



**IBM System Storage N series
Installation and Upgrade Guide
For Use with DataFabric Manager Server 4.0
Supporting Operations Manager, Protection
Manager, and Provisioning Manager**

Copyright and trademark information

Copyright information

Copyright ©1994 - 2010 NetApp, Inc. All rights reserved. Printed in the U.S.A.

Portions copyright © 2010 IBM Corporation. All rights reserved.

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

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

References in this documentation to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any of IBM's or NetApp's intellectual property rights may be used instead of the IBM or NetApp product, program, or service. Evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM and NetApp, are the user's responsibility.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S.A. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

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. A complete and current list of other IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

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

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

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

NetApp; the NetApp logo; the Network Appliance logo; Bycast; Cryptainer; Cryptoshred; DataFabric; Data ONTAP; Decru; Decru DataFort; FAServer; FilerView; FlexCache; FlexClone; FlexShare; FlexVol; FPolicy; gFiler; Go further, faster; Manage ONTAP; MultiStore; NearStore; NetCache; NOW (NetApp on the Web); ONTAPI; RAID-DP; SANscreen; SecureShare; SimulateONTAP; SnapCopy; SnapDrive; SnapLock; SnapManager; SnapMirror; SnapMover; SnapRestore; SnapValidator; SnapVault; Spinnaker Networks; Spinnaker Networks logo; SpinAccess; SpinCluster; SpinFlex; SpinFS; SpinHA; SpinMove; SpinServer; SpinStor; StorageGRID; StoreVault; SyncMirror; Topio; vFiler; VFM; and WAFL are registered trademarks of NetApp, Inc. in the U.S.A. and/or other countries. Network Appliance, Snapshot, and The evolution of storage are trademarks of NetApp, Inc. in the U.S.A. and/or other countries and registered trademarks in some other countries. The StoreVault logo, ApplianceWatch, ApplianceWatch PRO, ASUP, AutoSupport, ComplianceClock, DataFort, Data Motion, FlexScale, FlexSuite, Lifetime Key Management, LockVault, NOW, MetroCluster, OpenKey, ReplicatorX, SecureAdmin, Shadow Tape, SnapDirector, SnapFilter, SnapMigrator, SnapSuite, Tech OnTap, Virtual File Manager, VPolicy, and Web Filer are trademarks of NetApp, Inc. in the U.S.A. and other countries. Get Successful and Select are service marks of NetApp, Inc. in the U.S.A. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

NetApp is a licensee of the CompactFlash and CF Logo trademarks.

NetApp NetCache is certified RealSystem compatible.

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 on 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, N.Y. 10504-1785
U.S.A.

For additional information, visit the web at:
<http://www.ibm.com/ibm/licensing/contact/>

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.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

Table of Contents

	Preface	iii
Chapter 1	Before You Start.	1
	Installation and upgrade requirements	2
	Installation and upgrade issues	11
Chapter 2	Installing and Upgrading DataFabric Manager.	17
	Installing DataFabric Manager 4.0 on Windows	18
	Installing DataFabric Manager 4.0 on Linux	22
	Installing licenses	25
	Downloading and installing IBM N series Management Console	26
	Migrating the DataFabric Manager database from Solaris	30
	Migrating the DataFabric Manager data to a different location	34
	Configuring your database for Snapshot-based backups	35
Chapter 3	Configuring DataFabric Manager with MSCS	37
	Planning to install DataFabric Manager with MSCS	38
	Configuring DataFabric Manager with MSCS	48
	Configuring DataFabric Manager with MSCS manually.	52
	Managing DataFabric Manager with MSCS	60
	Upgrading cluster nodes with DataFabric Manager	67
Chapter 4	Configuring DataFabric Manager with VCS	69
	Planning to install DataFabric Manager with VCS.	70
	Configuring DataFabric Manager with VCS	76
	Managing DataFabric Manager with VCS	83
	Upgrading cluster nodes with DataFabric Manager	90

Appendix A	Microsoft Cluster Server Parameters	93
Appendix B	Veritas Cluster Server Parameters.	99
Appendix C	Related Documentation	101
	Index	103

Preface

About this guide

This document describes the installation and upgrade procedures for the following:

- ◆ DataFabric® Manager 4.0 server
This includes Operations Manager, the Web interface of DataFabric Manager.
- ◆ IBM® N series Management Console
This includes Performance Advisor, Protection Manager, and Provisioning Manager (see “[Downloading and installing IBM N series Management Console](#)” on page 26).
- ◆ DataFabric Manager High Availability with Microsoft® Cluster Server
- ◆ DataFabric Manager High Availability with Veritas® Cluster Server

DataFabric Manager software is optimized for IBM N series storage systems that operate with the Data ONTAP® operating system. The information in this document applies to all supported hardware and software, as indicated in Chapter 1, “[Before You Start](#),” on page 1.

This document does not include system or network administration topics. For administration information, see the IBM N series *Operations Manager Administration Guide*.

Audience

This document is for administrators and others interested in managing storage systems with DataFabric Manager.

This document assumes the following:

- ◆ You are familiar with the Data ONTAP operating system software.
- ◆ You are familiar with Microsoft Cluster Server, Veritas Cluster Server, or both, if installing DataFabric Manager on a cluster.
- ◆ You are familiar with the protocols (NFS, CIFS, or HTTP) you use for file sharing or transfers.
- ◆ You are familiar with the client-side operating systems (UNIX® or Windows®).

Supported features

IBM System Storage™ N series storage systems are driven by NetApp® Data ONTAP® software. Some features described in the product software documentation are neither offered nor supported by IBM. Please contact your local IBM representative or reseller for further details. Information about supported features can also be found at the following Web site:

www.ibm.com/storage/support/nas/

A listing of currently available N series products and features can be found at the following Web site:

www.ibm.com/storage/nas/

Getting information, help, and service

If you need help, service, or technical assistance or just want more information about IBM products, you will find a wide variety of sources available from IBM to assist you. This section contains information about where to go for additional information about IBM and IBM products, what to do if you experience a problem with your IBM N series product, and whom to call for service, if it is necessary.

Before you call

Before you call, make sure that you have taken these steps to try to solve the problem yourself:

- ◆ Check all cables to make sure that they are connected properly.
- ◆ Check the power switches to make sure that the system is turned on.
- ◆ Use the troubleshooting information in your system documentation and use the diagnostic tools that come with your system.

Using the documentation

Information about N series hardware products is available in printed documents and a documentation CD that comes with your system. The same documentation is available as PDF files on the IBM NAS support Web site:

www.ibm.com/storage/support/nas/

Data ONTAP software publications are available as PDF files on the IBM NAS support Web site:

www.ibm.com/storage/support/nas/

Web sites

IBM maintains pages on the World Wide Web where you can get the latest technical information and download device drivers and updates.

- ◆ For NAS product information, go to the following Web site:
www.ibm.com/storage/nas/
- ◆ For NAS support information, go to the following Web site:
www.ibm.com/storage/support/nas/
- ◆ For AutoSupport information, go to the following Web site:
www.ibm.com/storage/support/nas/
- ◆ For the latest version of publications, go to the following Web site:
www.ibm.com/storage/support/nas/

Accessing online technical support

For online Technical Support for your IBM N series product, visit the following Web site:

www.ibm.com/storage/support/nas/

Supported servers and operating systems

IBM N series products attach to many servers and many operating systems. To determine the latest supported attachments, follow the link to the Interoperability Matrices from the following Web site:

www.ibm.com/storage/support/nas/

Terminology

Storage systems and gateways that run Data ONTAP are sometimes referred to as *filers*, *appliances*, *storage appliances*, or *systems*. The terms used in DataFabric Manager reflect some of these common usages.

When the term *appliance* is used in DataFabric Manager, the information applies to all supported storage systems, near-line systems, and IBM N series storage systems.

When the term *filer* is used, it can refer to any supported storage system, including IBM N series storage systems or near-line systems.

DataFabric Manager provides infrastructure services for various applications through IBM N series Management Console. Examples of IBM N series Management Console applications are Performance Advisor, Protection Manager, and Provisioning Manager. The DataFabric Manager server is the Windows or Linux® system on which DataFabric Manager is installed. Services running on DataFabric Manager 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.

Protection Manager simplifies the managing and monitoring of SnapVault® and SnapMirror® data protection relationships. Provisioning Manager simplifies and automates provisioning and managing storage for NAS and SAN access.

Operations Manager is the Web-based user interface of DataFabric Manager, 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.

Path convention

In parenthetical references to paths leading to parts of Operations Manager, the greater-than symbol (>) is used to point to the next interface element connecting you to your final destination.

For example, File Systems > Views > Volume Growth means to select the File Systems tab, open the Views drop-down list and select Volume Growth.

Command, keyboard, and typographic conventions

This document uses command, keyboard, and typographic conventions that help you enter commands.

Command convention: In examples that illustrate commands executed on a UNIX workstation, the command syntax and output might differ, depending on your version of UNIX.

Keyboard conventions: The following list describes keyboard conventions used in this document:

- ◆ When describing key combinations, this document uses the hyphen (-) to separate individual keys. For example, “Ctrl-D” means pressing the “Control” and “D” keys simultaneously.
- ◆ This document uses the term “Enter” to refer to the key that generates the digital equivalent of a carriage return, although the key is named “Return” on some keyboards.

Typographic conventions: The following table describes typographic conventions used in this document.

Convention	Type of information
<i>Italic font</i>	Words or characters that require special attention. Placeholders for information you must supply. For example, if the guide says to enter the <code>arp -d hostname</code> command, you enter the characters “arp -d” followed by the actual name of the host. Book titles in cross-references.
Monospaced font	Command names, option names, keywords, and daemon names. Information displayed on the system console or other computer monitors. The contents of files.
Bold monospaced font	Words or characters you type. What you type is always shown in lowercase letters, unless you must type it in uppercase letters.

Special messages

This document might contain the following types of messages to alert you to conditions you need to be aware of:

Note

A note contains important information that helps you install or operate the system efficiently.

Attention

An attention notice contains instructions that you must follow to avoid a system crash, loss of data, or damage to the equipment.

How to send your comments

Your feedback is important in helping us provide the most accurate and high-quality information. If you have comments or suggestions for improving this document, send us your comments by e-mail to starpubs@us.ibm.com. Be sure to include the following:

- ◆ Exact publication title
- ◆ Publication form number (for example, GC26-1234-02)
- ◆ Page, table, or illustration numbers

- ◆ A detailed description of any information that should be changed

About this chapter

This chapter provides you with the following information about DataFabric Manager 4.0:

- ◆ [“Installation and upgrade requirements”](#) on page 2
- ◆ [“Installation and upgrade issues”](#) on page 11

You must read the information in this chapter before installing DataFabric Manager software.

Installation and upgrade requirements

About the requirements

Before installing or upgrading the DataFabric Manager software, you must ensure that you have met requirements in the following areas:

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

Hardware and software requirements

Your workstation or the DataFabric Manager server must meet the following requirements before you install DataFabric Manager 4.0.

These requirements are recommended for environments with 1 to 25 nodes. For more information about sizing guidelines, see the following section.

Windows Server 2008	
Hardware requirements	Software requirements
<ul style="list-style-type: none">◆ Intel®-based PC with single 2-GHz CPU (Xeon® or Pentium® 4)◆ 10 GB of free disk space minimum, 40 GB recommended◆ 2 GB of memory minimum	<ul style="list-style-type: none">◆ Windows 2008 server, 32-bit on x86◆ Windows 2008 server, 64-bit on x64 (in WOW64 mode)

Windows Server 2003	
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 of memory minimum 	<ul style="list-style-type: none"> ◆ Windows 2003 server, 32-bit on x86 ◆ Windows 2003 server, 64-bit on x64 (in WOW64 mode)
Windows Server 2008 on VMware® ESX 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) ◆ 10 GB of free disk space minimum, 40 GB recommended ◆ 2 GB of memory minimum 	<ul style="list-style-type: none"> ◆ Windows Server 2008, 32-bit and 64-bit (Standard and Enterprise editions)
Windows Server 2003 on VMware ESX 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 of memory minimum 	<ul style="list-style-type: none"> ◆ Windows Server 2003, 32-bit and 64-bit (Standard and Enterprise editions)

Linux workstation or server	
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 of memory minimum 	<ul style="list-style-type: none"> ◆ Oracle Enterprise Linux 4 for x86, 32-bit and 64-bit ◆ Oracle Enterprise Linux 5.x for x86, 32-bit and 64-bit ◆ Red Hat™ Enterprise Linux AS 4 (Update 3 or later) for x86, 32-bit and 64-bit ◆ Red Hat Enterprise Linux Advanced Platform 5.x 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
<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 of memory minimum 	<ul style="list-style-type: none"> ◆ Red Hat Enterprise Linux Advanced Platform 5 for x86, 32-bit and 64-bit ◆ Red Hat Enterprise Linux AS 4 (Update 3 or later) 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

Browser
See the NAS Interoperability Matrix at http://www.ibm.com/systems/storage/network/interophome.html .

Note

DataFabric Manager 4.0 is not supported on Windows NT® 4.0, Windows 2000, Windows XP, or distributions of Linux not listed in the preceding table.

DataFabric Manager 4.0 supports VMware VMotion™ and VMware High Availability features.

These requirements are for a DataFabric 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.

Sizing guidelines

You can view the Operations Manager *Sizing Guide* at <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.

License requirements

You must have a valid DataFabric Manager server license key to complete the DataFabric Manager installation. IBM provides your license key on the printed IBM System Storage N series Function Authorization form that you received from IBM with your software offering order. If you do not have the Function Authorization form, contact the IBM Manufacturing Quality Hotline.

After you complete the installation, you can enter additional license keys on the Options page in Operations Manager.

You can install (or upgrade to) DataFabric Manager 4.0 using the server 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 DataFabric 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.

The following table lists each DataFabric Manager feature, the license you must install to enable it, and the capabilities provided by it.

To use...	Install this license or application...	That enables these features...
Operations Manager	<ul style="list-style-type: none"> ◆ DataFabric Manager server license ◆ Operations Manager license <p>Note_____</p> <p>Required for all licensed Operations Manager installations. Sets the maximum number of storage systems that the DataFabric Manager server can monitor in this installation.</p> <p>_____</p>	<ul style="list-style-type: none"> ◆ Monitoring <ul style="list-style-type: none"> ❖ Reports ❖ Storage usage and availability, such as qtrees, volumes, aggregates, LUNs, and disks ❖ Storage systems ❖ vFiler® units ◆ Managing <ul style="list-style-type: none"> ❖ Storage system configuration ❖ Scripts ◆ Monitoring and managing active/active configurations on storage systems using Cluster Console ◆ Displaying historical and real-time performance data using Performance Advisor in IBM N series Management Console

To use...	Install this license or application...	That enables these features...
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 ◆ Secondary storage space management ◆ Backup Manager 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 mirror relationships of all systems enabled with SnapMirror ◆ When you install the Protection Manager license, the provisioning functions provided by Provisioning Manager are also enabled.

To use...	Install this license or application...	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 datasets ❖ Automated offline dataset and vFiler unit migration ❖ Automated online dataset and vFiler unit migration ❖ Deduplication to eliminate duplicate data blocks
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 	<ul style="list-style-type: none"> ◆ 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 	<ul style="list-style-type: none"> ◆ Tracking file system usage and capacity information

To use...	Install this license or application...	That enables these features...
Business Continuation	<ul style="list-style-type: none"> ◆ DataFabric Manager server license ◆ Operations Manager license ◆ Business Continuation Option (BCO) 	<ul style="list-style-type: none"> ◆ Backup Manager <ul style="list-style-type: none"> 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 mirror relationships of all systems enabled with SnapMirror <p>Note_____</p> <p>Because IBM has announced the end of availability for the BCO license, you cannot see Backup Manager or Disaster Recovery Manager in Operations Manager unless you install Protection Manager.</p>

Note_____

IBM has announced the end of availability for the BCO license.

