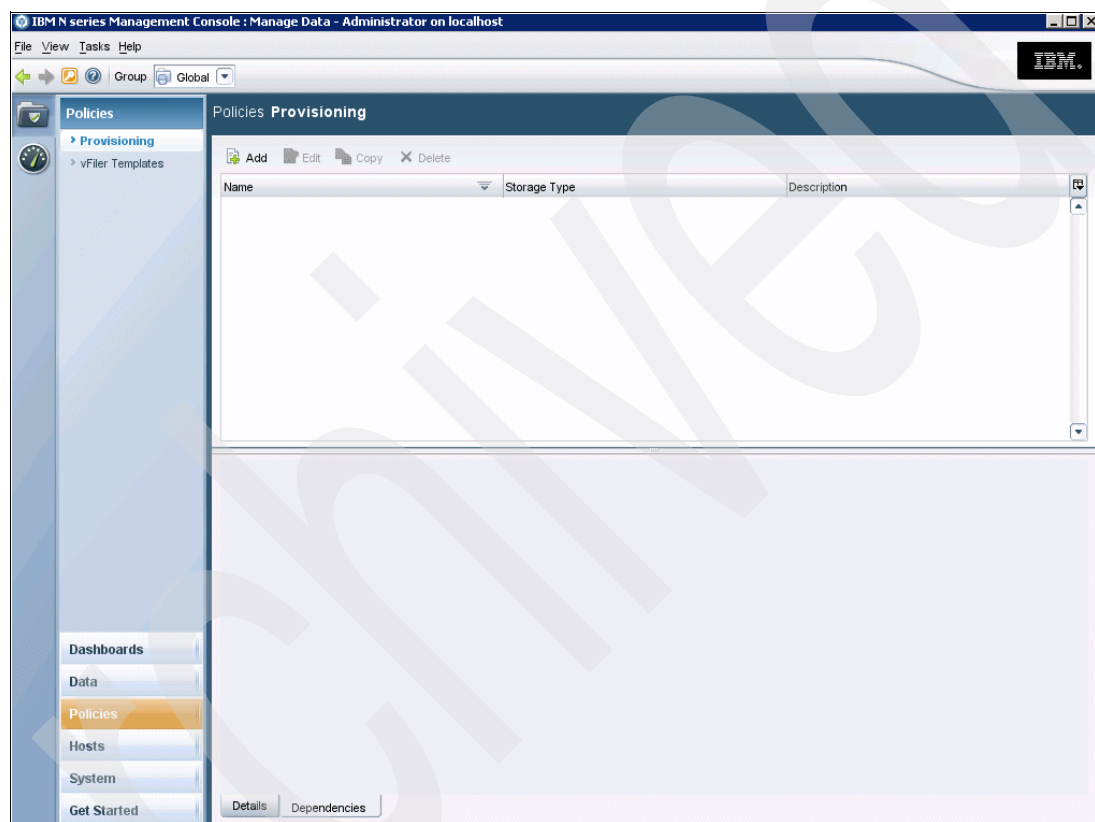


Figure 14-51 Provisioning Manager Policies Provisioning Dependencies

Policy provisioning configuration

This section provides the steps for the policy provisioning configuration.

Provisioning policy configuration involves the following steps:

1. Storage type: Type of storage to be provisioned. In our case, they type of storage will be NAS and the same steps will be repeated for a SAN storage for a second policy.
2. Availability properties: Storage reliability and availability settings (RAID properties and active/active controllers).
3. Container properties: Space and capacity attributes in NAS and SAN environments.
4. Space thresholds: Used space thresholds to report data set space status.

To do the policy provisioning configuration, do these steps:

1. Select **Policies** → **Provisioning** and click the **Add** button, as shown in Figure 14-52, to add new policy.

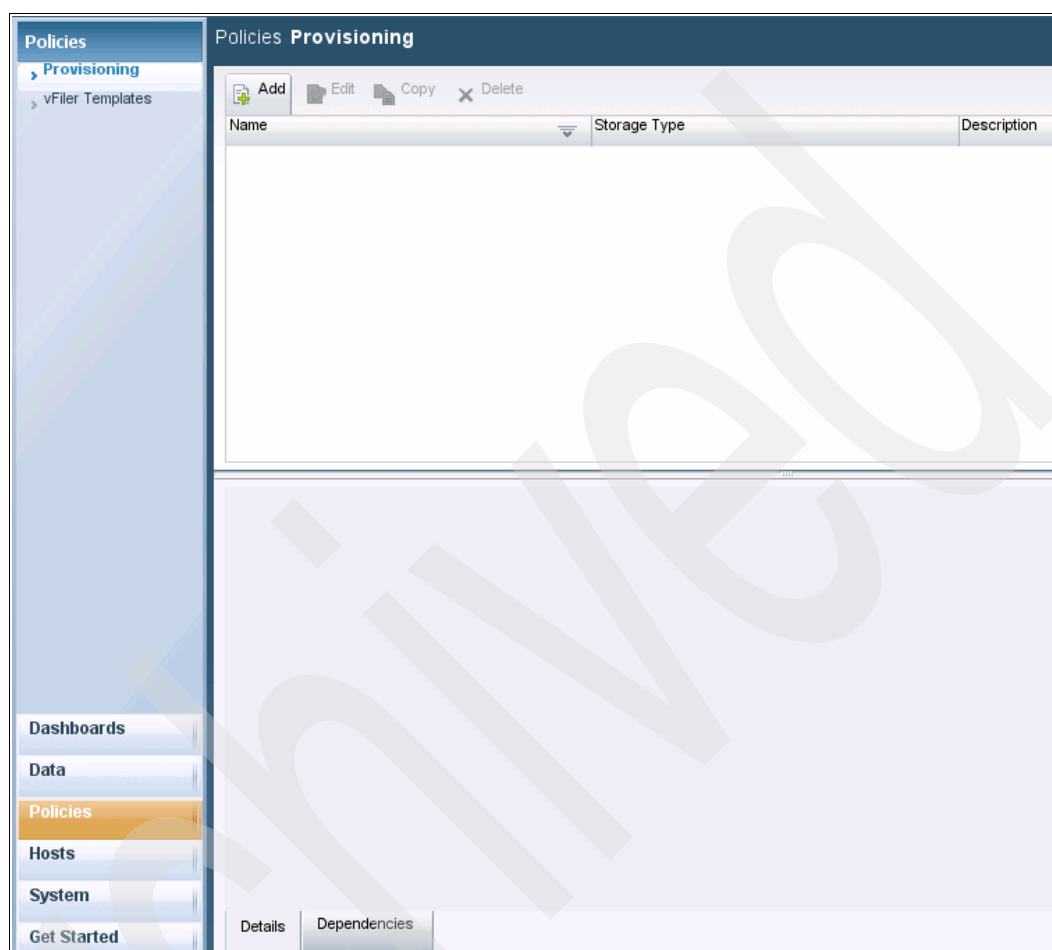


Figure 14-52 Policies Provisioning add new policy

The provisioning application monitors the provisioned space for conformance to the assigned provisioning policy, reports the status of provisioned storage conformance, and generates events when policy thresholds are nearly reached or exceeded. Based on actions configured in the policy, the provisioning application also attempts to correct nonconformance.

Policy-based provisioning conformance includes the following attributes of provisioned storage:

- Availability and reliability levels: A provisioning policy specifies the availability level required for storage provisioned for the associated data set.
- Actions taken when a volume runs out of space: A provisioning policy specifies whether the provisioning application increases the volume size, deletes old Snapshot copies, or increases quotas.
- Data access protocols: A provisioning policy specifies default values for the SAN (iSCSI or FCP) and NAS (NFS or CIFS) data access protocols.
- Thresholds: A provisioning policy specifies the space utilization Full and Nearly Full thresholds.

2. The Add Provisioning Policy Wizard window appears, as shown in Figure 14-53.

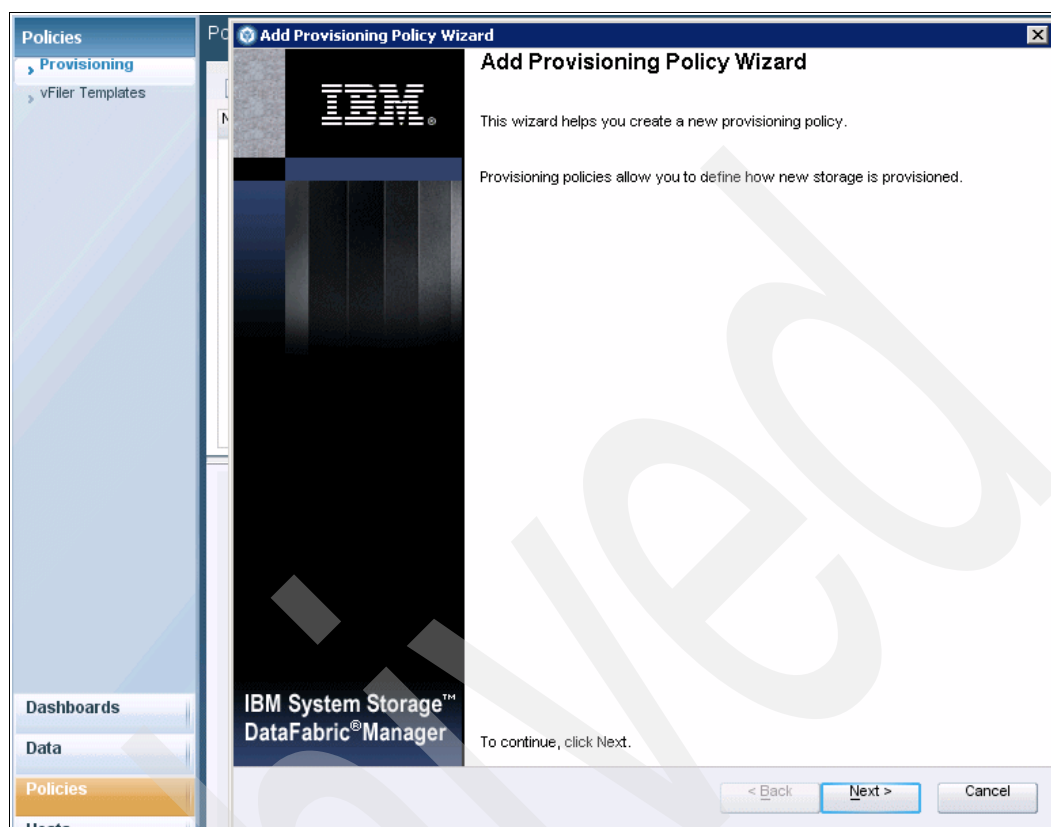


Figure 14-53 Provisioning Policy Wizard to add a new provisioning policy

Before creating a provisioning policy, you need to gather the information necessary to fill in the necessary fields. You need

- The name of the new policy. You can plan to use the name for the primary node.
- A description of the policy (optional).
- The type of storage you want to provision with this policy.
You can select the NAS option for the primary node policy and the secondary option for the backup and mirror node policy.
- For the level of storage availability that the data set requires, you can use the default option, double disks (RAID-DP™), or change to RAID 4.
- For the user and group quota settings, you can use the default settings of 0.
- For the space utilization settings, you can use the defaults:
 - Enable “Guarantee space for data and Snapshot copies” (space guarantee).
 - Enable “Reserve space for Snapshot copies”.
- For the space threshold settings, you can use the defaults:
 - Enable “Space utilization thresholds”.
 - Enable “Nearly Full threshold” at 80% (for the data set).
 - Enable “Full threshold” at 90% (for the data set).

Important: You can decide not to use a resource label or a provisioning script

Click **Next**.

3. in the General Properties window, we set up a policy for itsonas1 with storage type NAS to provision and export for (NFS or CIFS) access, as shown in Figure 14-54.

The screenshot shows the 'Add Provisioning Policy Wizard' window with the 'General Properties' tab selected. The window title is 'Add Provisioning Policy Wizard'. Below the title bar, it says 'General Properties' and 'You can name and describe the policy and select the storage type.' There are two text input fields: 'Name' with the value 'Data center1_itso1_policy1' and 'Description' with the value 'plocy for datacenter1'. Below these is a 'Storage type' section with three radio button options: 'NAS' (selected), 'SAN', and 'Secondary'. Each option has a description: 'NAS' is 'Provision and export storage for NAS (NFS or CIFS) access.', 'SAN' is 'Provision and export storage for SAN (FCP or iSCSI) access.', and 'Secondary' is 'Provision storage for backup or mirror destinations.' At the bottom of the window, there is a message 'To continue, click: Next.' and three buttons: '< Back', 'Next >', and 'Cancel'.

Figure 14-54 Provisioning Policy General Properties settings

You should use these settings:

- Name: The name assigned to the selected provisioning policy.
- Description: An optional description of the provisioning policy. It might include the intended use of the provisioning policy or some other common attribute that identifies why the provisioning policy was created.
- Storage Type: The type of storage that the policy is configured to support:
 - NAS: The policy provisions NAS storage and exports storage for NAS access. NFS and CIFS protocols are supported.
 - SAN: The policy provisions SAN storage and exports storage for SAN access. FCP and iSCSI protocols are supported.
 - Secondary: The policy provisions storage for secondary (backup or mirror) node. This type is available only if you have the protection license installed.

Click **Next**.

4. In the Availability Properties window, we set the Disk failure protection field. By default, RAID DP is enabled, but we change the setting to RAID 4, as shown in Figure 14-56 on page 491.

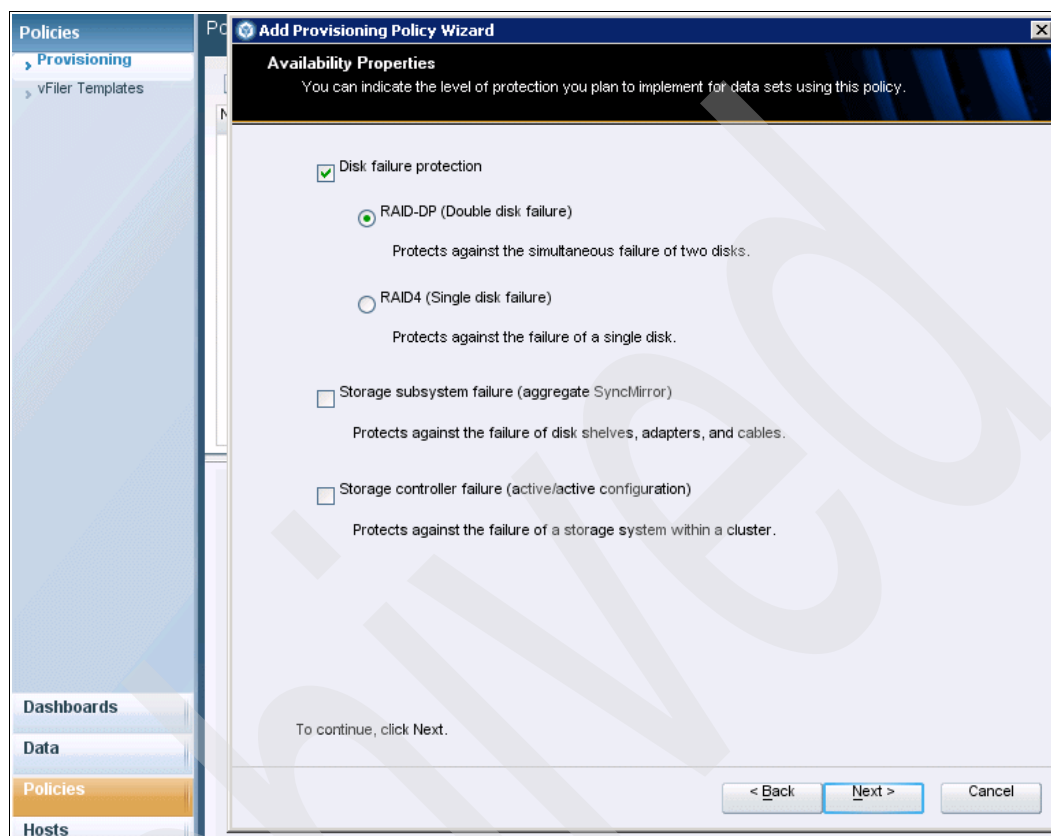


Figure 14-55 Provisioning Policy General Properties settings disk protection enabled

You can choose from the following options:

- RAID 0 (the Disk failure protection check box is unchecked): RAID 0 implements *striping*, which is a way of distributing reads and writes across multiple disks for improved disk performance. Striping reduces the overall load placed on each component disk in that different segments of data can be simultaneously read or written to multiple disks at once. The total amount of storage available is the sum of all component disks. Disks of different sizes may be used, but the size of the smallest disk will limit the amount of space usable on all of the disks. Data protection and fault tolerance is not provided by RAID 0, as none of the data is duplicated. A failure in any one of the disks will render the RAID unusable and data will have been lost.
- RAID 4: RAID 4 stripes data at the block-level and dedicates an entire disk for parity. RAID 4 is used on a limited basis due to the storage penalty and data corruption vulnerability of dedicating an entire disk to parity.
- RAID-DP: RAID-DP adds a second parity disk to each RAID group in a volume. Each traditional RAID 4 group has some number of data disks and one parity disk, with volumes containing one or more RAID 4 groups. While the parity disk in a RAID4 volume stores row parity across the disks in a RAID 4 group, the additional RAID-DP parity disk stores diagonal parity across the disks in a RAID-DP group. With these two parity stripes in RAID-DP, one the traditional horizontal, and the other diagonal, data protection is obtained even in the event of two disk failure events occurring in the same RAID group. In RAID 4, parity is at the block level.

Click **Next**.

5. In the Availability Properties, we check the **Disk failure protection** check box and selected **RAID 4**, check **Storage subsystem failure (aggregate SsyncMirror)**, and check **Storage controller failure (active/active configuration)**. Click Next.

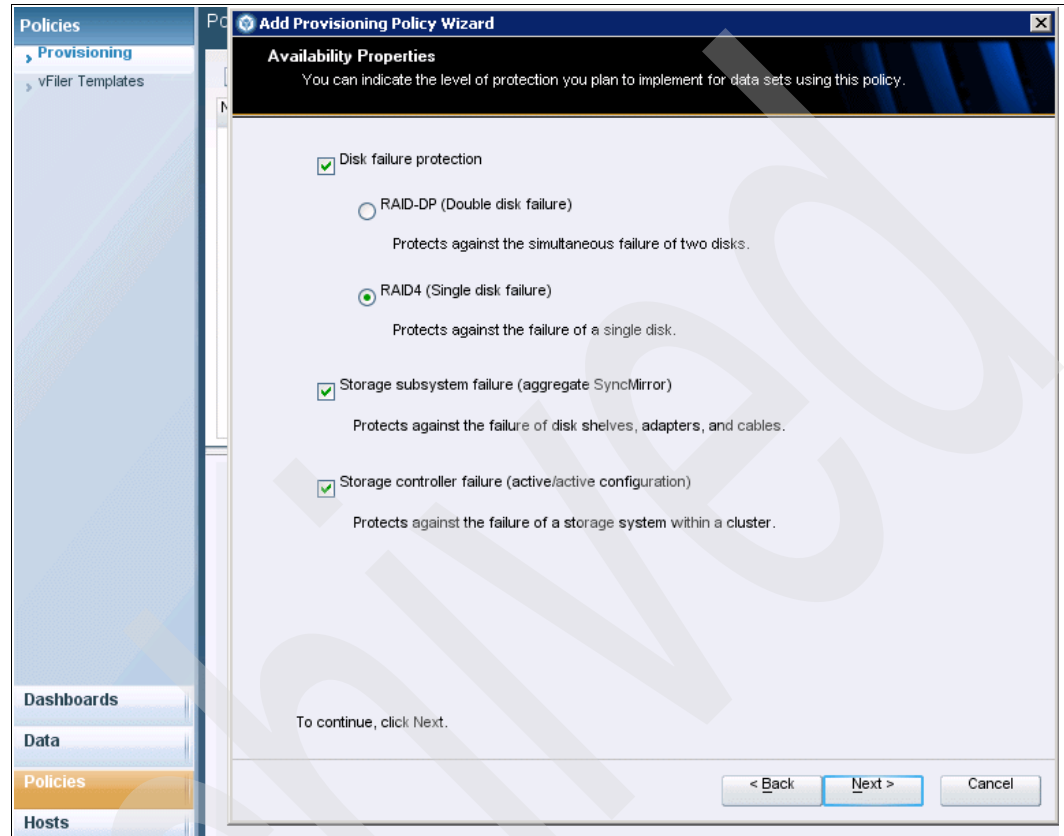


Figure 14-56 Provisioning Manager policy Availability Properties RAID 4 selected

6. The Resource Label window lists the preconfigured resource labels in the Resource labels field. In our example, we choose itso_itsonas1, as shown in Figure 14-57, which we have configured as the primary node of the cluster.

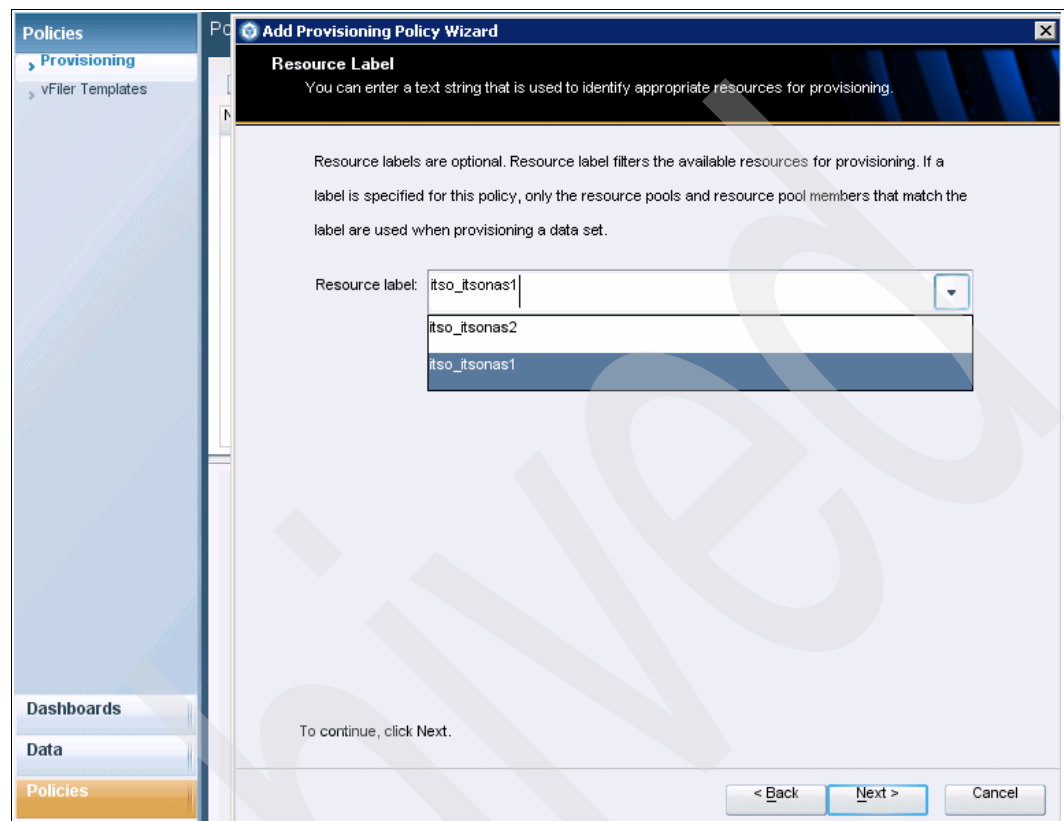


Figure 14-57 Provisioning Policy resource label

The resource label is an optional text string that is used in provisioning requests as a filter for specific resources. Only storage resources that match the label are used for provisioning.

Click **Next**.

7. The NAS Container Properties window appears, as shown in Figure 14-58.

The screenshot shows a software window titled "Add Provisioning Policy Wizard" with a sub-header "NAS Container Properties". Below the header is a descriptive text: "You can select space and quota properties and actions to be taken when a data set needs more space." The window is divided into two main sections: "Quota settings" and "Space utilization properties".

Quota settings: This section contains two rows of input fields. The first row is "Default user quota:" with a text box containing "0" and a dropdown menu set to "KB". The second row is "Default group quota:" with a text box containing "0" and a dropdown menu set to "KB".

Space utilization properties: This section contains two checked checkboxes: "Guarantee space for data and Snapshot copies" and "Reserve space for Snapshot copies". Below these is a text label: "Sample space breakout: (Current data space:100 MB,Snapshot space:25 MB)". A horizontal bar chart visualizes the space usage, with a green segment for "Guaranteed space" and a hatched segment for "Space for snapshot copies". A legend below the chart identifies these colors: green for "Guaranteed space" and orange for "Space allocated on demand". The hatched segment in the chart corresponds to the "Space for snapshot copies" in the legend.

At the bottom of the window, there is a text prompt: "To continue, click Next." and three buttons: "< Back", "Next >", and "Cancel".

Figure 14-58 Provisioning Policy NAS Container Properties

You can set the space and capacity settings for NAS storage in this window as follows:

- Guarantee space for data and Snapshot copies

If this option is enabled, the requested space is guaranteed for data and for Snapshot copies and guaranteed from the resource pool(s) associated with the data set. Writes to a specified FlexVol volume or writes to files with space reservations enabled do not fail due to a lack of available space in the containing aggregate.

If this option is disabled, the space is allocated on demand as data or Snapshot copies are written to data sets; some write requests might fail.

- Reserve space for Snapshot copies

If this option is enabled, an additional 20 percent of the requested space is provisioned from the resource pool(s) for Snapshot copies for every provisioned member in the data set. This guarantees that Snapshot copies do not fail because of a lack of disk space.

If this option is disabled, no additional Snapshot copy space is provisioned and Snapshot copies might fail if there is not enough space available.

- Quota settings (for NAS storage only): The default size of user quotas and group quotas. Quotas limit resource usage and provide notification when resource usage reaches specified levels

8. Figure 14-59 shows the NAS Container Properties with our settings. Click **Next**.

Add Provisioning Policy Wizard

NAS Container Properties
You can select space and quota properties and actions to be taken when a data set needs more space.