Data ONTAP requirements

You must be running Data ONTAP version 7.1 or later with DataFabric Manager 3.3.1 or later.

Note

You must have a DataFabric Manager plug-in for each version of Data ONTAP that you are running across your system. DataFabric 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 command line. You do not need to download a plug-in unless you are using a different version of Data ONTAP.

Installation and upgrade issues

About this section

Every release of DataFabric Manager involves changes in the system behavior of which you need to be aware. You might have to resolve issues before upgrading. This section provides you with information that you should review before installing or upgrading to DataFabric Manager 4.0.

rsh commands fail on Windows Server 2008

You cannot execute commands on storage systems by using the RSH protocol on Windows Server 2008. You can use either of the following workarounds to resolve this issue. However, Workaround 1 ensures greater security for your storage systems.

Workaround 1: Complete the following steps to set up DataFabric Manager server and storage systems to use SSH protocol:

Step	Action
1.	Ensure that you have SSH enabled on the storage system.
2.	Enter the following command on the DataFabric Manager server to enable SSH: <code>dfm option set hostLoginProtocol=ssh</code>

Workaround 2: Complete the following steps to enable rsh commands on the DataFabric Manager server:

Step	Action
1.	Install Subsystem for UNIX-based Applications (SUA) or Service for UNIX (SFU) on Windows Server 2008.
2.	Use the rsh binary from Microsoft SUA utilities and SDK by downloading Utilities and SDK for UNIX-based Applications_X86 (for Windows Server 2008).

Step	Action
3.	<p>Complete the following steps to configure the DataFabric Manager server:</p> <ol style="list-style-type: none"> <li data-bbox="575 326 1220 418">a. At the command prompt, enter the following command: <code>dfm option set rshBinary="posix /c <rsh_binary_path>\rsh"</code> <li data-bbox="575 440 1220 557">b. Enter the following command: <code>dfm service runas -u administrator -p <admin_password> http eventdmonitor scheduler server watchdog</code> <li data-bbox="575 579 1220 671">c. Stop the DataFabric Manager service by entering the following command: <code>dfm service stop</code> <li data-bbox="575 694 1220 786">d. Restart the DataFabric Manager service by entering the following command: <code>dfm service start</code>

Supported methods to upgrade from Windows 2000 to Windows Server 2003

DataFabric Manager 3.7 or later does not support Windows 2000. Therefore, you must upgrade your server operating system to Windows Server 2003 before you upgrade to DataFabric Manager 3.7 or later. If you have Windows Server 2003 already installed on another system, install DataFabric Manager 3.7 or later on it, and then migrate the database from the Windows 2000 system.

For information about database backup, see the section on setting up the DataFabric Manager database backup in the *Operations Manager Administration Guide*. For information about migrating the database, see [“Migrating the DataFabric Manager data to a different location”](#) on page 34.

Supported methods to upgrade from Solaris to Windows or Linux

DataFabric Manager 3.8 and later does not support Solaris. Therefore, you must migrate the DataFabric Manager database on Solaris to a server running Windows or Linux before you upgrade to DataFabric Manager 3.8 or later. For more information, see [“Migrating the DataFabric Manager database from Solaris”](#) on page 30.

Deploying DataFabric Manager software

You must deploy DataFabric Manager on a system that is running no other applications. When you purchase DataFabric Manager, you need to download the core license from the IBM Web site. You can install other licenses only after installing the Operations Manager license.

Deploying DataFabric Manager on a VMware server

DataFabric Manager deployed on a VMware server might cause DataFabric Manager database to hang or crash due to VMware. The virtual machine snapshot functionality locks the database transaction log and prevents Sybase iAnywhere database from writing to it.

Upgrading from DataFabric Manager 3.7.1

If you have created custom reports with GUILink and SecureGUILink as fields in DataFabric Manager 3.7.1 or earlier, upgrading to DataFabric Manager 3.8 or later causes the `dfm report view` command to fail. You must open the custom report in Operations Manager and save the report to view it.

Upgrading from DataFabric Manager 3.7 on Linux

If you are upgrading from DataFabric Manager 3.7 to DataFabric Manager 3.8 or later on Linux, the upgrade might fail with the following notification:

```
rpm: /opt/IBMdfrm/lib/libgcc_s.so.1: version `GCC_4.2.0' not found  
(required by /usr/lib/libstdc++.so.6)
```

You can resolve this issue by deleting the entry `/opt/IBMdfrm/lib` from the environment variable `LD_LIBRARY_PATH`.

Upgrading from DataFabric Manager 3.7 or earlier

If you are upgrading from DataFabric Manager 3.7 or earlier to DataFabric Manager 3.8 or later, you must delete the existing Data Source Name (DSN) entry for the Adaptive Server Anywhere 9.0 driver and create a new DSN entry for SQL Anywhere 10.

Upgrading from DataFabric Manager 3.5.1 or earlier

If you are upgrading from DataFabric Manager 3.5.1 or earlier to DataFabric Manager 3.6.1 or later, it takes a long time to upgrade the performance data files (data of 20 GB or more). The length of time depends on the platform used. The space used by the performance data files increases by about 65% during the upgrade.

Windows installation path for DataFabric Manager

Installing DataFabric Manager: The default installation path for DataFabric Manager 3.8 or later on Windows is as follows:

- ◆ 32-bit platform: C:\Program Files\IBM\DataFabric Manager
- ◆ 64-bit platform: C:\Program Files (x86)\IBM\DataFabric Manager

Upgrading to DataFabric Manager 3.8 or later: On upgrading to DataFabric Manager 3.8 or later, the default installation path remains unchanged.

Windows Server 2003 default browser security setting recommendations

After Windows Server 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 DataFabric Manager upgrade.

To ensure the best browsing experience when using DataFabric Manager, you might need to adjust your browser security setting to “medium.”

Resolving port conflicts

You might not be able to establish communication between the DataFabric Manager server and its storage systems, if port conflicts are not resolved. To ensure seamless communication, ensure that no application except DataFabric Manager uses the following ports: 8080, 8443, 8088, 8488, and 162. For more information, see the FAQ section at <http://www.ibm.com/storage/support/nas/>.

Installing DataFabric Manager in a custom Linux directory

If you use the `-d <new directory>` command to install DataFabric Manager in a custom directory, DataFabric Manager is installed in the new directory you specify; no additional IBMdfm directory is created.

Blocking access to Web servers

DataFabric Manager allows search engines to crawl and index the Web pages in Operations Manager. If you want to block access by search engines to the Web pages, create a file called `robots.txt` in the directory `<install-path>/DFM/web`. If the file contains only the following text, search engines do not crawl or index Operations Manager Web pages:

```
User-agent: *
```

```
Disallow: /
```

Viewing dynamic data

To use Disaster Recovery, the browser that you use to view Operations Manager must support Java applets.

About this chapter

The following topics are described in this chapter:

- ◆ [“Installing DataFabric Manager 4.0 on Windows”](#) on page 18
- ◆ [“Installing DataFabric Manager 4.0 on Linux”](#) on page 22
- ◆ [“Installing licenses”](#) on page 25
- ◆ [“Downloading and installing IBM N series Management Console”](#) on page 26
- ◆ [“Migrating the DataFabric Manager database from Solaris”](#) on page 30
- ◆ [“Migrating the DataFabric Manager data to a different location”](#) on page 34
- ◆ [“Configuring your database for Snapshot-based backups”](#) on page 35

Note

If you are upgrading from a previous version of DataFabric Manager, see the section on setting up the DataFabric Manager database backup in the IBM N series *Operations Manager Administration Guide*.

Installing DataFabric Manager 4.0 on Windows

About this section

This section provides you with the following information:

- ◆ [“Installing DataFabric Manager on Windows”](#) on page 18
- ◆ [“Uninstalling DataFabric Manager on Windows”](#) on page 20

Note

To achieve high availability, you should install Veritas Cluster Server (VCS) or Microsoft Cluster Server (MSCS) before installing DataFabric Manager.

Installing DataFabric Manager on Windows

To install DataFabric Manager 4.0 on your Windows server, complete the following steps:

Note

The installation software automatically detects and stops any DataFabric Manager services that are running on the system. Therefore, you do not need to manually stop DataFabric Manager services before starting the upgrade process.

Step	Action
1	<p>Ensure the following:</p> <ul style="list-style-type: none">◆ Your workstation meets the requirements described in “Installation and upgrade requirements” on page 2.◆ You have the information required to complete the installation, as described in “Installation and upgrade requirements” on page 2.”◆ You have Local Administrator login permission for the DataFabric Manager server. <hr/> <p>Note</p> <p>Before you install DataFabric Manager on Windows Server 2008, install Subsystem for UNIX-based Applications (SUA) or Service for UNIX (SFU) to support execution of rsh commands. For more information, see “Installation and upgrade issues” on page 11.</p> <hr/>

Step	Action	
2	If you are...	Then...
	Upgrading to DataFabric Manager 4.0	<p>Back up your existing DataFabric Manager database.</p> <ul style="list-style-type: none"> ◆ During the upgrade: The DataFabric Manager 4.0 install wizard includes a backup step. You can choose to back up your database and specify the type of backup. <p>For a Snapshot-based backup, you must enter the backup filename in the following format:</p> <p>-t sndb <backup name></p> <p>If the backup does not finish successfully, abort the install wizard, perform a manual backup, and then restart the installation.</p> <ul style="list-style-type: none"> ◆ Before the upgrade: If you choose to skip the backup step during the upgrade, you should back up the database before you start the install wizard. For instructions, see the section on setting up the DataFabric Manager database backup in the <i>IBM N series Operations Manager Administration Guide</i>.
	Installing DataFabric Manager 4.0	Go to Step 3 .
3	Get the DataFabric Manager 4.0 installer from http://www.ibm.com/storage/support/nas/ .	
4	Launch dfmsetup-4-0-win32.exe.	

Step	Action
5	<p>Follow the DataFabric Manager setup prompts to complete the installation.</p> <p>Note_____</p> <p>The installation process can take several minutes.</p> <p>_____</p> <p>During a new installation, you must specify the Operations Manager license key. If you are upgrading from an earlier licensed version of DataFabric Manager, you do not need a license key.</p> <p>The DataFabric Manager 4.0 installation and upgrade process automatically installs the AutoSupport feature with AutoSupport enabled and displays a message about how to disable the feature.</p>

After the installation is complete, Operations Manager launches automatically.

Uninstalling DataFabric Manager on Windows

To uninstall DataFabric Manager, complete the following steps:

Step	Action	
1	If you are using...	Then...
	Windows Server 2003	From the Start menu, select Settings > Control Panel > Add/Remove Programs.
	Windows Server 2008	From the Start menu, select Settings > Control Panel > Programs and Features.
2	From the list of applications, select IBM N series DataFabric Manager.	
3	Click the Remove button.	
4	Follow the prompts to uninstall the software.	

During the uninstallation process, DataFabric Manager saves the log files and the database file to a temporary directory named DFM-*<date, year, month, day, hour, minute, second>*, if it has permissions to create the directory. If DataFabric Manager does not have sufficient permissions, the database file and the log files are left in place.

Installing DataFabric Manager 4.0 on Linux

About this section

This section provides you with the following information:

- ◆ [“Installing DataFabric Manager on Linux”](#) on page 22
- ◆ [“Uninstalling DataFabric Manager on Linux”](#) on page 24

Installing DataFabric Manager on Linux

To install DataFabric Manager 4.0 on your Linux DataFabric Manager server, complete the following steps:

Step	Action
1	<p>Ensure the following:</p> <ul style="list-style-type: none">◆ Your DataFabric Manager server meets the requirements described in “Installation and upgrade requirements” on page 2.◆ You have the information required to complete the installation, as described in “Installation and upgrade requirements” on page 2.◆ You have the root privileges required to log in to the DataFabric Manager server.◆ The SELinux status is disabled if you are installing DataFabric Manager on Red Hat Enterprise Linux Advanced Platform 5.x.

Step	Action	
2	If you are...	Then...
	Upgrading to DataFabric Manager 4.0	<p>Back up your existing DataFabric Manager database.</p> <ul style="list-style-type: none"> ◆ During the upgrade: The DataFabric Manager 4.0 install wizard includes a backup step. You can choose to back up your database and specify the type of backup. <p>For a Snapshot-based backup, you must enter the backup filename in the following format:</p> <p>-t sndb <backup name></p> <p>If you choose this option, the DataFabric Manager SQL service must be running to perform the backup. If the backup does not complete successfully, abort the install wizard, perform a manual backup, and then restart the installation.</p> <ul style="list-style-type: none"> ◆ Before the upgrade: If you choose to skip the backup step during the upgrade, you should back up the database before you start the install wizard. For instructions, see the section on setting up the DataFabric Manager database backup in the IBM N series <i>Operations Manager Administration Guide</i>.
	Installing DataFabric Manager 4.0	Go to Step 3 .
3	Get the DataFabric Manager 4.0 installer from http://www.ibm.com/storage/support/nas/ .	
4	Launch dfmsetup-4-0-linux.sh.	

Step	Action
5	<p>Follow the DataFabric Manager setup prompts to complete the installation.</p> <p>Note_____</p> <p>The installation process can take several minutes.</p> <p>_____</p> <p>By default, DataFabric Manager installs itself in the /opt directory. To change the installation directory, use the <code>-d <new directory></code> command-line option. DataFabric Manager installs DataFabric Manager CLI wrappers at /usr/bin. To change this path, use the <code>-w <new path></code> command-line option.</p> <p>During a new installation, you must specify the Operations Manager license key. If you are upgrading from an earlier licensed version of DataFabric Manager, you do not need a license key.</p> <p>The DataFabric Manager 4.0 installation and upgrade process automatically installs the AutoSupport feature with AutoSupport enabled and displays a message about how to disable the feature.</p>

Uninstalling DataFabric Manager on Linux

To uninstall DataFabric Manager, complete the following step:

Note

You must log in to the DataFabric Manager server with root privileges.

Step	Action
1	<p>Use one of the following Linux commands:</p> <ul style="list-style-type: none"> ◆ <code>rpm -e IBMdfm</code> ◆ <code>rpm --erase IBMdfm</code>

During the uninstallation process, DataFabric Manager saves the log files and the database file to a temporary directory named IBMdfm- <year, month, day, hour, minute, second>, if it has permissions to create the directory. If DataFabric Manager does not have sufficient permissions, the database file and the log files are left in place.

Installing licenses

About license keys

Each DataFabric Manager installation requires a Core license or a DataFabric Manager server with a unique serial number. You must install the DataFabric Manager server license before you install any other licenses for DataFabric Manager features.

For DataFabric Manager features that require licenses, you must install the license key before you enable the feature. Each feature license need not have the same serial number as the DataFabric Manager server license. You can install these licenses in any order.

Installing license keys

You can install license keys when you install the Operations Manager license. You can also install them any time later, from the Options page on Operations Manager.

To install license keys from the Options page, complete the following steps:

Step	Action
1	From any page, select Options from the Setup menu. The Options page is displayed.
2	In the Edit Options list (in the left pane), click Licensed Features. The Licensed Features Options page is displayed.
3	In the New License Key field, enter the key for the new license and click Update.

Downloading and installing IBM N series Management Console

Function of IBM N series Management Console

IBM N series Management Console is a client software that contains a number of storage system management applications. IBM N series Management Console incorporates the following:

- ◆ 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.

IBM N series Management Console system requirements

IBM N series Management Console runs on Windows Server 2008, Windows Server 2003, Windows XP, Windows Vista, Red Hat Enterprise Linux, Oracle Enterprise Linux, and SUSE Linux Enterprise server platforms.

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

Note

IBM N series Management Console 3.0 is supported on DataFabric Manager 4.0 or later.

Minimum screen resolution: IBM N series Management Console applications require a minimum screen resolution of 1024 x 768 pixels.

Installing IBM N series Management Console

To install IBM N series Management Console on your Windows or Linux system, complete the following steps:

Step	Action	
1	From Operations Manager, select Download Management Console from the Setup menu.	
2	Click the link for the version of IBM N series Management Console required for your operating system: Windows or Linux.	
3	If you are installing on ...	Then ...
	Windows	Launch nmconsole-setup-3-0-win32.exe.
	Linux	Enter the following command to install nmconsole-setup-3-0-linux.rpm: <pre>rpm -i <rpm filename></pre> <p>Note_____</p> You must log in to Linux with root privileges.
4	Follow the setup prompts to complete the installation. <p>Note_____</p> The default installation path for IBM N series Management Console on Windows is C:\Program Files\IBM\N series Management Console. You can access IBM N series Management Console from the Start menu at Programs > IBM > N series Management Console. <p>The default installation path for IBM N series Management Console on Linux is /usr/lib/ibm/management_console/. You can access IBM N series Management Console from /usr/bin.</p>	

Note_____

You can install multiple versions of IBM N series Management Console by choosing a different install directory for each.

For instructions on how to start and use IBM N series Management Console, see the IBM N series *Operations Manager Administration Guide* and the Help.

Upgrading IBM N series Management Console

To upgrade IBM N series Management Console on your Windows system, complete the following steps:

Step	Action
1	Restart Windows.
2	Install IBM N series Management Console.

You can upgrade IBM N series Management Console on your Linux system using the following command:

```
rpm -U <rpm filename>
```

Alternatively, to upgrade IBM N series Management Console on your Linux system, complete the following steps:

Step	Action
1	Enter the following command to uninstall IBM N series Management Console: rpm -e nmconsole
2	Install IBM N series Management Console.

Locating information about the client configuration

From the Help menu in IBM N series Management Console, you can select the About option to get the following information about the software:

- ◆ Application name
- ◆ Build string
- ◆ Version string
- ◆ Copyright

In the About dialog box, click the Configuration Details button to get more specific information about the configuration of the client running IBM N series Management Console, such as the following:

- ◆ IBM N series Management Console installation directory
- ◆ Operating system architecture, name, and version
- ◆ Java version and associated memory
- ◆ DataFabric Manager version, serial number, host name, installation directory, and so on

In the About dialog box, click the Licenses button to display the installed data management products.

This type of information can be useful when you are troubleshooting problems or preparing to install software upgrades.

Migrating the DataFabric Manager database from Solaris

DataFabric Manager 3.8 or later is not supported on Solaris. However, you can migrate a DataFabric Manager database to a DataFabric Manager server on Windows or Linux.

Before you migrate the database, ensure the following:

- ◆ The installation directory has enough space.
The restore operation makes a temporary copy of the data in the backup. Therefore, you need at least twice as much space as the size of the backup in the new directory.
- ◆ The new server is on the same subnet as the old one.
In any case, all of the appliances that DataFabric Manager server is monitoring need to remain accessible.

To migrate the database on Solaris to a server running DataFabric Manager server 3.8 or later on Windows or Linux, complete the following steps:

Step	Action
1.	Enter the following command to create an archive-based backup on the Solaris server: <code>dfm backup create <backup_filename></code>

Step	Action	
2.	If you are migrating from...	Then...
	DataFabric Manager 3.4.1 or earlier on Solaris	<ul style="list-style-type: none"> a. Install DataFabric Manager 3.5 or later on Windows or Linux (excluding DataFabric Manager 3.8 or later). b. Copy the backup to the installed server. c. Enter the following command to restore the backup: dfm backup restore <backup_filename> d. Upgrade to DataFabric Manager 4.0 or later.
	DataFabric Manager 3.5 or later on Solaris	<ul style="list-style-type: none"> a. Install DataFabric Manager 4.0 on Windows or Linux. b. Copy the backup to the installed server. c. Enter the following command to restore the backup: dfm backup restore <backup_filename>

To ensure that you have the required operating system support for DataFabric Manager server, see the *Interoperability Matrix Tool* page at <http://www.ibm.com/storage/support/nas/>.

The following folders are not part of the archive-based backup:

- ◆ Reports
This folder contains the output of scheduled reports. You can use the `dfm options list reportsArchiveDir` command to locate the reports folder.
- ◆ Plug-ins
This folder contains storage system plug-ins. If you want these storage system plug-ins in DataFabric Manager server 4.0, then make sure that you

copy them manually to the new location. You can use the `dfm options list pluginsDir` command to locate the plug-ins folder.

◆ **Data**

This folder contains the DataFabric Manager database backups and the `monitordb.db` and `monitordb.log` files. You can use the `dfm options list databaseBackupDir` command to locate the plug-ins folder.

Note

You should not copy `monitordb.db` and `monitordb.log` files to DataFabric Manager server 4.0.

◆ **DataExport**

This folder contains the output of `dfm data export` command. You can use the `dfm options list dataExportDir` command to locate the plug-ins folder.

After you migrate, the administrators defined in the database might no longer be valid. Therefore, delete all the invalid DataFabric Manager users from the database, and re-add the valid users to the database on the migrated platform.

To delete all users using Operations Manager, complete the following steps:

Step	Action
1.	Log in to the administrator account.
2.	Click Setup > Administrative Users.
3.	Delete the users. Note You cannot delete root user on Linux.

To delete all users using the DataFabric Manager CLI, complete the following steps:

Step	Action
1.	Log in to the administrator account.
2.	Enter the following command to list all users: dfm user list

Step	Action
3.	<p data-bbox="490 234 981 269">Enter the following command to delete users:</p> <pre data-bbox="490 277 846 312">dfm user delete <user_name></pre> <p data-bbox="490 329 1229 407">Note _____ You cannot delete the root user on Linux. _____</p>

Note _____
On Windows, add Administrator as a user with the GlobalFullControl role.

Migrating the DataFabric Manager data to a different location

You can migrate the DataFabric Manager data to a different location using the `dfm datastore setup <dfm-data-dir>` command, where *dfm-data-dir* is the target location for the DataFabric Manager data. Besides configuring the DataFabric Manager database, this command copies the database, performance data, and script plug-in files in the specified target directory.

You can specify the target directories for the DataFabric Manager database, the script plug-ins, and the performance data by using the `dfm datastore setup -d <db-dir> -p <perf-dir> -s <script-plugins-dir>` command. The `-d`, `-p`, and `-s` options are used to set new locations for the DataFabric Manager database, performance data, and script plug-ins, respectively.

You can set up the DataFabric Manager server to use the data present at the new location by using the `dfm datastore setup -n <dfm-data-dir>` command. The `-n` option updates the configuration of the DataFabric Manager database to use the new data without copying the original data.

Example:

```
$ dfm datastore setup -n /mnt/dfm
```

Note

Do not run the `dfm` command while migrating the DataFabric Manager data. If commands are run, they can interfere with the migrate operation by locking the database tables and causing the operation to fail.

Configuring your database for Snapshot-based backups

You can quicken the backup process by using the Snapshot-based approach under the following conditions:

- ◆ The DataFabric Manager data resides on a LUN using either FC or iSCSI protocols.
- ◆ The volume holding the DataFabric Manager data is appropriately configured for SAN.

In this approach, you should disable the default Snapshot copies for the volume holding the DataFabric Manager data.

For more information about backups, see the section on setting up the DataFabric Manager database backup in the IBM N series *Operations Manager Administration Guide*.

To configure the DataFabric Manager database for Snapshot-based backups, complete the following steps:

Step	Action	
1	Install SnapDrive®.	
	If you are installing on....	Then use...
	Windows: <ul style="list-style-type: none"> ◆ Windows Server 2008 ◆ Windows Server 2003 SP1 or SP2 ◆ Windows Server 2003 R2 	<ul style="list-style-type: none"> ◆ SnapDrive 6.0 for Windows or later for Windows Server 2008. ◆ SnapDrive 4.1 for Windows or later for Windows Server 2003.
	Linux: <ul style="list-style-type: none"> ◆ SUSE <ul style="list-style-type: none"> ❖ SUSE Linux Enterprise Server 9 (SP3) ❖ SUSE Linux Enterprise Server 10 ◆ Red Hat <ul style="list-style-type: none"> ❖ Red Hat Enterprise Linux AS 4.0 Update 3 ❖ Red Hat Enterprise Linux Advanced Platform 5.0 	SnapDrive 2.2.1 for UNIX or later. Note _____ SnapDrive 3.0 for UNIX or earlier is not compatible with Red Hat Enterprise Linux Advanced Platform 5.0 and SUSE Linux Enterprise Server 10. _____
2	Create FC-based or iSCSI-based storage using SnapDrive.	
3	Enter the following command at the command line: dfm datastore setup <dfm-data-dir> <i>dfm-data-dir</i> is the target location for the DataFabric Manager data.	

About this chapter

This chapter describes how to configure DataFabric Manager 3.3.1 or later for high availability, using Microsoft Cluster Server (MSCS).

You must read the information in this chapter prior to installing DataFabric Manager software.

Note

Do not install DataFabric Manager until you install MSCS. For MSCS installation requirements and instructions, see the Microsoft documentation.

The following topics are described in this chapter:

- ◆ [“Planning to install DataFabric Manager with MSCS”](#) on page 38
- ◆ [“Configuring DataFabric Manager with MSCS”](#) on page 48
- ◆ [“Configuring DataFabric Manager with MSCS manually”](#) on page 52
- ◆ [“Managing DataFabric Manager with MSCS”](#) on page 60
- ◆ [“Upgrading cluster nodes with DataFabric Manager”](#) on page 67

Planning to install DataFabric Manager with MSCS

How DataFabric Manager with MSCS works

A cluster configured with DataFabric Manager consists of two nodes running DataFabric Manager 3.4.1 or later and configured for high availability using MSCS.

You configure DataFabric Manager services to be accessible through a network name and network address. DataFabric Manager and IBM N series Management Console can also use this network name or network address, so you do not need to add new network resources for the DataFabric Manager services. All DataFabric Manager data (database files, Performance Advisor files, and so on) are configured to be accessed from a shared data disk.

Cluster resources incorporate all the hardware and software components that exist in a cluster. This includes the DataFabric Manager services, the shared data disks, the network name, and the network address, among other components. At any time, these resources are all online on one of the two cluster nodes.

When any failure is detected, whether a node failure or a failure of one of the resources, all the resources are automatically moved, or failed over, to the partner node by MSCS.

This failover process is assisted by using a quorum resource on the cluster. The quorum resource is a physical storage device that can be accessed by both nodes in the cluster, although only one node has access to the quorum resource at any given time. The node that has access to the quorum resource is the node that has control of the cluster resource.

Overview of installing DataFabric Manager with MSCS

The overall process for using DataFabric Manager in a MSCS environment is as follows:

1. Install MSCS according to installation instructions provided by Microsoft and the following guidelines for DataFabric Manager:
 - ❖ See “[Supported configurations for DataFabric Manager with MSCS](#)” on page 39.
 - ❖ Modify the MSCS configuration for DataFabric Manager.
 - ❖ Configure MSCS with a shared data disk for a quorum resource, a network name, and a network address.
 - ❖ Place these cluster resources in a cluster resource group.

- ❖ See “[MSCS configuration requirements](#)” on page 40.
2. Configure the cluster nodes in preparation for the DataFabric Manager installation.
 - ❖ Set up new resources: domain users and shared data disks.
 - ❖ Add the new resources to the cluster resource group.
 - ❖ See “[Preparing to install DataFabric Manager with MSCS](#)” on page 41.
 3. Install DataFabric Manager.

See “[Installing DataFabric Manager with MSCS](#)” on page 43.
 4. Configure DataFabric Manager with configuration scripts or manually.

See “[Configuring DataFabric Manager with MSCS](#)” on page 48 or “[Configuring DataFabric Manager with MSCS manually](#)” on page 52.

Note

You should perform all cluster operations by using either Cluster Administrator or `cluster.exe`. Except where specifically indicated in installation and configuration procedures, do not use `dfm service start` and `dfm service stop`. These commands interfere with cluster operations.

Supported configurations for DataFabric Manager with MSCS

DataFabric Manager with MSCS is supported in the following configurations:

- ◆ Microsoft Windows servers running Windows Server 2003 Enterprise Edition or Data Center Edition, running the same patch versions on identical hardware

Note

MSCS is not supported on Windows Server 2008 and Windows Server 2008 R2. However, you can configure DataFabric Manager 3.8 or later for high availability on these platforms by using Failover Clustering. For more details, see <http://www.netapp.com/us/library/technical-reports/tr-3767.html>. This technical report contains information about NetApp products that IBM licenses and in some cases customizes. Technical reports might contain information about models and features that are not supported by IBM.

- ◆ Only two nodes in the cluster setup
- ◆ iSCSI-based storage for shared data disks with a IBM N series storage system as the storage back end
- ◆ Only the Single Quorum Device Cluster setup

- ◆ The same version of DataFabric Manager installed at the same path on both of the cluster nodes: for example, C:\Program Files\IBM\DataFabric Manager\DFM
- ◆ All DataFabric Manager administrators as domain users rather than local system users so that the user login works even when the DataFabric Manager services fail over to the partner node

Note

DataFabric Manager with MSCS is not supported on VMware.

MSCS configuration requirements

Before installing DataFabric Manager with MSCS, you must set up two Windows servers running on identical hardware platforms. Ensure that all of the requirements and guidelines for configuring cluster servers are followed, according to the MSCS documentation.

The configuration described in this document assumes that during the setup of MSCS you completed the following actions:

- ◆ Created a shared data disk to be used as a quorum resource
The shared quorum disk is used for storing the cluster configuration information.
- ◆ Created a network name resource and a network address resource
The network name and network address are used for managing the cluster server and DataFabric Manager.
- ◆ Added these cluster resources to a resource group
Putting the resources in a group makes all the resources available to both cluster nodes.

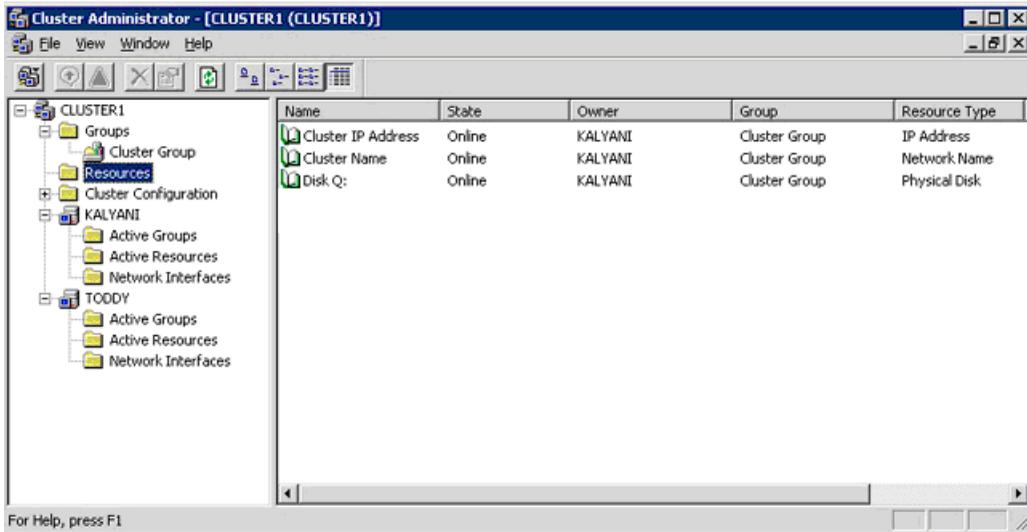
This is the configuration described in this document.