Quota settings

Default user quota: 0 KB
Default group quota: 0 KB

Space utilization properties

☒ Guarantee space for data and Snapshot
☒ Reserve space for Snapshot copies

Sample space breakout: (Current data space:100 MB,Snapshot space:25 MB)

Space usage: Space components:

☒ Guaranteed space
☐ Space allocated on demand
☒ Space for Snapshot copies

To continue, click Next.

< Back Next > Cancel

Figure 14-59 Provisioning Policy NAS Container space utilization properties

9. The Space Thresholds window appears, as shown in Figure 14-60.

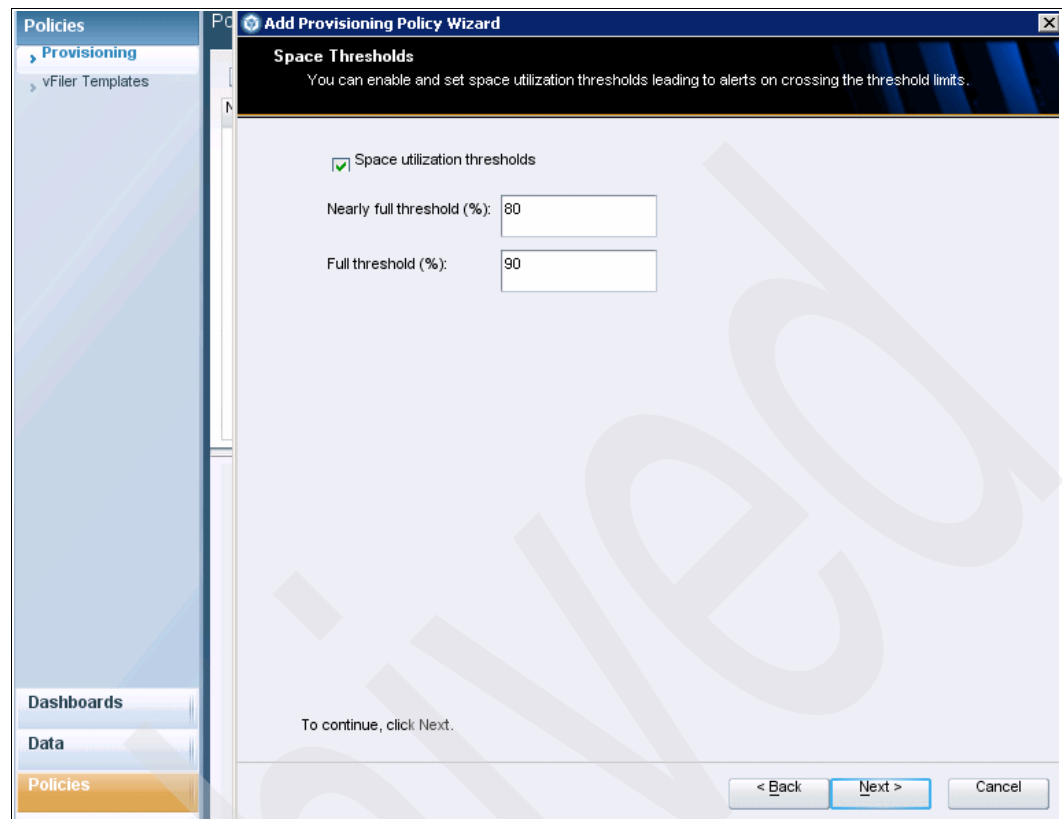


Figure 14-60 Provisioning Policy Space Thresholds window

You can set the properties for the Space Thresholds window as follows:

- Space utilization thresholds: This setting determines whether the licensed application generates space utilization events when a threshold is reached and, if so, what the threshold settings for the Full and Nearly Full thresholds are. The licensed application uses the thresholds to compute data set space status and generate events.
- Full threshold: The percentage of the maximum size of a data set at which a Full threshold event notification is generated.
- Nearly Full threshold: The percentage of the maximum size of a data set at which a Nearly Full threshold event notification is generated.

Click **Next**.

10. The Provisioning Script window appears, as shown in Figure 14-61.

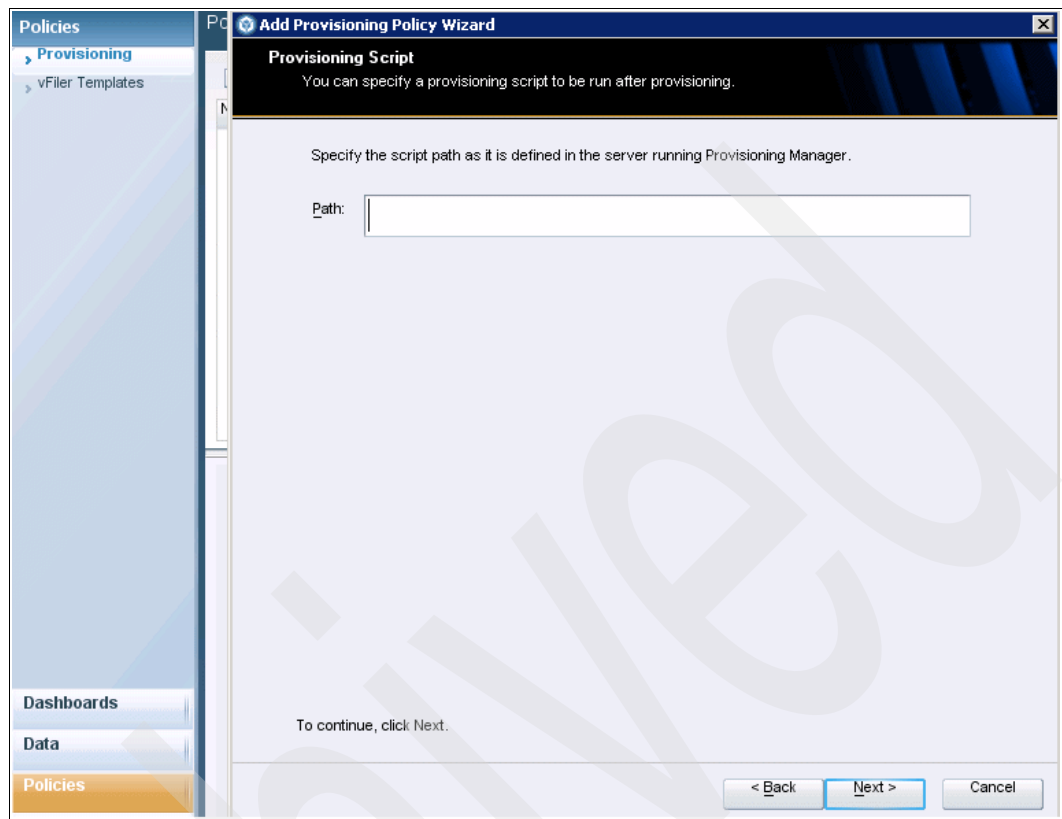


Figure 14-61 Provisioning Policy Provisioning Script window

The Provisioning Script field is where you enter the full Operation Manager server path of a provisioning script that performs custom tasks before or after storage is provisioned.

Click **Next**.

11. The Completing the Add Provisioning Policy Wizard window appears, as shown in Figure 14-62. Click **Finish** to complete the configuration.

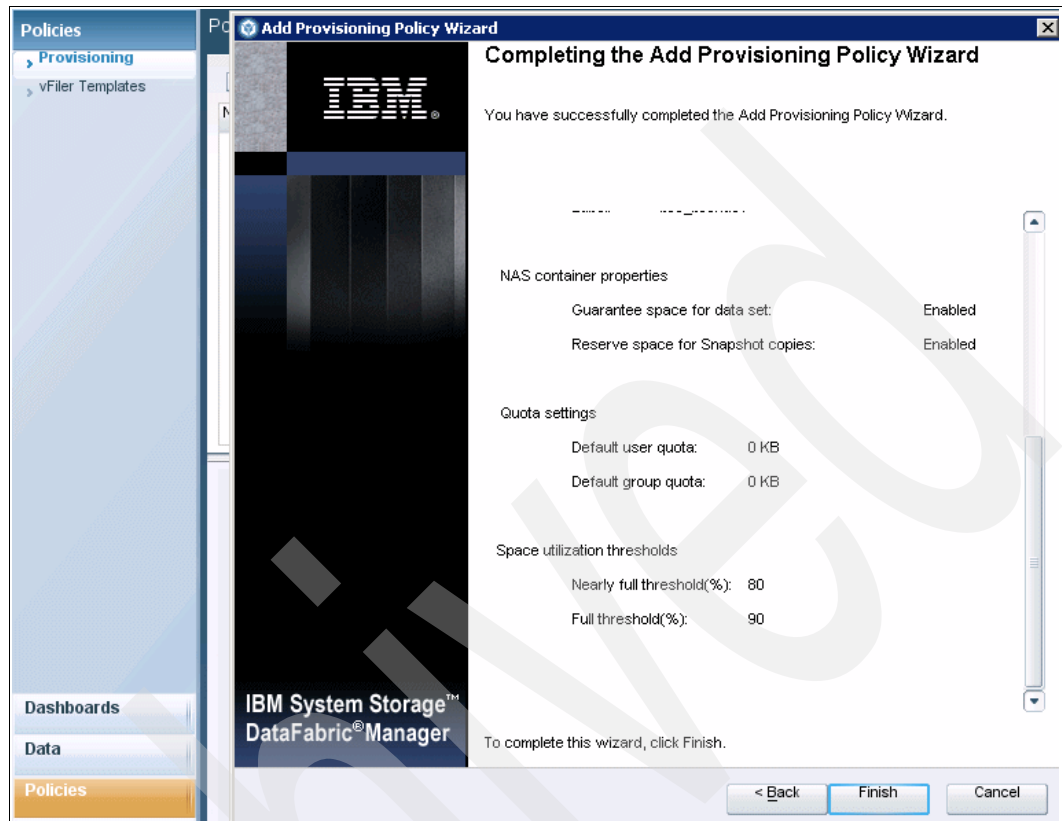


Figure 14-62 Completing the first provisioning policies configuration

12. Selecting **Policies** → **Provisioning** shows the new policy set for the NAS storage type, as shown in Figure 14-63.

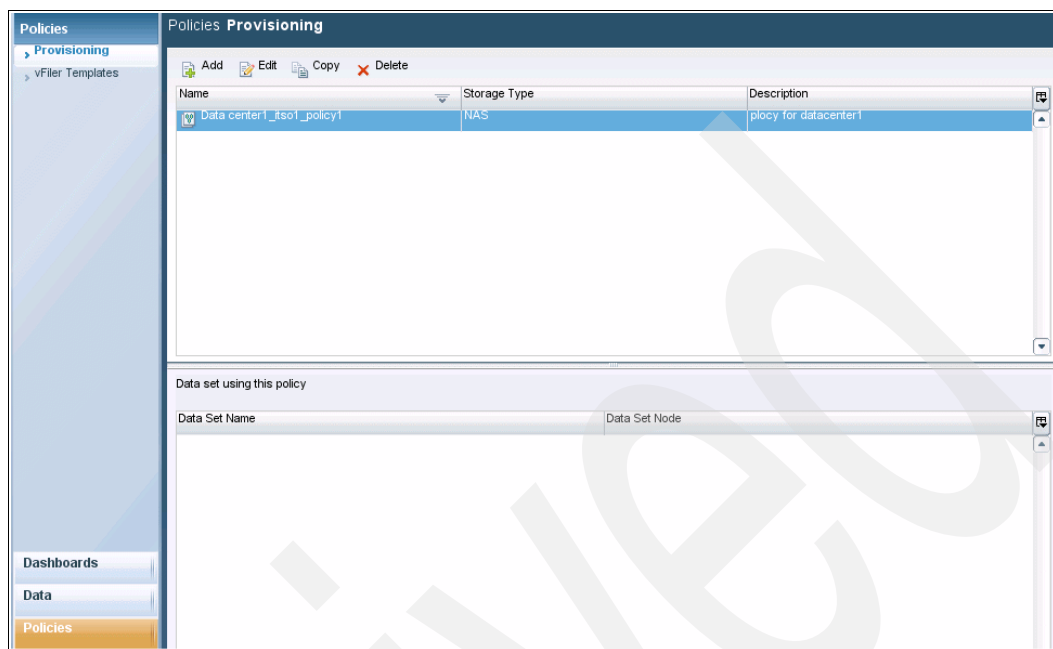


Figure 14-63 View of the provisioning policy

The same steps are used to set the policy for a SAN storage type with some modifications, as shown in Figure 14-64. You have the following tools to help you set the policy:

- ▶ **Add:** Starts the Add Provisioning Policy Wizard, which allows you to configure and add a new provisioning policy.
- ▶ **Edit:** Opens a window in which you can modify the properties of the selected provisioning policy.
- ▶ **Copy:** Copies the selected provisioning policy and adds the copy as a new provisioning policy.
- ▶ **Delete:** Opens a confirmation dialog box to verify whether you want to delete the selected policies. You can delete a provisioning policy only if it has no dependencies, that is, if the policy is not assigned to any data sets.

Do the following steps to set the policy for a SAN storage type:

1. Select **Policies** → **Provisioning**, click the **Add** button, and select **Provisioning Policy Wizard** → **General Properties**, which brings up the window shown in Figure 14-64. Here we set up a policy for itsonas2 with storage type SAN to provision and export storage for (FCP or iSCSI) access. Click **Next**.

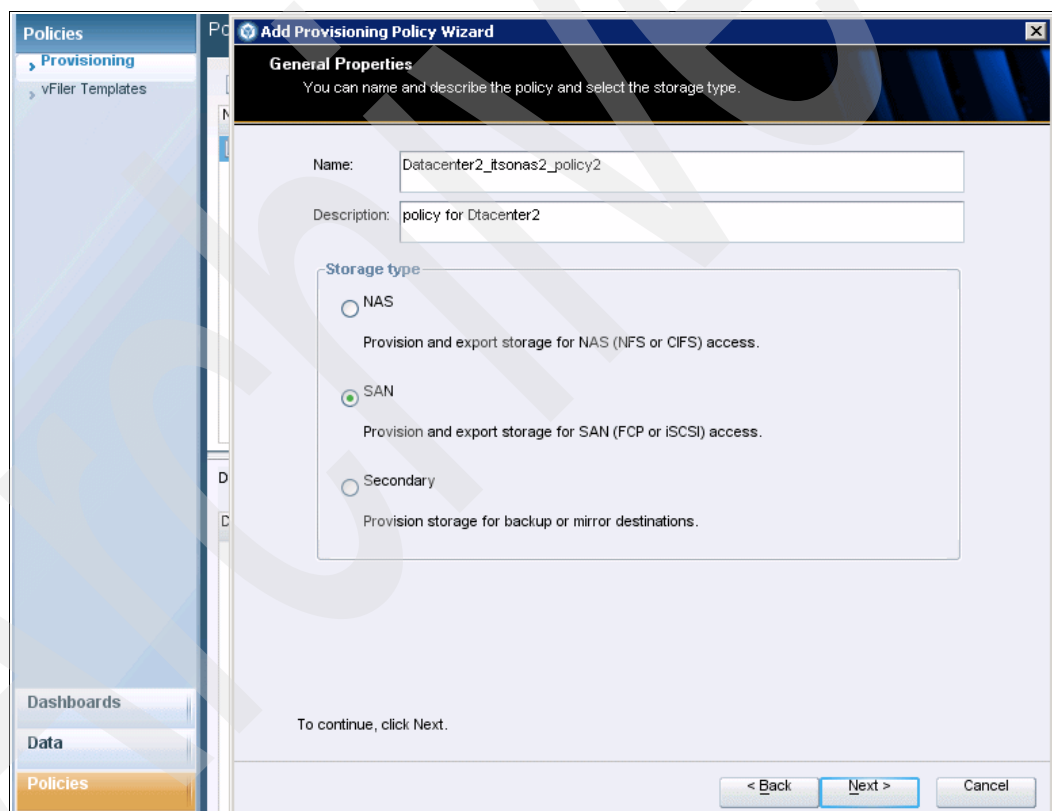


Figure 14-64 Provisioning Policy General Properties settings for the second policy

2. The Resource Label window, shown in Figure 14-65 lists the preconfigured resource labels under resource label configuration. In this example, we choose itso_itsonas2, which we configured for the secondary node of the cluster. Click **Next**.

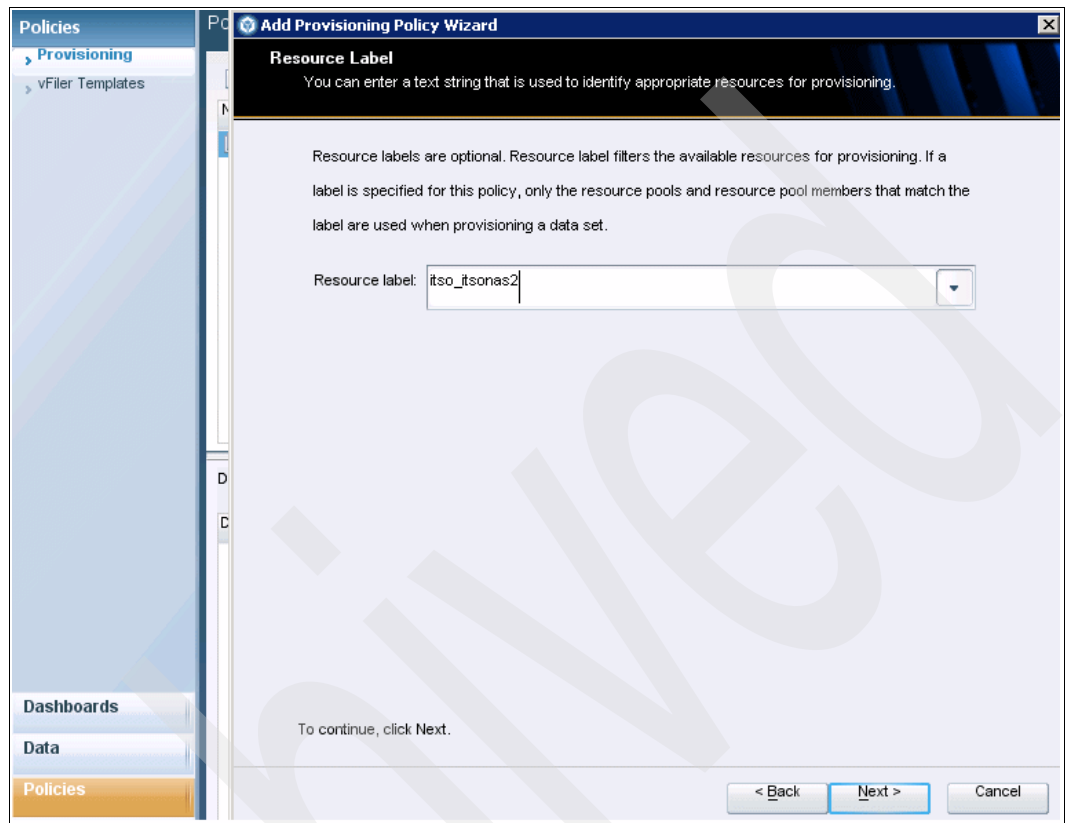


Figure 14-65 Provisioning Policy resource label for second policy

3. In the SAN Container Properties window, shown in Figure 14-66, we choose **LUN** and **Guarantee space for LUN and Snapshot copies**.

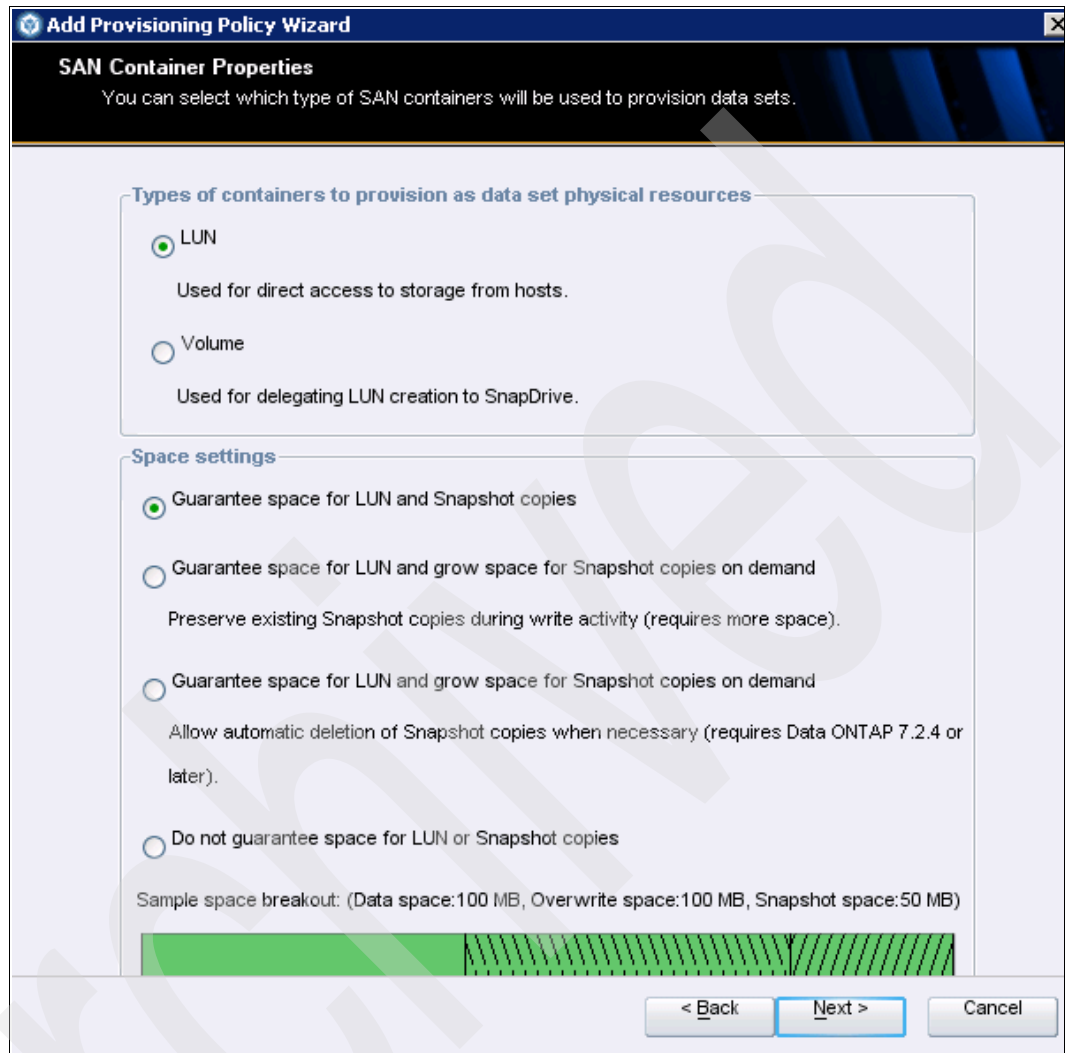


Figure 14-66 Provisioning Policy SAN Container Properties

The SAN Container Properties window (for SAN storage only) describes:

- ▶ Types of containers to provision as data set physical resources
The SAN container type to be used for provisioning: Volume or LUN
- ▶ Space settings
 - “Guarantee space for LUN and Snapshot copies” guarantees space for LUN and Snapshot copies.
 - “Guarantee space for LUN and grow space for Snapshot copies on demand” option is not set, as space might need to be allocated manually. This option normally preserves existing Snapshot copies during write activity, but requires more space.
 - “Guarantee space for LUN and grow space for Snapshot copies on demand, Allow automatic deletion of Snapshot copies when necessary (requires Data ONTAP 7.2.4. or later)” is not set.

If this option is enabled and a storage container needs more space, the Data ONTAP autosize option is used to automatically delete Snapshot copies to make more space available. To use this option, your storage must be using Data ONTAP V7.2.4 or later.

If this option is disabled, Snapshot copies are not automatically deleted when a storage container needs more space, therefore, you might need to delete them manually.

Click **Next**.

4. The SAN Container Properties window shown in Figure 14-67 has the same characteristics as the window shown in Figure 14-66 on page 501, but is for volumes instead of LUNs. “Guarantee space for Volume and Snapshot copies” is checked. Click **Next**.

Add Provisioning Policy Wizard

SAN Container Properties
You can select which type of SAN containers will be used to provision data sets.

Types of containers to provision as data set physical resources

☐ LUN
Used for direct access to storage from hosts.

☒ Volume
Used for delegating LUN creation to SnapDrive.

Space settings

☒ Guarantee space for Volume and Snapshot copies

☐ Guarantee space for Volume and grow space for Snapshot copies on demand
Preserve existing Snapshot copies during write activity (requires more space).

☐ Guarantee space for Volume and grow space for Snapshot copies on demand
Allow automatic deletion of Snapshot copies when necessary (requires Data ONTAP 7.2.4 or later).