Example of a MSCS initial setup

The following image presents an example of a Cluster Administrator display after initial setup of MSCS. The example shows the resources, nodes, and groups as they would appear, assuming the following configuration:

- ◆ Cluster name is Cluster 1.
- ◆ First node name is Kalyani.
- ◆ Second node name is Toddy.
- ◆ Physical disk, the shared disk quorum resource, is created with the name Disk Q: and is mapped to drive letter Q:.

- ◆ Network Name resource is called Cluster Name.
- ◆ Network IP Address resource is called Cluster IP Address.
- ◆ All of the above resources are added to a resource group named Cluster Group.



You can also set up DataFabric Manager to use a different network name and network address. You can add these resources, along with the DataFabric Manager service resources and shared data disk resource, to another resource group.

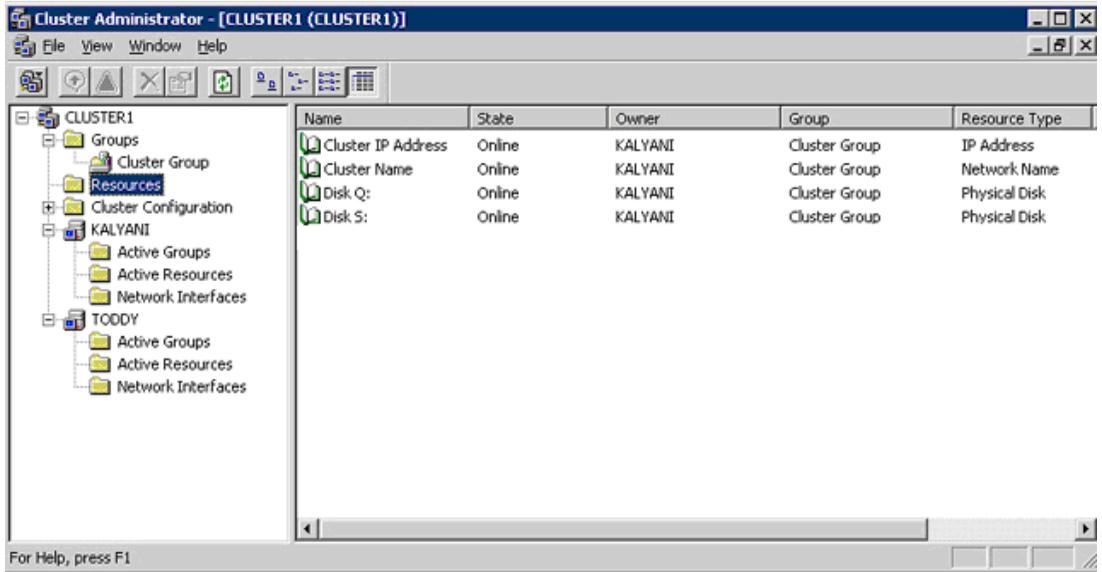
Preparing to install DataFabric Manager with MSCS

Before installing DataFabric Manager, you must set up users and shared resources, as indicated in the following procedure.

Step	Action
1	Select a domain user and add the domain user to the Administrators Group on both the cluster nodes. Example: domain\dfmuser

Step	Action
2	<p>Create a shared data disk:</p> <ul style="list-style-type: none"> a. Make the disk accessible to both the cluster nodes. b. Map the disk to a drive letter, such as drive S:. <p>Note _____ The data disk should be mapped to the same drive letter on both the cluster nodes.</p> <hr/> <ul style="list-style-type: none"> c. Add the shared data disk, as a physical disk resource, to the cluster server. <p>This disk is a resource for storing data specific to DataFabric Manager.</p>
3	<p>Verify that the resource group named Cluster Group can successfully fail over to the partner node.</p>

After completing these tasks, the Cluster Administrator displays the resources, nodes, and groups as shown in the following image. In addition to the content displayed after the initial setup, Cluster Administrator now shows a physical disk resource named Disk S:.



Installing DataFabric Manager with MSCS

For details about the DataFabric Manager installation process, see [“Installing and Upgrading DataFabric Manager”](#) on page 17.

Prerequisites:

- ◆ Microsoft Cluster Server must be properly installed and configured on both nodes of the cluster. Also, the pre-installation tasks discussed in "Preparing to install DataFabric Manager with MSCS" must be completed.
- ◆ Your workstation must meet the requirements described in "Installation and upgrade requirements."
- ◆ Ensure you have the information required to complete the installation, as described in "Installation and upgrade requirements."
- ◆ You must have Local Administrator login permission for the DataFabric Manager server.

To install DataFabric Manager on the cluster nodes, complete the following steps:

Step	Action
1	Log in to the first node of the cluster pair as a domain user with administrator privileges on the local system.

Step	Action	
2	<p>Ensure that this node owns the cluster resources.</p> <p>In the Cluster Administrator interface, check the Owner field of the Resources folder.</p>	
3	If you are...	Then...
	<p>Upgrading to DataFabric Manager 4.0</p>	<p>Back up your existing DataFabric Manager database.</p> <ul style="list-style-type: none"> ◆ During the upgrade: The DataFabric Manager 4.0 install wizard includes a backup step. You can choose to back up your database and specify the type of backup. <p>For a Snapshot-based backup, you must enter the backup filename in the following format:</p> <p>-t sndb <backup name></p> <p>If the backup does not finish successfully, abort the install wizard, perform a manual backup, and then restart the installation.</p> <ul style="list-style-type: none"> ◆ Before the upgrade: If you choose to skip the backup step during the upgrade, you should back up the database before you start the install wizard. For instructions, see the section on setting up the DataFabric Manager database backup in the <i>IBM N series Operations Manager Administration Guide</i>.
	<p>Installing DataFabric Manager 4.0</p>	<p>Go to Step 4.</p>
4	<p>Get the DataFabric Manager 4.0 installer from http://www.ibm.com/storage/support/nas/.</p>	
5	<p>Launch dfmsetup-4-0-win32.exe.</p>	

Step	Action
6	<p>Follow the DataFabric Manager setup prompts to complete the installation.</p> <p>Note the installation directory path for later reference.</p> <p>Example: C:\Program Files\IBM\DataFabric Manager\DFM</p> <p>Note_____</p> <p>The installation process can take several minutes.</p> <p>_____</p> <p>During a new installation, you must specify the Operations Manager license key. If you are upgrading from an earlier licensed version of DataFabric Manager, you do not need a license key.</p> <p>The DataFabric Manager 4.0 installation and upgrade process automatically installs the AutoSupport feature with AutoSupport enabled and displays a message about how to disable the feature.</p>
7	<p>When installation is complete, enter the following command to stop the DataFabric Manager services:</p> <pre>\$ dfm service stop</pre>
8	<p>Move the cluster resources to the second node by using the Move Group option.</p>
9	<p>Log in to the second node of the cluster pair, as a domain user with administrator privileges on the local system.</p> <p>You must log in with the same user name you used on the first node, in Step 1.</p>
10	<p>Install DataFabric Manager at the same directory path that you used on the first node.</p>
11	<p>Enter the following command to stop the DataFabric Manager services:</p> <pre>\$ dfm service stop</pre>

Step	Action
12	Disable the automatic startup of the DataFabric Manager server by entering the following command: <code>\$ dfm service enable -m</code>

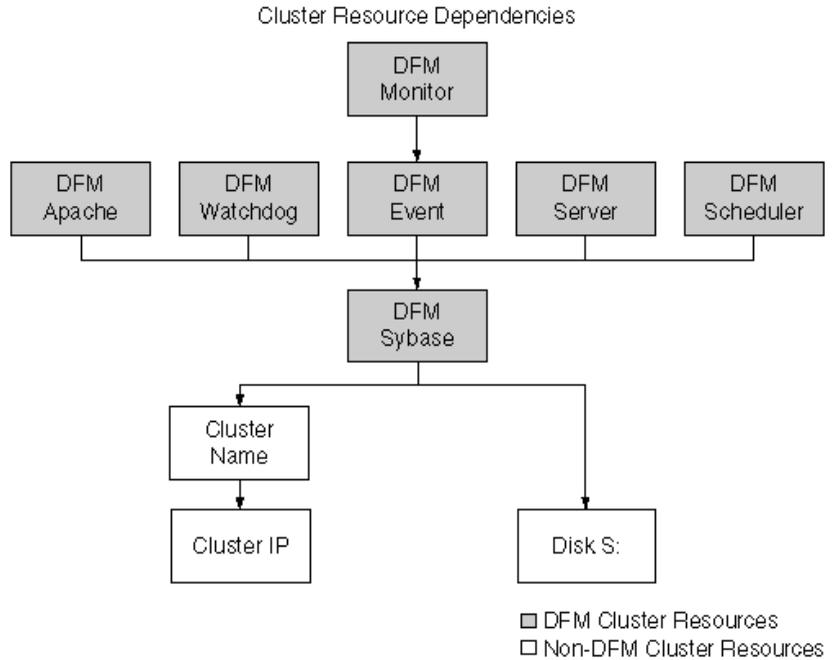
Note

Because the two DataFabric Manager nodes are configured to use the same database and to monitor the same set of nodes, you can use the same set of licenses for installation on both the nodes.

Block diagram of DataFabric Manager services and cluster resources

After completing the DataFabric Manager installations on both nodes of the cluster, you can configure them for high availability. You can perform setup configuration by using the configuration scripts that are provided with the installation, or you can perform the configuration manually.

The following block diagram describes the DataFabric Manager service resources that you need to add to MSCS and the dependencies among the various resources.



Configuring DataFabric Manager with MSCS

Configuring DataFabric Manager with MSCS

To configure DataFabric Manager with MSCS, complete the following steps:

Step	Action
1	Log in to the node that owns cluster resources.
2	Enter the following command to move the DataFabric Manager data onto the shared data disk: dfm datastore setup <drive name>
3	Stop DataFabric Manager services by entering the following command: dfm service stop Note _____This ensures that DataFabric Manager does not try to access the data disk to be moved to the secondary node. _____
4	Enter the following command: dfm service enable -m Note _____The option -m in the command ensures that the DataFabric Manager services do not start automatically. _____
5	Manually move the cluster group into the second node.
6	Verify that the secondary node owns the cluster resources.

Step	Action
7	Enter the following commands on the secondary node: <pre>dfm service enable -m dfm datastore setup -n <drive name></pre> <p>Note</p> <hr/> Ensure that the same drive letter is used for the secondary node as for first node. The option -n in the command ensures that the data is not copied again to the shared data disk. <hr/>
8	To add DataFabric Manager service to the MSCS cluster resource, go to the directory at C:\Program Files\IBM\DataFabric Manager\DFM\examples.
9	Run the following script to configure the DataFabric Manager services in MSCS: <pre>dfmcluster_add_resources.pl -g <cluster Group Name> -l <cluster IP resource Name> -n <Cluster Resource Name> -K <data disk Name></pre>

Adding a cluster resources setup script

You can use the following cluster setup script to add DataFabric Manager services as generic resources to the cluster. Run the script either on the first node or on the second node. Before running the script, ensure that you are logged in as the selected domain user. Verify that the current node owns all the cluster resources by reviewing the Resources details pane.

dfmcluster_add_resources.pl syntax:

```
perl dfmcluster_add_resources.pl <option> ...
```

Following is an example of the script:

```
$ perl dfmcluster_add_resources.pl -g "Cluster Group" -i "Cluster IP Address" -n "Cluster Name" -k "Disk S:"
```

Option	Description
-g <cluster-group>	Name of the cluster group to which the resources are added. This should be the name of the group in which the other resources already exist.
-i <cluster-ip-resource>	Name of the Cluster IP resource as it appears under the Name column in the Cluster Administrator.
-n <cluster-name-resource>	Name of the Cluster Name resource.
-k <data-disk-resource>	Name of the Data Disk resource.

dfmcluster_add_resources.pl operations:

This cluster setup script adds the following Generic Service resources to the MSCS configuration:

- ◆ DFM Sybase
- ◆ DFM Apache
- ◆ DFM Server
- ◆ DFM Event
- ◆ DFM Scheduler
- ◆ DFM Watchdog
- ◆ DFM Monitor

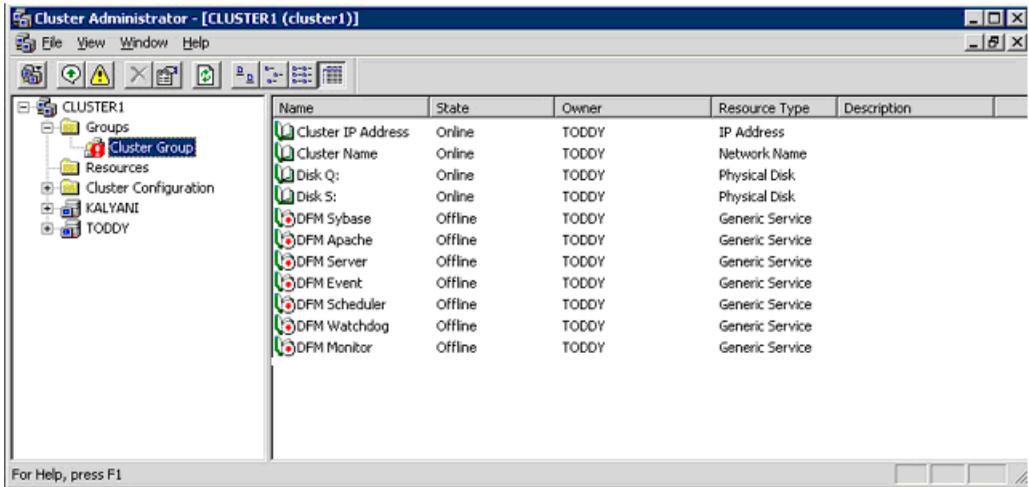
Bringing cluster services online

After you complete the setup scripts, DataFabric Manager with MSCS is ready, but the services are still offline.

To bring the cluster services online, complete the following step:

Step	Action
1	Select Cluster Group under the Groups folder and click File > Bring Online.

The following image is an example of Cluster Administrator displaying the Resource Types discussed previously, including the newly added Generic Services, shown as being offline.



Configuring DataFabric Manager to use the cluster name

By default, DataFabric Manager uses the local system name in the links in alert e-mail messages sent to administrators.

To make DataFabric Manager use the cluster name instead of the local system name, so that the URLs are always accessible, set the following option:

```
$ dfm option set localHostName=<fqdn-of-cluster>
```

Configuring DataFabric Manager with MSCS manually

You can use this section as a checklist for manually configuring DataFabric Manager with MSCS. The advantage of using the following checklist is that you can customize the way the cluster is set up. For example, you can move the different DataFabric Manager data files (database files, Performance Advisor files) to different shared data disks. You cannot move the data files using the configuration scripts described in the earlier section.