☐ Do not guarantee space for Volume or Snapshot copies

Sample space breakout: (Data space:100 MB, Overwrite space:100 MB, Snapshot space:50 MB)

< Back Next > Cancel

Figure 14-67 Provisioning Policy SAN Container Properties

5. The Completing the Add Provisioning Policy Wizard window appears, as shown in Figure 14-68. Click **Finish**.

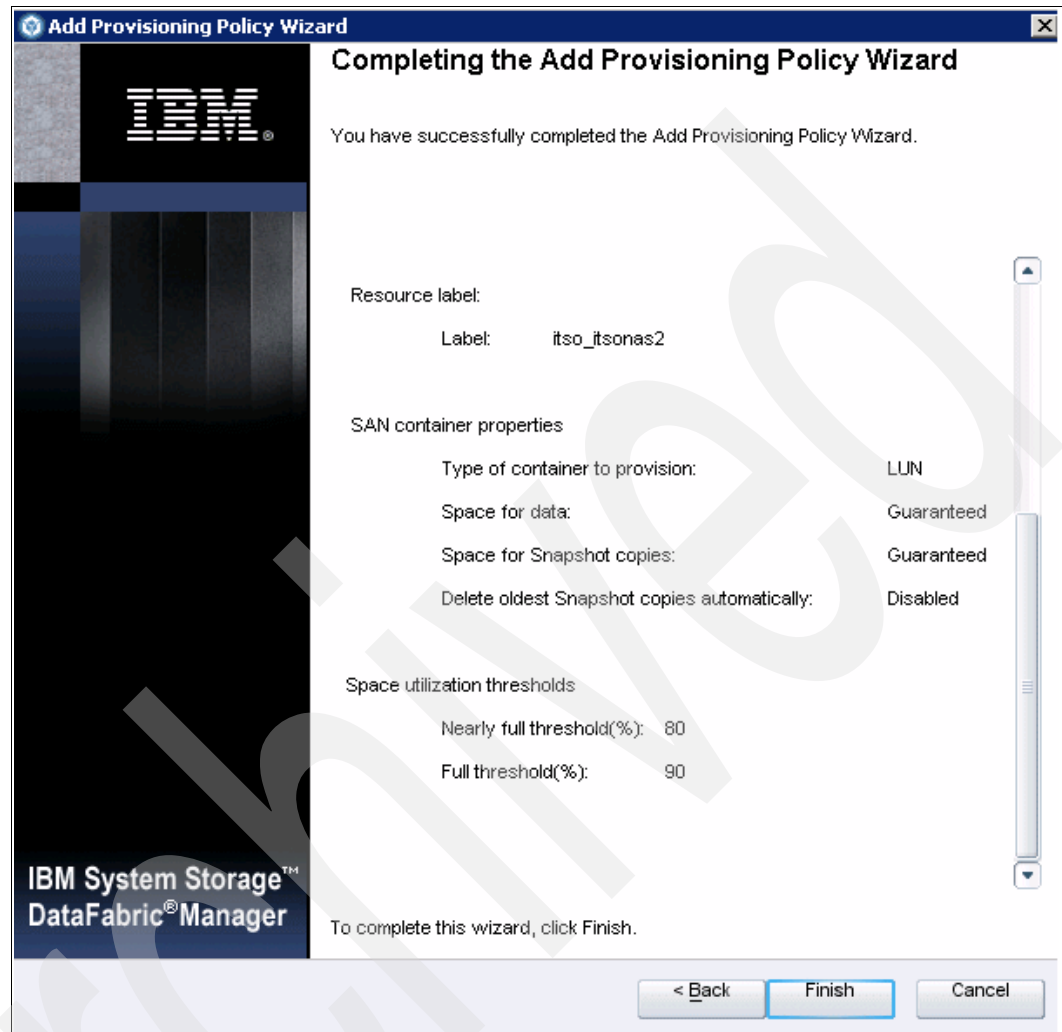


Figure 14-68 Completing the second provisioning policy configuration

Selecting **Policies** → **Provisioning** shows the new policy set for the SAN storage type, as shown in Figure 14-69.

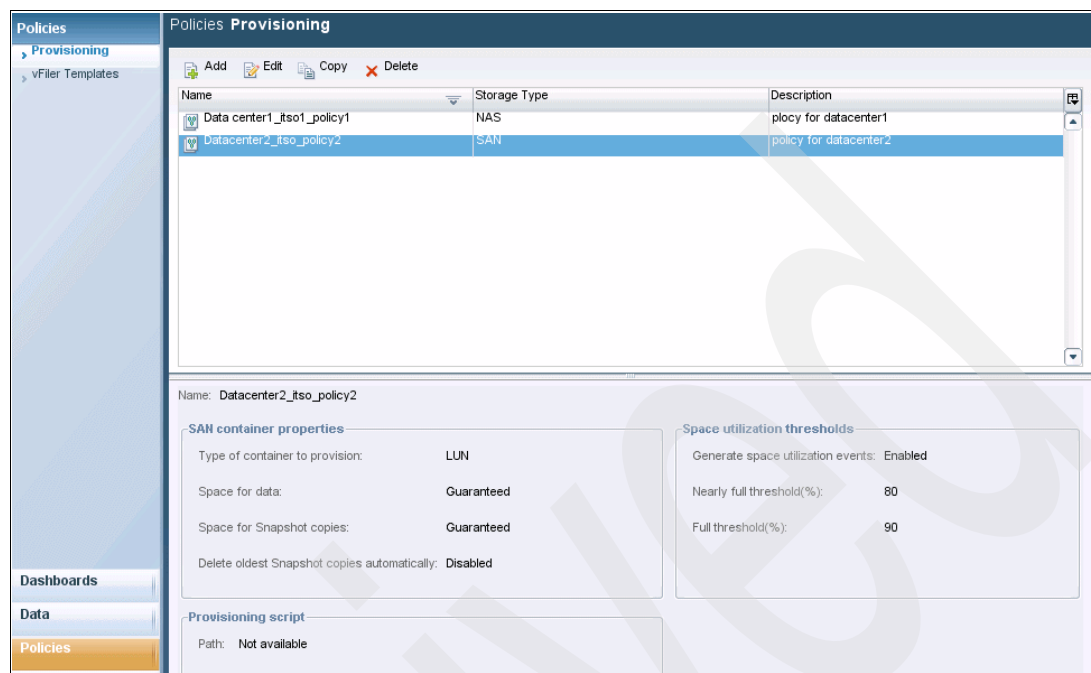


Figure 14-69 Provisioning Policy second policy configuration view

14.5.4 Policies vFilers Templates information

You can use the vFiler Templates window to view, add, edit, copy, or delete vFiler templates and for configuration. Refer to “Policies vFilers Templates configuration” on page 506 for more information.

The provisioning application provides vFiler templates to help manage the exporting of storage through vFiler units by configuring NAS and SAN access to those vFiler units. A vFiler template is a set of vFiler configuration settings, including the corresponding CIFS, DNS, NIS, and administrative host configuration settings that you want to use as default settings for one or more vFiler units you plan to add as hosts. You can configure as many vFiler templates as you need.

To open the vFile Templates window, Select **Policies** → **vFiler Templates**, as shown in Figure 14-70 on page 505.

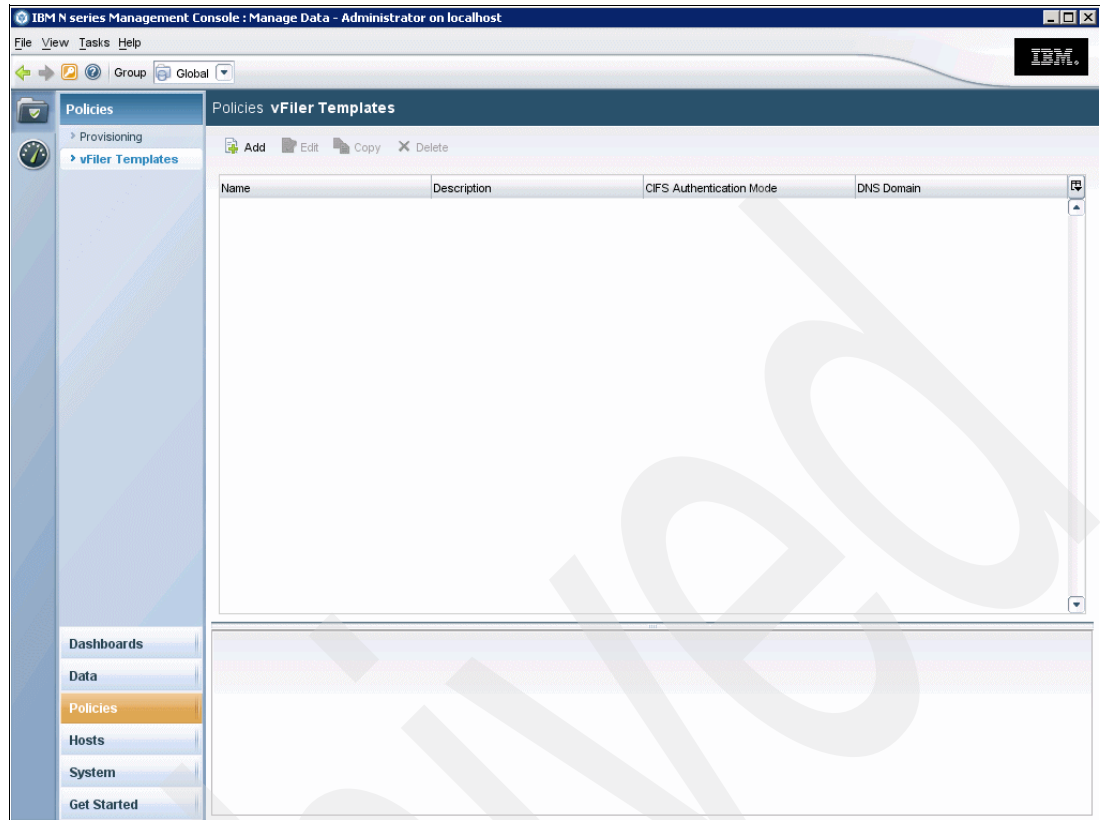


Figure 14-70 Provisioning Manager Policies vFiler Templates

You can do the following actions from this window:

- ▶ **Add:** Starts the Add vFiler Template Wizard, which allows you to configure and add a new vFiler template.
- ▶ **Edit:** Opens a property sheet in which you can modify the properties of the selected vFiler template.
- ▶ **Copy:** Copies the selected vFiler template and adds the copy as a new vFiler template.
- ▶ **Delete:** Deletes the selected vFiler templates.

The vFiler Templates list displays a list of the currently configured vFiler templates. The list is updated dynamically when the status changes. You can customize the display as follows:

- ▶ You can select a vFiler template to see the configuration details for that vFiler template.
- ▶ You can use the sort arrows in the column header to specify the sort order of the entries (you can click the column header to display the sort arrows).
- ▶ You can click the upper-right corner of the list to select which columns you want to be displayed.
- ▶ You can drag the bottom of the list area up or down to resize that area.

Policies vFilers Templates configuration

This section provides the steps to configure the vFilers Templates:

1. Select **Policies** → **vFile Templates** and press the **Add** button, as shown in Figure 14-71.

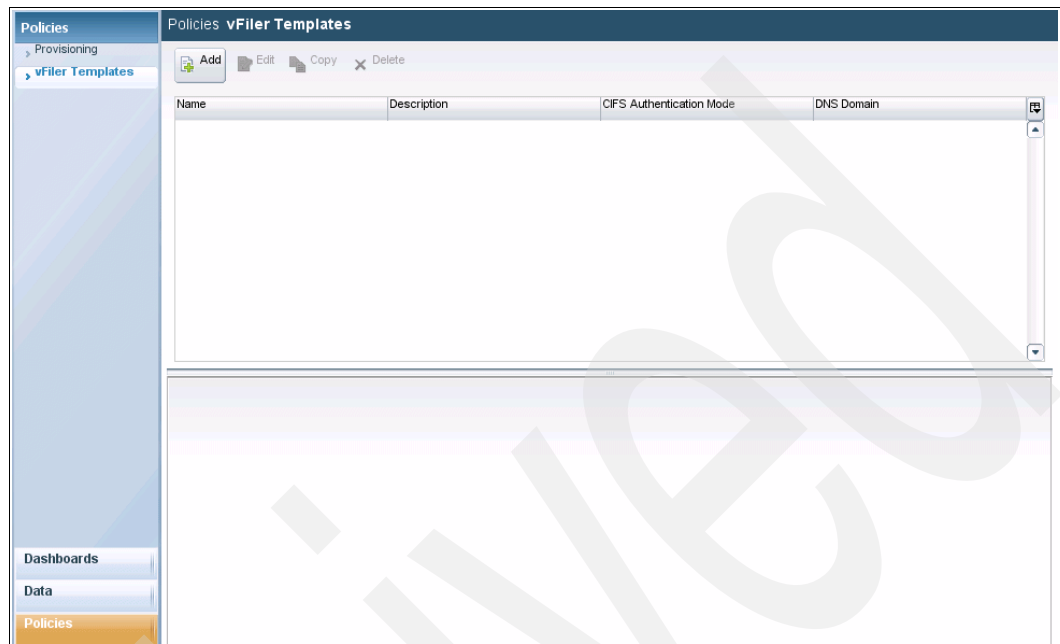


Figure 14-71 Provisioning Manager Policies vFile Templates Add window

2. The Welcome to the Add vFiler Template Wizard window appears, as shown in Figure 14-72. Click **Next**.

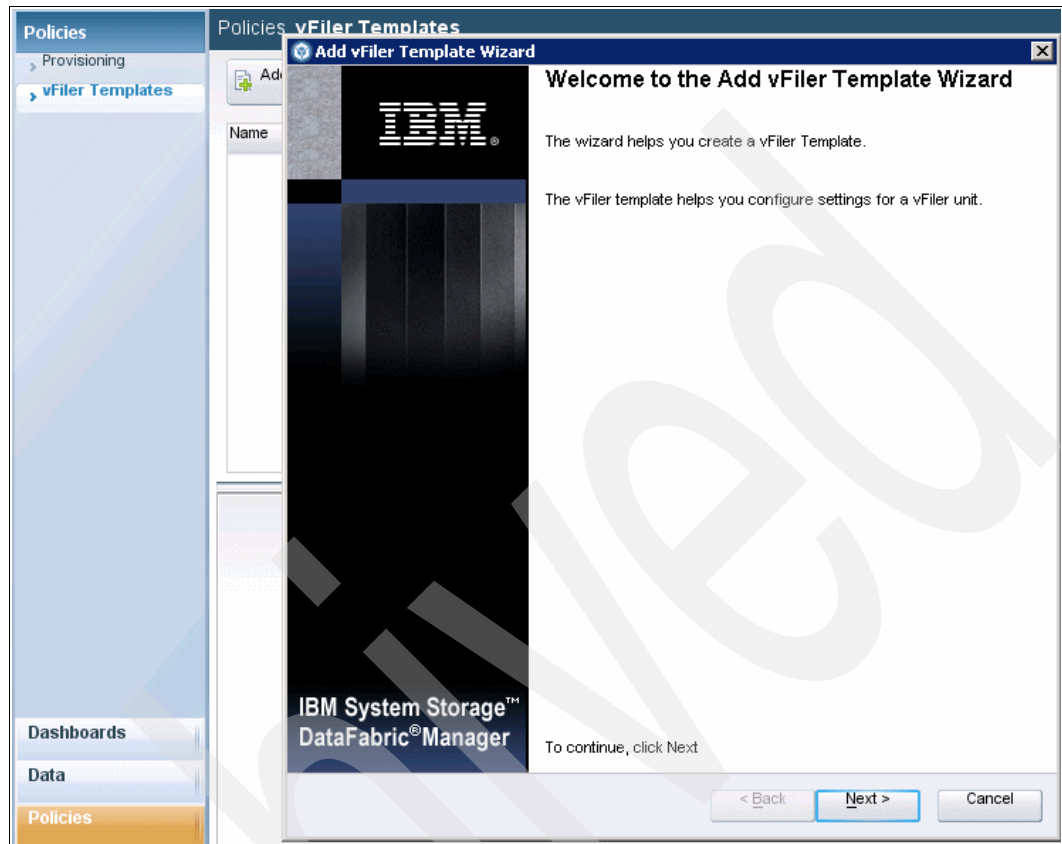


Figure 14-72 Welcome window for Provisioning Manager Policies vFiler Templates

3. The vFiler Template Overview window appears, as shown in Figure 14-73.

The screenshot shows the 'vFiler Template Overview' window within the 'Policies vFiler Templates' section. The window has a title bar that says 'Add vFiler Template Wizard'. Below the title bar, there's a subtitle 'vFiler Template Overview' and a note 'You may specify the details and network settings for the vFiler template.' The main area contains three input fields: 'Name' with the value 'ITSO4vFiler', 'Description' with the value 'vFiler for ITSO4vFiler', and 'Administrative Host' with the value '9 . 11 . 218 . 238'. At the bottom, there's a prompt 'To continue, click Next' and three buttons: '< Back', 'Next >', and 'Cancel'. The 'Next >' button is highlighted with a blue border.

Figure 14-73 Provisioning Manager Policies vFiler Templates Overview

You need to fill in the following fields:

- Name: The name of the selected vFiler template.
- Description: The description provided when the vFiler template was configured.
- Administrative Host: The host name or the IP address of the host that has root access to the files needed to manage a vFiler unit using the selected vFiler template.

Click **Next**.

4. The DNS and NIS Settings window appears, as shown in Figure 14-74.

The screenshot shows the 'Add vFiler Template Wizard' window with the 'DNS and NIS Settings' tab selected. The window is divided into two main sections: DNS and NIS. The DNS section has a 'Domain Name' field with the value 'Domain Controller' and a 'Domain Servers' list containing '9.11.218.102'. The NIS section has a 'Domain Name' field and a 'Domain Servers' list. At the bottom, there is a message 'Max of 3 DNS domain servers can be added' and a 'To continue, click Next' instruction. The left sidebar shows the 'Policies' menu with 'vFiler Templates' selected.

Policies vFiler Templates

Add vFiler Template Wizard

DNS and NIS Settings
You may specify the DNS and NIS settings.

DNS

Domain Name: Domain Controller

Domain Servers: 9.11.218.102 Add Delete

9.11.218.102

NIS

Domain Name:

Domain Servers: 0.0.0.0 Add Delete

Max of 3 DNS domain servers can be added

To continue, click Next

< Back Next > Cancel

Figure 14-74 Provisioning Manager Policies vFiler Templates DNS and NIS Settings

You need to fill in the following fields:

- DNS Domain: The name of the DNS domain to which a vFiler unit using the selected vFiler template belongs.
- NIS Domain: The name of the NIS domain to which a vFiler unit using the selected vFiler template belongs.

Click **Next**.

5. The CIFS Settings window appears, as shown in Figure 14-75.

The screenshot shows the 'Add vFiler Template Wizard' window with the 'CIFS Settings' tab selected. The window has a sidebar on the left with 'Policies' and 'vFiler Templates' sections. The main area contains three fields: 'Security Protocol' set to 'Multiprotocol', 'Authentication Mode' set to 'Active Directory', and 'Domain Name' set to 'Domain Controller'. At the bottom, there are buttons for '< Back', 'Next >', and 'Cancel'. A message at the bottom says 'To continue, click Next'.

Field	Value
Security Protocol	Multiprotocol
Authentication Mode	Active Directory
Domain Name	Domain Controller

Figure 14-75 Provisioning Manager Policies vFiler Templates CIFS Settings

The CIFS Authentication Mode field is the name of the CIFS domain to which a vFiler unit using the selected vFiler template belongs. Valid values are Active Directory, Windows Workgroup, DNS Domain, the name of the domain, or the IP address of the PDC.

6. The CIFS Security Protocol setting window appears, as shown in Figure 14-76.

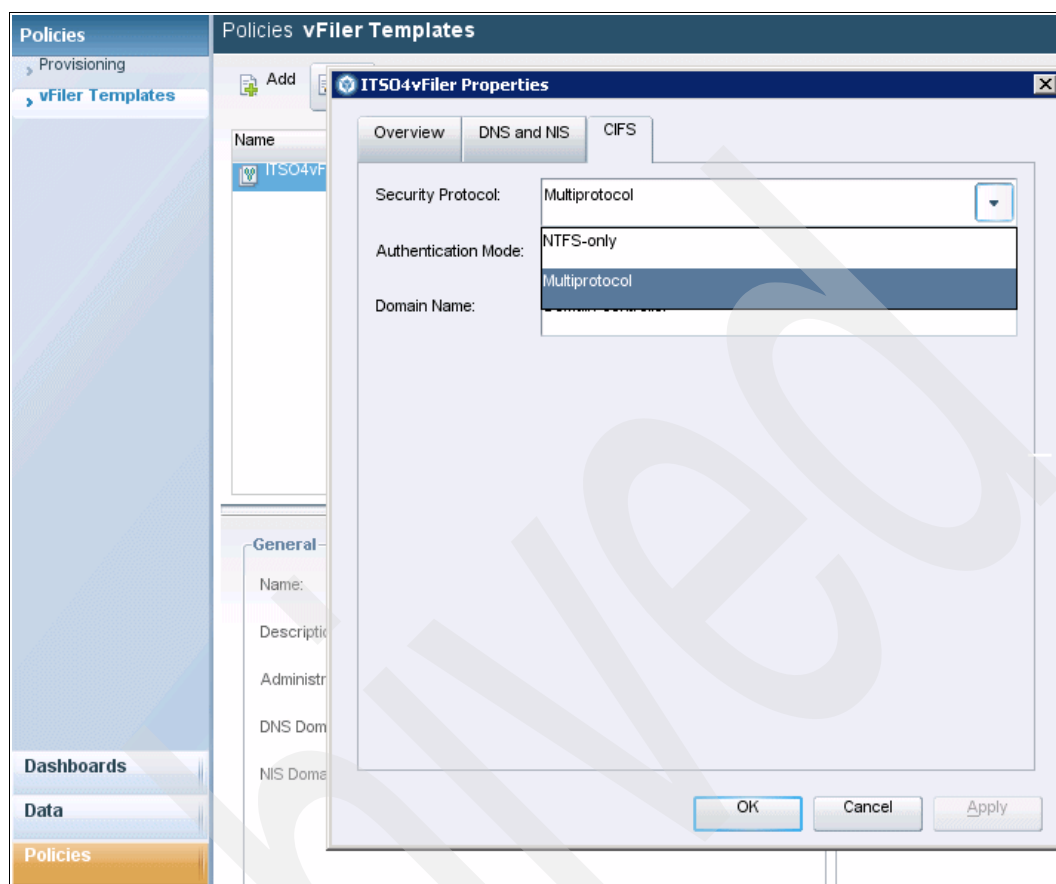


Figure 14-76 Provisioning Manager Policies selected vFiler Templates properties

You need to fill in the following fields:

- Security Protocol: The security protocol to be used by a vFiler unit using this vFiler template. Valid values are NTFS-only or Multiprotocol. We use Multiprotocol.
- Authentication Mode: The mode used to authenticate data requests to or from a vFiler unit using this vFiler template. Valid values are Active Directory or Windows workgroup. We use Active Directory.
- CIFS Domain: The name of the CIFS domain to which a vFiler unit using the selected vFiler template belongs to. We use Active Directory.

Click **Apply**, as shown in Figure 14-77.

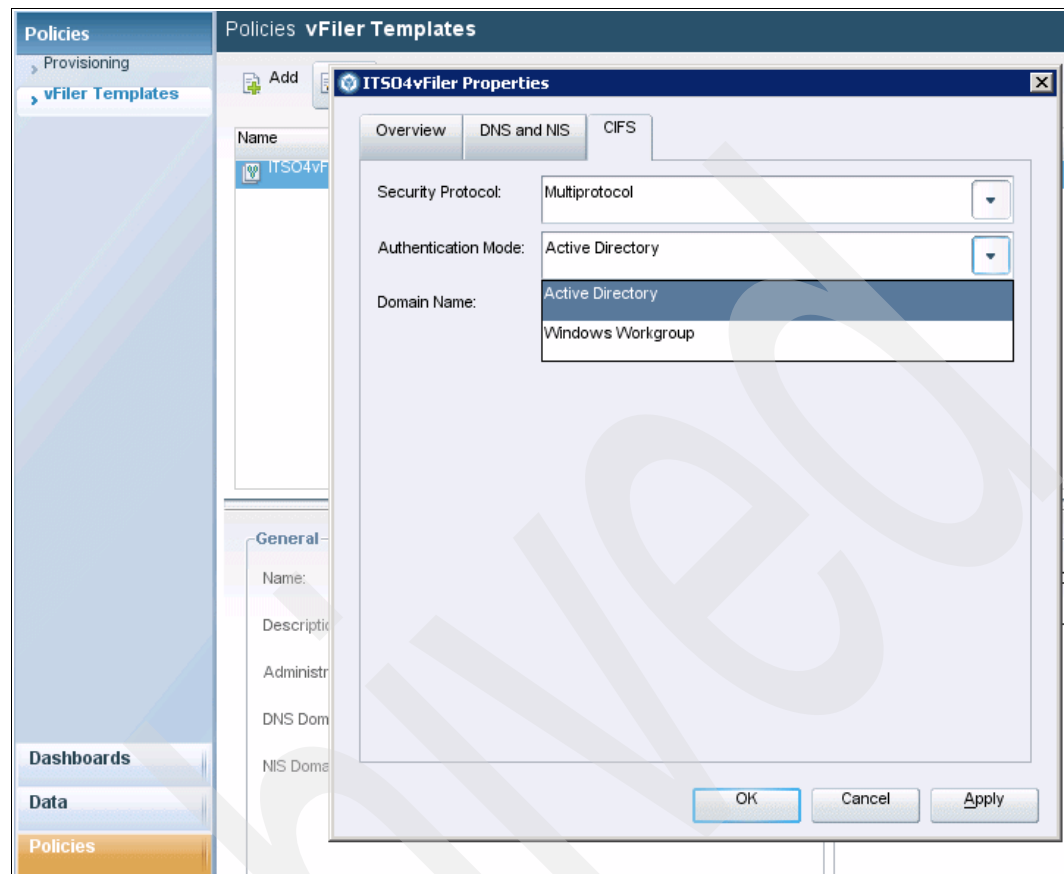


Figure 14-77 Provisioning Manager Policies selected vFiler Templates properties CIFS settings

7. The Completing the Add vFiler Template Wizard window appears, as shown in Figure 14-78. Click **Finish**.

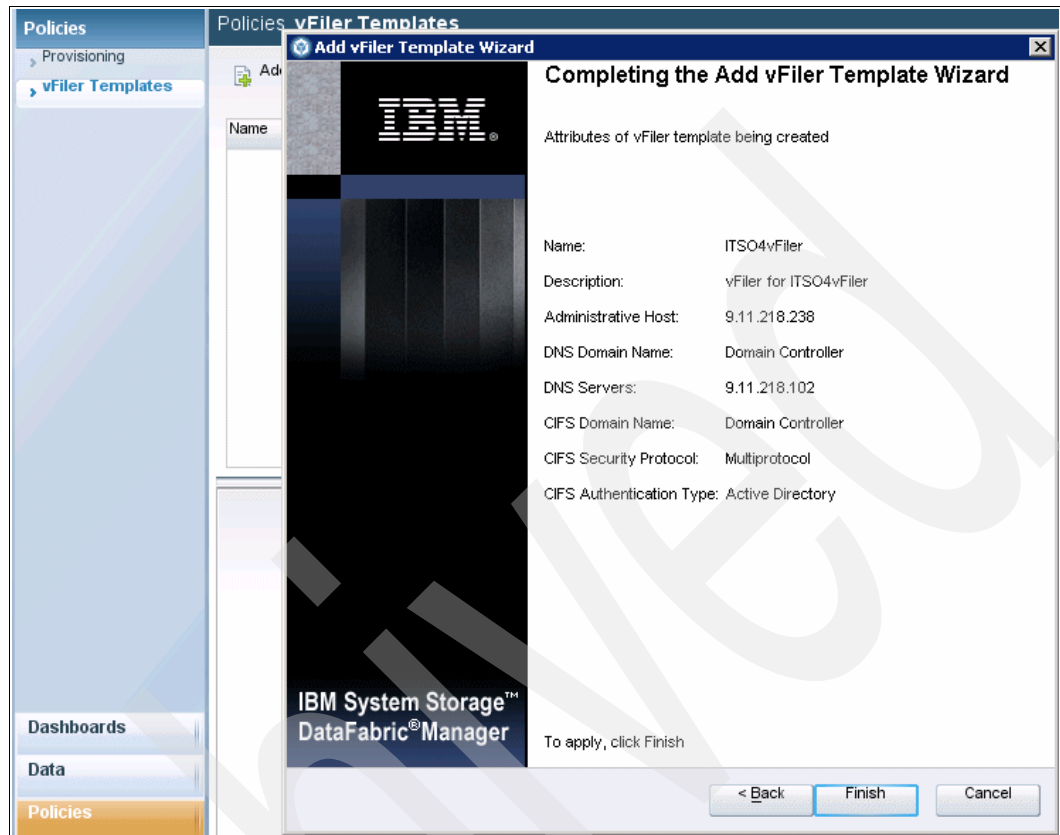


Figure 14-78 Completing Provisioning Manager policies vFiler Templates wizard

Select **Policies** → **vFile Templates** to view the vFile details after a successful configuration, as shown in Figure 14-79.