Tasks involved of manual setup

The manual setup of DataFabric Manager with MSCS involves the following tasks:

1. Configure the first node.
2. Configure the second node.
3. Configure the cluster resources.

Configuring the first node to access all data files from a shared disk

You can access the shared data drive by logging in as a domain user with the appropriate privileges.

Note

By default, DataFabric Manager services run using the local system account. This default does not provide access to the shared drive where the database files and other files reside.

To configure the DataFabric Manager services to log in as a domain user with administrative privileges on the local system, complete the following steps:

Step	Action
Log in	
1	Log in to the first node as the selected domain user with administrator privileges on the local system.
2	<p>Ensure that the first node currently owns the resource group named Cluster Group.</p> <p>In the Cluster Administrator interface, check the Owner field of the Resources folder.</p>
3	<p>Enter the following command to stop DataFabric Manager services:</p> <pre>\$ dfm service stop</pre>
Specify the account	
4	<p>Specify a Logon As user account:</p> <ol style="list-style-type: none"> a. From the Start menu, select Control Panel. b. In the Control Panel, select Administrative Tools > Services. c. Double-click DFM Sybase ASA. d. In the General tab, change the Startup type option to Manual. <p>Note _____ Setting the Startup type to manual is very important, because the services should not be started automatically on a node reboot. All service starting and stopping should be done only by the cluster server.</p>
5	Click Apply.
6	Click the Log On tab.
7	Enter the name of the domain user account you want to use to access the DataFabric Manager service from the shared drive.
8	Click OK.

Step	Action
Repeat the service, and perform other configuration tasks	
9	<p>Repeat Step 4 through Step 8 for each additional service.</p> <p>Each time, replace the service name in Step 4c with one of the following services:</p> <ul style="list-style-type: none"> ◆ DFM Apache ◆ DFM Server ◆ DFM Scheduler ◆ DFM Event ◆ DFM Watchdog ◆ DFM Monitor <p>Note_____</p> <p>Step 4 through Step 8 must be performed for all services, for the cluster to work properly.</p> <p>_____</p>
10	<p>Move the database files to a nonroot folder in the shared data drive.</p> <p>An example of a shared data folder is S:\dfm\data.</p> <p>The default location for database files is <installation-directory>\data.</p> <p>Note_____</p> <p>This step is required only on the first node you configure to be part of a cluster.</p> <p>_____</p>
11	<p>Point DataFabric Manager to the relocated database.</p> <p>Example:</p> <pre>\$ dfm database set dbDir=S:\dfm\data \$ dfm database set dbLogDir=S:\dfm\data</pre>
12	<p>Enter the following command to verify that all the services are still stopped:</p> <pre>\$ dfm service list</pre>

Step	Action
13	<p>Move all other data files to folders in the shared disk.</p> <ul style="list-style-type: none"> ◆ Move the Performance Advisor data files to a folder in the shared data drive. The default location for performance data files is <installation directory>\perfdata (for example, S:\dfm\perfdata). ◆ Move the Script plug-in files to a folder in the shared data drive. The default location for performance data files is <installation directory>\script-plugins (for example, S:\dfm\script-plugins). ◆ Move the Configuration Management plug-in files to a folder in the shared data drive. The default location for plug-in files is <installation directory>\plugins (for example, S:\dfm\ plugins). ◆ Move the archived reports to a folder in the shared data drive. The default location for the archived reports is <installation directory>\reports (for example, S:\dfm\reports).
14	<p>Enter the following command to start the SQL service:</p> <pre>\$ dfm service start sql</pre>
15	<p>Set options to point DataFabric Manager to the new location of the relocated files.</p> <pre>\$ dfm option set perfArchiveDir=S:\dfm\perfdata \$ dfm option set pluginsDir=S:\dfm\plugins \$ dfm option set scriptDir=S:\dfm\script-plugins \$ dfm option set reportsArchiveDir=S:\dfm\reports</pre> <p>Note_____</p> <p>Do not start the services after setting each option as suggested by the messages printed by the <code>dfm option set</code> command.</p> <p>_____</p>
16	<p>Enter the following command to stop the SQL service:</p> <pre>\$ dfm service stop sql</pre>

Configuring the second node to access all data files from the shared disk

You can access the shared data drive by logging in as a domain user with the appropriate privileges.

To configure the second node to log in as a domain user with administrative privileges on the local system, complete the following steps:

Step	Action
1	Log in to the second node as the selected domain user with administrator privileges on the local system.
2	Ensure that the second node currently owns the resource group named Cluster Group. In the Cluster Administrator interface, check the Owner field of the Resources folder.
3	Repeat Step 2 through Step 9 , and then Step 11 , as described in “ Configuring the first node to access all data files from a shared disk ” on page 52.
4	After you complete the previous steps, enter the following command to ensure that all the services are stopped: <code>\$ dfm service list</code>

Configuring DataFabric Manager services as cluster resources

After you install the DataFabric Manager services, you need to configure them as cluster resources and make them available to both nodes.

Prerequisite: Determine any dependencies that exist between various cluster resources. A dependency requires that one service be running before its associated service can be brought online. For example, most services cannot function unless Sybase ASA is already running.

To configure DataFabric Manager services as cluster resources, complete the following steps:

Step	Action
Log in, open Cluster Administrator, and select Resource	
1	Log in to the node as the selected domain user with administrator privileges on the local system.
2	Ensure that the node currently owns all the cluster resources. In the Cluster Administrator interface, check the Owner field of the Resources folder.

Step	Action
3	Open Cluster Administrator.
4	In the console tree, double-click the Groups folder.
5	In the details pane, click the group named Cluster Group.
6	<p>On the File menu, select New > Resource.</p> <p>Select or type the appropriate options on each UI page, for each resource.</p>
Install the cluster services, starting with the DFM Sybase ASA service	
7	<p>On the New Resource page, complete the following steps:</p> <ul style="list-style-type: none"> a. Enter the name of the resource in the Name field. <p style="margin-left: 40px;">Example:</p> <p style="margin-left: 40px;">Name = DFM Sybase ASA</p> <ul style="list-style-type: none"> b. Choose Generic Service as Service Type. c. Select Cluster Group as the group. d. Click Next.
8	<p>On the Possible Owners page, complete the following steps:</p> <ul style="list-style-type: none"> a. Add both nodes as the possible owners of the resource. b. Click Next.

Step	Action
9	<p>On the Dependencies page, complete the following steps:</p> <ul style="list-style-type: none"> a. Add Dependencies related to the new service. <p>Example:</p> <p>DFM Sybase ASA Dependencies = Cluster IP Address, Cluster Name, Disk S:, Data Disk</p> <ul style="list-style-type: none"> b. Click Next.
10	<p>On the Generic Service Parameters page, complete the following steps:</p> <ul style="list-style-type: none"> a. Set Service Name. <p>Example:</p> <p>Service Name = DFM Sybase</p> <ul style="list-style-type: none"> b. Leave the "Use Network name for computer name" option unchecked.
11	<p>On the Registry Replication page, click Finish.</p> <p>There is no registry replication required.</p>
Add the remaining cluster service	
12	<p>Repeat Step 7 through Step 11 to add the remaining cluster services.</p> <p>Note_____</p> <p>Step 7 through Step 11 must be performed for all services, for the cluster to work properly.</p> <p>_____</p>

The following table indicates Resource Name, Dependencies, and Service Name to enter for each new cluster service.

Resource Name field	Dependencies field	Service Name field
DFM Sybase ASA	Cluster I, Cluster Name, Shared Data Disk, Data Disk	DFMSybase
DFM Apache	DFM Sybase ASA	DFMApache
DFM Scheduler	DFM Sybase ASA	DFMScheduler
DFM Watchdog	DFM Sybase ASA	DFMWatchdog
DFM Server	DFM Sybase ASA	DFMServer
DFM Event	DFM Sybase ASA	DFMEvent
DFM Monitor	DFM Event	DFMMonitor

Bringing cluster services online (manual setup)

After you complete the manual setup, all the DataFabric Manager services are listed under the Cluster Group resource group. DataFabric Manager with MSCS is ready, but the services are still offline.

To bring the cluster services online, complete the following step:

Step	Action
1	Select Cluster Group under the Groups folder and click File > Bring Online.

Configuring DataFabric Manager to use the cluster name (manual setup)

By default, DataFabric Manager uses the local system name in the links in alert e-mail messages sent to administrators.

To make DataFabric Manager use the cluster name instead of the local system name so that the URLs are always accessible, set the following option:

```
$ dfm option set localHostName=<fqdn-of-cluster>
```

Managing DataFabric Manager with MSCS

Ways to start and stop DataFabric Manager services

After you set up DataFabric Manager with MSCS, do not use the `dfm service start` and `dfm service stop` commands, except where specifically indicated in installation and configuration procedures. These commands interfere with the working of MSCS. Instead, you should perform all operations by using either the Cluster Administrator or the `cluster.exe` command.

Also, do *not* change the Service startup type to Automatic in the Service Control Manager on any of the nodes. Keep this option set to “manual.”

Backing up the DataFabric Manager database

The database backup improvements in DataFabric Manager 3.4.1 and later releases do not require you to stop the database service while backing up the database. Therefore, there are no additional steps to be performed to back up the database.

Restoring the DataFabric Manager database

To restore from an existing backup of the database, complete the following steps:

Step	Action
1	Log in to the node that currently owns the cluster resources.
2	Using Cluster Administrator, take the services offline: <ul style="list-style-type: none">a. Right-click the DFM Sybase service.b. Select Take offline. This also takes the other services offline.
3	Enter the following command to restore the database: <code>\$ dfm backup restore</code>
4	Enter the following command to stop all the services: <code>\$ dfm service stop</code>

Step	Action
5	Using Cluster Administrator, bring the DataFabric Manager services online: <ul style="list-style-type: none"> a. Select Cluster Group under the Groups folder. b. Click File > Bring Online.

Configuring DataFabric Manager Web UI to use HTTPS

To set up DataFabric Manager Web UI to use HTTPS, complete the following steps:

Step	Action
1	Log in to the first node in the cluster.
2	Using Cluster Administrator, take the DataFabric Manager services offline: <ul style="list-style-type: none"> a. Right-click the DFM Sybase service. b. Select Take offline. <p>This also takes the other services offline.</p>
3	On the command line, enter the following command to start the SQL service: <pre>\$ dfm service start sql</pre>
4	Enter the following command to set up HTTPS, by creating an SSL certificate: <pre>\$ dfm ssl server setup</pre> <p>This creates two files, server.crt and server.key, in the <install-dir>conf folder. Copy these files to the second node before starting the services on that node.</p>
5	Enter the following command to enable HTTPS, by setting the DataFabric Manager option: <pre>\$ dfm option set httpsEnabled=yes</pre>

Step	Action
6	<p>Enter the following command to start the HTTP service:</p> <pre>\$ dfm service start http</pre> <p>Starting the service using <code>dfm service start</code> re-creates the <code>httpd.conf</code> file with the changed options.</p>
7	<p>Enter the following command to stop all the services:</p> <pre>\$ dfm service stop</pre> <p>Note_____</p> <p>Ensure that DataFabric Manager services are still in the offline status before proceeding. If the services are online, take them offline and then proceed to Step 8. Otherwise the HTTP service fails to come up on the other node because the configuration is not complete.</p>
8	<p>Using Cluster Administrator, move the cluster group to the second node by using the Move Group option.</p>
9	<p>Log in to the second node in the cluster.</p>
10	<p>Copy the <code>server.crt</code> and <code>server.key</code> files created on the first node to the <code><install-dir>\conf</code> folder.</p>
11	<p>From the command line, enter the following command to start the services and verify they are operating properly:</p> <pre>\$ dfm service start</pre> <p>Starting the service by using <code>dfm service start</code> re-creates the <code>httpd.conf</code> file with the changed options.</p>
12	<p>Enter the following command to stop the services:</p> <pre>\$ dfm service stop</pre>
13	<p>Using Cluster Administrator, bring the DataFabric Manager services online:</p> <ol style="list-style-type: none"> a. Select Cluster Group under the Groups folder. b. Click File > Bring Online.

Changing HTTP options

To change the HTTP options `httpEnabled`, `httpPort`, `httpsEnabled`, `httpsPort`, complete the following steps:

Step	Action
1	Log in to the first node in the cluster.
2	Using Cluster Administrator, take the DataFabric Manager services offline: <ol style="list-style-type: none"> a. Right-click the DFM Sybase service. b. Select Take offline. This also takes the other services offline.
3	From the command line, enter the following command to start the SQL service: <pre>\$ dfm service start sql</pre>
4	Enter the following command to set the required HTTP option: <pre>\$ dfm option set <option-name>=<option-value></pre> Example: <pre>\$ dfm option set httpsPort=443</pre>
5	Enter the following command to start the HTTP service: <pre>\$ dfm service start http</pre> Starting the service by using <code>dfm service start</code> re-creates the <code>httpd.conf</code> file with the changed options.
6	Enter the following command to stop all the services: <pre>\$ dfm service stop</pre> Note Ensure that DataFabric Manager services are still in the offline status before proceeding. If the services are online, take them offline and then proceed to Step 8 . Otherwise the HTTP service fails to come up on the other node because the configuration is not complete.
7	Using Cluster Administrator, move the cluster group to the second node by using the Move Group option.
8	Log in to the second node in the cluster.

Step	Action
9	From the command line, enter the following command to start the services: <code>\$ dfm service start</code>
10	Enter the following command to stop the services: <code>\$ dfm service stop</code>
11	Using Cluster Administrator, bring the DataFabric Manager services online: <ul style="list-style-type: none"> a. Select Cluster Group under the Groups folder. b. Click File > Bring Online.

Monitoring DataFabric Manager with MSCS

Microsoft provides Server Clusters Management Pack as part of Microsoft Operations Manager. This provides the capability to monitor the cluster server, and to report node status, resource status, and alerts.

DataFabric Manager does not provide any additional cluster monitoring or alerting functionality.

Data that can be shared by DataFabric Manager server cluster nodes

The following table describes what data can be shared by the DataFabric Manager cluster nodes.

You should configure DataFabric Manager nodes to access these files from a shared disk. If each node uses its own local copy of files, updates to files might not be accessible to the other nodes after a failover. For example, if a new storage system configuration management plug-in is installed on one node, it is accessible to only that node.

Default path	Description
<install-dir>\data	Sybase database files
<install-dir>\perfdata	Performance Advisor data files
<install-dir>\script-plugins	Installed script plug-ins and related files

Default path	Description
<install-dir>\plugins	Storage system configuration plug-ins
<install-dir>\reports	Archived reports
<install-dir>\dataExport	DataFabric Manager and Performance Advisor data

Data that is not shared by DataFabric Manager server cluster nodes

The following table describes what data is not shared by the DataFabric Manager cluster nodes.

Default path	Description
<install-dir>\bin	Executable files
<install-dir>\conf	Configuration files
<install-dir>\docs	Third-party licenses
<install-dir>\examples	Cluster configuration scripts, and so on.
<install-dir>\log	Log files
<install-dir>\misc	Configuration files
<install-dir>\sbin	Third-party executables
<install-dir>\scripts	Sybase_install.sql
<install-dir>\src	Storage system configuration plug-ins
<install-dir>\web\clients	Performance Advisor clients
<install-dir>\web\com	JAR files for applets
<install-dir>\web\help	Help files
<install-dir>\web\man	Manual (man) pages
<install-dir>\web\media	Images used on Web interfaces
<install-dir>\web\scripts	Java script files
<install-dir>\web\styles	CSS style sheets

Default path	Description
<install-dir>\perfExport	Exported performance counter data for specified objects

Uninstalling DataFabric Manager from a cluster

To uninstall DataFabric Manager from both the cluster nodes and disable the cluster setup, complete the following steps:

Step	Action
1	Using the Cluster Administrator, delete all the DataFabric Manager services: <ul style="list-style-type: none"> a. Right-click the DFM Sybase resource. b. Select Delete. This also deletes all other DataFabric Manager services.
2	Log in to either cluster node.
3	Select Uninstall from the Add/Remove Programs tool.
4	Complete the procedure for removing the program from your system.
5	Repeat Step 2 through Step 4 for the other cluster node.

Upgrading cluster nodes with DataFabric Manager

Upgrading DataFabric Manager cluster nodes

If you want to upgrade the cluster nodes, you should upgrade all of the nodes together.

To upgrade all the nodes together, complete the following steps:

Step	Action
1	Using Cluster Administrator, take the DataFabric Manager services offline: <ul style="list-style-type: none">a. Right-click the DFM Sybase service.b. Select Take offline. This takes the other services offline also.
2	To upgrade the first node, ensure that the first node currently owns all the cluster resources. In the Cluster Administrator interface, check the Owner field of the Resources folder.
3	Upgrade the DataFabric Manager installation on this node: <ul style="list-style-type: none">a. Back up your existing DataFabric Manager database.b. Get the DataFabric Manager 4.0 installer from http://www.ibm.com/storage/support/nas/, and launch dfmsetup-4-0-win32.exe.c. Follow the DataFabric Manager setup prompts to complete the installation. See “ Installing DataFabric Manager 4.0 on Windows ” on page 18 for details.
4	Enter the following command to stop all the DataFabric Manager services: \$ dfm service stop

Step	Action
5	Enter the following command to disable the automatic service startup during reboot: \$ dfm service enable -m
6	To upgrade the second node, ensure that the second node currently owns all the cluster resources. In the Cluster Administrator interface, check the Owner field of the Resources folder.
7	Upgrade the DataFabric Manager installation on the second node: <ol style="list-style-type: none"> a. Back up your existing DataFabric Manager database. b. Get the DataFabric Manager 4.0 installer from http://www.ibm.com/storage/support/nas/, and launch dfmsetup-4-0-win32.exe. c. Follow the DataFabric Manager setup prompts to complete the installation. See “Installing DataFabric Manager 4.0 on Windows” on page 18 for details.
8	Enter the following command to stop all the DataFabric Manager services: \$ dfm service stop
9	Disable the automatic service startup during reboot by entering the following command: \$dfm service enable -m
10	Using Cluster Administrator, bring the DataFabric Manager services online: Select Cluster Group under the Groups folder and click File > Bring Online.

About this chapter This chapter describes how to configure DataFabric Manager 4.0 server with Veritas Cluster Server (VCS) for high availability. VCS provides automatic failover for shared storage. In case of application failures or host failures, you can access the DataFabric Manager services quickly from another host with minimum or no interruption.

Additional reading For more information about the VCS requirements, guidelines, and management, see the following documentation:

- ◆ *Veritas Storage Foundation and High Availability Solutions 5.0 Getting Started Guide*
- ◆ *Veritas Cluster Server 5.0 Installation Guide*
- ◆ *Veritas Cluster Server 5.0 User's Guide*

Sections in this chapter

This chapter describes the following tasks:

- ◆ [“Planning to install DataFabric Manager with VCS”](#) on page 70
- ◆ [“Configuring DataFabric Manager with VCS”](#) on page 76
- ◆ [“Managing DataFabric Manager with VCS”](#) on page 83
- ◆ [“Upgrading cluster nodes with DataFabric Manager”](#) on page 90

Planning to install DataFabric Manager with VCS

How VCS works with DataFabric Manager server

A cluster configured with DataFabric Manager server consists of two nodes running DataFabric Manager 3.7 or later. Using VCS, the cluster is configured for high availability.

You can make DataFabric Manager services accessible by configuring a network address. DataFabric Manager and IBM N series Management Console can also use this network name or network address, so you do not need to add new network resources for DataFabric Manager services.

All DataFabric Manager server data (database files, Performance Advisor files, and so on) are configured to be accessed from a shared data disk.

The hardware and software components that exist in a cluster are called "cluster resources." This includes the DataFabric Manager services, the shared data disks, the network name, and the network address, among other components. These resources remain online on one of the two cluster nodes at any instance.

When a resource of node failure is detected, all the resources are automatically moved, or failed over, to the partner node by VCS.

To know more about the VCS failover process, see *Veritas Cluster Server 5.0 User's Guide*.

Overview of installing DataFabric Manager server with VCS

An overview of the process for using DataFabric Manager server in a VCS environment is as follows:

1. Install VCS according to the instructions given in the *Veritas Cluster Server 5.0 Installation Guide*.
2. Ensure that the correct configurations exist on your system.
For more information, see "[Configuration requirements for DataFabric Manager server with VCS](#)" on page 71.
3. Install DataFabric Manager server.
For more information, see "[Installing DataFabric Manager server on cluster nodes](#)" on page 73.
4. Configure DataFabric Manager server on the first node to use the data from a shared disk.

For more information, see “[Configuring DataFabric Manager server on the first node](#)” on page 76.

5. Configure DataFabric Manager server using configuration scripts or manually.

For more information, see “[Using configuration scripts to configure DataFabric Manager services for clustering](#)” on page 77 or “[Manually configuring DataFabric Manager Server for clustering](#)” on page 80.

6. Configure DataFabric Manager server on the second node to use the data from a shared disk.

For more information, see “[Configuring DataFabric Manager server on the second node](#)” on page 78.

Note

Perform all cluster operations using Cluster Manager. Other than in installation and configuration procedures, do not use the commands `dfm service start` and `dfm service stop`. These commands interfere with cluster operations.

Configuration requirements for DataFabric Manager server with VCS

The following configuration requirements should be met before configuring DataFabric Manager with VCS:

- ◆ Both cluster nodes are running on the supported version of the operating system at the suggested patch level.

The minimum supported versions for Linux are Red Hat Enterprise Linux 4 Update 3 and SUSE Linux Enterprise Server 9 with SP3.

- ◆ DataFabric Manager server is using Veritas File System and Volume Manager on Native ext3 File System and Logical Volume Manager (LVM) on Linux.
- ◆ You are running Veritas Storage Foundation™ and High Availability Solutions 5.0 with Maintenance Pack 1 (MP 1).
- ◆ You have FC connectivity to storage systems having DataFabric Manager server data.

Also ensure that the FC link is active and that LUNs created on the storage systems are accessible to both cluster nodes. The shared data disk should have sufficient space to accommodate DataFabric Manager server database, performance data, and script plug-in folders.

- ◆ There is a minimum of two network interfaces set up on each system.

At least two network interfaces should be set up: one for node-to-node communication and the other for node-to-client communication. The name

of the network interface used for node-to-client communication should be the same on both the systems.

- ◆ The same version of DataFabric Manager server is installed at the same path on both cluster nodes.
- ◆ A heartbeat link is established between cluster nodes.

Note

DataFabric Manager server supports only a two-node cluster setup.

Configuration limitations with VCS

The following are configuration limitations with VCS:

- ◆ DataFabric Manager server with VCS is not supported on VMware.
- ◆ DataFabric Manager server supports a two-node cluster setup only.
- ◆ An FC-based storage for shared data disks is initially required as the storage back end; however, iSCSI can also be used as a storage option after the qualification is complete.

Steps for configuring VCS and creating shared storage

Ensure that all of the requirements and guidelines for configuring cluster servers are followed according to the VCS documentation.

To configure VCS, complete the following steps:

Step	Action
1	Install Veritas Storage Foundation and High Availability Solutions 5.0. On Linux, Veritas File System and Veritas Volume Manager components are not required because Native ext3 File System and LVM are used.
2	Configure VCS: a. Enter the following command: \$ installvcs -configure b. Provide the network address (virtual IP address).
3	Use SnapDrive for UNIX to create file systems and logical volumes.

Installing DataFabric Manager server on cluster nodes

Before you install DataFabric Manager server on cluster nodes, ensure the following:

- ◆ VCS is properly installed and configured on both nodes of the cluster.
- ◆ The same version of DataFabric Manager server is installed on each of the two nodes.
- ◆ The installation directory is the same on both nodes: for example, /opt/IBMdsm.
- ◆ The first node owns the cluster resources.
- ◆ You have the root privileges required to log in to the DataFabric Manager server.

To install DataFabric Manager server on the cluster nodes, complete the following steps:

Step	Action	
1	Log in to the first node of the cluster pair.	
2	If you are...	Then you should...
	Upgrading to DataFabric Manager 4.0	<p>Choose one of the following backup options:</p> <ul style="list-style-type: none"> ◆ During the upgrade: The DataFabric Manager 4.0 install wizard includes a backup step. You can choose to back up your database and specify the type of backup. For a Snapshot-based backup, you must enter the backup filename in the following format: -t sndb <backup name> <p>If you choose this option, the DataFabric Manager SQL service must be running to perform the backup. If the backup does not finish successfully, abort the install wizard, perform a manual backup, and then restart the installation.</p> <ul style="list-style-type: none"> ◆ Before the upgrade: If you choose to skip the backup step during the upgrade, you should back up the database before you start the install wizard. For instructions, see the section on setting up the DataFabric Manager database backup in the IBM N series <i>Operations Manager Administration Guide</i>.
	Installing DataFabric Manager 4.0	Go to Step 3 .
3	Get the DataFabric Manager 4.0 installer from http://www.ibm.com/storage/support/nas/ .	

Step	Action
4	For Linux, launch <code>dfmsetup-4-0-linux.sh</code> .
5	Follow the DataFabric Manager server setup prompts to complete the installation. Note _____ The installation process takes few minutes to complete. _____
6	By default, DataFabric Manager server installs itself in the directory <code>/opt</code> . To change the installation directory, enter the following command: -d <new directory>
7	For a new installation rather than an upgrade from a previously licensed version of DataFabric Manager server, specify the DataFabric Manager server license key.
8	If you want to disable AutoSupport, which is automatically installed and enabled as part of the DataFabric Manager 4.0 installation, enter the following command: \$ dfm option set autosupportEnabled=no
9	When installation is complete, stop the DataFabric Manager server services by entering the following command: \$ dfm service stop
10	Disable the automatic startup of the DataFabric Manager server by entering the following command: \$ dfm service enable -m

Note _____

Because the two DataFabric Manager server nodes are configured to use the same database and to monitor the same set of nodes, the same set of licenses for installation is used on both the nodes.

Configuring DataFabric Manager with VCS

You can configure DataFabric Manager server with VCS by completing the following steps:

- ◆ Configure DataFabric Manager server on the first node.
- ◆ Configure DataFabric Manager services for clustering, either manually or using configuration steps.
- ◆ Configure DataFabric Manager server on the second node.
- ◆ Bring cluster services online.

These steps are discussed in the following sections.

Configuring DataFabric Manager server on the first node

To configure DataFabric Manager server on the first node, complete the following steps:

Step	Action
1	Ensure that the first node owns all the cluster resources, such as Mount, NIC, and Volume.
2	Move all the shared data, such as database files and performance data files, to the shared storage system by entering the following command: <code>\$ dfm datastore setup</code>

The following table describes the data shared by the DataFabric Manager cluster nodes.

Default data location	Details
<code>+/opt/IBMdfm/data</code>	Sybase database files
<code>+/opt/IBMdfm/perfdata</code>	Performance Advisor data files
<code>+/opt/IBMdfm/script-plugins</code>	Installed script plug-ins
<code>+/opt/IBMdfm/reports</code>	Reports Archive directory
<code>+/opt/IBMdfm/plugins</code>	Storage system configuration
<code>+/opt/IBMdfm/web/cms</code>	CMS TOCs, ACLs, and error files

Note

Log and configuration files are not shared by cluster nodes.

Using configuration scripts to configure DataFabric Manager services for clustering

DataFabric Manager server can be configured either manually, or using the Perl configuration script.

Before running the script, ensure that the node is the current owner of the cluster resources, and all the DataFabric Manager services are stopped on all the nodes.

The Perl script is located at \$ InstallDir/examples. The syntax for running the script is as follows:

```
$ perl dfmcluster_add_resources.pl [-t cluster type] [option] ...
```

The following table lists and describes all of the options you can use with this command.

Option	Description
-t [cluster type]	Cluster solution used for high availability. Possible values are vcs and mscs . Default value is vcs on UNIX.
-h [cluster-node1] [cluster-node2]	Nodes used for cluster setup, separated by space.
-g [cluster-group name]	Name of the cluster service group to which the resources are added.
-e [nic resource name]	Name of the Network Interface Card. This must be same on both cluster nodes.
-i [cluster-ip resource name]	Name of the cluster IP resource.
-n [cluster-name resource name]	Host name of the cluster (mapped to the cluster virtual IP address).
-f [mount-point resource name]	Name of the mount resource.
-v [volume resource name]	Name of the volume resource that contains the file-system represented by mount-point resource.

Option	Description
-d [disk-group resource name]	Name of the disk group that contains the volume represented by volume resource.
-m [netmask]	Netmask associated with the cluster IP address.
-l [install-dir]	DataFabric Manager installation directory. Default values is /opt/IBMdfm.

For instructions on manual configuration, see “[Configuring DataFabric Manager server on the second node](#)” on page 78.

Configuring DataFabric Manager server on the second node

Follow these steps to configure the second cluster node:

Step	Action
1.	Perform manual failover so that the second node is the owner of all the resources like disks and virtual IP addresses: <ol style="list-style-type: none"> a. At the command prompt, type the <i>hagui</i> command to open Cluster Manager. b. Right-click the service group, click Switch To and select the second cluster node for failover.
2.	Bring the services offline on the second node.
3.	Configure the node to use the shared data by entering the following command: <pre>\$ dfm datastore setup -n</pre>

Note

Ensure that the -n option is enabled when you configure the second node. This ensures that DataFabric Manager server uses the data that was copied during the configuration of the first node.

Bringing cluster services online

After you have configured DataFabric Manager on both nodes and also configured DataFabric Manager for clustered services, you must still bring the services online.

To bring the cluster services online, complete the following steps:

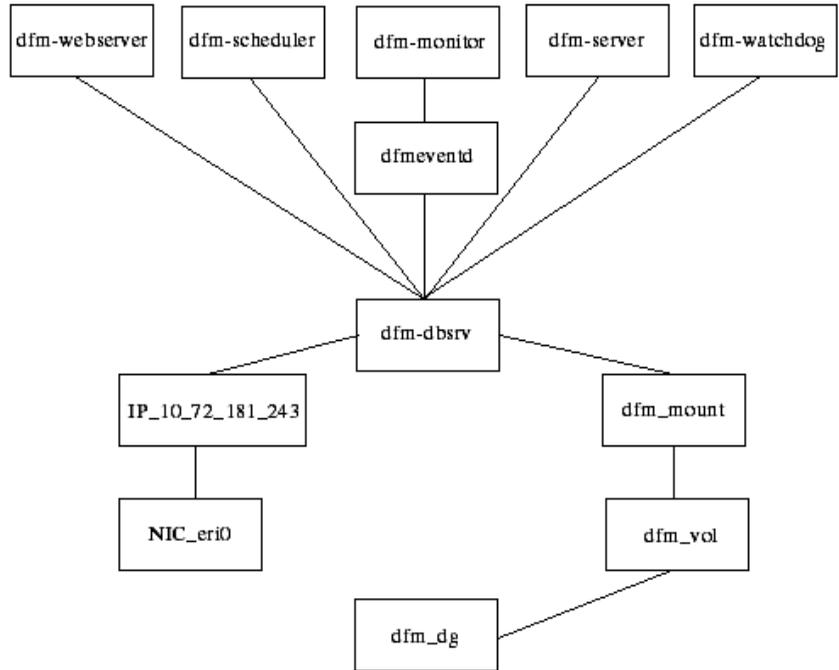
Step	Action
1	Run <i>hagui</i> in the command prompt to open Cluster Manager.
2	Right-click the service group.
3	Click Online.
4	Select the first cluster node.

Note

After every reboot of the server, use the `snapdrive config prepare luns -count <n>` command to make the LUNs visible to host.

Block diagram of DataFabric Manager server services and cluster resources

The following block diagram describes the DataFabric Manager server service resources that you need to add to VCS and the dependencies among the various resources.



Note
dfm_dg is not created on Linux.

Manually configuring DataFabric Manager Server for clustering

The advantage of configuring DataFabric Manager server with VCS manually is that you can customize the cluster.

You can manually configure DataFabric Manager server by using the VCS Application Configuration wizard, described next.

Step	Action
1	On the node where the VCS is set up, run the hawizard application at the command prompt. The VCS Application Configuration Wizard opens.

Step	Action
2	Select Create Application Service Group and click Next.
3	Enter Service Group Name and select the cluster from Available Cluster systems.
4	Click Next.
5	Enter the application details. <ul style="list-style-type: none"> Specify /usr/bin/dfm as the path in Start Program. Select root as user.
6	Click Next.
7	Select the processes monitor string. The string should match the output of the <code>ps -u root -o args</code> command. For more information, see the text following this table.
8	Click Next.
9	Configure the mount resources and click Next.
10	Configure IP and NIC resources and click Next.
11	Repeat Steps 5 through 10 for each DataFabric Manager process.
12	Open the Cluster Manager GUI to configure remaining cluster resources.
13	Select service group dfm_sg in the left pane.
14	In the Resource tab, right-click Resource View.
15	Enter the details for each Resource Type. Note _____ On Linux, you should select only LVMLLogicalVolume and Mount as the Resource Types. The FSType attribute should be set to "ext3" for Mount. _____
16	Select NIC from the Resource Type list.
17	Right-click the added resources and select Link.

Step	Action
18	Create a dependency tree and bring all services online.
19	To make DataFabric Manager server use the cluster name instead of the local system name, enter the following command: \$ dfm option set localHostName=<fqdn-of-cluster>

The DataFabric Manager server uses the name of the local system to send e-mail alerts to administrators.

The following table provides the DataFabric Manager server monitor strings required by VCS monitor process.

Process	Process monitor string
+ dfmmonitor	/opt/IBMdfm/sbin/dfmmonitor
+ dfmserver	/opt/IBMdfm/sbin/dfmserver
+ dfmscheduler	/opt/IBMdfm/sbin/dfmscheduler
+ dfmeventd	/opt/IBMdfm/sbin/dfmeventd start
+ database server	/opt/IBMdfm/sbin/dbsrv10 @/opt/IBMdfm/conf/sybase.conf
+ Apache server	/opt/IBMdfm/sbin/httpd -f /opt/IBMdfm/conf/httpd.conf

Note

If DataFabric Manager server is not installed at /opt/IBMdfm, the entries should be made accordingly.

Managing DataFabric Manager with VCS

Ways to start and stop DataFabric Manager services

After you set up DataFabric Manager server with VCS, do not use the `dfm service start` and `dfm service stop` commands, except where specifically indicated in installation and configuration procedures. These commands interfere with the working of VCS. Instead, perform all operations by using Cluster Manager.

Also, disable the DataFabric Manager init scripts after installation on both the cluster nodes. Do *not* change the service startup type to Automatic in Service Control Manager on any of the nodes. DataFabric Manager server reactivates these scripts during an upgrade and then disables them again when the upgrade is complete.

Backing up the DataFabric Manager server database

Database backup improvements in DataFabric Manager server 3.3 and later releases remove the requirement that you to stop the database service while backing up the database. Therefore, there are no additional steps to be performed to back up the database. Because the DataFabric Manager data is stored in storage systems, you can also create a Snapshot-based backup.

Restoring the DataFabric Manager server database

To restore an existing backup of the database, complete the following steps:

Step	Action
1	Disable DataFabric Manager service through Cluster Manager, by right-clicking the name of the service group and then clicking Offline.
2	Select the first cluster node in which the services are online.
3	Ensure that one of the nodes still owns the cluster resources (like the mount point), by following these steps: <ul style="list-style-type: none">a. Select the service group <code>dfm_sg</code>.b. In the Resources tab, right-click Resource View.c. Right-click the resource Mount and click Online.

Step	Action
4	In the node owning the Mount resource, enter the following command to restore the database: \$ dfm backup restore
5	Using Cluster Manager, restore the DataFabric Manager services by right-clicking the service group, clicking Online, and selecting the first cluster node used for backup restore.

Configuring DataFabric Manager to use HTTPS

To set up DataFabric Manager to use HTTPS, complete the following steps:

Step	Action
1	Using Cluster Manager, make the DataFabric Manager server services offline (except the service dfm-dbsrv): <ul style="list-style-type: none"> a. Right-click the service group dfm-sg, and click Offline. b. Select the first cluster node where the services are online. c. In the Resources tab, right-click Resource View. d. Right-click the resource dfm-dbsrv and click Online.
2	Enter the following command to create an SSL certificate: \$ dfm ssl server setup
3	Copy the files server.crt and server.key (located in the <install-dir>/conf folder) to the second node before starting the services on that node.
4	Enter the following command to enable HTTPS, by setting the following DataFabric Manager option to yes: \$ dfm option set httpsEnabled=yes
5	Enter the following command to start the HTTP service: \$ dfm service start http This re-creates the httpd.conf file with the changed options.

Step	Action
6	Enter the following command to stop all the services: <pre>\$ dfm service stop</pre> <p>Note_____</p> <p>Ensure that DataFabric Manager server services are still in the offline status before proceeding. If the services are online, take them offline and then proceed to Step 7. Otherwise the HTTP service fails to come up on the other node, as the configuration is not complete.</p>
7	Using Cluster Manager, move the cluster group to the second node by using the option Switch To.
8	Log in to the second node in the cluster.
9	Copy the files server.crt and server.key created on the first node to the folder <install-dir>/conf.
10	From the command line, enter the following command to start the services and verify that they are operating properly: <pre>\$ dfm service start</pre> <p>This creates the httpd.conf file with the changed options.</p>
11	Enter the following command to stop the services: <pre>\$ dfm service stop</pre>
12	Using Cluster Manager, reenable the DataFabric Manager services.

Changing HTTP options

To change the HTTP options httpEnabled, httpPort, httpsEnabled and httpsPort, complete the following steps:

Step	Action
1	<p>Using Cluster Manager, take the DataFabric Manager server services offline except dfm-dbsrv resource. To bring the services offline, follow these steps:</p> <ol style="list-style-type: none"> a. Right-click the service group, and click Offline. b. Select the first cluster node where the services are online. c. In the Resources tab, right-click Resource View. d. Right-click the resource dfm-dbsrv and select Offline.
2	<p>Enter the following command to set the required HTTP option:</p> <pre>\$ dfm option set <option-name>=<option-value></pre> <p>Example:</p> <pre>\$ dfm option set httpsPort=443</pre>
3	<p>Enter the following command to re-start the HTTP service:</p> <pre>\$ dfm service start http</pre> <p>This re-creates the file httpd.conf with the changed options.</p>
4	<p>Enter the following command to stop all the services:</p> <pre>\$ dfm service stop</pre> <p>Note _____</p> <p>Ensure that DataFabric Manager server services are still in the offline status before proceeding. If the services are online, take them offline and then proceed to Step 6. Otherwise the HTTP service fails to come up on the other node because the configuration is not complete.</p> <p>_____</p>
5	<p>Using Cluster Manager, move the cluster group to the second node by using the option Switch To.</p>
6	<p>Log in to the second node in the cluster.</p>
7	<p>From the command line, enter the following command to start the services:</p> <pre>\$ dfm service start</pre>

Step	Action
8	Enter the following command to stop the services: <code>\$ dfm service stop</code>
9	Using Cluster Manager, bring the DataFabric Manager services online.

Data that can be shared by DataFabric Manager server cluster nodes

The following table describes which data can be shared by the DataFabric Manager server cluster nodes.

Default path	Description
<install-dir>/data	Sybase database files
<install-dir>/perfdata	Performance Advisor data files
<install-dir>/script-plugins	Installed script plug-ins and related files
<install-dir>/plugins	Storage system configuration plug-ins
<install-dir>/reports	Archived reports
<install-dir>/dataExport	DataFabric Manager and Performance Advisor data

Configure DataFabric Manager server nodes to access files from a shared disk. If each node uses its own local copy of files, updates to files are not accessible to the other nodes after failover. For example, if a new storage system configuration management plug-in is installed on one node, it is accessible to only that node.

Data that is not shared by DataFabric Manager server cluster nodes

The following table describes which data is *not* shared by the DataFabric Manager server cluster nodes.

Default path	Description
<install-dir>/bin	Executable files
<install-dir>/conf	Configuration files

Default path	Description
<install-dir>/docs	Third-party licenses
<install-dir>/examples	Cluster configuration scripts, and so on.
<install-dir>/log	Log files
<install-dir>/misc	Configuration files
<install-dir>/sbin	Third-party executables
<install-dir>/scripts	Sybase_install.sql
<install-dir>/src	Storage system configuration plug-ins
<install-dir>/web/clients	Performance Advisor clients
<install-dir>/web/com	JAR files for applets
<install-dir>/web/help	Help files
<install-dir>/web/man	Manual (man) pages
<install-dir>/web/media	Images used on Web interfaces
<install-dir>/web/scripts	Java script files
<install-dir>/web/styles	CSS style sheets

Uninstalling DataFabric Manager server from a cluster

To uninstall DataFabric Manager server from both the cluster nodes and to disable the cluster setup, complete the following steps:

Step	Action
1	Using the Cluster Manager, delete all the DataFabric Manager services: <ul style="list-style-type: none"> a. Right-click the service group dfm-sg. b. Select Delete.
2	Log in to either of the cluster nodes.

Step	Action
3	For Linux, use either of the following Linux commands: <pre>\$ rpm -e IBMdfm</pre> <pre>\$ rpm --erase IBMdfm</pre>
4	Repeat Steps 2 and 3 for the other cluster node.

Upgrading cluster nodes with DataFabric Manager

Upgrading DataFabric Manager server cluster nodes

To upgrade the cluster nodes, you should upgrade all the nodes together.

To upgrade all the nodes together, complete the following steps:

Step	Action
1	Using Cluster Manager, make all the DataFabric Manager services offline. To make the services offline, follow these steps: <ul style="list-style-type: none">a. Right-click the service group, and then click Offline.b. Select the first cluster node where the services are online.
2	To ensure that the DataFabric Manager server resource Mount and its dependant resources are online, complete the following steps: <ul style="list-style-type: none">a. Select the service group dfm_sg.b. In the Resources tab, right-click Resource View.c. Right-click the resource Mount and select Online.
3	Upgrade DataFabric Manager server installation on this node.
4	Enter the following command to stop all of the DataFabric Manager server services: <code>\$ dfm service stop</code>
5	Enter the following command to disable the automatic service startup during reboot: <code>\$ dfm service enable -m</code>
6	Switch the mount resource and its dependents to the second node by selecting the option Switch To.
7	Ensure that the second node owns all the cluster resources.
8	Upgrade DataFabric Manager server installation on this node.

Step	Action
9	Stop all of the DataFabric Manager server services by entering the following command: \$ dfm service stop
10	Disable the automatic service startup during reboot by entering the following command: \$ dfm service enable -m
11	Make all DataFabric Manager server services online using Cluster Manager.

To run DataFabric Manager with MSCS efficiently, it is important to understand and to set the cluster parameters appropriately. This section describes some of the parameters that are particularly useful to understand.

Generic service restart and failover parameters

DataFabric Manager services are added to the cluster as resources of the type Generic Service. For a Generic Service resource type, MSCS defines the following default parameters of the service on the Advanced tab of the Properties view.

Restart Action: Specifies the action to be taken by the MSCS service if the resource fails. You can set this property to Do Not Restart or Restart.

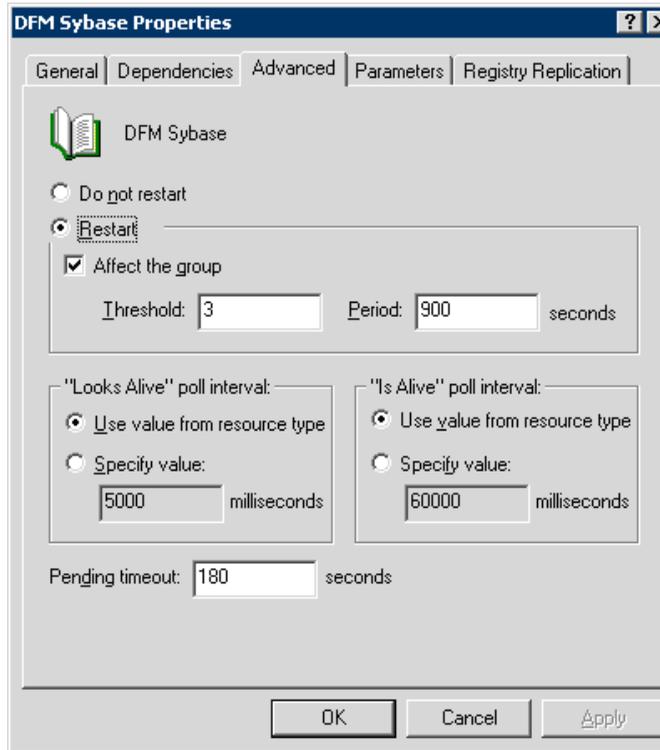
- ◆ **Affect the Group** property specifies whether the group should be failed over when this resource fails.
- ◆ **Restart Threshold** specifies the maximum number of restart attempts that can occur on a resource, within the interval defined by the Restart Period property. If the threshold and restart period are met, the MSCS service initiates the action specified by the Restart Action property (Do Not Start or Restart).
- ◆ **Restart Period** defines an interval of time, in milliseconds, during which a specified number of restart attempts can be made on a nonresponsive resource.

Looks Alive Poll Interval: Provides the recommended interval, in milliseconds, at which the MSCS service should poll the resource to determine whether it appears operational. Looks Alive polling performs basic polling to see whether the resource appears to be online.

Is Alive Poll Interval: Provides the recommended interval, in milliseconds, at which the MSCS service should poll the resource to determine whether it is operational. Is Alive polling performs thorough polling to see whether the resource is online and functioning.

Pending Timeout: Sets the number of seconds that a Resource Monitor waits for a resource in Online Pending or Offline Pending state before terminating the resource. For services that have long startup duration, you should specify a larger pending timeout value.

The following image is an example of the DFM Sybase Properties view, and it indicates the default values used for Generic Service. DataFabric Manager services use the same default values.



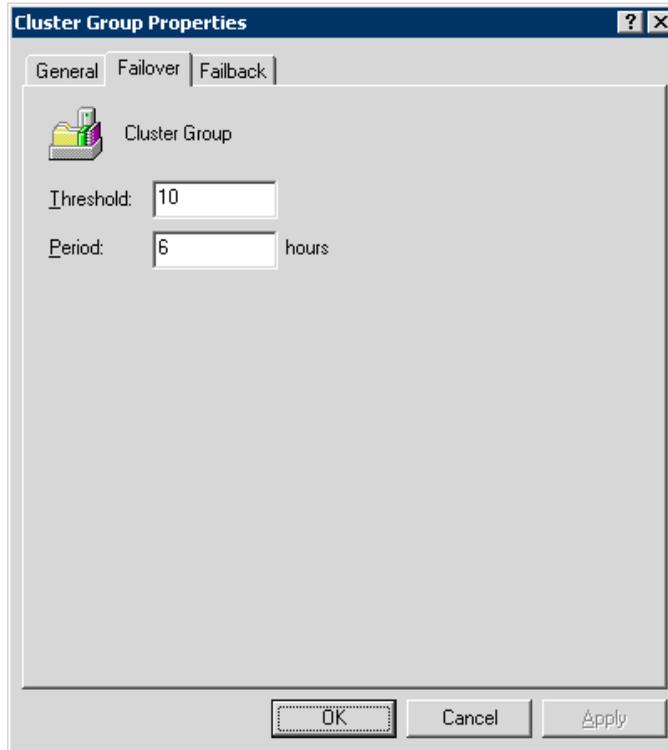
A summary of system behavior, based on the example in the image:

MSCS service waits for 180 seconds for a service to start up. After the service is online, MSCS performs Looks Alive polling every 5 seconds and Is Alive polling every 60 seconds. For a Generic Service resource, both these operations are performed by contacting the Service Control Manager. If a failure is detected for any service, the MSCS service tries to restart it. If the MSCS service has to restart the service three times in the duration of 15 minutes, it is considered a critical failure. Remaining services are brought offline and a controller failover is initiated.

Repeated service failure parameters

In some conditions it is possible that the DataFabric Manager resource group can fail again, after the resource group fails over to the new node. This situation can result in a failover loop. The cluster group properties define when such repeated failovers should stop.

The following image is an example, and it shows the default settings for Cluster Group failover limits.

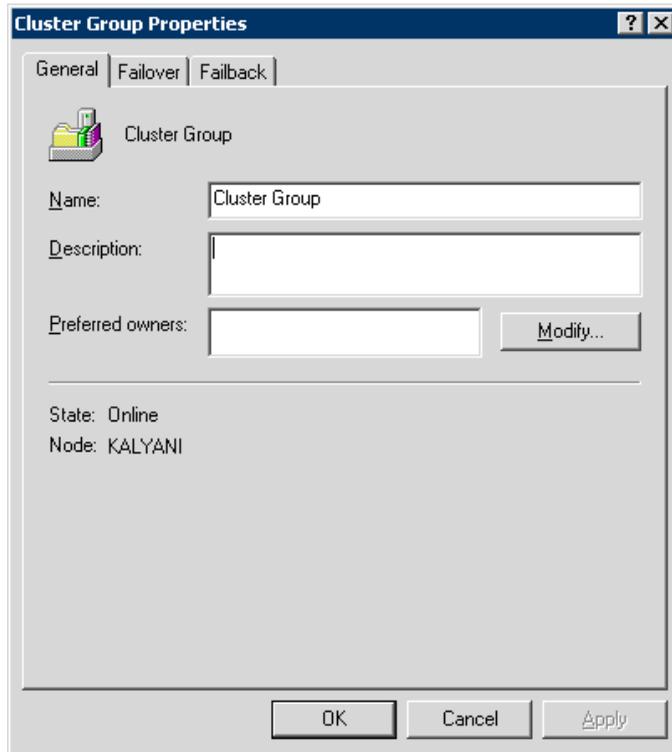


A summary of system behavior based on the example in the image:

Since the group failover Threshold is set to 10 and its failover Period is set to 6, the MSCS service fails over the group, for a maximum of 10 times within a 6-hour period. The eleventh time a resource in the group fails, the cluster service fails all other resources in the group and leaves the entire group offline, instead of failing over the group.

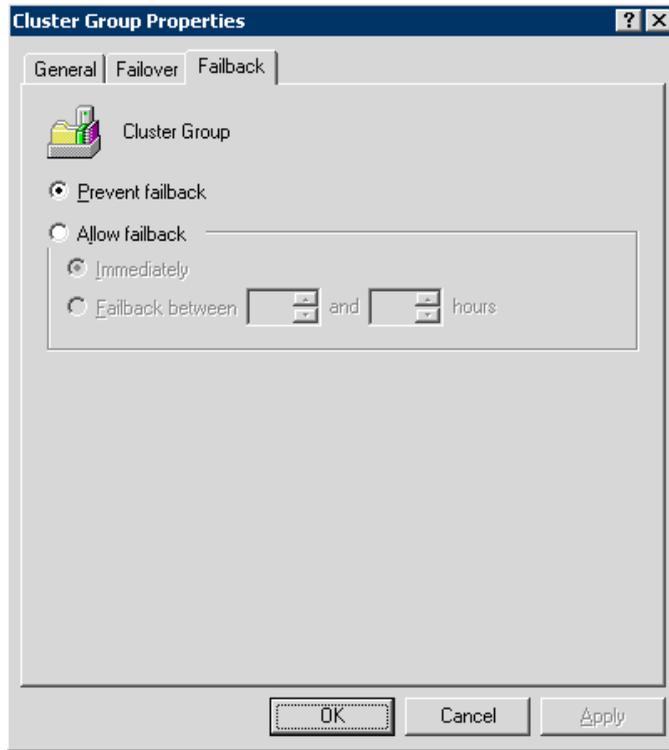
Node failback parameters

Preferred owners: For each resource group, you can specify a node as the preferred owner. You can specify the preferred owner node from the General tab of the Cluster Group Properties view, as shown in the following example (image). By default, the Preferred owners parameter is left empty.



Failback: Failback is the process by which the MSCS service moves a resource group back to its preferred node after the preferred node has failed and has come back online. You can set the failback parameter on the Failback tab of the Cluster Group Properties view, as shown in the following example (image). The options are Prevent Failback and Allow Failback. By default, groups are set to Prevent Failback.

If you want to enable failback, you can specify the failback to happen immediately after the preferred node comes online or during specific time periods (such as off-peak hours). For a failback to happen, you must specify a preferred node.



Additional reading

For more information about the MSCS requirements, guidelines, and management, see Microsoft's documentation at the following locations:

- ◆ The Microsoft Web site (www.microsoft.com):
 - ❖ Microsoft Operations Manager documentation
 - ❖ *Exchange Server 2003 Administration Guide*
 - ❖ Frequently Asked Questions
- ◆ The Microsoft TechNet Web site (technet.microsoft.com):
 - ❖ *Server Clusters Technical Reference*
 - ❖ *Windows Server 2003 Quick Start Guide for Server Clusters*
 - ❖ Frequently Asked Questions

To run DataFabric Manager server with VCS efficiently, it is important to understand and to set the cluster parameters appropriately. These parameters are used to control the behavior of VCS at the application source level. This section describes the parameters related to generic service restart and failover of clusters.

Generic service restart and failover parameters

For the resource type Generic Service, VCS defines the following default parameters:

- ◆ **RestartLimit:** If a resource fails, the number of attempts VCS makes to restart it, before sending an error message. Although the default value is 0 for VCS, you should change this value to 3 during manual configuration of the cluster. Using the configuration scripts automatically updates this value.
- ◆ **MonitorInterval:** The interval between two consecutive monitor calls for an online or transitioning resource, measured in seconds. The default MonitorInterval is 60 seconds. If many resources of the same type exist, a low value for this parameter impacts the performance of the cluster; however, a high value delays the detection of a faulted resource.
- ◆ **ToleranceLimit:** The number of times the monitor entry point goes offline, before it declares the resource faulty. The default value for ToleranceLimit is zero.
- ◆ **OnlineTimeout:** The maximum time, in seconds, that the online entry point has to complete before it is terminated. The default value for OnlineTimeout is 300 seconds.
- ◆ **ConfInterval:** Interval, in seconds, that a resource remains online. The agent ignores previous faults and restart attempts. The default value is 600 seconds.

To edit these parameters, click **Application > Properties**. DataFabric Manager server services are added to the cluster as resources of the type Generic Service.

Related Documentation

Available documentation

The following list describes available DataFabric Manager user documentation.

Help:

After you install or upgrade the DataFabric Manager software, you can access the following Help:

- ◆ IBM N series Operations Manager Help:
 - Click the Help icon on any Operations Manager page.
- ◆ IBM N series Management Console Help:
 - ❖ Install IBM N series Management Console.
 - ❖ Click the Help icon.

Administration Guides:

The administration guides are available in the printed and electronic form. Both forms are in the software and documentation package you received with your product. If you did not receive a software and documentation package, you can order a package through your sales representative.

You can also find the guides at <http://www.ibm.com/storage/support/nas/>.

Additional information available about DataFabric Manager

The following resources provide additional information about Operations Manager, File SRM, Business Continuance Option, and DataFabric Manager:

- ◆ IBM N series *Operations Manager Administration Guide*
- ◆ Man pages, which are available in the Help
- ◆ Frequently asked questions (FAQs)
- ◆ IBM N series *Performance Advisor Administration Guide*
- ◆ IBM N series *Provisioning Manager and Protection Manager Administration Guide*
- ◆ <http://www.ibm.com/storage/support/nas/>

Accessing IBM N series Operations Manager Help

By default, the *compact* Help window opens when you select Help and click About This Page. The compact window displays only the Help topic, while the full window includes the Table of Contents. To switch between the full and compact Help windows, click Show.

Using the FAQs

The Help includes FAQs that provide answers to many typical questions regarding software installation and features. It is also available at <http://www.ibm.com/storage/support/nas/>.

Index

- A**
 - additional license keys
 - Options page in Operations Manager 5
 - appliance requirements
 - Data ONTAP 10
- B**
 - backing up
 - DataFabric Manager with MSCS database 60, 83
 - browser
 - default security setting 14
 - Java applet support 15
 - requirements 5
- C**
 - cluster nodes, upgrading 67, 90
 - cluster. *See* DataFabric Manager with MSCS or Microsoft Cluster Server (MSCS)
 - configuring DataFabric Manager with MSCS
 - manually 52
 - using scripts 76
- D**
 - data not shared in a cluster 65, 87
 - Data ONTAP requirements 10
 - data shared in a cluster 64, 87
 - database backup
 - DataFabric Manager with MSCS 60, 83
 - DataFabric Manager
 - removing from a cluster 66, 88
 - upgrade issues 13
 - DataFabric Manager cluster
 - quorum 38, 70
 - set up using manual steps 52
 - upgrading nodes 67, 90
 - DataFabric Manager database
 - migration from Solaris 30
 - DataFabric Manager sizing guidelines 5
 - DataFabric Manager with MSCS
 - add resources setup script 49
 - additional reading 69, 97
 - bringing online
 - after manual setup 59
 - after setup scripts 50, 83
 - configuration scripts 76
 - configuring name 51, 82
 - data that can be shared 64, 87
 - data that is not shared 65, 87
 - database backup 60, 83
 - first node setup script 76
 - first node, manual configuration 52
 - hardware requirements 39, 71
 - HTTP options 63, 85
 - HTTPS setup 61, 84
 - installation diagram 47, 80
 - installation requirements 40
 - installation, understanding 38, 70
 - installing 43, 73
 - management 60, 83
 - manual configuration
 - first node 52
 - resources 56
 - second node 55
 - monitoring 64
 - name 51, 82
 - preparing to install 41
 - resource dependencies 47, 80
 - resources, manual configuration 56
 - second node, manual configuration 55
 - shared data 64, 87
 - software requirements 39, 71
 - start and stop services 60, 83
 - understanding 38, 70
 - uninstalling 66, 88
 - using cluster name
 - after manual setup 59
 - after script setup 51, 82
 - Web UI setup 61, 84
 - dependencies
 - DataFabric Manager with MSCS resources 47,

80

- dfm option set localHostName, cluster 51, 59, 82
- dfm service start, cluster 60, 83
- dfm service stop, cluster 60, 83
- dfmcluster_add_resources.pl
 - operations 50
 - syntax 49
- dfmcluster_firstnode.pl
 - syntax 76
- diagram
 - DataFabric Manager with MSCS installation 47, 80
- Disaster Recovery feature 15
- documentation available 101

G

- generic service restart and failover 93, 99
- guidelines
 - sizing for DataFabric Manager server 5

H

- hardware requirements
 - DataFabric Manager with MSCS 39, 71
 - Linux 4
 - Windows 3
- Help 101
- HTTP options
 - DataFabric Manager with MSCS 63, 85
- HTTPS
 - DataFabric Manager with MSCS 61, 84

I

- IBM N series Management Console
 - installing 26, 27, 28
 - system requirements 26
- installation
 - DataFabric Manager cluster, preparation 41
- installation requirements
 - Data ONTAP 10
 - DataFabric Manager with MSCS 40
 - IBM N series Management Console 26
 - Microsoft Cluster Server 40

- Windows 3
- installing DataFabric Manager
 - Linux 22
 - Windows 18
- installing DataFabric Manager with MSCS
 - understanding 38, 70
- installing IBM N series Management Console 26, 27, 28
- installing licenses 25
- Is Alive Poll Interval 93

J

- Java applet support 15

L

- licenses
 - installing 25
 - license keys 5
 - requirements 5
- Linux
 - installing DataFabric Manager 22
 - requirements 4
 - uninstalling DataFabric Manager 24
- Looks Alive Poll Interval 93

M

- management
 - Microsoft Cluster Server 67, 93
- Microsoft Cluster Server
 - configuration requirements 40
 - management 67, 93
 - parameters
 - Is Alive Poll Interval 93
 - Looks Alive Poll Interval 93
 - node failback 95
 - Pending Timeout 93
 - repeated service failure 94
 - Restart Action 93, 99
 - restart and failover 93, 99
 - setup example 40
- monitoring
 - DataFabric Manager with MSCS 64

N

node failback 95

O

Operations Manager
 definition iii
 license requirements 5

P

Pending Timeout 93
Performance Advisor
 installing 26
Protection Manager
 installing 26
 license 7

Q

quorum 38, 70

R

removing. *See* uninstalling DataFabric Manager
 from a cluster
 repeated service failure 94
 requirements
 browser 5
 Data ONTAP 10
 IBM N series Management Console 26
 license keys 5
 Linux 4
 Windows 3
Restart Action 93, 99
rpm command (Linux) 24

S

scripts, DataFabric Manager with MSCS 76
 first node setup 76
search engine access, blocking 14
sizing guidelines for DataFabric Manager 5
software requirements
 browser 5
 Data ONTAP 10
 DataFabric Manager with MSCS 39, 71
 Linux 4
 Windows 3
special messages vii

U

uninstalling DataFabric Manager
 from a cluster 66, 88
 Linux 24
 Windows 20
upgrade
 issues 11
 nodes, DataFabric Manager with MSCS 67, 90

V

viewing dynamic data 15

W

Windows
 browser
 recommendation 14
 default browser security setting 14
 installing DataFabric Manager 18
 requirements 3
 uninstalling DataFabric Manager 20



NA 210-04811_A0, Printed in USA

GC26-7892-08