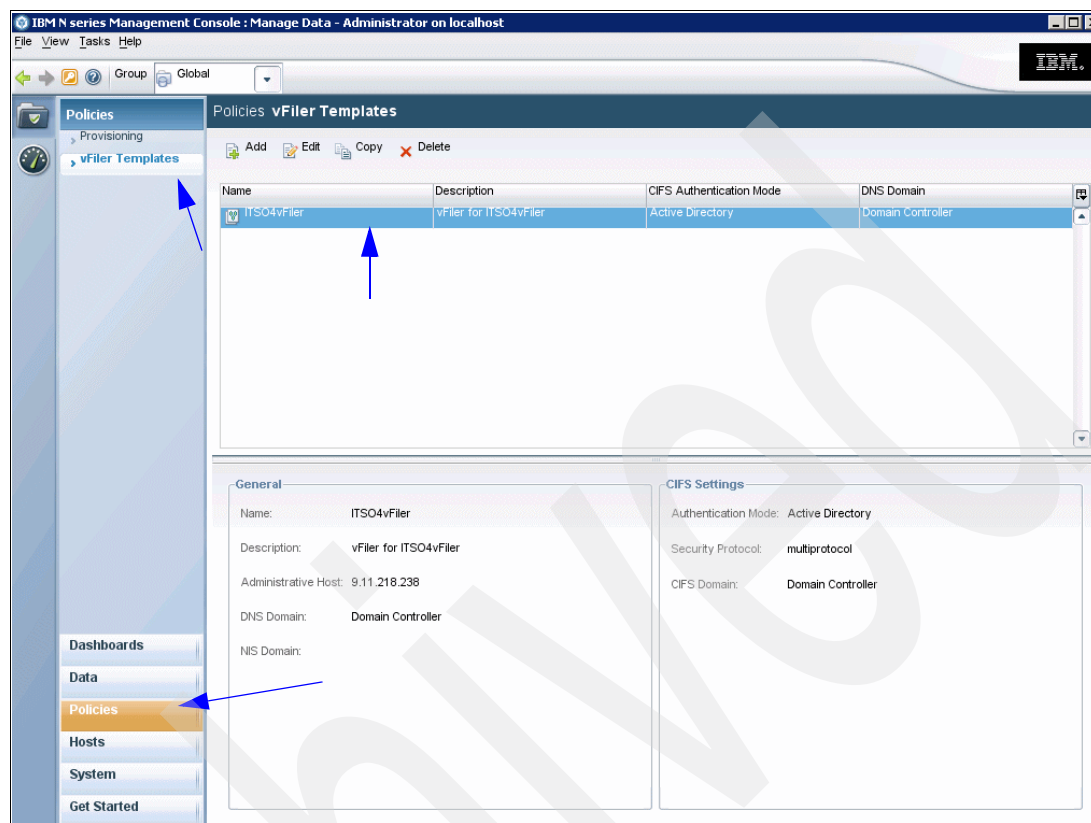


Figure 14-79 Provisioning Manager policies vFile Templates view

14.5.5 Host storage system information

You can use the Hosts Storage Systems window to view detailed information about storage systems discovered by Operations Manager. From this window, you can add a storage system to Management Console, edit the properties on existing storage, and diagnose a storage system's configuration. You can also manage Data ONTAP service licenses.

For information about configuration, refer to “Hosts Storage Systems configuration” on page 516.

Select **Hosts** → **Storage Systems** and click the **Details** tab to view the list of the N series storage systems from Provisioning Manager through Operations Manager, as shown in Figure 14-80.

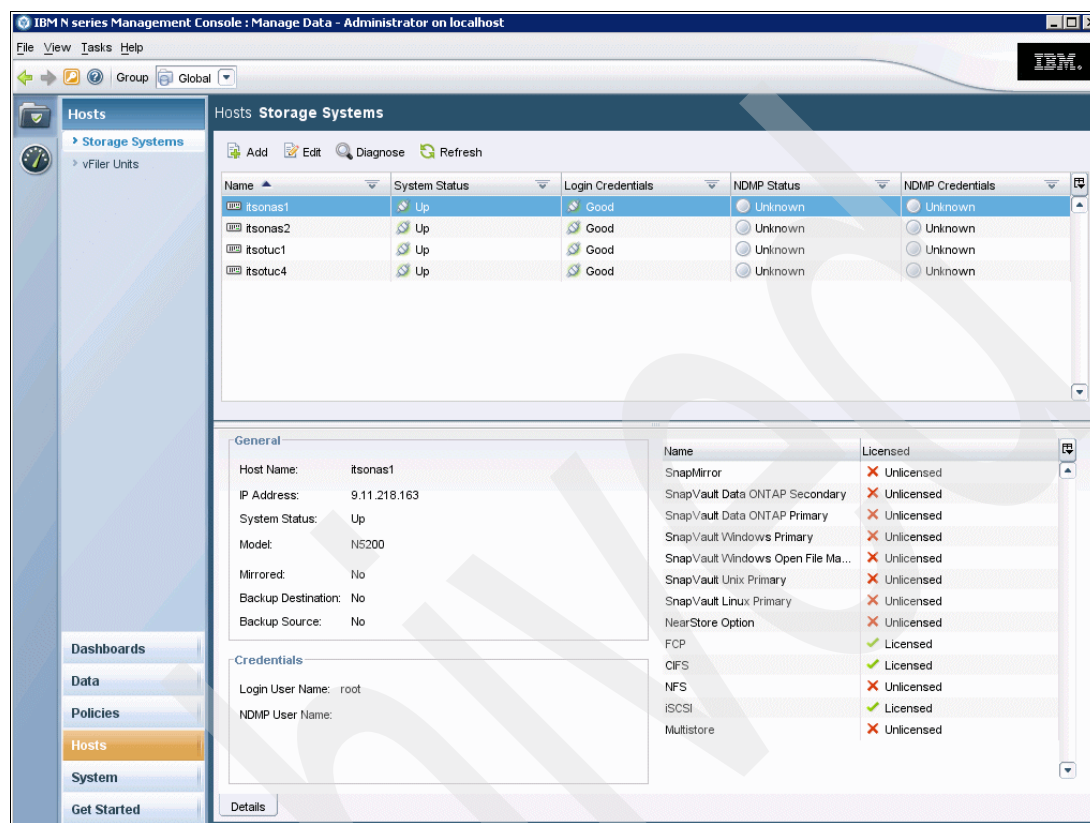


Figure 14-80 Provisioning Manager Hosts Storage Systems details

This window contains the following items:

- ▶ **Name:** Displays the name of the storage system or vFiler unit as it appears in the Operations Manager database.
- ▶ **Details tab:** The Details tab has three areas: General, Credentials, and Licenses. You can see the following items in the General pane:
 - **Host Name:** The name of the currently selected storage system.
 - **IP Address:** The IP address associated with the selected storage system.
 - **System Status:** Whether the status of the host is Up, Down, Unknown, or Not Applicable (N/A).
 - **Model:** Displays the model number of this storage system.
 - **Mirrored:** Indicates whether the SnapMirror license is enabled on this host. Possible values are Yes or No.
 - **Backup Destination:** Indicates whether a SnapVault Secondary license is enabled on this host, making the host a potential destination for backups. Possible values are Yes or No.
 - **Backup Source:** Indicates whether the SnapVault Data ONTAP Primary license is enabled on this host, making the host a potential source of backups. Possible values are Yes or No.

Note: Other hosts might also be potential backup sources, but the SnapVault Primary licenses for Linux, UNIX, and Windows are installed on the secondary system. This list confirms only whether the SnapVault Data ONTAP Primary license is enabled on the host.

Hosts Storage Systems configuration

This section provides the steps for configuring Hosts Storage Systems:

1. Select **Hosts** → **Storage Systems**, select itsonas1, and click the **Diagnose** button, as shown in Figure 14-81.

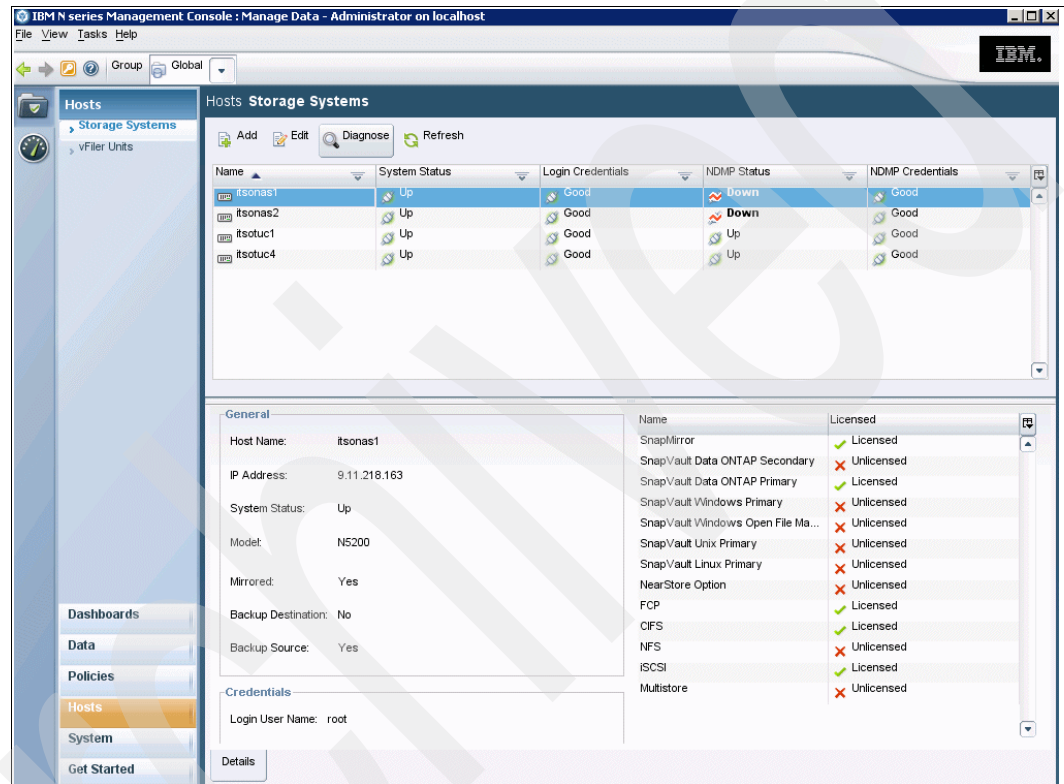


Figure 14-81 Provisioning Manager Hosts Storage Systems diagnose

You can see the following systems in Figure 14-81:

- itsonas1 and itsonas2 (in the Active/Active Cluster)
- itsotuc1
- itsotuc4

We show the configuration for itsonas1 first and will repeat the same steps for itsonas2 with some modifications.

Note the following options at the top of the window:

- Add: Starts the Add Storage System Wizard that allows you to set up storage system hosts.
- Edit: Opens a window in which you can modify the properties of the selected host.
- Diagnose: Starts the Diagnose Storage System wizard, which allows you to modify some aspects of a storage system's configuration. This button is disabled if more than one host is selected in the hosts list. This is the option we are using in this configuration.

- Refresh: Pulls updated data into the host list for the selected host only.
2. The Diagnose Storage System Welcome window appears, as shown in Figure 14-82. Click **Next**.

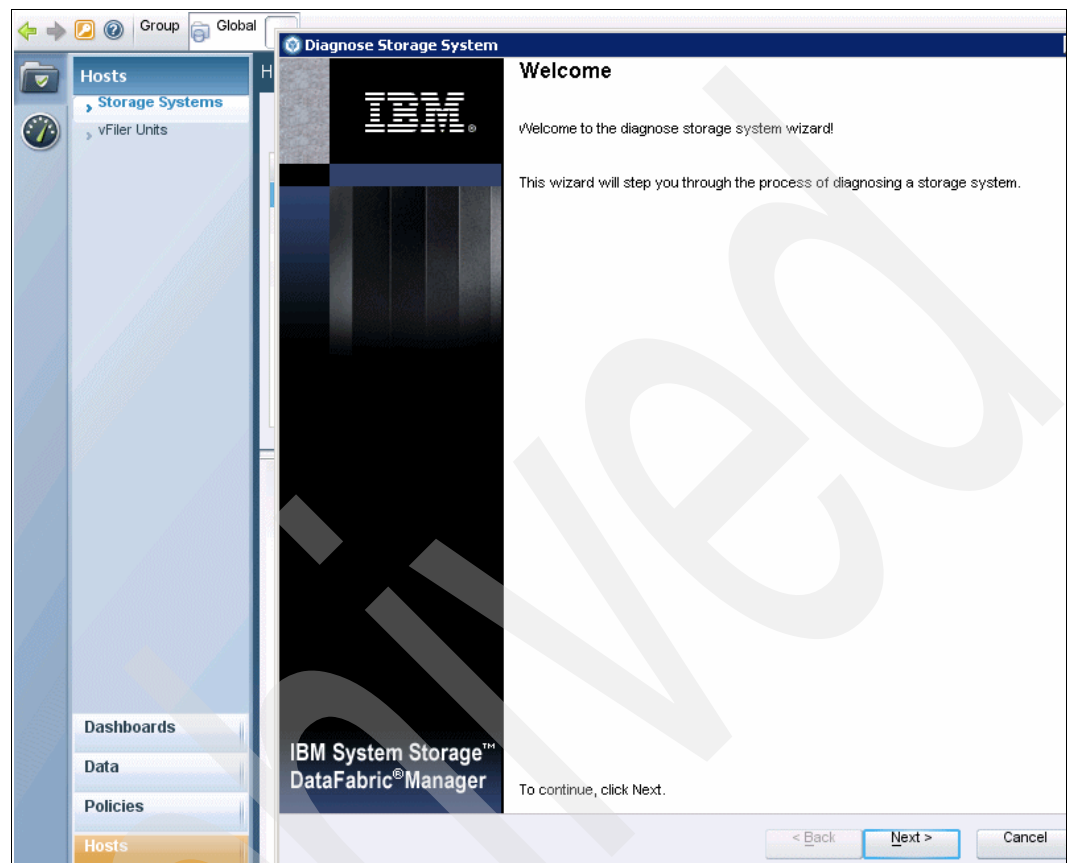


Figure 14-82 Provisional Manager Host Storage Systems diagnose storage system welcome window

3. The System Storage Information window appears, as shown in Figure 14-83. This window displays information about the vFile, such as Host Name, Domain Name, IP Address, Host Status, Host Login Credentials, NDMP Enabled/Disabled, NDMP Status, NDMP Credentials, Enabled Licenses, SnapVault Access Control List, and SnapMirror Access Control List.

We are interested in the Host Status field, which displays the current status of the storage system. Possible values are Up, Down, and Unknown. The default monitoring interval is one minute. The interval is specified as the Ping Monitoring Interval in the DFM server.

The screenshot shows a web-based interface for managing storage systems. On the left is a navigation pane with links for Hosts, Storage Systems, vFile Units, Dashboards, Data, Policies, and Hosts. The main window is titled 'Diagnose Storage System' and contains a 'Storage System Information' section. This section displays various attributes for a storage system identified as 'itsonas1'. The attributes and their values are: Host Name (itsonas1), Domain Name (itsonas1), IP Address (9.11.218.163), Host Status (Up), Host Login Credentials (Good), NDMP Enabled (No), NDMP Status (Down), NDMP Credentials (Good), Enabled Licenses (SnapMirror), SnapVault Access Control List (none), and SnapMirror Access Control List (none). At the bottom of the window, there is a prompt 'To continue, click Next.' and three buttons: '< Back', 'Next >', and 'Cancel'.

Field	Value
Host Name	itsonas1
Domain Name	itsonas1
IP Address	9.11.218.163
Host Status	Up
Host Login Credentials	Good
NDMP Enabled	No
NDMP Status	Down
NDMP Credentials	Good
Enabled Licenses	SnapMirror
SnapVault Access Control List	none
SnapMirror Access Control List	none

Figure 14-83 Provisioning Manager Hosts Storage Systems Storage System Information window

Click **Next**.

4. The Login Credentials window appears, as shown in Figure 14-84. Specify the User Name and the Password.

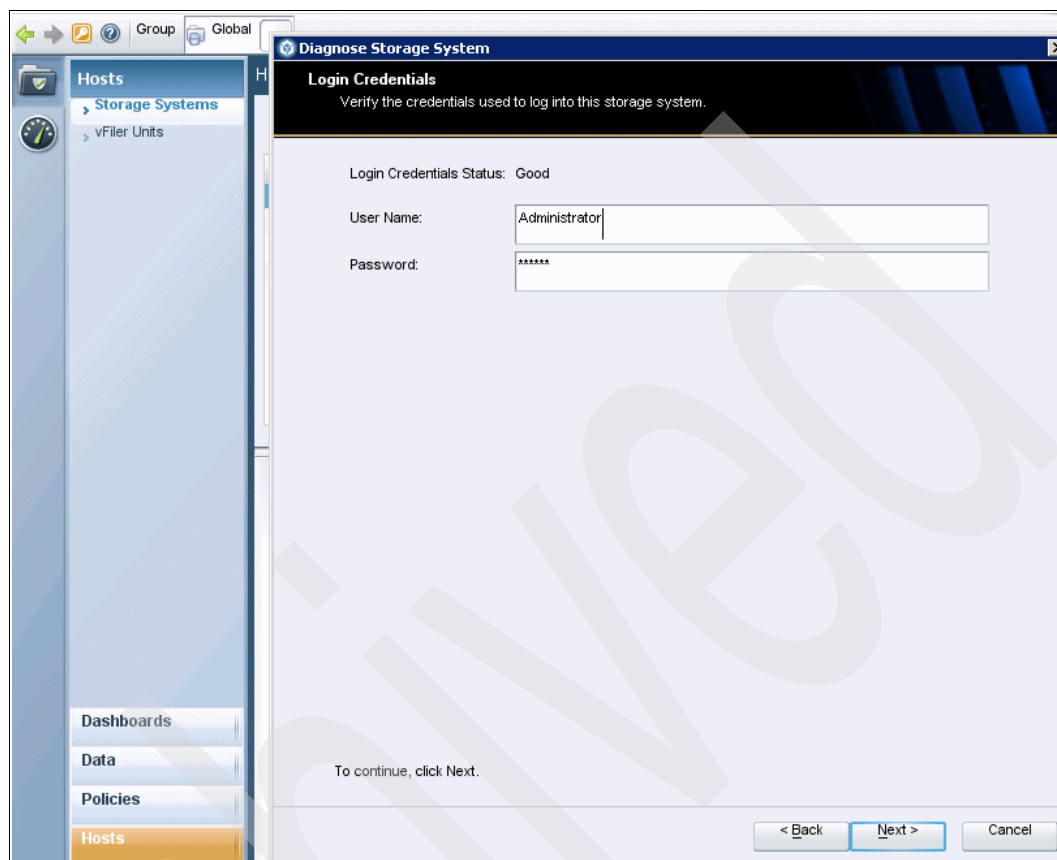


Figure 14-84 Provisioning Manager Hosts Storage Systems Login Credentials

You should note the following information about the fields in this window:

- Login Credentials Status: Displays the current status of the login credentials that Operations Manager uses to log in to the host. Possible values are Good, Bad, Read Only, Unknown, and Not Applicable. NDMP credentials for vFilter units are designated Not Applicable because Operations Manager uses the credentials of the hosting system.
- Login User Name: The name that Operations Manager uses to log into the selected host ITSONAS1.
- NDMP User Name: The name that Operations Manager uses to log in to the selected host by using NDMP.

Click **Next**.

5. The Licences window appears, as shown in Figure 14-85. You must make sure that the SnapMirror and SnapVault Data ONTAP primary licenses are enabled. If both licenses are not enabled, you can add a new license in the New License window.

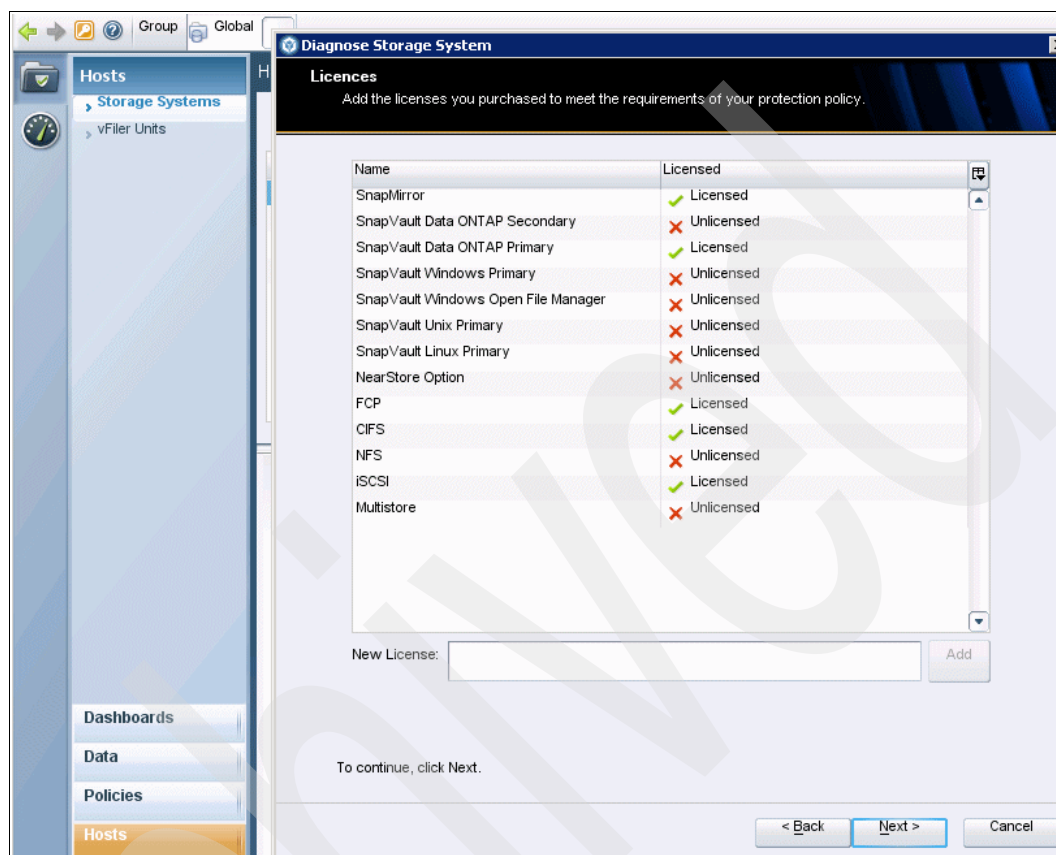


Figure 14-85 Provisioning Manager Hosts Storage Systems primary storage system licenses

The licenses shown in this window each enable a service of Data ONTAP. You must enter a software license code on the storage system to enable the corresponding service. You do not need to indicate which license the code enables. The code is matched automatically to the appropriate service license.

- SnapMirror: You should add the SnapMirror licenses on both the source and destination storage systems for the mirrored data. If the source and destination volumes are on the same system, only one license is required. SnapMirror replicates data to one or more networked storage systems. SnapMirror updates the mirrored data to keep it current and available for disaster recovery, offloading tape backup, read-only data distribution, testing on nonproduction systems, online data migration, and so on. You can also enable the SnapMirror license to use Qtree SnapMirror for backup.
- SnapVault Data ONTAP Secondary: You install the SnapVault Secondary license on storage systems hosting the backups of protected data. SnapVault creates backups of data stored on multiple primary storage systems and copies the backups to a secondary storage system. If data loss or corruption occurs, backed up data can be restored to a primary or open storage system with little of the downtime and uncertainty associated with conventional tape backup and restore operations.
- SnapVault Linux Primary: You install the SnapVault Linux Primary license on a secondary storage system, in addition to the SnapVault Secondary license, to support a Linux-based primary storage system running the Open Systems SnapVault agent. A

Linux-based primary storage system running the Open Systems SnapVault agent does not require a SnapVault license.

- NearStore Option: The NearStore license enables your storage system to use transfer resources as conservatively as though it were optimized as a backup system. This approach is useful when the storage system on which you want to store backed up data is not a system optimized for storing backups, and you want to minimize the number of transfer resources the storage system requires.

Storage systems using the NearStore license must meet the following criteria:

- The storage system must be a N5000, N6000, or N7000 series system.
- The version of Data ONTAP software must be Version 7.1 or later.

If you are using the SnapVault service, the storage system must have a SnapVault secondary license enabled.

- FCP: Fibre Channel Protocol (FCP) is a licensed service on the storage system that enables you to export LUNs and transfer block data to hosts using the SCSI protocol over a Fibre Channel fabric.
- CIFS: The Common Internet File System (CIFS) protocol is a licensed service for remote file access that runs over TCP/IP on the Windows operating system. It enables application access and file sharing across the Internet.
- NFS: Network File System (NFS) is client/server application that runs over TCP/IP on the UNIX operating system. It enables application access and file sharing across the Internet.
- iSCSI: The iSCSI protocol is a licensed service on the storage system that enables you to transfer block data to hosts using the SCSI protocol over TCP/IP. iSCSI supports Gigabit Ethernet and is often used in a SAN environment.
- MultiStore: The Data ONTAP MultiStore license enables you to partition the storage and network resources of a single storage system so that it appears as multiple storage systems, called vFilers, in the network. Be sure the host on which you intend to install the MultiStore license is running Data ONTAP Version 7.1 or later.

Click **Next**.

6. The SnapVault Settings window appears, as shown in Figure 14-86.

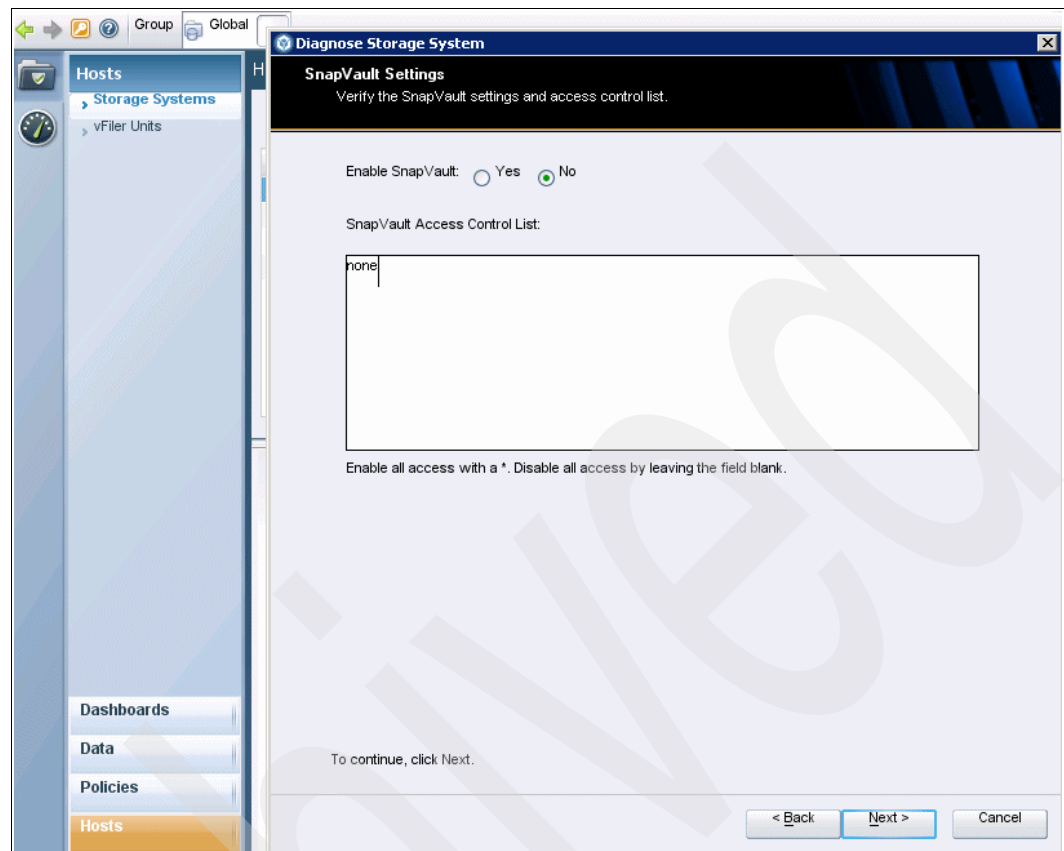


Figure 14-86 Provisioning Manager Hosts Storage Systems SnapVault Settings window

7. In this SnapVault Settings window, make sure that the Enable SnapVault setting is set and that SnapVault Access Control List is set to host-itsonas2, as shown in Figure 14-87. Click **Next**.

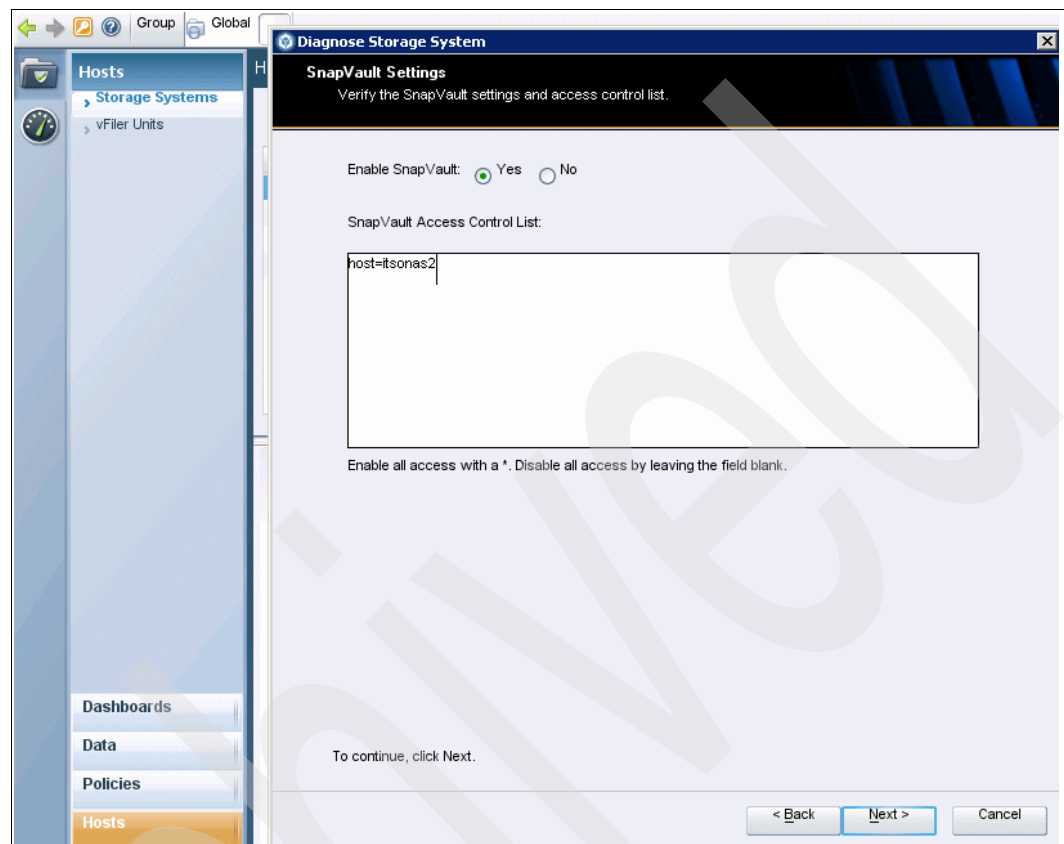


Figure 14-87 Provisioning Manager Hosts Storage Systems SnapVault enabled

8. The SnapMirror Settings window appears, as shown in Figure 14-88.

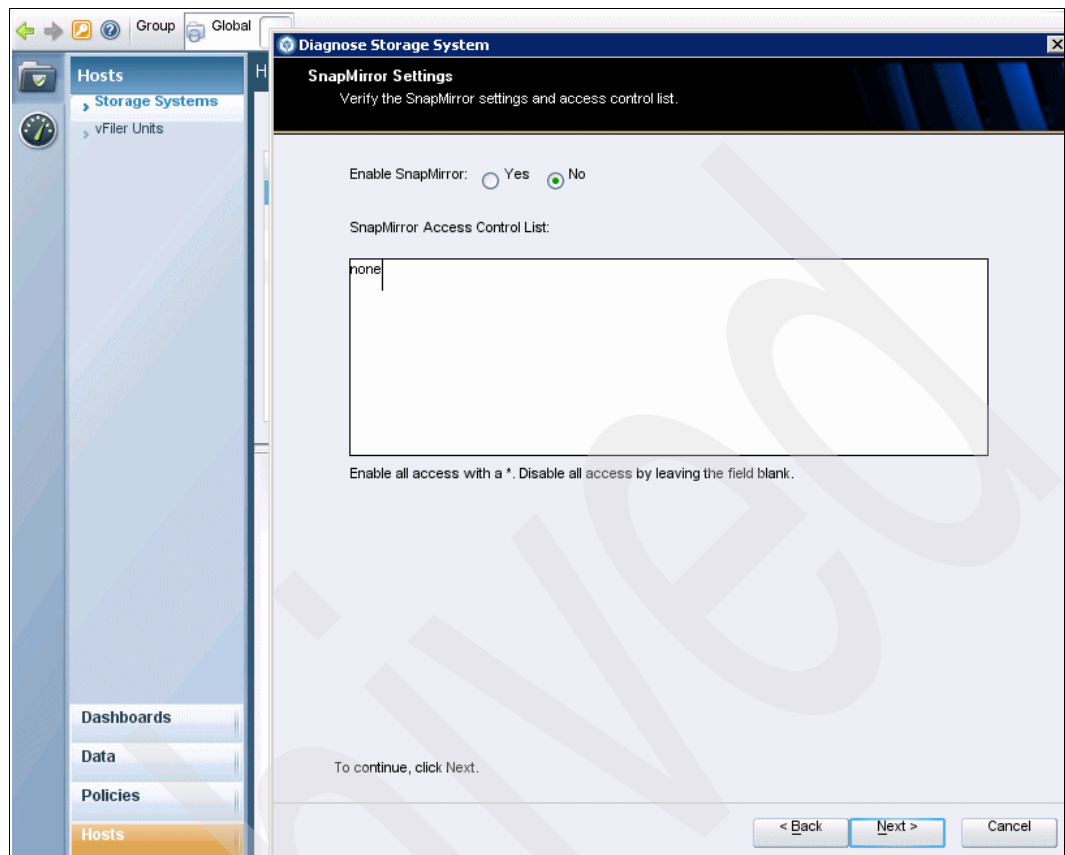


Figure 14-88 Provisioning Manager Hosts Storage Systems SnapMirror Settings

9. In the SnapMirror Settings window, make sure SnapMirror is enabled and that the SnapMirror Access Control List is set to host-itsonas2, as shown in Figure 14-89. Click **Next**.

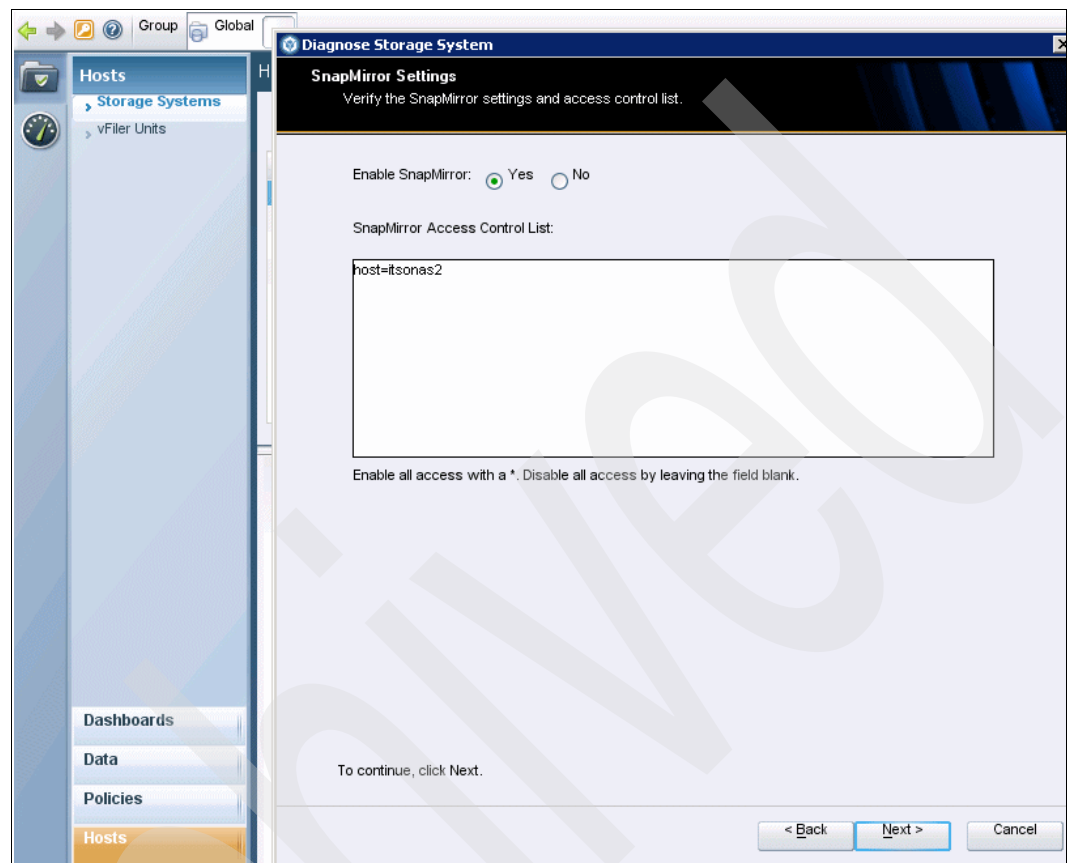


Figure 14-89 Provisioning Manager Hosts Storage Systems SnapMirror enabled

10. The NDMP Settings window appears, as shown in Figure 14-90.

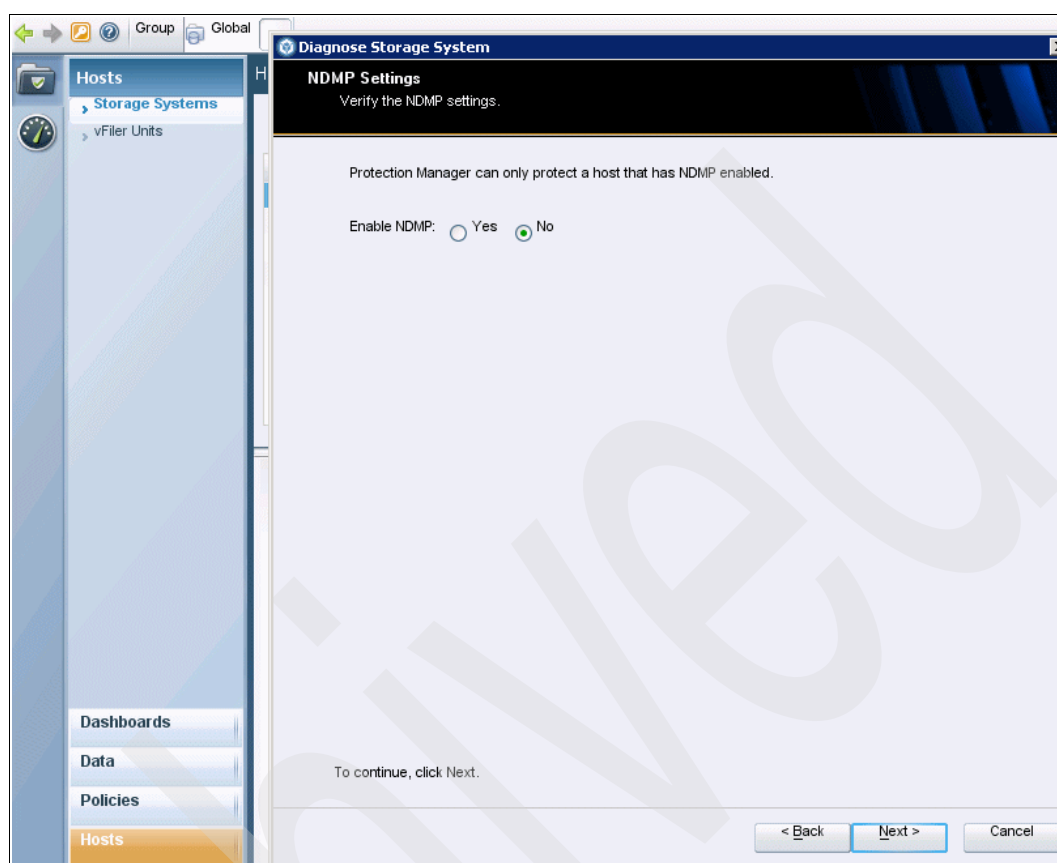


Figure 14-90 Provisioning Manager Hosts Storage Systems NDMP Settings

11. Make sure that NDMP is enabled, as shown in Figure 14-91. Click **Next**.

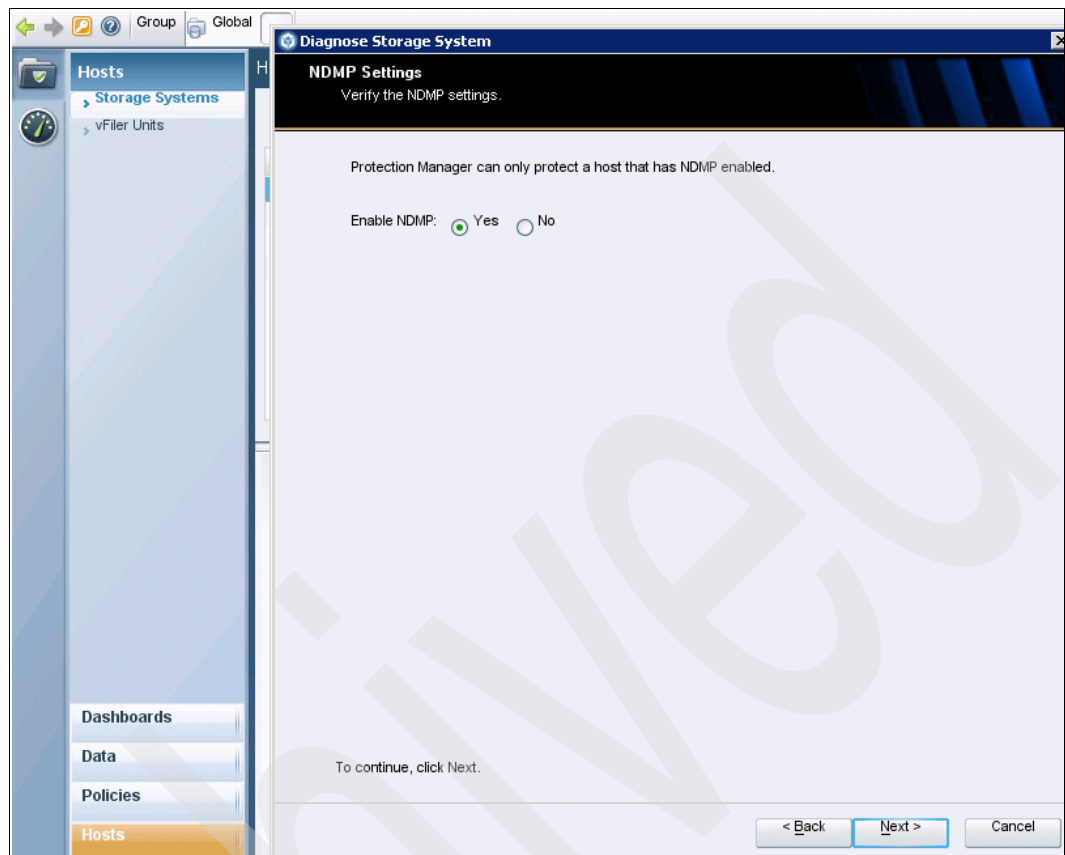


Figure 14-91 Provisioning Manager Hosts Storage Systems NDMP enabled

12.The NDMP Credentials window appears, as shown in Figure 14-92.

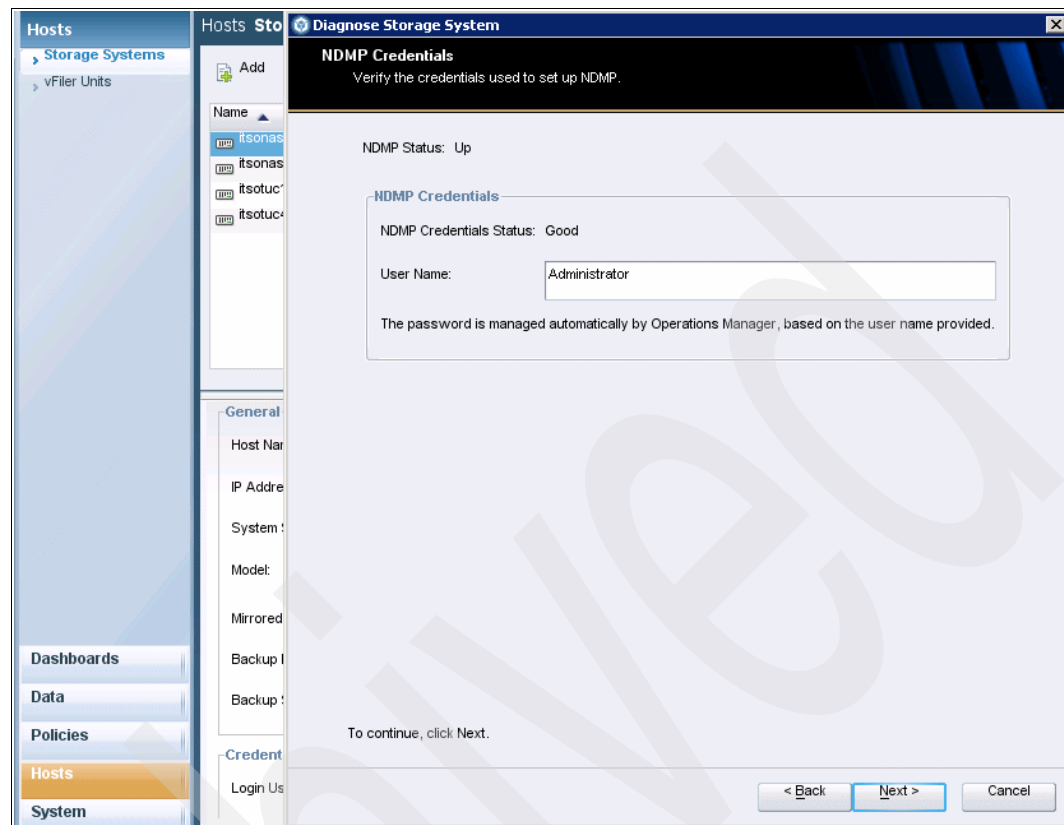


Figure 14-92 Provisioning Manager Hosts Storage Systems NDMP Credentials

Note the information in the following fields:

- NDMP Status: This is the Network Data Management Protocol (NDMP) availability of the storage system as of the most recent NDMP monitoring check. Possible values are Up, Down, and Unknown. The default interval for NDMP monitoring is 30 minutes. You can use the Backup Discovery Options window in Operations Manager to change the NDMP monitoring interval.
- NDMP Credentials: Displays the current status of the Network Data Management Protocol credentials that Operations Manager uses to communicate with the host. Possible values are Good, Bad, Unknown, and Not Applicable. NDMP credentials for vFiler units are always designated Not Applicable because Operations Manager uses the credentials of the hosting system. You can use the NDMP Credentials window in Operations Manager to edit the credentials for NDMP discovery.

Click **Next**.

13. The Completing the Diagnose Storage System Wizard window appears, as shown in Figure 14-93) Click **Finish** to complete the configuration.

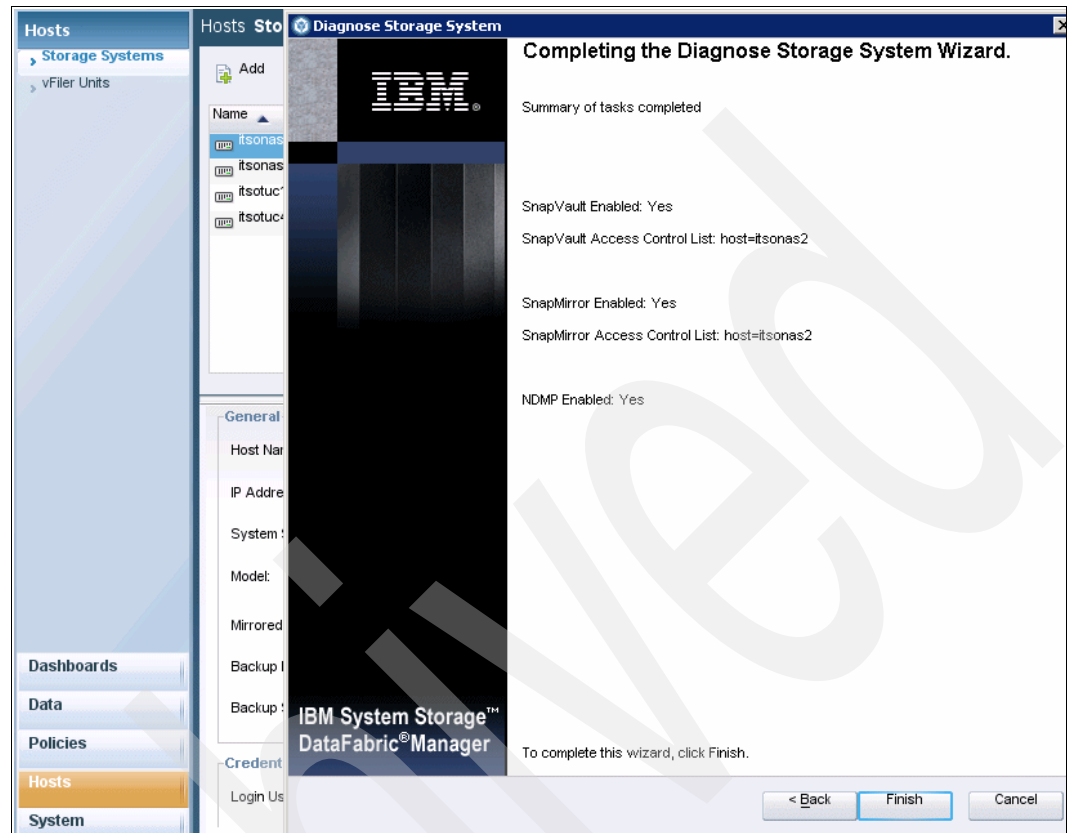


Figure 14-93 Completing the Diagnose Storage System Wizard window

To diagnose itsonas2, select **Hosts** → **Storage Systems**, select itsonas2, and click **Diagnose**, which will start the Diagnose Storage System Wizard, as shown in Figure 14-94. The steps to diagnose itsonas2 is the same as for itsonas1, except for the changes noted in the following steps.

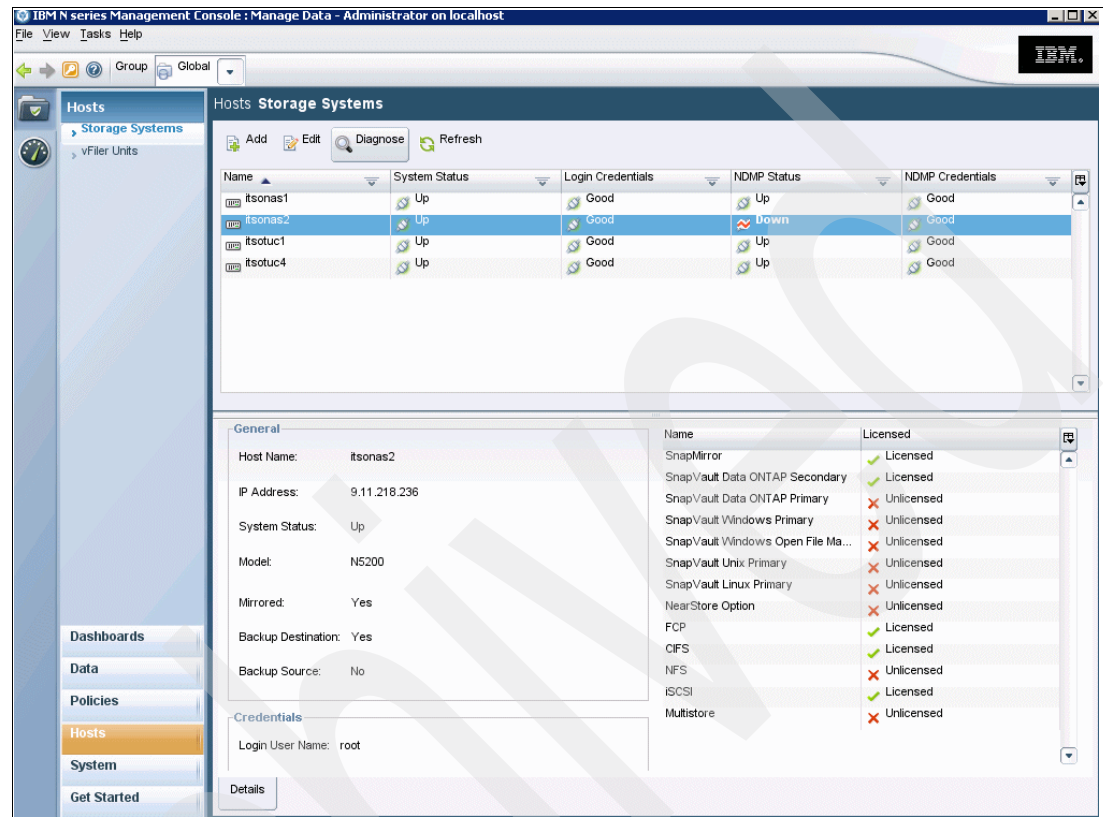


Figure 14-94 Provisioning Manager Hosts Storage Systems Diagnose button

The changed steps are as follows:

1. In the Licences window, shown in Figure 14-95, a SnapMirror license and SnapVault Data ONTAP Secondary license are required. If both licenses are not enabled, you have the option to add new license using the New License field. Click **Next**.

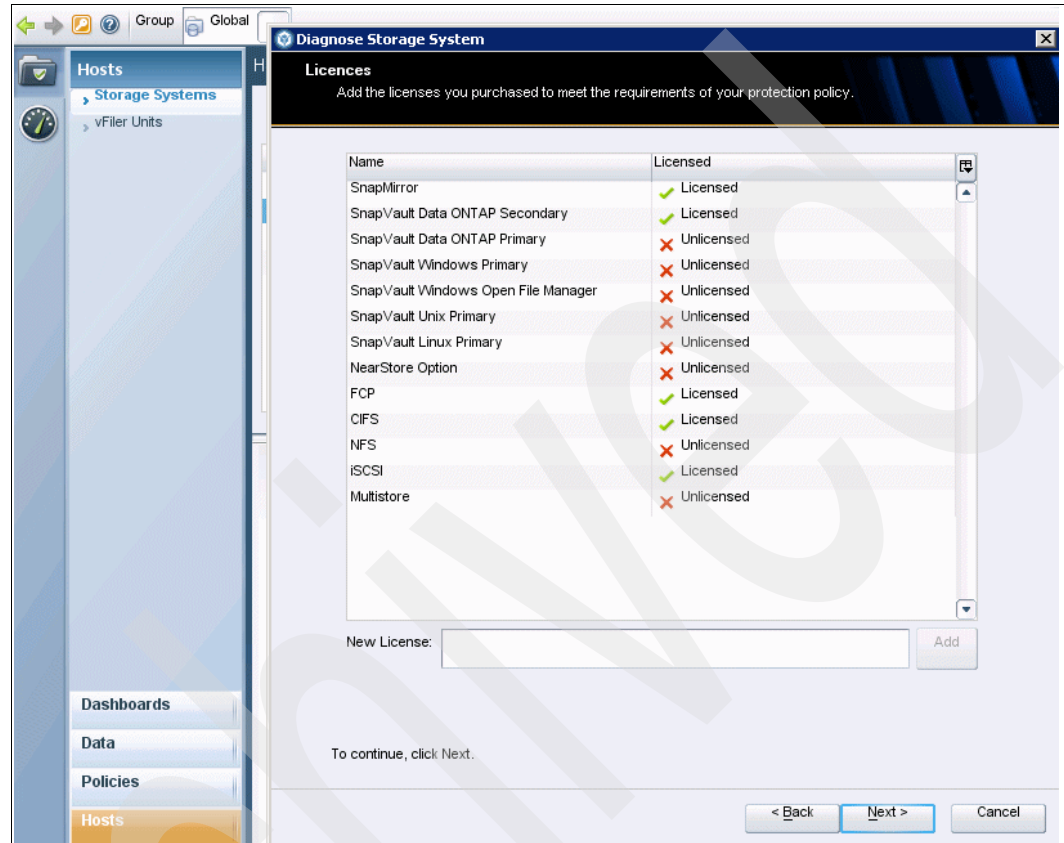


Figure 14-95 Provisioning Manager Hosts Storage Systems Licences for secondary storage

2. In the SnapVault Settings window, SnapVault is enabled and the SnapVault Access Control List is set to host=itsonas1, as shown in Figure 14-96. Click **Next**.

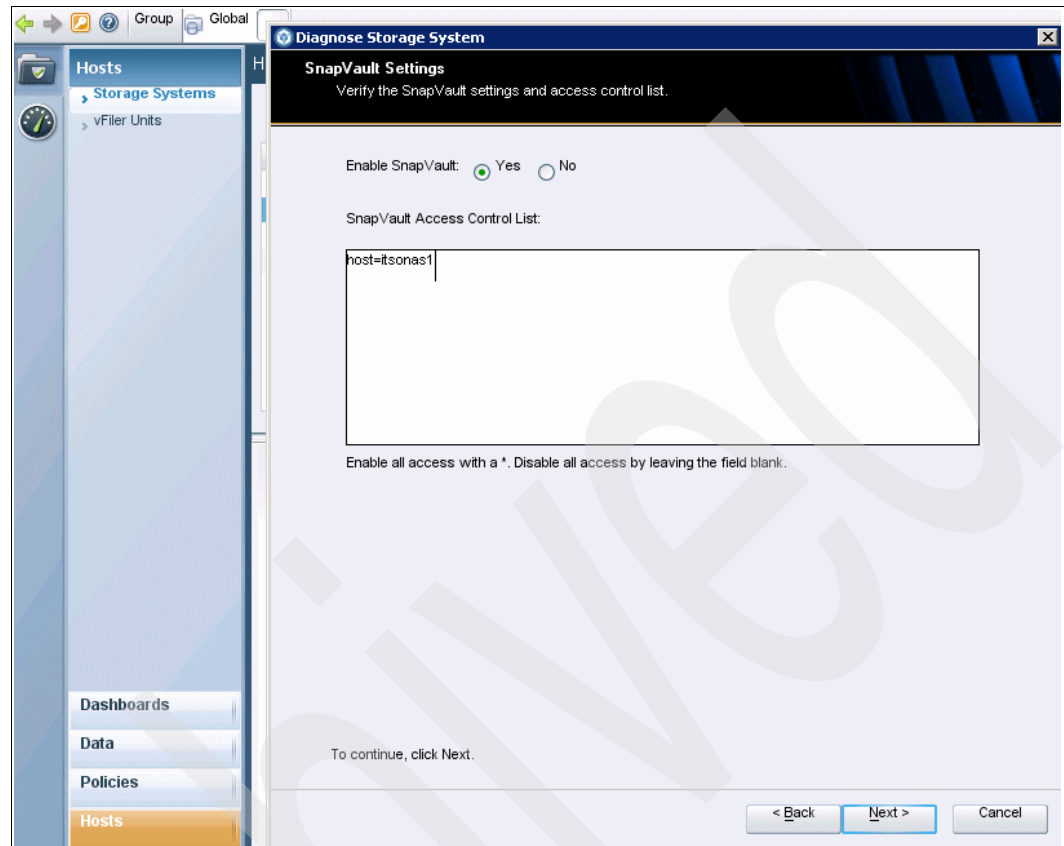


Figure 14-96 Provisioning Manager Hosts Storage Systems SnapVault enabled for secondary storage

3. In the SnapMirror Settings window, shown in Figure 14-97, SnapMirror is enabled and the SnapMirror Access Control List is set to host=itsonas1. Click **Next**.

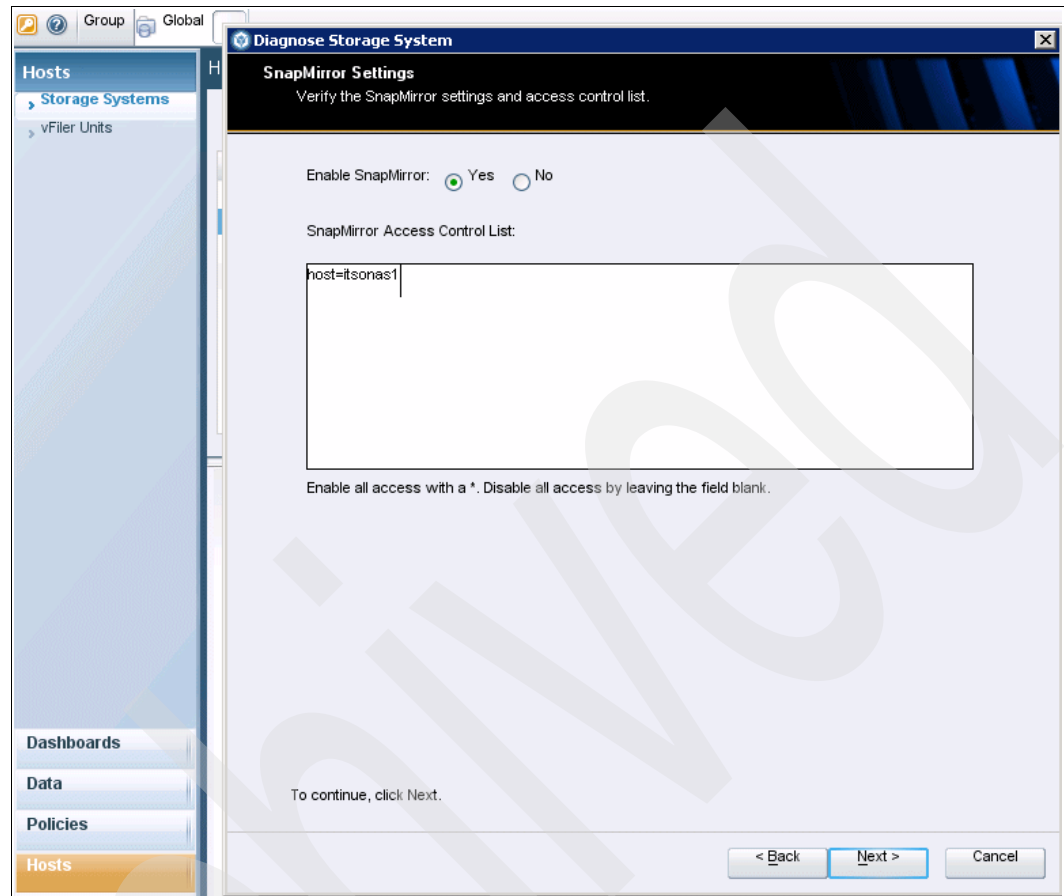


Figure 14-97 Provisioning Manager Hosts Storage Systems SnapMirror enabled for secondary storage

4. The Completing the Diagnose Storage Wizard window appears, as shown in Figure 14-98. Click **Finish**.

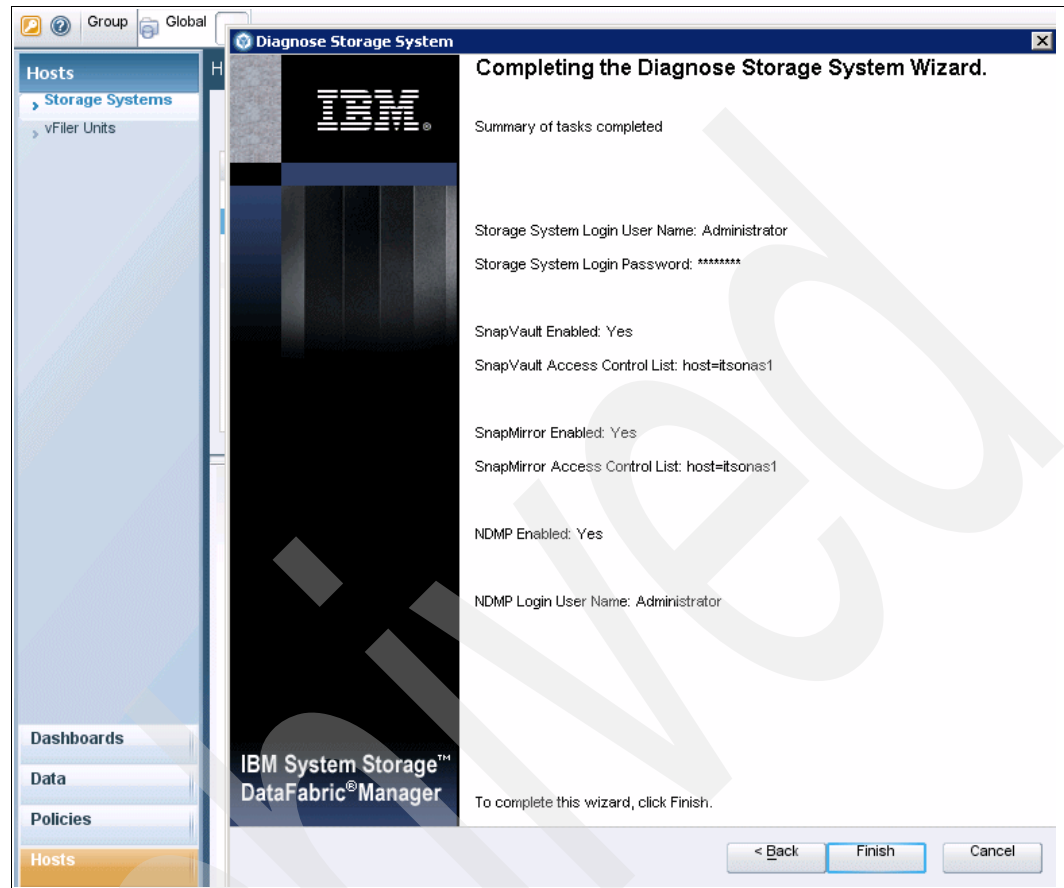


Figure 14-98 Completing the Diagnose Storage System Wizard for secondary storage

Select **Hosts** → **Storage Systems**, and you can see that itsonas1 and the itsonas2 are up, as shown in Figure 14-99.

Hosts Storage Systems

Add Edit Diagnose Refresh

Name	System Status	Login Credentials	NDMP Status	NDMP Credentials
itsonas1	Up	Good	Up	Good
itsonas2	Up	Good	Up	Good
itsotuc1	Up	Good	Up	Good
itsotuc4	Up	Good	Up	Good

General

Host Name: itsonas2

IP Address: 9.11.218.236

System Status: Up

Model: NS200

Mirrored: Yes

Backup Destination: Yes

Backup Source: No

Credentials

Login User Name: Administrator

Details

Name	Licensed
SnapMirror	✓ Licensed
SnapVault Data ONTAP Secondary	✓ Licensed
SnapVault Data ONTAP Primary	✗ Unlicensed
SnapVault Windows Primary	✗ Unlicensed
SnapVault Windows Open File Ma...	✗ Unlicensed
SnapVault Unix Primary	✗ Unlicensed
SnapVault Linux Primary	✗ Unlicensed
NearStore Option	✗ Unlicensed
FCP	✓ Licensed
CIFS	✓ Licensed
NFS	✗ Unlicensed
iSCSI	✓ Licensed
MultiStore	✗ Unlicensed

Navigation: Dashboards, Data, Policies, **Hosts**, System, Get Started

Figure 14-99 Provisioning Manager Hosts Storage Systems view for secondary storage

14.5.6 Host vFiler Units information

This section shows the options for Host vFiler Units configuration.

To open the vFiler Units details, select **Hosts** → **vFilers Units** and click the **Details** button, as shown in Figure 14-100.

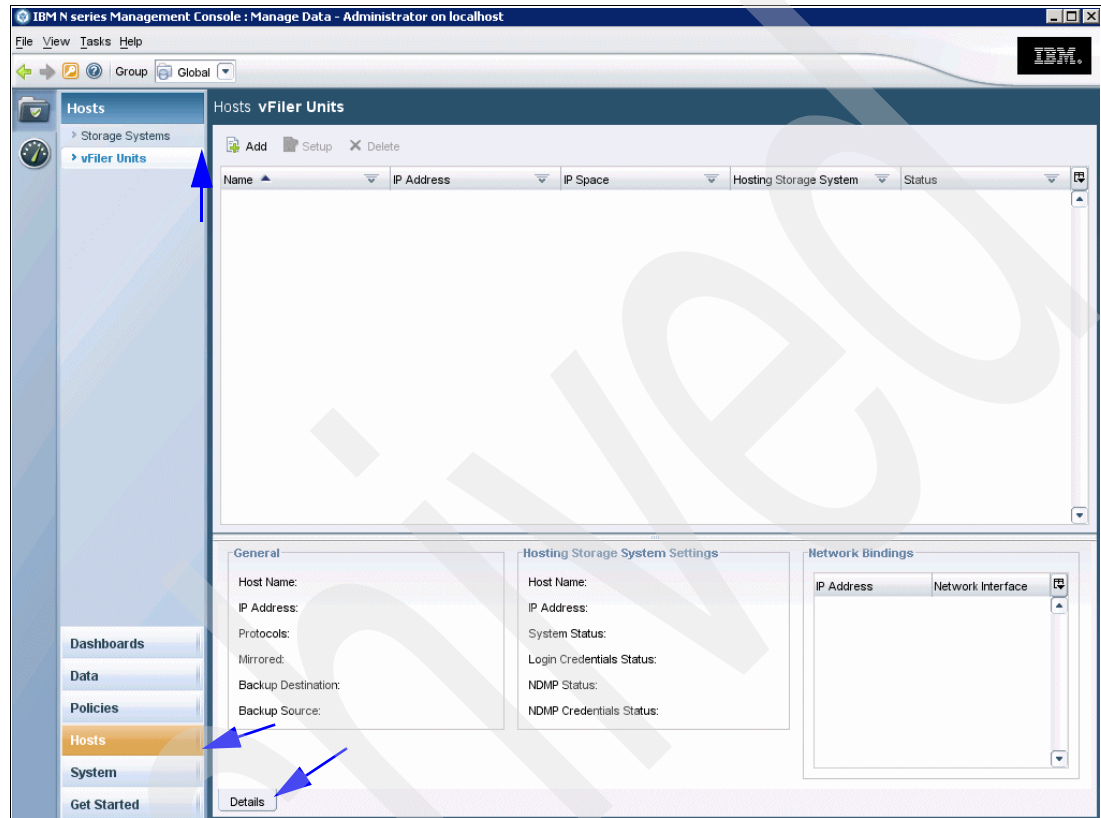


Figure 14-100 Provisioning Manager Hosts vFilers Units details

You can use the vFiler Units window to view detailed information about vFiler units discovered by Operations Manager. From this window, you can add vFiler units to or delete them from the N series Management Console host list. You can also access the Details tab, Paths tab, and the Relationships tabs for vFiler units.

You can use the following options by using the buttons at the top of the pane:

- ▶ **Add:** Starts the Add vFiler Unit Wizard, which allows you to create and configure vFiler units. You can choose to create a vFiler unit and configure it later by using the Setup button.
- ▶ **Setup:** Starts the Setup vFiler Unit Wizard, which allows you to configure or modify an existing vFiler unit.
- ▶ **Delete:** Deletes the selected vFiler unit from the N series Management Console vFiler list.

Here are descriptions of the fields and panes in this window:

- ▶ **vFiler host list**
 - **Name:** Displays the name of the vFiler unit as it appears in the Operations Manager database.
 - **IP Address:** The IP address associated with the selected vFiler unit.

- IP Space: The name of the IP space, if any, assigned to the vFiler unit.
- Hosting Storage System: Displays the name of the storage system that hosts the vFiler unit.
- Status: Displays the current status of the vFiler unit. Possible values are Up, Down, and Unknown. The default monitoring interval is five minutes. The interval is specified as the Ping Monitoring Interval in the DFM server.
- Details tab: The Details tab has two areas: General and Hosting Storage System Settings.
 - General
 - Host Name: The name of the currently selected vFiler unit.
 - IP Address: The IP address associated with the selected vFiler unit.
 - Protocols: The protocol type associated with the vFiler unit, either CIFS, NFS, or iSCSI.
 - Mirrored: Indicates whether the SnapMirror license is enabled on the hosting storage system. Possible values are Yes or No.
 - Backup Destination: Indicates whether a SnapVault Secondary license is enabled on this host, making the host a potential destination for backups. Possible values are Yes or No.
 - Backup Source: Indicates whether the SnapVault Data ONTAP Primary license is enabled on the hosting storage system, making the host a potential source of backups. Possible values are Yes or No.

Note: Other hosts might also be potential backup sources, but the SnapVault Primary licenses for Linux, UNIX, and Windows are installed on the secondary system. This list confirms only whether the SnapVault Data ONTAP Primary license is enabled on the host.

- Hosting Storage System Settings: These settings apply to the storage system that is hosting the vFiler unit that you selected in the vFiler list.
 - Host Name: The name that Operations Manager uses to log into the storage system.
 - IP Address: The IP address of the hosting storage system that is associated with the selected vFiler unit.
 - System Status: Whether the status of the storage system is Up, Down, Unknown, or Not Applicable (N/A).
 - Login Credentials Status: Displays the current status of the login credentials that Operations Manager uses to log in to the hosting storage system. Possible values are Good, Bad, Read Only, Unknown, and Not Applicable.
 - NDMP Status: The Network Data Management Protocol (NDMP) availability of the storage system as of the most recent NDMP monitoring check. Possible values are Up, Down, and Unknown. The default interval for NDMP monitoring is 30 minutes. You can use the Backup Discovery Options window in Operations Manager to change the NDMP monitoring interval.
 - NDMP Credentials Status: Displays the current status of the Network Data Management Protocol credentials that Operations Manager uses to communicate with the hosting storage system. Possible values are Good, Bad, Unknown, and Not Applicable. You can use the NDMP Credentials window in Operations Manager to edit the credentials for NDMP discovery.

14.5.7 Provisioning Manager configuration help

This Provisioning Manager help includes information for all applications installed on N series Management Console. By using the table of contents and index, you can find information about application features and how to use them.

You can access the help menu by selecting **Get Started** → **Overview**, as shown in Figure 14-101.

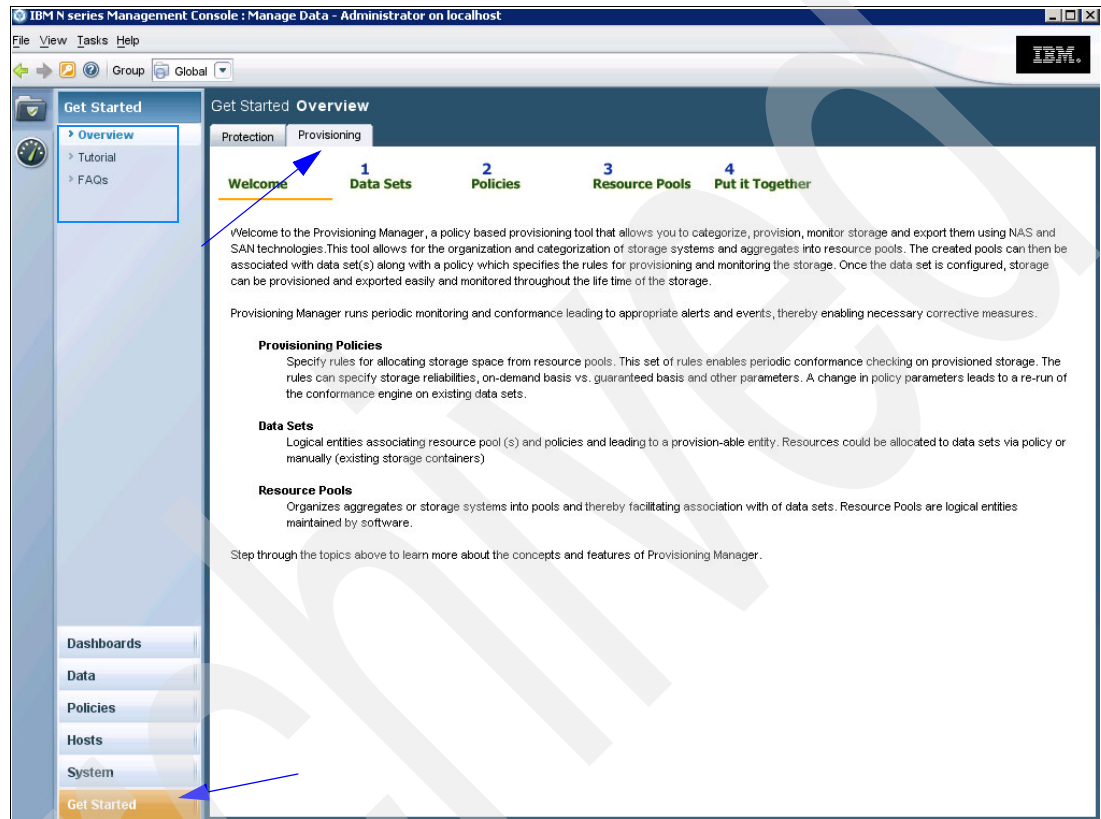


Figure 14-101 N series Management Console Provisioning Manager Get Started Overview

This help is structured as follows:

- General information about N series Management Console and tools common to all applications installed on N series Management Console, such as the dashboard, alarms, and events
- Information about Protection Manager
- Information about Provisioning Manager
- Information about Performance Advisor

The Get Started section of the user interface provides an overview of Provisioning Manager and a list of frequently asked questions (FAQs) for Provisioning Manager.

This concludes our chapter on the Provisioning Manager setup.



Operations Manager with VMware ESX server

This chapter will discuss solutions for VMware Server, VMware Workstation, and VMware Infrastructure with the IBM System Storage N series Operations Manager.

15.1 Introduction

In this publication, we have covered the features and functions of DataFabric Manager Server and the components of Operations Manager. As you have seen, there are many benefits of utilizing this product in your data center. However, how would you handle this product in a virtualized data center? For example, you might be running out of real estate, or adding more hardware is not possible, so it becomes necessary to utilize the resources you have to accomplish more tasks. Virtualization is a real solution that needs to be considered.

When you think of data center virtualization, one of the first applications that comes to mind is VMware ESX Server. While this chapter focuses on the benefits of utilizing a VMware guest to run Operations Manager, and a discussion of the benefits of virtualization may not be necessary, we offer the following information for your consideration.

Virtualization helps you take back control of your infrastructure. Virtualization enables you to see and manage your computing resources in ways that offer more flexibility because you are not restricted by implementation, location, or physical packaging. With virtualization, you have a logical rather than a physical view of data, computing power, storage capacity, and other resources. By gaining greater control of your infrastructure, you can improve cost management, as shown in Figure 15-1.

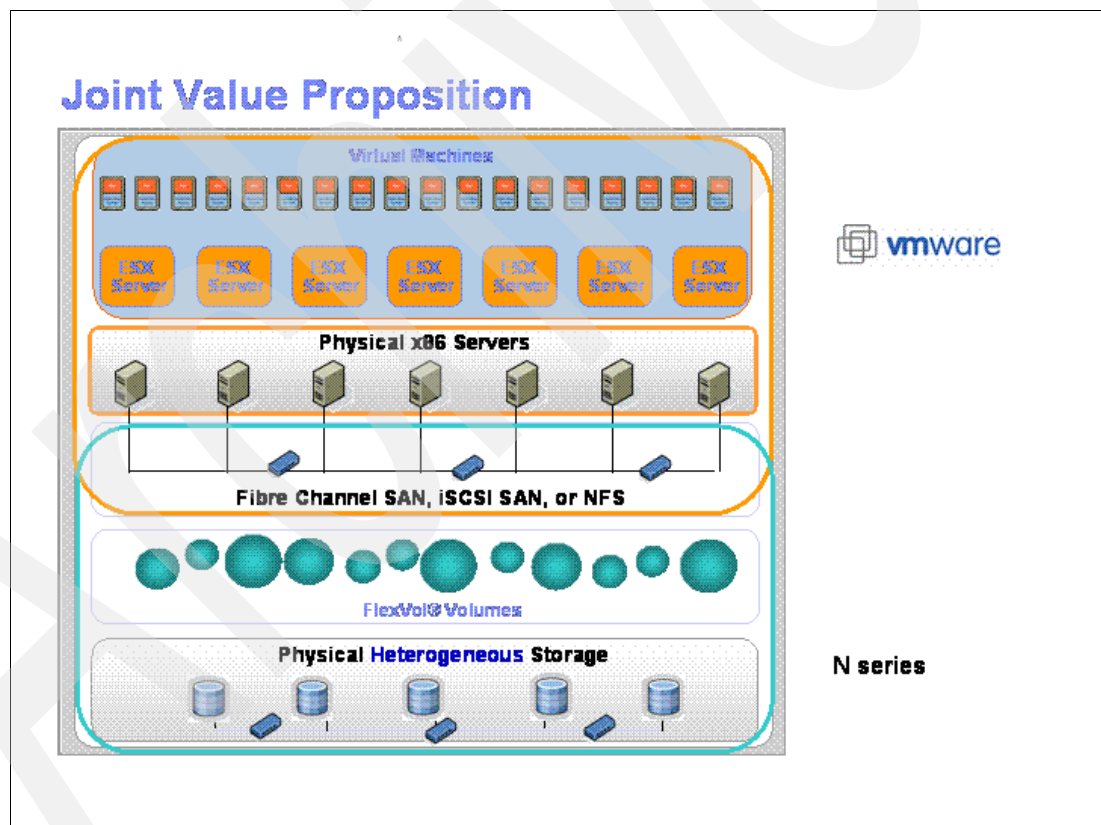


Figure 15-1 Virtualization with N series

Businesses are pursuing financial savings through both server and storage consolidation. This consolidation is achieved using virtualization. Virtualization is the abstraction of a physical resource into a virtual resource that is decoupled from the underlying hardware. Consolidation of server and storage hardware using virtualization offers return on investment for the business.

While cost saving is a primary driver for initial virtualization deployment, the full value of virtualization lies in its ability to offer:

- Improved total cost of ownership (TCO)

By decreasing management costs and increasing asset utilization, you can experience a rapid return on investment (ROI) with virtualization. In addition, by virtualizing resources and making them easier to migrate or fail over to other physical devices or locations, you can enhance system availability and help lower the cost and complexity of disaster recovery solutions.

- Increased flexibility

Virtualization supports the pooling of resources that can be managed centrally through an enterprise hub to better support changing business requirements dynamically.

- Enabled access through shared infrastructure

Virtualization provides a resilient foundation and shared infrastructure that enables better access to infrastructure and information in support of business applications and service-oriented architectures (SOA).

Companies of all sizes are aggressively adopting virtualization solutions to help with the following items:

- Infrastructure simplification

Virtualization can help control infrastructure sprawl through the deployment of virtual servers and storage that run securely across a shared hardware environment. Virtualization not only helps with server consolidation, but also server containment when deploying new systems. Consolidating to a virtual infrastructure can enable you to increase server utilization rates from 5% to 15% to over 70%, thus helping improve ROI. In addition, a simplified infrastructure can help lower management costs with a common management platform and tooling.

- Rapid application deployment

Virtualization can help enable rapid infrastructure provisioning (for example, minutes compared to days). It can help developers speed application test and deployment, enhance collaboration, and improve access to the infrastructure. The ease and flexibility of creating and reconfiguring guest operating systems means that development and test environments can realize significant benefits from virtualization.

- Business resiliency

Virtualization can help IT managers secure and isolate application workloads and data within virtual servers and storage devices for easier replication and restoration. This added resiliency can provide IT managers with greater flexibility to maintain a highly available infrastructure while performing planned maintenance and to configure low-cost disaster recovery solutions.

Virtualization technologies solve many traditional backup issues, because they decouple the bindings between the operating system (with the application and data) and the underlying hardware. For example, you could have a different hardware topology in the recovery site, both in terms of numbers of servers and configuration of those, and still be able to boot all your guests on the two different data centers.

- Green DataCenter

Though it has been mentioned above, virtualization allows you to reduce the hardware footprint in your data center. In many businesses and organizations, power consumption has hit a wall. Data center managers are confronted with the dilemma of spending precious dollars on either data center infrastructure or processing power. Though the choice would be for more processing power, the shortage of power affects the growth of

the business because you are forced to do more with less. Not only is the data center manager faced with the prospect of running out of power, but now additional cooling has to be considered as well. While the process of delivering these infrastructure upgrades goes on, processing power remains the same causing a lag in the growth of the company.

Virtualization offers a way to consolidate servers and downsize the processing footprint, thus reducing or delaying significantly the need for more infrastructure investments.

15.2 Operations Manager for VMware ESX Server host using N series storage

Storage virtualization software, similar in concept to server virtualization, abstracts the storage hardware volumes of data into a logical or virtual view of the volume. Using N series hardware with storage virtualization gives the data center a method to support storage provisioning in a manner that is independent of the underlying storage hardware, as shown in Figure 15-2. Storage virtualization can enable data sharing, data tiering, improved storage hardware utilization, improved availability, and disaster recovery capabilities. Storage virtualization software separates the representation of the storage to the operating system from the physical device. Utilization rates of storage are likely to be improved when moving towards network based storage that is virtualized.

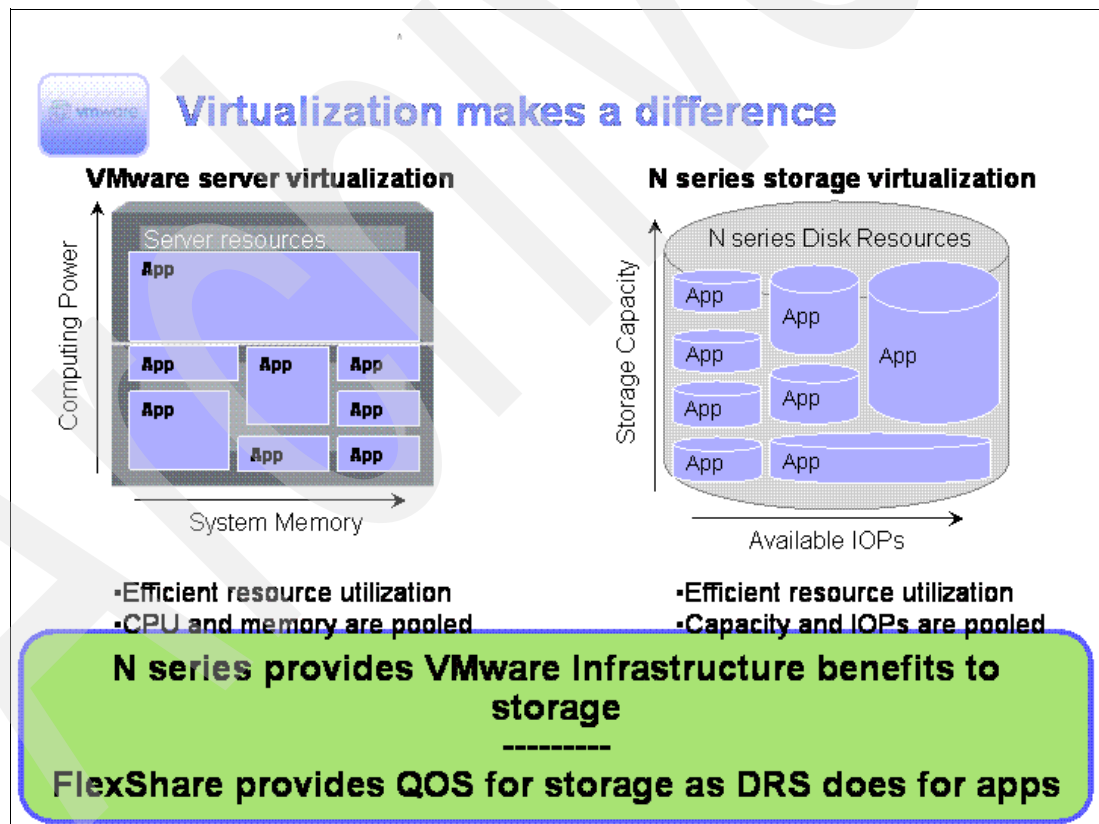


Figure 15-2 Virtualization benefits

Virtualization essentially lets one computer do the job of multiple computers, by sharing the resources of a single computer across multiple environments. Virtual servers and virtual desktops let you host multiple operating systems and multiple applications on a single computer.

The benefits of using VMware ESX Server with N series products has been discussed at length in *IBM System Storage N series with VMware ESX Server*, SG24-7636. The benefits of running Operations Manager on a VMware ESX Server host are similar in that you are able to get the same level of management as you would from running Operations Manager on a separate server platform.

15.3 Running a DataFabric Manager Server on a VMware guest

The very first consideration for running Operations Manager on a VMware guest is to be sure that you have properly sized your VMware guest machine. Though the N series Management Console and Host Agent will not require a large amount of resources, Operations Manager needs to have enough for good performance and growth. Refer to *IBM System Storage N series Operations Manager Sizing and Installation Guide*, REDP-4270 for detailed guidance about sizing your Operations Manager server.

In order for a virtual server to provide the desired application performance, it is important to ensure that all applications, on all guest virtual machines, are accounted for in regards to the following requirements: CPU, memory, networking bandwidth, and I/O capabilities. While the specific applications for meeting the business needs are at the forefront of discussions when virtualization projects are started, do not forget to account for the needs of often overlooked, yet required, applications such as the guest operating system itself, third-party firewall programs, antivirus programs, and so on. One missed supporting application may not be too big of a problem when you are working with a standard, single server. However, when you virtualize your servers into a single hardware platform where you may have 20 or more virtual systems running, a 1 MB memory oversight quickly turns into a 20 MB or more memory shortage that could cause performance impacts to the primary applications. Hardware shortages are often masked when all VMware ESX servers are running and applications are well distributed. However, if a server goes down due to a problem, and the applications that were running on that server are distributed to the surviving systems, a small shortage of headroom can become a critical business problem that manifests as poor performance.

Table 15-1 and Table 15-2 on page 544 shows what the hardware and software requirements are for running DataFabric Manager Server on a VMWare guest on a Windows 2003 server and a Linux server, respectively.

Table 15-1 Windows 2003 server

Windows 2003 Servers VMware ESX Server V3.0.1 or later	
Hardware requirements	Software requirements
<ul style="list-style-type: none"> ▶ Intel-based PC with single 2 GHz CPU (Xeon or Pentium 4) ▶ 4 GB of free disk space minimum, 8 GB recommended ▶ 1 GB Memory minimum 	<ul style="list-style-type: none"> ▶ Windows 2003 server 32-bit (Standard and Enterprise editions)

Table 15-2 Linux server

Linux Servers on VMware ESX server 3.0.1 or later	
Hardware requirements	Software requirements
<ul style="list-style-type: none"> ▶ Intel-based PC with single 2 GHz CPU (Xeon or Pentium 4) ▶ 4 GB of free disk space minimum, 8 GB recommended ▶ 1 GB Memory minimum 	Red Hat Enterprise Linux AS 4 (Update 3 or later) for x86, 32-bit and 64-bit

Note: Operations Manager V3.7 does not support VMware VMotion and VMware High Availability features.

Remember that these are minimum requirements and support up to 25 nodes. These requirements are for a Operations Manager installation with only basic system monitoring enabled. If you enable additional features and monitor additional objects, a more powerful platform is probably required. Examples of objects and features that might require a more powerful platform include additional storage systems, qtrees, user quotas, and use of the Storage Resource Management, Performance Advisor, Business Continuity Option, Provisioning Manager, or Protection Manager features.

The installation of Operations Manager requires the same steps outlined in Chapter 2, “Installing Operations Manager: Windows 2003 32-bit operating system” on page 21 and Chapter 4, “Installing Operations Manager: Linux” on page 55.

Note: Operations Manager deployed on a VMware server might cause the Operations Manager database to hang or crash. The virtual machine’s Snapshots functionality locks the database transaction log and prevents the Sybase iAnywhere database from writing to it. To prevent the Operations Manager database from hanging or crashing, complete the following steps:

1. Before you take a VMware Snapshot, stop the DFM SQL service using the `dfm service stop sql` command.
2. Restart the SQL service using the `dfm service start sql` command after the Snapshot is taken.

15.4 Running N series Management Console on a VMware guest

The N series Management Console is a small piece of code and does not have a significant impact on the amount of storage or resource required to run on a VMware guest. Refer to Chapter 5, “Host Agent installation for Windows 2003” on page 79 for the procedure to install the host agent on Windows 2003 Server.

As stated earlier in publication, the N series Management Console must be loaded on the system being used to monitor Operations Manager. You cannot access the DataFabric Manager Server using a browser.

15.5 Host Agent

The Operations Manager Host Agent is a small piece of code that does not have a significant impact on the amount of storage or resource required to run on a VMware guest. Refer to the Windows (Chapter 5, “Host Agent installation for Windows 2003” on page 79) or Linux Host Agent (Chapter 6, “Host Agent installation for Linux” on page 91) installation for the steps required to install the software.

Table 15-3 shows the systems that will run Operations Manager Host Agent V2.6. Table 15-4 shows the systems that will run Operations Manager Host Agent V2.7

Table 15-3 Operations Manager Host Agent V2.6

Operations Manager Host Agent 2.6 VMware ESX Server, Standard or Enterprise Edition, Version 3	
Guest	Windows 2003 Server (32-bit)
Guest	Red Hat Enterprise Linux AS, Version 4

Table 15-4 Operations Manager Host Agent V2.7

Operations Manager Host Agent 2.7 VMware ESX Server, Standard or Enterprise Edition, Version 3	
Guest	Windows 2003 Server (32-bit)
Guest	Windows XP, Windows Vista
Guest	Windows Server 2008 (32 and 64-bit)
Guest	Red Hat Enterprise Linux ES v4 (Update 3 or later) on both 32-bit and 64-bit x86
Guest	Red Hat Enterprise Linux AS v4 (Update 3 or later) on both 32-bit and 64 bit x86
Guest	SUSE Linux Enterprise Server 9 (SP2 or later) on 32-bit and 64-bit x86
Guest	SLES 10 on both 32-bit and 64-bit x86

Archived

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this book.

IBM Redbooks

For information about ordering these publications, see “How to get Redbooks” on page 548. Note that some of the documents referenced here may be available in softcopy only.

- ▶ *IBM System Storage N series MultiStore Discovery, Monitoring, and Reporting in Operations Manager*, REDP-4277
- ▶ *IBM System Storage N series Operations Manager Sizing and Installation Guide*, REDP-4270
- ▶ *IBM System Storage N series Reporting With Operations Manager*, SG24-7464
- ▶ *IBM System Storage N series with VMware ESX Server*, SG24-7636
- ▶ *Operations Manager 3.7 for IBM System Storage N series*, REDP-4457

Other publications

These publications are also relevant as further information sources:

- ▶ *Data ONTAP Network Management Guide*, GC52-1280
- ▶ *Installation and Upgrade Guide for Use with DataFabric Manager Server 3.7*, GC26-7892
- ▶ *Operations Manager Administration Guide for Use With DataFabric Manager Server 3.7*, GC26-7889
- ▶ *Operations Manager Host Agent 2.6 Installation and Administration Guide*, GC26-7894
- ▶ *Performance Advisor Administration Guide for Use with DataFabric Manager Server*, GC26-7897

Online resources

These Web sites are also relevant as further information sources:

- ▶ Supported features:
<http://www.ibm.com/storage/support/nas/>
- ▶ Product information:
<http://www.ibm.com/storage/nas/>
- ▶ Product publications:
<http://www.elink.ibm.link.ibm.com/public/applications/publications/cgibin/pbi.cgi/>

How to get Redbooks

You can search for, view, or download Redbooks, Redpapers, Technotes, draft publications and Additional materials, as well as order hardcopy Redbooks publications, at this Web site:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

Symbols

.doc 272
.JPG 289

A

Access Control 137
Active Directory 163
active/active configuration 167, 173, 182–183
Active/Active Configuration Status window 182
active/active controller configurations 173
adapter hardware 193
adapter speed 193
Add a FSRM path 304
Add a New Group option 186
Add a schedule 304
Add paths 284
Add SRM Path 300
Add SRM Path menu 305
Add Threshold Summary window 264
Add Threshold wizard 253
adding SRM host agents 292
adding SRM paths 299
Additive license 58
additive license 23, 58, 131, 309
administration access 298
administration transport 196
administration user name 297
administrative access 159, 298, 303
 for host agents 298
 passwords, host agents 296
administrative traffic 174
administrative users 162, 166
administrator 160, 162–164, 198
administrator account 160, 162
administrator accounts 160–161
administrator group 162
Administrator id 122
administrators 186
Administrators window 197–198, 303
agent administration access 274
agents 134
Agents Total section 294
aggr0 222
aggr0 disks 223
aggregate 144, 171–172, 211, 256, 323–324
Aggregate column 171
Aggregate Full 156
aggregate names 324
Aggregate Resource group 138
aggregates 144, 250, 256, 323–325
alarm 153, 157, 240–241, 249
alarm notification 157, 265
alarm setup 241
Alarms window 157

Alert Trap Received 155
alerting mechanisms 325
Appliance CPU 175
Appliance Details window
 cluster console, accessing 181
 description 172
 tasks 173
appliance discovery 13
Appliance Discovery Options section 291
Appliance group 139
appliance management. *See* appliances, managing
appliance name 172
Appliance Resource group 138
Appliance Tools 174
appliances
 Appliance Details window 172–173
 configuration changes 184
 console, connecting 180
 grouping 168
 managing
 DataFabric Manager settings 174
 management tasks 167
 modifying the appliance IP address 174
 SNMP requirement 135
 viewing
 Details window 172
Appliances tab
 Appliance Details window 172
applicable threshold 265
application suites 326
application tier 316
array of counters 258
assign roles 162
Audit Log options 12
authentication method 181
authorization key 29
autodiscovery 132
automated process 132
Automatic provisioning 17
Autosupport 75
Autosupport feature 75
Autosupport Notice 46
Autosupport options 12
average latency counter 250

B

Backup category 141
backup node 320
backup targets 325
base Operations Manager 309
buffer time 236
business application 325
Business Continance 40, 130
Business Continance Option 10

business recovery Service Level Agreements 319

C

- capacity reports 304
- catalog 150
- catalogs 149–150
- Certificate Authority 96
- chargeback 273
- Chargeback options 12
- Check Cluster Configuration 183
- Check Cluster Configuration tool 182
- CIFS 19, 196, 273, 292, 300
- CIFS systems 300
- CIFS, SRM requirements 300
- CLI 14, 141, 149–150, 185, 324
- CLI output 149
- Cloning 186
- cluster console 13, 173, 181
 - accessing 181
 - colors of storage system icons 182
 - information displayed on 181
 - requirements for using 181
 - status indicator colors 182
 - tools
 - Check Cluster Configuration 182
 - Giveback 184
 - Takeover 184
- Cluster Console window 182
- cluster partner 195
- colors
 - storage system icons, cluster console 182
- command-line interface 101
- command-line tools 245
- commands by group 137
- communication options 298
- complex storage environment 128
- configuration file 186
- configuration files 186–187
- configuration level 127
- configuration metrics 150
- configuration policies 314
- configuration resource group 138, 186–187
- configuration resource groups 186
- configuration settings 186
- Confirm Password parameter 164
- conformance checker 314
- Connect Device to Console tool 180
- console 267
- console connection 180
- Control Center 100, 304
- Control Center screen 133, 136, 251
- controlled access 241
- controller hardware 182
- controllers 182–183
- core license 129
- Core license key 58
- core license key 58, 309
- counter 253
- counter name 253
- counter value 250

- CPU 57, 175, 308
- CPU usage 153, 208
- CPU utilization 201, 208
- Create a LUN wizard 196
- create alarms 241
- Create new SRM Path 284
- Create Threshold Template 266
- creating configuration files 187
- Credential Cache options 12
- Critical severity level 156
- Critical Trap Received 155
- CSV format. 141
- custom alarm 157
- custom comment fields 167
- custom protection policy 326
- custom reports 150
- custom views 201, 228–229
- Customer Information window 115

D

- Dashboard 123, 317
- dashboard panel 205
- Dashboard window 205, 207–208
- data infrastructure 199, 209
- data migration 16
- Data ONTAP 24, 38, 41, 59, 64, 163, 183, 192, 216, 317
- Data ONTAP CLI 9
- Data ONTAP LUN 272
- Data ONTAP requirements 25
- Data ONTAP. 59
- Data ONTAP® version 7.1 59
- data protection 314, 316–317, 326
- data protection agents 325
- data restoration 318
- data service 182
- data set 138, 314, 316–317, 320, 322, 325–326
- Data Set Lags summary pane 322
- data set members 316
- Data Set Protection Status summary pane 322
- data sets 240
- Data Suite 16
- Database Backup option 13
- database servers 314
- DataFabric Manager 5, 20, 34, 41, 57, 141
 - acronyms 19
 - appliance settings 174
 - CLI 271
 - console connection 180
 - grouping objects 168
 - plug-in 24, 41, 59
- DataFabric Manager Host Agent software 190
 - administration access 298
 - capabilities 190
 - passwords 296
 - SAN host discovery requirement 188
- DataFabric Manager plug-in 59
- DataFabric Manager Server 6, 11, 14–16, 18–20, 22, 24, 38, 44, 59, 80, 82, 92, 154, 156, 160, 190, 199–200, 216, 268, 271, 273, 308, 547
 - command line 24, 44

- DataFabric Manager Server Data ONTAP plug-in 24
- DataFabric Manager Server license 23, 58, 309
- DataFabric Manager Server license key 58, 309
- DataFabric Manager Server software 11
- Default Thresholds options 13
- default view 222
- delete a host agent 283
- delete SRM path walk schedules 289
- delete SRM paths 288
- deleting
 - SRM host agents 295
- deploy Operations Manager 25
- Deploying Operations Manager software 25
- desktop 112
- Destination Folder 116
- destination mirror site 318
- detail view 258
- Details window 172
- DFM 5, 20
- dfm command 65, 310
- DFM Host Agent 20
- dfm host diag command 182
- dfm licenses command 310
- dfm option set 152
- dfm option set command 152
- DFM plug-ins 59, 64–65
- DFM Server 56, 59, 63–64, 67–69, 75, 92, 95–96, 100, 106, 310–312, 317, 321, 323–324, 326
- DFM Server package 64
- dfmDataExportEnabled 152
- Diagnose Connectivity 196
- Diagnose Connectivity tool 176, 182
- different login name 299
- Directory Details window 302
- directory paths 285
- directory structures 271
- directory-level data 289
- Disable field 157
- disaster 319
- disaster recovery 141
 - capability 314
 - concepts 314
 - configuration 314
 - license 318
 - operations 14
- discoveries 134
- discovery 128, 132–133, 136–137
 - DataFabric Manager Host Agent software 188
 - host 132
 - methods 132
 - SAN hosts 188, 190
- discovery function 132
- discovery methods 135
- Discovery option 13, 291
- Discovery Options window 134, 291
- discovery process 132
- disks 250
- Display options 13
- Distribution ACL options 13
- DNS 63

- downtime 156
- DR 327
- DR storage system 319
- drop-down menu 299, 304

E

- edit
 - SRM host agent settings 294–295
- Edit Agent Logins link 304
- Edit Appliance Settings window 174, 182
- Edit Groups link 186
- Edit Groups Membership window 139
- Edit Host Agent Settings window 293
- Edit Host Agents window 292, 295
- Edit Logins 295
- Edit Logins for Host Agents window 294, 298
- Edit Options list 295
- Edit Options pane 291
- Edit Options section 304
- Edit Settings window 180
- Edit SRM paths 288
- Edit SRM Paths window 299
- edit thresholds 259
- editing
 - SRM host agent settings 292
- element 219
- e-mail 162
- Emergency severity level 156
- Emergency Trap Received 155
- environment 128
- Error severity level 156
- Error Trap Received 155
- Ethernet port 81, 92
- event name 244
- events 153, 156
 - severity levels 156
- Events and Alerts options 13
- Events and Jobs queues 317
- Events window 124, 238
- export data 152
- exporting data 152
- Extended Details link 302
- external management applications 271

F

- fabric 193
- failover 183
- failover instructions 318
- Failover Status 182
- FC topology 193
- FCP targets
 - FC topology 193
- FCP Targets details window 193
- Fibre Channel SAN 319
- File SRM 81, 92, 271, 274, 303, 306
- File SRM license 81, 92
- File SRM paths 299
- File SRM tab 273, 304–305
- File SRM tab window 302

- File Storage Resource Manager 10, 18, 40, 59, 130, 269, 271
- file system 272, 284
- File System Resource group 138
- File Systems report 143
- File-level and directory-level statistics collection. 13
- FilerView 167, 169, 180, 182, 184–186, 192
- FilerView GUI 186
- Firefox plug-ins 74
- Firefox Web browser 69, 100
- flexible volumes 144
- FlexVol volume 324
- FlexVol volumes 144, 324
- formats 157
- free space 147
- FSRM 15, 19, 81–82, 270–271, 273–274, 298
 - capabilities 80
 - features 81, 92, 300, 304
 - functions 296–297
 - host 81
 - host agents 82, 274
 - host monitoring 92
 - information 272
 - license 81–82, 92
 - monitoring 273
 - path 305
 - path walk feature 299
 - paths 92, 285
 - prerequisites 303
 - tasks 80, 82
- FSRM summary window 272
- FSRM-generated file system data 80, 271
- Full threshold 320

G

- generating data 167
- giveback 182, 184
- giveback activities 184
- giveback tool in cluster console 184
- Global Backup 162
- Global Default 299
- Global group 140
- global Host Agent Options section 295
- global monitoring options 154
- global options 298
- global password 299
- global roles 163
- global values 174
- GlobalPerfManagement role 201
- GlobalRead role 163, 201
- GlobalSRM role 303
- GlobalWrite 201, 204
- graphical user interface 128
- graphs
 - SAN hosts 195
- group hierarchy 140
- group level 163
- group members 140
- Group Membership 186
- group role 163

- Group Selection window 243
- Group Summary View 220
- group summary windows 173
- Group the FSRM paths 304
- groups
 - advantages 168
- Groups pane 167–168
- Groups section 306
- Group-Top Objects 220

H

- HBA 14–15, 80, 257, 271
- HBA information 271
- HBA ports 191, 197
- high-availability storage 316
- Home link 305–306
- host administration access 274
- Host Agent 14–16, 18, 88–89, 128, 194–195, 271, 299
- host agent 93, 96, 275, 284, 286, 289–290, 297, 299
- host agent access 296
- host agent administration 278
- Host Agent code 59
- host agent communication 303
- Host Agent Details window 194, 292
- Host Agent Discovery 291
- Host Agent file 83
- Host Agent for Linux 91
- host agent global settings 281
- Host Agent Installation 77
- Host Agent link 296–297
- Host Agent Login option 284, 298, 304
- host agent login settings 304
- Host Agent Login=guest 298
- Host Agent Management Password 304
- Host Agent management tasks 274
- Host Agent Monitoring Password 296
- Host Agent options 13
- Host Agent Options window 292, 298
- host agent passwords 284
- host agent settings 275
- Host Agent software 15, 19, 93, 271, 274, 285, 298
- Host Agent's configuration UI 297
- host agents 13, 92, 270, 274, 281, 292, 294, 298, 304
 - administration access 298
 - passwords 298
- Host Agents menu option 194
- Host Agents, SRM window 290
- host agents, SRM. *See* SRM hosts
- host discovery 132
- host name 94
- Host Name parameter 164
- host ports 94
- HTTP port 89, 95
- HTTP_ops counter 214
- HTTPS certificate 96
- HTTPS port 95

I

- IBM Autosupport authorization Information 27

- IBM ID 60
- IBM Redbooks publication 128, 325
- IBM storage system 318
- IBM Support 60, 75
- IBM System Storage N series 3, 6, 11, 128, 308
- IBM System Storage N series Operations Manager 128
- IBM System Storage N series Operations Manager Host Agent software 79
- IBM System Storage N series SnapVault technology 318
- IBM System Storage N series storage systems 6, 11, 59
- IBM System Storage N series Volume SnapMirror technology 319
- IBM System Storage support 93
- IBM System Storage support Web site 60, 64
- IBM System Storage N series 5200 139
- IBM System x3655 38
- IBM® N series 56
- ICMP 135
- imported aggregates 321
- individual groups 140
- individual members 320
- information fields 150
- Information severity level 156
- Information Trap Received 155
- infrastructure 132
- initialization scripts 317
- initiator group
 - viewing LUNs mapped to 192
- Installation of Operations Manager 26
- installation of the N series Management Console 114
- interval 157
- IP address 68, 80–81, 89, 105, 122, 157
- IP network 216
- iSCSI 214
- IT operations staff 100
- ITSO Sample 234

J

- Java 100
- Java applets 128
- JDBC 151

K

- key management 96

L

- lag time 319, 322
- LDAP Directory 106
- LDAP options 13
- license 219
- license key
 - SRM 274
- License Key for Operations Manager 29
- Licensed Features options 13
- licenses 130
- Linux 271
- Linux DFM Server 65
- Linux downloads 110

- Linux host 67, 80, 96
- Linux host agent 16
- Linux system 308
- Linux workstation 92
- list of counters 252–253
- local backup operations 318
- Local User 166
- localhost 122
- logical objects 142
- Logical Objects category 141–142
- login and password
 - SAN hosts 196
- login protocol 177
- logs
 - syslog messages 182
- LUN Path Tools 192
- LUN Resource group 138
- LUNs 147, 188, 191–192, 197–198, 250, 258, 284, 299, 320, 326
 - creating 196
 - destroying 192
 - expanding 192
 - information 190–191
 - initiator groups mapped to 192
 - management 190
 - names 324
 - status of 191

M

- main panes 321
- managed storage systems 155
- Management Console 37, 110, 202, 241, 249, 260
- Management Console icon 217
- Management Console window 264
- Management Station License 238
- Management Station Node Limit Reached 238
- managing administrator access
 - Local Users
 - adding 164
- Maximum Password 164
- menu selections 234
- meta data 271, 299
- MIB 155
- Minimum Password 164
- mirror 327
- mirror operations 319
- mirror targets 319
- mix-types objects 140
- Monitor Events window 205
- monitor intervals
 - editing 295
- Monitored Devices field 194
- monitored objects 167
- monitored systems 169
- monitoring agents 324
- monitoring cycle time 175
- monitoring cycles 175
- monitoring interval 154
- monitoring intervals
 - See* options

- monitoring options 13
 - guidelines for changing 154
 - location to change 154
- monitoring passwords, host agents 296
- Mozilla 75
- MSCS 195
- multiple data sets 317
- multiple host agents 299
- multiple operations levels 163

N

- N series 17, 132, 163
 - arrays 68
 - products 252
 - storage network 5
 - technology 326
- N series Management Console 6, 11, 17, 20, 34, 75, 97, 100, 103, 105–106, 108, 110–112, 115–122, 204, 215, 217, 238–239, 241, 250, 268, 312, 319, 321
 - prerequisites 100
 - setup file 112
- N series Management Console Welcome window 114
- N series Operations Manager 56
- N series Operations Manager Interoperability Matrix 38
- N series Protection Manager 313
- N series storage systems 59, 93, 129, 132, 152–153, 163–164, 186, 214, 317, 324
- N7600 cluster 222
- NAS 6, 11, 17, 108, 147
- NAS volume 147
- navigation pane 206
- nconsole file 103
- nconsole link 104
- NetApp 127
- NetCache appliances
 - changing DataFabric Manager settings 174
 - IP address 174
 - threshold interval values 175
 - threshold values 174
- network 153, 178, 234
- network administration tasks 167
- Network Configuration field 183
- Network Credentials window 136
- network interface 250
- network traffic 154
- networking 183
- networks 132
- Networks to Discover field 134
- new alarm 242
- New Alarm Wizard 242
- New Group menu 305
- New Host Agent section 292
- New Schedule section 305
- New SRM Path section 305
- NFS 81, 92, 273
- node count 23, 131, 309
- Nodes and Connections tab 328
- Non-ASCII 315
- Normal severity level
 - description 156
- notification 157
- notification of events 241
- Notification Trap Received 155

O

- object types 256
- online help 304
- Open Systems SnapVault 132, 323, 326
- operating system version 168
- Operation Manager 12
- Operation Manager licenses 13
- Operational status 193
- Operations Manager 3–6, 9, 13, 15, 19–27, 31–34, 37–39, 41–42, 45, 48, 53–60, 63, 68, 75, 79–82, 91–92, 96, 122–124, 128–129, 132, 135, 138, 140–141, 144, 146, 148–158, 160–163, 167, 170–172, 174, 177, 180–182, 184–185, 201, 205, 207, 216, 218, 232, 238–239, 241, 245, 250, 258, 273, 285–286, 289, 293, 296–297, 304–305, 308–309, 317, 319, 321, 324 38
- Operations Manager 3.7 38, 57, 308
- Operations Manager 3.7.1 128
- Operations Manager Administration Guide 127, 135, 163, 269, 547
- Operations Manager Administrator 160, 164
- Operations Manager agents 75
- Operations Manager application 129
- Operations Manager commands 152
- Operations Manager components 39, 56
- Operations Manager Control Center 169
- Operations Manager Core 19
- Operations Manager data export 152
- Operations Manager database 128, 151–152, 163
- Operations Manager default settings 158
- Operations Manager environment 75
- Operations Manager events 167
- Operations Manager Events tab 239
- Operations Manager files 51
- Operations Manager FSRM feature 18, 92
- Operations Manager functions 162
- Operations Manager Groups summary 194
- Operations Manager GUI 82, 241
- Operations Manager Host Agent 80, 82–83, 89, 92, 271
- Operations Manager Host Agent 2.5 95
- Operations Manager Host Agent software 271
- Operations Manager Host Agent Version 2.6 84
- Operations Manager infrastructure 128, 194
- Operations Manager installation 57, 127
- Operations Manager installation files 60
- Operations Manager installations 23, 58
- Operations Manager interface 100, 184, 271, 324
- Operations Manager license 10, 13, 39
- Operations Manager logs 167
- Operations Manager main window, 167
- Operations Manager multiple-storage system remote configuration feature 184
- Operations Manager on Linux 57
- Operations Manager operation 54
- Operations Manager performance monitoring server 211
- Operations Manager resources 163

- Operations Manager Setup 25
- Operations Manager setup file 26
- Operations Manager software 25, 160
- Operations Manager storage systems 208
- Operations Manager User Interface 9–10, 14, 19, 180
- Operations Manager Web Application Interface 60
- Operations Manager window 159, 174
- operations per second 207, 212
- optional reports 142
- options
 - monitoring options
 - changing 154
 - guidelines for changing 154
- options ssh command 179
- Options window 174–175, 291, 295, 299
- OS Version field 183
- OSSV 16, 318

P

- pager address 162
- parent group 163
- Password field 122
- Password=userspecified value 297
- passwords for host agents 296
- Path field 305
- path management tasks 274
- path walk 285
- path walk schedule 274, 305
- path walk times 286
- pattern-matching expression 245
- per-file and per-directory data 285
- Performance Advisor 6, 11, 108, 152, 199–204, 208–209, 215–216, 227–228, 234, 237, 239, 252, 268
- Performance Advisor application 109, 200
- Performance Advisor data export 152
- Performance Advisor Events window 239
- Performance Advisor instances 201
- Performance Advisor options 13
- Performance Advisor view 241
- Performance Advisor window 229
- Performance Advisor windows 205
- performance blockage 223
- performance chart 211
- performance counter 211
- performance data 218
- performance information 201
- Performance Manager 217
- performance metrics 150
- performance monitoring 201
- performance monitoring server 215
- performance objects 209, 211
- performance tiers 326
- performance views 201
- performance-monitoring server 200
- physical level 257
- physical objects 256
- Physical Objects category 141
- physical resources 320
- physical selections 256
- physical storage 320

- plug-ins 37, 41, 317
- Point-To-Point 194
- policies 314
- policy parameters 317
- Policy-based provisioning 14
- policy-based provisioning 14
- potential blockages 250
- power 241
- power failure 241
- predefined condition 238
- predefined roles 160
- prerequisites 215
 - managing
 - clusters 181
- primary storage 317
- Protected Data Summary pane 322
- protecting data sets 325
- Protection 238
- Protection Manager 6, 11, 14, 16, 40, 58, 108, 129, 308–309, 314, 317–321, 323–326
- Protection Manager 3.7.1 315
- Protection Manager application 109
- Protection Manager components 321
- Protection Manager Dashboard 321
- Protection Manager DR license 319
- Protection Manager Option 14
- Protection Manager protection policies 317
- Protection Manager with Disaster Recovery 59, 130
- protection policies 317, 319, 324–326
- protection policy settings 317
- protection requirements 316
- protection scripts 325
- protection strategy 325
- provisioned storage 18
- provisioned volumes 320
- provisioning 314, 316, 320–321
- provisioning applications 320
- Provisioning Manager 6, 11, 17, 38, 40, 57, 59, 108–109, 130, 308, 316
- Provisioning Manager option 14
- Provisioning policies 17
- provisioning policies 14, 320, 326
- provisioning policy 320
- provisioning request 320
- provisioning requirements 316

Q

- qtrees 9, 13, 16, 38, 167, 315–316
- Qtrees rows 322
- quotas 81
- Quotas subtab 92
- Quotas tab renamed to SRM 274

R

- RBAC 9, 11, 163
- Real Time Data 235–236
- real-time status 181–182
- Recipients window 246
- Recovery Point Objectives (RPO) 319

- Red Hat DFM Server 70
- Red Hat Enterprise Linux 57
- Red Hat Enterprise Linux Advanced Platform 57, 308
- Red Hat Enterprise Linux AS 57, 308
- Red Hat Enterprise Linux AS 4 57, 308
- Redbooks Web site 548
 - Contact us xi
- Refresh Monitoring Samples 192, 196
- regular expressions. 157
- remote backup operations 318
- Remote backup protection 318
- remote backup site 318
- remote backups 314
- remote configuration feature 186
- Remote LAN Module 173
- remote secondary 318
- remote systems 13
- repeat notification 157
- Repeat Notify field 247
- report catalog 149
- Report drop-down list 306
- report status 154
- Reports options 13
- reserved space 147
- resiliency characteristics 320
- resizing volumes 18
- resource management 314
- resource pool 138, 314, 320–322, 326
- resource pools 320
- Resource Pools summary pane 322
- restoration mechanism 325
- restore operations 163
- Restore roles 163
- Review SRM path details 288
- Review SRM path walk schedule details 288
- RLM 177
- role management 162
- roles
 - global 163
 - group 163
- Roles list 198
- root 105
- root privileges 95
- RSH 177
- RSH Enabled 182

S

- SAN 15, 80, 271
- SAN capabilities 81
 - 15
- SAN HBA information 292
- SAN host 194–195
- SAN host reports 194
- SAN hosts 15, 80, 191, 197, 271
 - administration transport 196
 - discovery 188
 - editing settings for 196
 - port for communication 196
- SAN hosts agent software 188
- SAN objects 198
- SAN tasks 271
- SAN volumes 146–147
- Schedule field 305
- Schedule link 305
- Schedule Template list 305
- scheduled replication 317
- Script plug-ins options 13
- secondary storage 316, 318
- secondary storage site 318
- secureadmin addcert 180
- secureadmin command 180
- secureadmin disable command 180
- secureadmin enable 180
- secureadmin setup 180
- secureadmin setup ssh 179–180
- secureadmin status 179–180
- security options 13
- separate server 122
- serial numbers 131
- Set Up Alarm window 241
- setting passwords 297
- Setup drop-down menu 291
- Setup menu 159, 161, 295
- severity 157
- severity levels 156, 239
- severity of event 167
- severity type 238, 240
- simultaneous SRM path walks 286
- SLAs 319, 322
- SnapDrive 284
- SnapDrive software 190
- SnapMirror 10, 314, 317, 327
- SnapMirror protection 318
- SnapMirror technology 318
- Snapshot copies 147, 314
- Snapshot copy protection 318
- Snapshot reserve 146
- Snapshots 145–146
- SnapVault 10–11, 14, 108–109
- SnapVault operations 319
- SnapVault relationships 146
- SNMP 38, 135, 137, 153, 155, 246
 - requirement 135
 - storage system discovery on a specific network using v3 136
 - traps 154–156
- SNMP queries 153–154, 182
- SNMP trap 154
- SNMP trap global options 155
- SNMP trap listener 154–155
- SNMP trap listener configuration 155
- SNMP trap listener options 13
- SNMP version 137
- SNMPv1 135, 137
- SNMPv3 135–137
- Software Licenses menu 183
- Solaris 8 308
- Solaris host 80
- Space management 14
- specific host. 299

- specific privileges 158
- speedometer 204
- SRM
 - license key 274
 - passwords, setting 298
 - path walks
 - recommendations 286
 - paths
 - CIFS requirements 300
 - CLI quoting conventions 300
 - UNC requirements 300
 - valid path formats 300
 - paths, managing
 - quick reference 284
 - viewing details, directory-level 301
 - viewing details, file-level 300
 - quick reference for tasks 274
 - set up, steps for 274
 - Summary window, finding 274
 - viewing data 289
- SRM feature 300
- SRM Files, Least Recently Accessed report 306
- SRM Host Agent information 282
- SRM Host Agent monitoring 282
- SRM host agents 290
- SRM Host drop-down list 299, 305
- SRM Host window 83
- SRM hosts
 - administration access
 - configuration settings 298
 - deleting 295
 - disabling discovery 291
 - editing monitor intervals 295
 - editing settings 292, 294
 - enabling administration access 298
 - identifying
 - passwords, types of 296
 - quick reference of management tasks 275
- SRM options 13
- SRM path 300, 305–306
- SRM path group 138
- SRM Path Tools list 302
- SRM path walk 287–288
- SRM Path Walk Schedule Times window 305
- SRM paths 284, 305
- SRM Paths Details window 300–301
- SRM Paths, All menu 305
- SRM path-specific data 289
- SRM report tasks 274
 - quick reference 289
- SRM Summary window 272, 274, 290, 292, 294–295, 298–301, 304–305
- SSH 178, 180, 183
- SSH protocol 178–179, 183
- standard reports 143, 151
- statistical data 298
- Storage Area Network 41, 130
- storage array 316
- storage capacity 314
- Storage Capacity heading 190
- storage container 320
- storage device 177–178, 222
- storage infrastructure 6
- storage management 144–145
- storage object 299
- Storage Resource Management 57, 308
- Storage Resource Management Option 13
- Storage Resource Management option 13
- storage system 153–154, 167, 169–170, 173–174, 178, 191, 193, 205, 208–209, 232, 265, 316, 318–320, 323, 326
 - threshold interval value 175
 - threshold values 174
- storage system CLI 184
- storage system configuration 173
- storage system configuration settings 186
- storage system groups 167
- Storage System GUI 170
- storage system infrastructure 9
- storage system IP address 174
- storage system load 154
- storage systems 131–132, 134–135, 142, 152, 154, 167–168, 177–178, 184–187, 199, 209, 211, 216, 272–273, 304, 308
- Storage systems sub menu 323
- storage technology 326
- streaming protocols 13
- subsystem names 209
- Summary view 221
- summary view 258
- Summary View window 205–208
- Sun Java Runtime Environment 69
- Sun JRE 69–70, 73
- supported operating systems 271
- SUSE Linux Enterprise Server 10 57, 308
- SUSE Linux Enterprise Server 9 308
- SUSE® Linux Enterprise Server 57
- Synchronize Time Axis 223
- Synchronous 315
- Synchronous SnapMirror relationships 317
- syslog messages 182
- System Events 239
- System Events window 321
- System ID key 96
- system management applications 108
- System-ID variable 96

T

- takeover 181
- Takeover tool in cluster console 184
- target. *See* FCP targets.
- Targets Report 193
- TCP/IP 81, 274
- telnet
 - appliance console 180
- template 265, 267
- templates 326–327
- terminal server 180
- terminal window 102
- tertiary storage 318

- text strings 245
- threshold breach 259
- threshold creation 264
- threshold events 156
- threshold interval 175
- threshold intervals 175, 250
 - appliances 175
 - modifying 175
- threshold levels 263
- threshold template 265
- Threshold Templates view 267
- threshold values 174–175
- thresholds 238, 250–251, 262
 - modifying 174
- Throughput in Blocks icon 225
- throughput per second 205
- time 328
- time zone 328
- top five events 321
- Top Five Events Summary pane 321
- Top Performance Events 205
- Top Storage Systems by CPU Utilization 208
- Top Storage Systems by Network Throughput 205
- Top Storage Systems by Total OPs 207
- total network load 232
- total node count 131
- total_ops counter 212
- traditional volumes 315
- trap 154
- traps 157
- tree of items 218
- trend analysis 199

U

- UNC format 300
- UNIX 300
- UNIX host 284
- unpacking options 44
- Unprotected Data pane 323
- Unprotected Data summary pane 322
- Unprotected Data window 325
- upgrade 128
- usable space 171
- usage parameters 102
- USB drive 111
- user account 105
- user interface window 284, 289
- user name 122
- User Name parameter 164
- User name=admin value 297
- user quotas 167
- Users options 13

V

- vertical axis 208
- vFiler 9, 163, 172, 228
- vFiler Details window 173
- vFiler settings 174
- vFiler unit 164, 172–174, 200–201, 250

- vFiler unit name 185
- vFiler units 200
- View Cluster Console 183
- view permission 160
- View SRM reports 290
- visible directory paths 274
- VMware ESX Server 92
- volume 191
- volume data 129, 146
- volume Snapshot reserve space 147

W

- warning events 325
- Warning severity level 156
- Warning Trap Received 155
- Web browser 67
- Web servers 314
- Web-based UI 182
- Windows 200, 271, 289
- Windows 2000 57
- Windows 2003 25, 79, 108
- Windows 2003 64-bit operating system 45
- Windows 2003 64-bit server 37, 83
- Windows 2003 server 38
- Windows 64-bit host 79–80
- Windows system name 115
- Windows XP 108
- workflow 328
- WoW64 mode 45



Redbooks

Managing Unified Storage with IBM System Storage N series Operation Manager

(1.0" spine)

0.875" <-> 1.498"

460 <-> 788 pages



Managing Unified Storage with IBM System Storage N series Operation Manager



Learning about N series Provisioning Manager

Using N series Performance Advisor

Using N series Protection Manager to protect your data

IBM System Storage N series with Operations Manager software offers comprehensive monitoring and management for N series enterprise storage and content delivery environments.

Operations Manager is designed to provide alerts, reports, and configuration tools from a central control point, helping you keep your storage and content delivery infrastructure in-line with business requirements for high availability and low total cost of ownership.

We focus especially on Protection Manager, which is designed as an intuitive backup and replication management software for IBM System Storage N series unified storage disk-based data protection environments. The application is designed to support data protection and help increase productivity with automated setup and policy-based management.

This IBM Redbooks publication demonstrates how Operation Manager manages IBM System Storage N series storage from a single view and remotely from anywhere. Operations Manager can monitor and configure all distributed N series storage systems, N series gateways, and data management services to increase the availability and accessibility of their stored and cached data. Operations Manager can monitor the availability and capacity utilization of all its file systems regardless of where they are physically located. It can also analyze the performance utilization of its storage and content delivery network. It is available on Windows, Linux, and Solaris.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks