

Lotus. software

IBM

Lotus Domino 6



Administering Domino Clusters

Disclaimer

THIS DOCUMENTATION IS PROVIDED FOR REFERENCE PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS DOCUMENTATION, THIS DOCUMENTATION IS PROVIDED "AS IS" WITHOUT ANY WARRANTY WHATSOEVER AND TO THE MAXIMUM EXTENT PERMITTED, IBM DISCLAIMS ALL IMPLIED WARRANTIES, INCLUDING WITHOUT LIMITATION THE IMPLIED WARRANTIES OF MERCHANTABILITY, NONINFRINGEMENT AND FITNESS FOR A PARTICULAR PURPOSE, WITH RESPECT TO THE SAME. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION, DIRECT, INDIRECT, CONSEQUENTIAL OR INCIDENTAL DAMAGES, ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS DOCUMENTATION OR ANY OTHER DOCUMENTATION. NOTWITHSTANDING ANYTHING TO THE CONTRARY, NOTHING CONTAINED IN THIS DOCUMENTATION OR ANY OTHER DOCUMENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF THE APPLICABLE LICENSE AGREEMENT GOVERNING THE USE OF THIS SOFTWARE.

Copyright

Under the copyright laws, neither the documentation nor the software may be copied, photocopied, reproduced, translated, or reduced to any electronic medium or machine-readable form, in whole or in part, without the prior written consent of IBM, except in the manner described in the documentation or the applicable licensing agreement governing the use of the software.

© Copyright IBM Corporation 1985, 2002

All rights reserved.

Lotus Software
IBM Software Group
One Rogers Street
Cambridge, MA 02142

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

List of Trademarks

1-2-3, cc:Mail, Domino, Domino Designer, Freelance Graphics, iNotes, Lotus, Lotus Discovery Server, Lotus Enterprise Integrator, Lotus Mobile Notes, Lotus Notes, Lotus Organizer, LotusScript, Notes, QuickPlace, Sametime, SmartSuite, and Word Pro are trademarks or registered trademarks of Lotus Development Corporation and/or IBM Corporation in the United States, other countries, or both. AIX, AS/400, DB2, IBM, iSeries, MQSeries, Netfinity, OfficeVision, OS/2, OS/390, OS/400, S/390, Tivoli, and WebSphere are registered trademarks of International Business Machines Corporation in the United States, other countries, or both. Pentium is a trademark of Intel Corporation in the United States, other countries, or both. Microsoft, Windows, and Windows NT are registered 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. Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

All other trademarks are the property of their respective owners.

Contents

Preface	vii	3 Planning a Cluster	3-1
1 Cluster Benefits and Requirements	1-1	Determining how many servers to include in a cluster	3-1
What is a Domino cluster?	1-1	Determining the number and placement of replicas in a cluster	3-3
How do clusters help you?	1-1	How many replicas to create	3-4
Clustering requirements	1-3	Distributing databases in a cluster	3-7
2 How Domino Clustering Works ..	2-1	Distributing mail databases	3-7
Clustering basics	2-1	Distributing application databases	3-9
How failover works	2-2	Determining whether to create a private LAN for your cluster	3-10
When failover occurs	2-3	Clustering over a wide area network	3-11
How Domino finds a replica during failover	2-5	Fault recovery in a cluster	3-11
How workload balancing works	2-5	Examples of cluster configurations	3-12
The cluster components	2-7	Example of clustering two servers for mail and applications	3-13
The Cluster Manager	2-7	Example of clustering three servers for mail and applications	3-14
The Cluster Database Directory	2-8	Example of clustering six servers for mail and applications	3-15
The Cluster Database Directory Manager	2-9	Examples of clustering hub servers ...	3-16
The Cluster Administrator	2-9	Example of clustering for disaster preparedness	3-18
The Cluster Replicator	2-9	Example of clustering partitioned servers	3-19
How replication works in a cluster	2-10	Example of clustering passthru servers	3-21
Replication history in a cluster	2-11	Example of using a Domino cluster with an operating system cluster	3-22
Private folder replication in a cluster ...	2-11		
Mail failover in a cluster	2-12		
How calendars work in a cluster	2-13		
How operating system clusters work	2-14		
Benefits of using OS clusters with Domino clusters	2-17		

4 Setting Up a Cluster	4-1	Making sure Domino is using the private LAN	4-25
Before you create a cluster	4-1	Making sure that cluster replication is using the private LAN	4-25
Creating a cluster	4-2	NOTES.INI settings for setting up a cluster	4-26
Verifying that the cluster was created correctly	4-4	5 Managing and Monitoring a Cluster	5-1
Using Cluster Analysis to check the cluster configuration	4-5	Monitoring a cluster	5-1
Running Cluster Analysis	4-6	Displaying a list of cluster members and their availability	5-2
Managing user access to databases	4-7	Enabling statistic reporting in the Monitoring Results database	5-3
Setting up database ACLs in a cluster	4-7	Enabling statistic report generation	5-3
Controlling other settings that restrict database access	4-9	Creating a server statistic collection	5-3
Creating replicas in a cluster	4-9	Starting the Statistic Collector task	5-4
Setting up mail in a cluster	4-10	Viewing Cluster Manager events and statistics	5-5
Changing the mail routing failover setting	4-11	Viewing failover and workload balancing events	5-5
Using shared mail in a cluster	4-11	Viewing Cluster Manager statistics	5-6
Creating mail database replicas in a cluster during user registration	4-12	Viewing cluster replication events and statistics	5-7
Using standard replication in a cluster	4-13	Viewing cluster replication events	5-8
Scheduled replication in a cluster	4-13	Viewing cluster replication statistics	5-9
Replicating with all servers in a cluster	4-15	Using Tell commands to display cluster replication information	5-11
Enabling the display of cluster replication status messages	4-17	About the Tell commands for cluster replication	5-13
Obeying database size quotas during cluster replication	4-17	Viewing the cluster replication information	5-14
Using directory assistance in a cluster	4-19	Monitoring all the servers in a cluster at the same time	5-15
Setting up roaming in a cluster	4-19	Balancing the workload in a cluster	5-17
Setting up new users for roaming in a cluster	4-20	Limiting the workload of a server	5-18
Setting up existing users for roaming in a cluster	4-20	Setting the maximum number of users on a server	5-24
Setting up cluster access for mobile users	4-21	Redistributing replicas	5-26
Using the Server Web Navigator in a cluster	4-23		
Setting up a private LAN for a cluster	4-23		

Using the IBM Tivoli Analyzer for Lotus Domino in a cluster	5-26	Moving a server from one cluster to another	5-49
Managing failover in a cluster	5-27	Viewing information in the Cluster Database Directory	5-49
Causing failover to occur	5-27	Creating a new Cluster Database Directory	5-50
Replacing a cluster server with a different server	5-29	NOTES.INI settings for managing a cluster	5-51
Other ways to manage failover	5-29	6 Clustering Domino Servers That Run Internet Protocols	6-1
Managing replication in a cluster	5-30	How the Internet Cluster Manager works	6-1
Configuring and controlling cluster replication	5-30	Generating URLs that refer to the ICM	6-3
Using multiple Cluster Replicators	5-31	Planning to use the ICM	6-4
Disabling cluster replication for an entire server	5-33	Planning the cluster	6-4
Disabling cluster replication for specific databases	5-35	Planning the location of the ICM	6-4
Pausing cluster replication	5-37	Example of a single ICM outside the cluster	6-5
Resuming cluster replication	5-38	Example of multiple ICMs outside the cluster	6-6
Forcing the Cluster Replicator to retry failed replications immediately	5-39	Example of a single ICM inside the cluster	6-7
Forcing the Cluster Replicator to update the Cluster Database Directory information immediately	5-40	Example of multiple ICMs inside the cluster	6-8
Forcing the Cluster Replicator to log immediately	5-40	Example of one ICM outside the cluster and one ICM inside the cluster	6-9
Creating configuration settings for all servers in a cluster	5-41	Configuring the ICM	6-9
Managing database availability in a cluster	5-42	To configure the ICM	6-10
Making a database unavailable for user access	5-42	Setting up a separate IP address for the ICM	6-12
Making an “out of service” database available again	5-43	Starting the ICM	6-13
Deleting databases from a cluster	5-44	Failover and workload balancing	6-14
Adding a server to a cluster	5-45	When a server fails	6-15
What occurs when you add a server to a cluster	5-46	Security	6-16
Removing a server from a cluster	5-47	Managing and monitoring the ICM	6-16
What occurs when you remove a server from a cluster	5-48	Viewing the log file	6-16
		Viewing ICM statistics	6-17
		Monitoring all the ICM servers at the same time	6-17

Compatibility with previous releases of Domino	6-19
Using an IP sprayer with Domino for HTTP and POP3 failover	6-19
Setting up failover for IMAP	6-21
Using an IP sprayer with Domino for LDAP failover	6-22
Appendix A Cluster Statistics	A-1
Appendix B Troubleshooting	
Cluster Problems	B-1
Index	Index-1

Preface

The documentation for IBM Lotus Notes, IBM Lotus Domino, and IBM Lotus Domino Designer is available online in Help databases and, with the exception of the Notes client documentation, in print format.

License information

Any information or reference related to license terms in this document is provided to you for your information. However, your use of Notes and Domino, and any other IBM program referenced in this document, is solely subject to the terms and conditions of the IBM International Program License Agreement (IPLA) and related License Information (LI) document accompanying each such program. You may not rely on this document should there be any questions concerning your right to use Notes and Domino. Please refer to the IPLA and LI for Notes and Domino that is located in the file LICENSE.TXT.

System requirements

Information about the system requirements for Lotus Notes and Domino is listed in the *Release Notes*.

Printed documentation and PDF files

The same documentation for Domino, and Domino Designer that is available in online Help is also available in printed books and PDF files.

You can order printed books from the IBM Publications Center at www.ibm.com/shop/publications/order.

You can download PDF files from the IBM Publications Center and from the Documentation Library at the Lotus Developer Domain at www-10.lotus.com/ldd.

Related information

In addition to the documentation that is available with the product, other information about Notes and Domino is available on the Web sites listed here.

- IBM Redbooks are available at www.redbooks.ibm.com.

- A technical journal, discussion forums, demos, and other information is available on the Lotus Developer Domain site at www-10.lotus.com/ldd.

Table of conventions

This table lists conventions used in the Notes and Domino documentation.

<i>Convention</i>	<i>Description</i>
<i>italics</i>	Variables and book titles are shown in italic type.
<code>monospaced type</code>	Code examples and console commands are shown in monospaced type.
file names	File names are shown in uppercase, for example NAMES.NSF.
hyphens in menu names (File - Database - Open)	Hyphens are used between menu names, to show the sequence of menus.

Structure of Notes and Domino documentation

This section describes the documentation for Notes, Domino, and Domino Designer. The online Help databases are available with the software products. Print documentation can be downloaded from the Web or purchased separately.

Release Notes

The Release Notes describe new features and enhancements, platform requirements, known issues, and documentation updates for Lotus Notes 6, Lotus Domino 6, and Lotus Domino Designer 6. The Release Notes are available online in the Release Notes database (README.NSF). You can also download them as a PDF file.

Documentation for the Notes client

The Lotus Notes 6 Help database (HELP6_CLIENT.NSF) contains the documentation for Notes users. This database describes user tasks such as sending mail, using the Personal Address Book, using the Calendar and Scheduling features, using the To Do list, and searching for information.

Documentation for Domino administration

The following table describes the books that comprise the Domino Administration documentation set. The information in these books is also found online in the Lotus Domino Administrator 6 Help database (HELP6_ADMIN.NSF).

The book *Installing Domino Servers* ships with Domino. The other books are available for purchase, or for free download as PDF files.

<i>Title</i>	<i>Description</i>
<i>Upgrade Guide</i>	Describes how to upgrade existing Domino servers and Notes clients to Notes and Domino 6. Also describes how to move users from other messaging and directory systems to Notes and Domino 6.
<i>Installing Domino Servers</i>	Describes how to plan a Domino installation; how to configure Domino to work with network protocols such as Novell SPX, TCP/IP, and NetBIOS; how to install servers; and how to install and begin using Domino Administrator and the Web Administrator.
<i>Administering the Domino System, Volumes 1 and 2</i>	Describes how to register and manage users and groups, and how to register and manage servers including managing directories, connections, mail, replication, security, calendars and scheduling, activity logging, databases, and system monitoring. This book also describes how to use Domino in a service provider environment, how to use Domino Off-Line Services, and how to use IBM Tivoli Analyzer for Lotus Domino.
<i>Administering Domino Clusters</i>	Describes how to set up, manage, and troubleshoot Domino clusters.

Documentation for Domino Designer

The following table describes the books that comprise the Domino Designer documentation set. The information in these books is also found online in the Lotus Domino Designer 6 Help database (HELP6_DESIGNER.NSF) with one exception: *Domino Enterprise Connection Services (DECS) Installation and User Guide* is available online in a separate database, DECS User Guide Template (DECSDOC6.NSF). The printed documentation set also includes Domino Objects posters.

In addition to the books listed here, the *Domino Designer Templates Guide* is available for download in NSF or PDF format. This guide presents an in-depth look at three commonly used Designer templates: TeamRoom, Discussion, and Documentation Library.

<i>Title</i>	<i>Description</i>
<i>Application Development with Domino Designer</i>	Explains how to create all the design elements used in building Domino applications, how to share information with other applications, and how to customize and manage applications.
<i>Domino Designer Programming Guide, Volume 1: Overview and Formula Language</i>	Introduces programming in Domino Designer and describes the formula language.

continued

<i>Title</i>	<i>Description</i>
<i>Domino Designer Programming Guide, Volumes 2A and 2B: LotusScript/COM/OLE Classes</i>	Describes the LotusScript/COM/OLE classes for access to databases and other Domino structures.
<i>Domino Designer Programming Guide, Volume 3: Java/CORBA Classes</i>	Provides reference information on using the Java and CORBA classes to provide access to databases and other Domino structures.
<i>Domino Designer Programming Guide, Volume 4: XML Domino DTD and JSP Tags</i>	Describes the XML and JSP interfaces for access to databases and other Domino structures.
<i>LotusScript Language Guide</i>	Describes the LotusScript programming language.
<i>Domino Enterprise Connection Services (DECS) Installation and User Guide</i>	Describes how to use Domino Enterprise Connection Services (DECS) to access enterprise data in real time.
<i>Lotus Connectors and Connectivity Guide</i>	Describes how to configure Lotus Connectors for use with either DECS or IBM Lotus Enterprise Integrator for Domino (LEI). It also describes how to test connectivity between DECS or LEI and an external system, such as DB2, Oracle, or Sybase. Lastly, it describes usage and feature options for all of the base connection types that are supplied with LEI and DECS. This online documentation file name is LCCON6.NSF.
<i>Lotus Connector LotusScript Extensions Guide</i>	Describes how to use the LC LSX to programmatically perform Lotus Connector-related tasks outside of, or in conjunction with, either LEI or DECS. This online documentation file name is LSXLC6.NSF.
<i>IBM Lotus Enterprise Integrator for Domino (LEI) Installation Guide</i>	Describes installation, configuration, and migration information and instructions for LEI. The online documentation file names are LEIIG.NSF and LEIIG.PDF. This document is for LEI customers only and is supplied with LEI, not with Domino.
<i>IBM Lotus Enterprise Integrator for Domino (LEI) Activities and User Guide</i>	Provides information and instructions for using LEI and its activities. The online documentation file names are LEIDOC.NSF and LEIDOC.PDF. This document is for LEI customers only and is supplied with LEI, not with Domino.

Chapter 1

Cluster Benefits and Requirements

This book describes clustering with IBM® Lotus® Domino™ 6 servers and tells you how to set up, configure, and use Domino clusters. This chapter explains the benefits of clusters and the hardware and software requirements for setting up and using clusters.

What is a Domino cluster?

A Domino cluster is a group of two or more servers that provides users with constant access to data, balances the workload between servers, improves server performance, and maintains performance when you increase the size of your enterprise. The servers in a cluster contain replicas of databases that you want to be readily available to users at all times. If a user tries to access a database on a cluster server that is not available, Domino opens a replica of that database on a different cluster server, if a replica is available. Domino continuously synchronizes databases so that whichever replica a user opens, the information is always the same.

IBM Lotus Notes® clients can access all Domino cluster servers. HTTP clients (Internet browsers) can access only Domino Web servers in a Domino cluster.

How do clusters help you?

The main benefits of clusters are:

- **High availability of important databases**

When a hardware or software problem occurs, clustered servers redirect database open requests to other servers in the cluster to provide users with uninterrupted access to important databases. This process is called failover. Clusters provide failover for business-critical databases and servers, including passthru server failover to other servers in the cluster. Failover also lets you perform server maintenance, such as hardware and software upgrades, with little negative effect on users.

- **Workload balancing**

When users try to access databases on heavily used servers, Domino can redirect the user requests to other cluster servers that aren't as busy so that the workload is evenly distributed across the cluster. Workload balancing of cluster servers helps your system achieve optimum performance, which leads to faster data access.
- **Scalability**

As the number of users you support increases, you can easily add servers to a cluster to keep server performance high. You can also create multiple database replicas to maximize data availability, and you can move users to other servers or clusters as you plan for future growth. As your enterprise grows, you can distribute user accounts across clusters and balance the additional workload to optimize system performance within a cluster.
- **Data synchronization**

A key to effective clustering is setting up replicas on two or more cluster servers so that users have access to data when a server is down or is being used heavily. Cluster replication ensures that all changes, whether to databases or to the cluster membership itself, are immediately passed to other databases or servers in the cluster. Thus, databases are continuously synchronized to provide high availability of information.
- **Analysis tools**

Using the cluster analysis tools, as well as the log file, the Monitoring Configuration and Monitoring Results databases, and the server monitor, you can analyze cluster activity and make any changes necessary to improve performance.
- **Ease of changing operating systems, hardware, or versions of Domino**

When you want to change your hardware, operating system, or Domino release, you can mark the clustered server as RESTRICTED so that requests to access a database on the server fail over to other cluster servers that contain replicas. This lets you make changes without interrupting the productivity of your users.
- **Data backup and disaster planning**

You can set up a cluster server as a backup server to protect crucial data. You can prevent users from accessing the server, but cluster replication keeps the server updated at all times. You can even do this over a WAN so that the backup is in a different geographical location.

- **Easy administration**

You can create a cluster with a few keystrokes. You can also add servers to a cluster, remove servers from a cluster, and move servers between clusters with a few keystrokes. In addition, you can drag and drop databases into a cluster and specify which cluster servers should receive replicas. You can also create multiple mail replicas and roaming file replicas for users when you register them, and you can monitor all the servers in a cluster simultaneously.

- **Use of any hardware and operating system that Domino supports**

You can set up a cluster using the same hardware you use for your Domino servers. You do not need to use special hardware to create a Domino cluster. In addition, the cluster can contain servers that use any operating system that Domino supports.

Clustering requirements

Server requirements

- All servers in a cluster must run one of the following: the Lotus Domino 6 Enterprise server, the Lotus Domino 6 Utility server, the Domino Release 5 or Domino Release 4.62 Enterprise server, or the Domino Release 4.6 or Domino Release 4.5 Advanced Services server.

Note Earlier releases of Domino do not have access to features that are new in Domino 6.

- All servers in a cluster must be connected using a high-speed local area network (LAN) or a high-speed wide area network (WAN). You can also set up a private LAN for cluster traffic.
- All servers in a cluster must use TCP/IP and be on the same Notes named network
- All servers in a cluster must be in the same Domino domain and share a common Domino Directory.
- You must specify an administration server for the Domino Directory in the domain that contains the cluster. If you do not specify an administration server, the Administration Process cannot change cluster membership. The administration server does not have to be a member of a cluster.
- Each server in the cluster must have a hierarchical server ID. If any servers have flat IDs, you must convert them to hierarchical IDs to use them in a cluster.
- A server can be a member of only one cluster at a time.

- Each server must have adequate disk space to function as a cluster member. Because clusters usually require more database replicas, servers in clusters require more disk space than unclustered servers.
- Each server must have adequate processing power and memory capacity. In general, clustered servers require more computer power than unclustered servers.

For more information on determining the resources you need to set up a cluster, see the chapter “Planning a Cluster.”

Client requirement

- Notes clients must run Notes Release 4.5 or later to take advantage of the cluster failover feature.

Chapter 2

How Domino Clustering Works

This chapter explains how Domino clustering works for Notes clients and Domino servers. This chapter also describes operating system clusters and how Domino clusters work in conjunction with operating system clusters. For an explanation of how Domino clustering works with Internet clients and Internet protocols, see the chapter “Clustering Domino Servers That Run Internet Protocols.”

Clustering basics

All the servers in a Domino cluster continually communicate with each other to keep updated on the status of each server and to keep database replicas synchronized. Each server in a cluster contains cluster components that are installed with the Lotus Domino 6 Enterprise server or the Lotus Domino 6 Utility server. These components, and the Administration Process, perform the cluster management and monitoring tasks that run the cluster and let you administer the cluster. The components keep replica databases synchronized, and they communicate with each other to ensure that the cluster is running efficiently and smoothly. They also let you set limits for workload balancing, track the availability of servers and databases, and add servers and databases to the cluster.

To take advantage of failover and workload balancing, you distribute databases and replicas throughout the cluster. You do not need a replica of every database on every server. The number of replicas you create for a database depends on how busy the database is and how important it is for users to have constant access to that database. For some databases, you may not need to create any replicas; for others, you may need to create multiple replicas.

For information about deciding how many replicas to create, see the chapter “Planning a Cluster.”

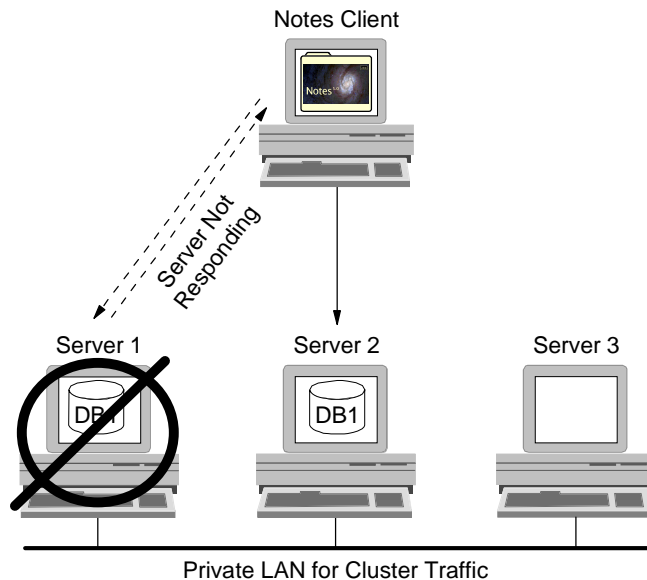
How failover works

A cluster's ability to redirect requests from one server to another is called failover. When a user tries to access a database on a server that is unavailable or in heavy use, Domino directs the user to a replica of the database on another server in the cluster.

The Cluster Manager on each cluster server sends out probes to each of the other cluster servers to determine the availability of each server. The Cluster Manager also checks continually to see which replicas are available on each server. When a user tries to access a database that is not available, the user request is redirected to a replica of the database on a different server in the cluster. Although the user connects to a replica on a different server, failover is essentially transparent to the user.

Example

This example describes the process that Domino uses when it fails over. This cluster contains three servers. Server 1 is currently unavailable. The Cluster Managers on Server 2 and Server 3 are aware that Server 1 is unavailable.



1. A Notes user attempts to open a database on Server 1.
2. Notes realizes that Server 1 is not responding.
3. Instead of displaying a message that says the server is not responding, Notes looks in its cluster cache to see if this server is a member of a cluster and to find the names of the other servers in the cluster. (When a Notes client first accesses a server in a cluster, the names of all the

servers in the cluster are added to the cluster cache on the client. This cache is updated every 15 minutes.)

4. Notes accesses the Cluster Manager on the next server listed in the cluster cache.
5. The Cluster Manager looks in the Cluster Database Directory to find which servers in the cluster contain a replica of the desired database.
6. The Cluster Manager looks in its server cluster cache to find the availability of each server that contains a replica. (The server cluster cache contains information about all the servers in the cluster. Cluster servers obtain this information when they send probes to the other cluster servers.)
7. The Cluster Manager creates a list of the servers in the cluster that contain a replica of the database, sorts the list in order of availability, and sends the list to Notes.
8. Notes opens the replica on the first server in the list (the most available server). If that server is no longer available, Notes opens the replica on the next server in the list. In this example, Server 2 was the most available server.

When the Notes client shuts down, it stores the contents of the cluster cache in the file CLUSTER.NCF. Each time the client starts, it populates the cluster cache from the information in CLUSTER.NCF.

When failover occurs

Failover occurs when users cannot access the server that contains the database they want or they cannot access the database itself. The following table describes reasons that users may not be able to access the server or the database they want.

<i>Category</i>	<i>Cause of failover</i>
Unable to access server	The server is down
	There are network connectivity problems
	The server has reached the maximum number of users allowed (set with Server_MaxUsers in NOTES.INI)
	The server is restricted by the administrator (using the Server_Restricted setting in NOTES.INI)
	The server is BUSY because it has reached the maximum load allowed (the server availability threshold)
Unable to access database	The database is marked "Out of Service" in the Cluster Database Directory
	The database is marked "Pending Delete" in the Cluster Database Directory

When a server or database is not available, failover occurs when a user attempts to use Notes to perform certain actions. The following table describes the actions that trigger failover.

<i>Category</i>	<i>Action that triggers failover</i>
Database open operations	Opening a database from a bookmark
	Clicking a document link, a view link, or a database link
	Using Domain Search when a clustered indexing server is unavailable
	Accessing roaming files when a clustered roaming server is unavailable
	Activating a field, action, or button that contains @command ((FileOpenDatabase))
	Running a LotusScript® routine that contains the OpenWithFailover method of the NotesDatabase class
	Using Java™ that contains the OpenDatabase method of the DbDirectory class
	Replicating with a database on a cluster server that is not running or not reachable on the network
Mail server related operations	Sending mail
	Name lookups
	Type-ahead
	Routing mail messages
	Mail predelivery agents
	Meeting invitations
	Free time lookups
	Server lookups
Web server operations	Selecting the Open URL icon
	Clicking a URL hotspot
	Accessing a URL with a Web browser

When failover does not occur

Failover does not occur in the following cases:

- When a server becomes unavailable while a user has a database open

Note The user can reopen the database, which causes failover to a different replica, if one exists in the cluster. If the user was editing a document when the server became unavailable, the user can copy the document to the replica.

- When a user chooses File - Database - Properties or File - Database - Open
- When the router attempts to deliver mail when you've disabled mail routing failover or set MailClusterFailover to 0
- When the template server is unavailable while creating a new database
- When a server fails while running agents, other than the mail predelivery agents
- When a server fails while running the Administration Process
- When replicating with a server that is restricted by the administrator or has reached the maximum number of users or the maximum usage level set by the administrator. Also, when replicating with a database marked "Out of Service." Replication occurs regardless of such restrictions, so there is no need for failover to occur.

How Domino finds a replica during failover

When a user tries to open a database that is not available, the Cluster Manager looks in the Cluster Database Directory for a replica of that database. To find the replica, the Cluster Manager looks for a database that has the same replica ID as the original database. It also makes sure that the replica is not marked "Out of Service" or "Pending Delete."

If there is more than one replica of a database on a server, the Cluster Manager assumes that selective replication is used to replicate these databases. To be sure it fails over to the correct replica, the Cluster Manager selects a replica that has the same path as the original database. Therefore, if you put multiple replicas on a server, be sure that all replicas in the cluster that use the same selective replication formulas have the same path. Otherwise, users may fail over to a different replica.

When failover occurs on a server, Domino generates a failover or workload balance event in the log file.

How workload balancing works

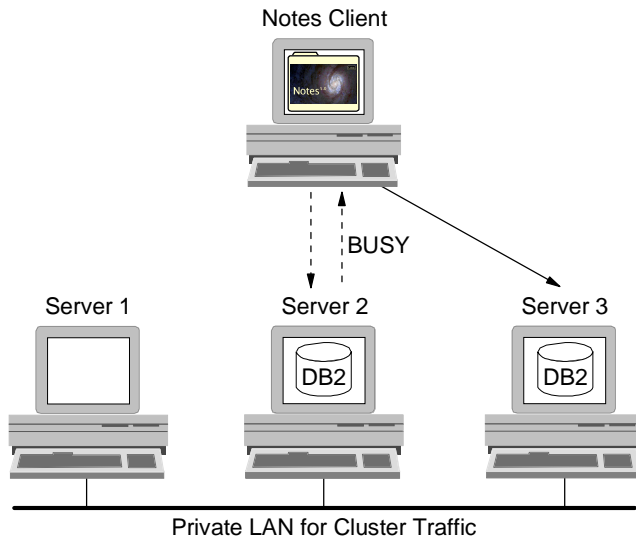
By distributing databases throughout the cluster, you balance the workload in the cluster so that no server is overloaded. In addition, there are several NOTES.INI variables you can set to help balance the workload. For example, you can specify a limit on how busy a server can get by specifying an availability threshold. When the server reaches the availability threshold, the Cluster Manager marks the server BUSY. When a server is BUSY, requests to open databases are sent to other servers that contain replicas of the requested databases. You can also specify the maximum number of

users you want to access a server. When the server reaches this limit, users are redirected to another server. This keeps the workload balanced and keeps the server working at optimum performance.

When a user tries to open a database on a BUSY server, the Cluster Manager looks in the Cluster Database Directory for a replica of that database. It then checks the availability of the servers that contain a replica and redirects the user to the most available server. If no other cluster server contains a replica or if all cluster servers are BUSY, the original database opens, even though the server is BUSY.

Example

This example describes how Domino performs workload balancing. This cluster contains three servers. Server 2 is currently BUSY because the workload has reached the availability threshold that the administrator set for this server. The Cluster Managers on Server 1 and Server 3 are aware that Server 2 is BUSY.



1. A Notes user attempts to open a database on Server 2.
2. Domino sends Notes a message that the server is BUSY.
3. Notes looks in its cluster cache to find the names of the other servers in the cluster.
4. Notes accesses the Cluster Manager on the next server listed in the cluster cache.
5. The Cluster Manager looks in the Cluster Database Directory to find which servers in the cluster contain a replica of the desired database.

6. The Cluster Manager looks in its server cluster cache to find the availability of each server that contains a replica.
7. The Cluster Manager creates a list of the servers in the cluster that contain a replica of the database, sorts the list in order of availability, and sends the list to Notes.
8. Notes opens the replica on the first server in the list (the most available server). If that server is no longer available, Notes opens the replica on the next server in the list.

For information about deciding how many replicas to create, see the chapter “Planning a Cluster.”

The cluster components

There are several components that work together to make clustering function correctly. These include:

- The Cluster Manager
- The Cluster Database Directory
- The Cluster Database Directory Manager
- The Cluster Administrator
- The Cluster Replicator
- The Internet Cluster Manager

These components are described in the following sections, except the Internet Cluster Manager, which is described in the chapter “Clustering Domino Servers That Run Internet Protocols.”

The Cluster Manager

A Cluster Manager runs on each server in a cluster and tracks the state of all the other servers in the cluster. It keeps a list of which servers in the cluster are currently available and maintains information about the workload on each server.

When you add a server to a cluster, Domino automatically starts the Cluster Manager on that server. As long as the server is part of a cluster, the Cluster Manager starts each time you start the server.

Each Cluster Manager monitors the cluster by exchanging messages, called probes, with the other servers in the cluster. Through these probes, the Cluster Manager determines the workload and availability of the other cluster servers. When it is necessary to redirect a user request to a different replica, the Cluster Manager looks in the Cluster Database Directory to determine which cluster servers contain a replica of the requested database.

The Cluster Manager then informs the client which servers contain a replica and the availability of those servers. This lets the client redirect the request to the most available server that contains a replica.

The tasks of the Cluster Manager include:

- Determining which servers belong to the cluster. It does this by periodically monitoring the Domino Directory for changes to the ClusterName field in the Server document and the cluster membership list.
- Monitoring server availability and workload in the cluster.
- Informing other Cluster Managers of changes in server availability.
- Informing clients about available replicas and availability of cluster servers so the clients can redirect database requests based on the availability of cluster servers (failover).
- Balancing server workloads in the cluster based on the availability of cluster servers.
- Logging failover and workload balance events in the server log file.

When it starts, the Cluster Manager checks the Domino Directory to determine which servers belong to the cluster. It maintains this information in memory in the server's Cluster Name Cache. The Cluster Manager uses this information to exchange probes with other Cluster Managers. The Cluster Manager also uses the Cluster Name Cache to store the availability information it receives from these probes. This information helps the Cluster Manager perform the functions listed above, such as failover and workload balancing.

To view the information in the Cluster Name Cache, type "show cluster" at the server console.

The Cluster Database Directory

A replica of the Cluster Database Directory (CLDBDIR.NSF) resides on every server in a cluster. The Cluster Database Directory contains a document about each database and replica in the cluster. This document contains such information as the database name, server name, path, replica ID, and other replication and access information. The cluster components use this information to perform their functions, such as determining failover paths, controlling access to databases, and determining which events to replicate and where to replicate them to.

The Cluster Database Directory Manager

The Cluster Database Directory Manager on each server creates the Cluster Database Directory and keeps it up-to-date with the most current database information. When you first add a server to a cluster, the Cluster Database Directory Manager creates the Cluster Database Directory on that server. When you add a database to a clustered server, the Cluster Database Directory Manager creates a document in the Cluster Database Directory that contains information about the new database. When you delete a database from a clustered server, the Cluster Database Directory Manager deletes this document from the Cluster Database Directory. The Cluster Database Directory Manager also tracks the status of each database, such as databases marked “Out of Service” or “Pending Delete.”

When there is a change to the Cluster Database Directory, the Cluster Replicator immediately replicates that change to the Cluster Database Directory on each server in the cluster. This ensures that each cluster member has up-to-date information about the databases in the cluster.

For information about marking a database “Out of Service” or “Pending Delete,” see the chapter “Managing and Monitoring a Cluster.”

The Cluster Administrator

The Cluster Administrator performs many of the housekeeping tasks associated with a cluster. For example, when you add a server to a cluster, the Cluster Administrator starts the Cluster Database Directory Manager and the Cluster Replicator. The Cluster Administrator also starts the Administration Process, if it is not already running. When you remove a server from a cluster, the Cluster Administrator stops the Cluster Database Directory Manager and the Cluster Replicator. It also deletes the Cluster Database Directory on that server and cleans up records of the server in the other servers’ Cluster Database Directories.

The Cluster Replicator

The Cluster Replicator constantly synchronizes data among replicas in a cluster. Whenever a change occurs to a database in the cluster, the Cluster Replicator quickly pushes the change to the other replicas in the cluster. This ensures that each time users access a database, they see the most up-to-date version. The Cluster Replicator also replicates changes to private folders that are stored in a database. Each server in a cluster runs one Cluster Replicator by default, although you can run more Cluster Replicators if there is a lot of activity in the cluster.

The Cluster Replicator looks in the Cluster Database Directory to determine which databases have replicas on other cluster members. The Cluster Replicator stores this information in memory and uses it to replicate

changes to other servers. Periodically (every 15 seconds by default), the Cluster Replicator checks for changes in the Cluster Database Directory. When the Cluster Replicator detects a change in the Cluster Database Directory — for example, an added or deleted database or a database that now has Cluster Replication disabled — it updates the information it has stored in memory.

The Cluster Replicator pushes changes to servers in the cluster only. The standard replicator task (REPLICA) replicates changes to and from servers outside the cluster.

How replication works in a cluster

Cluster replication is event-driven, rather than schedule-driven. When the Cluster Replicator learns of a change to a database, it immediately pushes that change to other replicas in the cluster. If there is a backlog of replication events, the Cluster Replicator stores these in memory until it can push them to the other cluster servers. If a change to the same database occurs before a previous change has been sent, the Cluster Replicator pools these changes and sends them together to save processing time.

Because Domino stores replication events in memory only, both the source and destination servers must be available for the replication to complete successfully. If a destination server is not available, the Cluster Replicator continues to store the events in memory until the destination server becomes available. The Cluster Replicator attempts periodically to push these replication events to the destination server. The interval between these attempts starts at one hour and increases over time to a maximum of one day.

If the source server shuts down before the replication completes, the replication events in memory are lost. For this reason, you should use standard replication (the REPLICA task) to perform immediate replication with all members of the cluster whenever you restart a cluster server. It is also a good idea to schedule replication between cluster servers on a regular basis, such as several times per day, to ensure that databases remain synchronized.

When the Cluster Replicator logs replication events to the log file, any replication events that are awaiting a retry are also recorded. This lets you see which databases are not currently synchronized and see what errors are preventing replication. After the errors are corrected and successful replication is completed, the error information is no longer recorded.

For more information about performing immediate replication and scheduling replication, see the book *Administering the Domino System*.

The Cluster Replicator leaves the processing of replication formulas to the standard replicator. Because these formulas can use a lot of processing power, they are not processed by the Cluster Replicator in order to minimize the overhead of using cluster replication. If you use selective replication, therefore, a database may temporarily include documents that do not match the selection formula. Domino deletes these documents when you run standard replication.

In addition, the Cluster Replicator does not honor the settings on the Advanced panel in the Replication Settings dialog box. Therefore, you cannot disable the replication of specific elements of a database, such as the ACL, agents, and design elements. The Cluster Replicator always attempts to make all replicas identical so that users who fail over do not notice that they failed over.

Caution Standard replication cannot automatically remove changes to specific database elements, such as the ACL, agents, or design elements. If limiting the replication of these items is important for a database, consider using only standard replication, not cluster replication, with that database.

Replication history in a cluster

Because replication events occur so frequently in a cluster, the Cluster Replicator does not read from or write to the replication history of a database each time it replicates the database. When replication is successful, the history information is stored in memory. Each subsequent successful replication event adds to the history information kept in memory. The Cluster Replicator transfers the history information to the databases approximately once an hour.

For more information about replication history, see the book *Administering the Domino System*.

Private folder replication in a cluster

During standard replication, private folders and their contents do not replicate, except when replicating with the client of the folder owner. Within clusters, however, private folders do replicate to other replicas within the cluster. This behavior ensures that when clients fail over, no matter which replica they access, the database contents are identical. Both cluster replication and standard replication support replicating private folders and their contents within a cluster.

Private folders can be accessed only by the creator of the folder or a server within the cluster. Only servers defined as user type “server” or “server group” in a database access control list (ACL) can access and replicate private folders within a database. Servers that are not explicitly included in the ACL cannot replicate private folders.

Mail failover in a cluster

If you create replicas of mail databases in a cluster, failover occurs in the following instances:

- When a user tries to open a mail database that is unavailable
Failover for mail works the same as for any database.
- When a user tries to send a message after the user's mail server has become unavailable

If a user is composing a message when the mail server becomes unavailable, the user can still send the message. The delivery fails over to another cluster server, where Notes deposits the message in the outgoing mailbox. Saving the message does not fail over, however, so this message is not saved in the sender's Sent folder.

- When the router tries to deliver mail to a server that is unavailable
If the server that contains the mail database is unavailable, the router delivers the mail to a cluster server that contains a replica mail database. The router uses the following process to find the correct mail database.

First, the router checks to see if mail routing failover is enabled for the local server and if the user's mail server is in a cluster. If the local server is in the same cluster and has a replica of the user's mail database, the router delivers the mail to that database. Otherwise, it asks an available cluster member for a server that contains a replica of the user's mail database, and delivers the mail to that database. If there is no replica available, the router tries again to deliver the mail to the user's mail server without using failover.

- When the user is using shared mail
Shared mail works the same on a cluster server as it does on a non-clustered server. When a user opens a shared message in a replica of the user's mail database, Domino opens the message header from the user's mail database and opens the message body from the shared mail database on that server.

When the Cluster Replicator replicates a new message to another replica of the mail database, the Cluster Replicator combines the message header and message body together as one note and sends the note to the cluster server where the replica resides. The receiving server determines if shared mail is being used. If shared mail is being used, the server deposits the message header in the replica of the user's mail database and deposits the message body in the shared mail database on the server. If shared mail is not being used, the server deposits the entire message in the replica of the user's mail database.

For more information about shared mail, see the book *Administering the Domino System*.

How calendars work in a cluster

Domino supports clustering of calendars and the Free Time database. When clustered, the scheduling system works a little differently behind the scenes than when it is not clustered. However, these differences are not noticeable to users.

When not in a cluster, each server contains a database that includes scheduling information for all users who use that server as their mail server. This database is named BUSYTIME.NSF and is known as the Free Time database. In a cluster, there is a single Free Time database for everyone whose mail server is in the cluster. This database is named CLUBUSY.NSF, and it contains all the information that was in all the Free Time databases on all the servers in the cluster. Every server in the cluster contains a replica of this database.

When you add a server to the cluster, the Schedule Manager deletes the BUSYTIME.NSF database on that server and creates the CLUBUSY.NSF database, which then replicates with the other servers in the cluster. When a user in the cluster looks for free time, the server looks in its own CLUBUSY.NSF first to find information for every user in the cluster. For users whose mail servers are outside the cluster, a request is sent to those servers for the free time information. When a user outside the cluster makes a request for information about a user in the cluster, the request fails over to another server in the cluster if the user's mail server is unavailable. Whenever there is a change to CLUBUSY.NSF on any server in the cluster, the Cluster Replicator replicates the change to the other servers in the cluster.

When you remove a server from a cluster, the Schedule Manager deletes CLUBUSY.NSF from that server and creates BUSYTIME.NSF on the server. The Schedule Manager on each server in the cluster removes the deleted information from its replica of CLUBUSY.NSF.

Note If there are Domino Release 4.5 or 4.6 servers in a cluster, these servers maintain their BUSYTIME.NSF databases. These databases are not converted to CLUBUSY.NSF. Calendaring for these servers works the same as in a non-clustered environment.

For more information about how calendaring works in Domino, see the book *Administering the Domino System*.

How operating system clusters work

The method of clustering that Domino uses is called “application clustering.” Domino, which is an application, monitors the cluster and determines when failover and workload balancing should occur, based on parameters that you set.

Another form of clustering is “operating system clustering.” In this form of clustering, the operating system monitors the cluster and determines when failover should occur. When failover occurs, the server (called a node) to which you fail over takes over the resources of the failed node, accesses the storage space the failed node was accessing, and runs the applications the failed node was running.

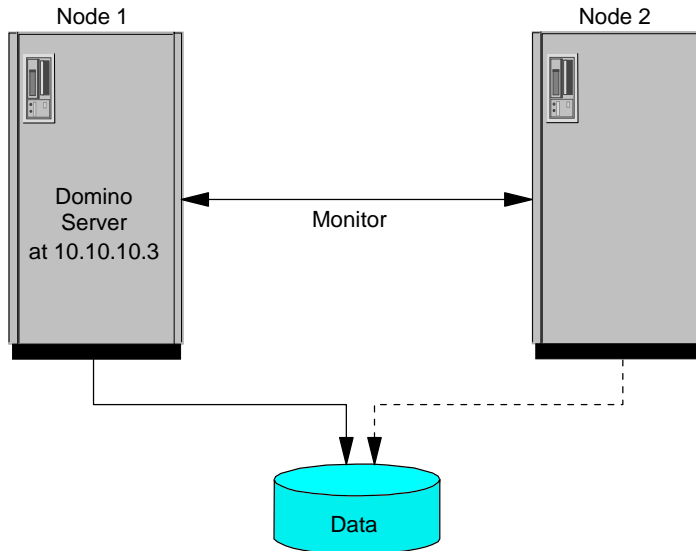
There are two basic methods of running operating system clusters, active-passive and active-active. In an active-passive cluster, passive nodes do not run their own applications but instead stand by to take over if active nodes fail. In an active-active cluster, the nodes all run their own applications but are also available to take over if other nodes in the cluster fail. In addition, you can configure an operating system cluster to fail over only when there is a hardware failure or to fail over when there is either a hardware failure or a software failure.

Because Domino uses application clustering, this section does not give detailed information about the various methods and configurations that operating system clusters use. However, because you can run Domino in conjunction with several operating system clusters, including IBM High Availability Cluster Multi-Processing (HACMP), Microsoft® Cluster Server (MSCS), and Sun(TM) Cluster, this section describes basic information about operating system clusters.


Note For information about configuring your operating system cluster software to run with Domino, see the documentation that came with your operating system cluster.

Operating system clusters provide failover that is transparent to users. Because the receiving node takes over the resources of the failed node, the user sees the same server name and same network address as on the original server. Unlike many operating system clusters, Domino clustering does intelligent failover. When a server fails, Domino checks its cluster cache to find the server that is most available in the cluster. Domino also lets you actively control workload balancing, which operating clusters may not offer. In addition, Domino clustering lets you set up clusters of servers that run different operating systems, while operating system clusters require that all nodes run the same operating system.

The following example shows a basic configuration for active-passive operating system clustering. A Domino server runs on Node 1. Node 2 monitors Node 1 and waits for a failure to occur.




When a failure occurs, Node 2 picks up the resources of Node 1 and takes over running the Domino server. Node 2 uses the same disk set and the same IP address for the Domino server that Node 1 used.




To run Domino in an active-active cluster, you must use Domino partitioned servers on the nodes. Doing so lets each node take over the tasks of the other node while also maintaining its own tasks.

The following example shows a basic configuration for active-active operating system clustering. Node 1 and Node 2 each have Domino running in the first partition. The second partition on each node duplicates the resources of the first partition on the other node. Each node has its own disk set, but both nodes have access to both disk sets in case failover occurs. The nodes monitor each other.



When Node 1 fails, Node 2 picks up the resources of Node 1 and runs the Domino servers for both nodes.



To use an active-active configuration, you must be sure that each node can handle the load of the other node if failover occurs.

Benefits of using OS clusters with Domino clusters

When you use an operating system cluster in conjunction with a Domino cluster, the few things that do not fail over in a Domino cluster will fail over in the operating system cluster. Here are a few examples:

- Most Domino agents do not fail over, so when a server fails over in a Domino cluster, agents that were running do not continue running on the new server. If these agents are configured to run on a specific server, they will not be able to run on another server after Domino failover occurs. In an operating system cluster, however, the same server name is used after failover occurs. Therefore, the agents can run on this server. In an operating system cluster, agents that were running on a schedule when failover occurred will restart the next time they are scheduled to run.
- If you have applications that use hard-coded server names, the applications will not work if they fail over to a different Domino server. These applications will run after failing over in an operating system cluster, however, because the server name is still the same.
- If a user is editing a document when the server fails, the user can't save the document in a Domino cluster. The user has to paste the document into a replica on the new server. In an operating system cluster, however, users can save documents that they were editing when the server failed.
- The Administration Process does not fail over in Domino. Therefore, it is useful to set up operating system clusters for your administration servers.

For these features, it is a good idea to set up an active-passive operating system cluster to run in conjunction with the Domino cluster.

Chapter 3

Planning a Cluster

When planning a cluster, it is important to consider the performance and ability of your hardware. The cluster must have enough CPU power, memory, and disk space to handle the cluster traffic and the number of databases and replicas required.

Planning a cluster also includes the following:

- Determining the number of servers to include in the cluster
- Determining the number and placement of replicas in a cluster
- Distributing databases across cluster servers
- Determining whether to create a private LAN for cluster traffic
- Deciding whether to create a cluster over a wide area network (WAN)
- Deciding whether to use fault recovery in a cluster
- Determining the cluster configuration

After the cluster is up and running, you can further balance the workload by setting a maximum number of users for each server and setting the availability threshold in a way that does not allow any server to become overloaded. You can track the cluster statistics to determine whether you need to make any changes to the cluster setup.

For more information about balancing the workload, see the chapter “Managing and Monitoring a Cluster.”

Determining how many servers to include in a cluster

In general, adding servers to the cluster increases the cluster’s ability to balance the workload so that no server becomes overloaded and performance stays high. However, if servers use too many server tasks, CPU-intensive applications, and replicas, adding servers can decrease performance because of the additional amount of cluster traffic required to keep databases synchronized on all servers.

If your organization is small, you can start with two servers and add servers as your enterprise grows without affecting the performance to your users. Keep in mind that each server you add creates additional network traffic

when it probes other cluster servers to find out their status and when it does cluster replication. Therefore, do not add servers to a cluster until you need the additional capacity or additional redundancy.

In a larger organization, you must decide whether to create large clusters or small clusters. A larger cluster is better able to absorb the workload when a cluster server fails. If you have a cluster with only two servers, for example, if one of the servers fails, the other server must absorb 100% of the failed server's workload. That means that you could run each server at only 50% of its capacity so that it has enough capacity available to absorb the workload of the other server. If the cluster has six servers, however, each of the remaining five servers must absorb only 20% of the failed server's workload. That means you could run each server at 80% of capacity, and they would still be able to absorb the workload if a server goes down. (Of course, there are other factors that determine how the workload of a failed server is absorbed, such as the way you have distributed replicas across the cluster servers.)

Hardware considerations

The number of servers you decide to include in a cluster can be affected by the amount of disk space and the processing power of each server. Keep the following in mind as you decide which hardware to use in your cluster:

- The more replicas you create, the more disk space you need and the more processing power you need for cluster replication.
- The Cluster Database Directory requires approximately 2MB of disk space plus an additional 1MB for each 2,000 databases in the cluster.
- The more servers in the cluster, the more processing power each server uses to communicate with the other cluster servers.
- The more server tasks and CPU-intensive applications you run on a server, the more processing power you need.
- Each server needs adequate processing power for the databases it contains and for any databases that might fail over to the server.
- Clustered servers require more memory than nonclustered servers. The actual amount you will require depends on the level of activity on the server.

To see if you need additional memory or processing power on your computer, check the Platform statistics.

For information about the Platform statistics, see the book *Administering the Domino System*.

- When you have a large cluster or a cluster with a heavy workload, you might need to use multiple Cluster Replicators to improve the performance of cluster replication. Check the Cluster Replicator statistics to see if there is a large queue of databases waiting to be processed. If so, add Cluster Replicators one at a time until the statistics improve adequately.

Because Cluster Replicators use system resources, the overall performance of the system could decrease as you add Cluster Replicators. Therefore, do not add more Cluster Replicators than you need.

For more information about using multiple Cluster Replicators, see the chapter “Managing and Monitoring a Cluster.”

Determining the number and placement of replicas in a cluster

There are two major reasons to create a replica for a database in a cluster — to provide constant availability of the data and to distribute the workload between multiple servers. Before you create replicas in a cluster, consider how frequently users access a database and their need for data redundancy. If a database is extremely busy or its availability is extremely important, you may want to create multiple replicas and locate them on your most reliable servers. For databases that are not very busy and whose constant availability is not important, you may not want to create any replicas at all. A server log file, for example, does not need to have a replica on another server.

In general, the more replicas of a database, the more accessible the data. Creating too many replicas, however, can add unnecessarily to the overhead of maintaining a system and affect performance. As you plan your cluster strategy, try to create a balance between your users’ requirements for data availability and the physical ability of each server in your cluster to manage additional workload. More than three replicas of a database may not provide you with significant incremental availability. If users can adequately access a database from one or two servers, do not increase the number of replicas in the cluster.

When users require the constant availability of a specific database, consider placing replicas on every server in the cluster if you have adequate disk space and resources. If you are a public service provider, this configuration provides the highest possible redundancy of data.

In addition, try to distribute the busiest databases to different servers so that no server contains too many busy databases. If the servers in the cluster all have a similar amount of processing power, you can have an equal load on each server, including the processing power reserved for failover. If a server

has significantly more or less processing power than the other servers, consider changing the number of databases on the server and the number of databases that can fail over to the server. Also, distribute mail files across a cluster, or set up separate servers or separate clusters for mail.

Because busy databases in a cluster can create a lot of replication events, it is a good idea to install these replicas on the fastest disk hardware available in the cluster. If possible, place these replicas where other processes are not in contention — for example, on a partition other than the one that contains the operating system swap file.

To view which databases and replicas already exist in the cluster, open the Cluster Database Directory (CLDBDIR.NSF). It contains a document that stores information about each database and replica in a cluster.

Note Selective replication formulas work differently in a cluster.

For more information about selective replication in a cluster, see the chapter “How Domino Clustering Works.”

How many replicas to create

The following list describes some factors to consider when determining how many replicas to create.

- The number of replicas of a database you create depends on how important the availability of that database is and the amount of use the database receives.
- You should create at least one replica of a database for which you want data redundancy. If a database becomes unavailable, users can then fail over to the replica.
- If you want to be sure that a database is available at all times, you can create more than one replica. The more important availability is, the more replicas you should create. Add multiple replicas for very important databases only. Unneeded replicas can diminish cluster and network resources.
- For most databases, a single replica is adequate. Rarely are more than three replicas needed, unless a database is truly mission-critical.

- Consider the power and bandwidth of your system when creating replicas. The busier a database is, the more network traffic and processing power it takes to keep replicas updated. If you have systems with limited power and bandwidth, you may want to create fewer replicas of busy databases than you would if you had more power and bandwidth, or you may want to add more processors and other resources to the servers. In a cluster with limited resources, creating replicas of busy databases can be counterproductive because of the additional resources needed for cluster replication. (Clustering is not a solution for inadequate resources.) The less busy a database is, however, the less overhead it takes to keep that database updated.
- If you aren't sure how many replicas to create, start with one and track the cluster statistics. If the statistics show that the server becomes unavailable or that performance becomes a problem, increasing the number of replicas may solve the problem.
- Do not create replicas of databases for which availability or workload balancing is not one of your goals.

Analyzing databases to determine the number of replicas

There are many factors to consider when deciding how many replicas to create. Some factors suggest creating more replicas, and some suggest creating fewer replicas. Below is a list of factors and how they might affect your cluster traffic and performance.

Prior to distributing databases in a cluster, it can be helpful to create a table of information about the databases and the cluster hardware. You can use the table to determine how important specific databases are and how adequate your resources are. You can include some or all of the following in the table:

- **Titles of the databases**
This identifies each database.
- **Size of each database**
Large databases consume a lot of disk space. Depending on your disk capacity, you may want to create fewer replicas of larger databases to preserve disk space.
- **Number and distribution of database users**
If you have a large number of users, they will probably experience better performance if usage is spread across multiple servers. This requires multiple replicas. If the number of users is small, they probably won't notice a performance improvement from additional replicas.

- How often user transactions take place
If the transaction rate is high, creating multiple replicas may improve performance.
To find out the rate of activity for a database, look in the Notes log file.
- Expected volume of new data
If you expect a large amount of new data in the database, additional replicas may slow down performance because cluster replication will cause a lot of additional traffic. If you have powerful servers and a lot of bandwidth, this may not create a problem.
- Capacity of Domino server hardware
The more powerful the servers and the more disk space they have, the more active replicas you can create without significantly affecting performance.
- Type of network connections between servers
Cluster replication can create a bottleneck on a network that does not have enough bandwidth. Therefore, the greater the bandwidth, the more replicas you can create.
- How critical the database is to the functioning of your business
For databases that are mission-critical, you should create multiple replicas. For databases where availability is less important, create fewer replicas or none at all.

Example table

When you create a table of database information, include the factors that are most important to you. The following table uses a subset of the preceding information to determine the number of replicas needed.

<i>Database title</i>	<i>Size</i>	<i>Maximum concurrent users</i>	<i>Transaction rate</i>	<i>Growth rate</i>	<i>Need for availability</i>	<i>Suggested number of replicas</i>
Product Discussion	4GB	600	High	High	High	2
Sales Tracking	1GB	200	Medium	High	Critical	2 or more
Company Research	2GB	20	Low	Medium	Medium	0 or 1
Classified Ads	1GB	50	Medium	Medium	Low	0

This table helps identify which databases require high availability, which databases are busiest, and how much additional disk space you will need in the future. In this example, two databases are very important and are growing rapidly. You should be sure that there are enough replicas of these

databases so that they are always available. You should also be sure there is adequate disk space for growth on every server that contains a replica of these databases. One database is of medium importance, not growing as quickly, and not very active. You should provide no more than one replica of this database, unless it would affect your business negatively if the database was not available for a while. One database is not very important and does not require a replica in the cluster.

The number of concurrent users helps you determine the need for workload balancing. In this example, two databases are very busy and both are very important. Therefore, you should consider placing these databases on different servers to balance the workload. You should also be sure that workload balancing parameters are set on the servers that contain these databases so that users will fail over to another server when these databases become busy.

For more information about managing workload balancing, see the chapter “Managing and Monitoring a Cluster.”

Distributing databases in a cluster

The way you distribute databases significantly affects workload balancing, as well as the performance of your equipment.

Distributing mail databases

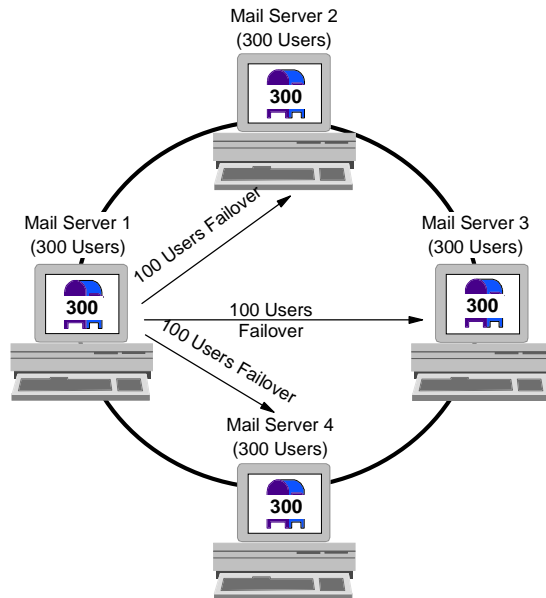
Clustering mail databases provides high availability to users. Some companies set up a cluster for mail only. This is not required but is one way to arrange your organization. Because mail is an important application, it is a good idea to have your mail databases in a cluster, whether or not the cluster is dedicated to mail servers only.

When you create a mail cluster, distribute the replicas for each server among the other servers in the cluster. This ensures that all the other servers share the load when a server fails, thus balancing the workload and keeping performance as high as possible.

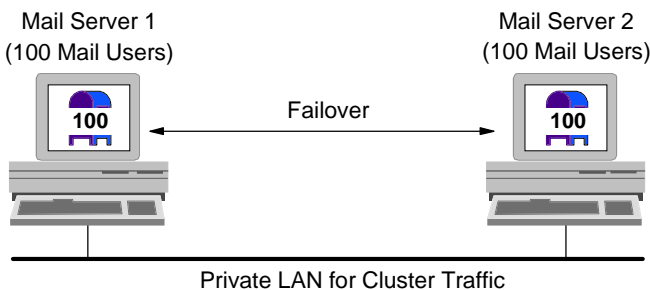
It is a good idea to distribute replicas evenly across servers, as long as the servers are fairly equal in resources. If your mail cluster contains four servers and 1,200 users, for example, place 300 mail databases on each server, and place 100 replicas from each server onto each of the other three servers. When a server fails, 100 mail users fail over to each of the other three servers,

increasing each server's work load by 33%. You might be tempted to place all 300 replicas from Server 1 onto Server 2 and all the replicas for Server 3 onto Server 4. In such a case, however, if Server 1 fails, all 300 users fail over to Server 2, increasing the workload on Server 2 by 100% but not increasing the workload on Server 3 or Server 4 at all.

The following figure shows a mail cluster that contains four servers with 300 mail databases on each server. Replicas of the mail databases are evenly distributed among all the other servers in the cluster, keeping the workload of the other servers as low as possible, even when failover occurs.



The following figure shows a mail cluster that contains two servers with 100 mail files on each server. Because there are only two servers, each server must fail over to the other server. Therefore, each server contains replicas of all the mail databases on the other server.



Since users often open mail databases once a day and leave them open, distributing the mail databases is usually adequate for workload balancing. You do not usually have to use separate workload balancing settings, especially if you dedicate servers to mail only.

After failing over to a replica mail database, users automatically return to the mail database on their mail server the next time they start their Notes clients, as long as the Location document that points to that mail database is the current Location document.

Note If you do not create a dedicated mail cluster, you should distribute mail databases equally among the cluster servers, if the cluster servers are approximately equal in power. If some servers are more powerful than others, distribute more databases to the more powerful servers. This distribution helps to keep the workload balanced.

Caution If you plan to create a cluster that includes some Domino 6 servers and some Domino Release 5 or Domino Release 4.6 or 4.5 servers, keep the following in mind:


- The Domino 6 mail template does not work properly on Domino Release 5 or Domino Release 4.6 or 4.5 servers. If a user has a Domino 6 mail database, do not create a replica on a Domino Release 5 or a Domino Release 4.6 or 4.5 server.
- Because the Cluster Replicator always replicates the template design between replicas, a user's mail replicas should all use the same template — the Domino 6 mail template, the Domino Release 5 mail template, or the Domino Release 4.6 or 4.5 mail template.

Distributing application databases


When clustering applications, some applications need a higher level of availability than others. In addition, some databases are used much more frequently than others, and some require more computing resources to run than others. You should consider all of these factors when distributing application databases. With application databases, workload balancing is more important than with mail databases.

As with mail, you should try to distribute workload evenly across the cluster. However, take into consideration the power of each server. If some servers are more powerful than others, be sure to increase their workload appropriately. Also, take into consideration the amount of usage for each database.

The following figure shows a cluster with four servers of varying amounts of power. The databases in the cluster are distributed in a way that takes advantage of the resources of each server.



The following figure shows a cluster with four servers that are equal in power. The databases in this example all receive a similar amount of use. DB1 is a critical database, so each server contains a replica.



Determining whether to create a private LAN for your cluster

To make a busy cluster more efficient, you can create a private network for your cluster. To do so, you install an additional network interface card in each cluster server and connect these network interface cards through a private hub or switch.

The main reason for creating a private LAN for your cluster is to separate the network traffic a cluster creates when it uses cluster replication and server probes. A private LAN can leave more bandwidth available on your primary LAN. If you anticipate a lot of cluster replication activity, you should create a private LAN.

You can also consider creating a private LAN for intra-cluster communication to ensure that cluster servers remain in communication with each other, even when certain network problems occur. By adding a private network, all servers in the cluster are connected by at least two distinct LAN segments. So if a network board or a cable on one LAN segment fails, there is still network connectivity between all servers in the cluster. This assures you that cluster servers remain in contact with each other and that cluster replication continues to keep databases synchronized.

If you create a private LAN for your cluster, all cluster members must be connected to both the private LAN, for intra-cluster communication, and the primary LAN, for client access.

For information about setting up a private LAN, see the chapter “Setting Up a Cluster.”

Clustering over a wide area network

A cluster over a wide area network (WAN) works the same way as a cluster on a LAN. However, if you have a low-speed WAN, you should consider disabling cluster replication. Instead, use scheduled replication more frequently than usual, such as every hour. This reduces WAN traffic, bottlenecks, and the cost of continual transmission.

Also keep in mind that Domino fails over to the most available server in the cluster, not the closest server. For example, if you have three servers — one in Boston, one in New York, and one in Hong Kong — the Boston server would fail over to the Hong Kong server if it is more available than the New York server. You can control this behavior to some extent by changing the server availability thresholds on the cluster servers.

For information about the server availability threshold, controlling failover, and disabling cluster replication, see the chapter “Managing and Monitoring a Cluster.”

Using a cluster over a WAN is a good idea for disaster planning. Having emergency backup servers at different locations is a good way to ensure that necessary data is always available when you need it.

Fault recovery in a cluster

Fault recovery is the ability of a Domino server to clean up and restart itself after a failure. Fault recovery works well in a Domino cluster. If there is no Domino server to fail over to, fault recovery still ensures that users will have constant access to their data. Even if users fail over to another cluster server, fault recovery increases availability because the failed server becomes available again. In addition, depending on the workload balancing parameters you’ve set, some users will fail back to the original server when they open new databases.

If you are using an operating system cluster in conjunction with a Domino cluster, the decision about whether or not to use fault recovery depends on how you configured the operating system cluster. If you configured the operating system cluster to fail over on a hardware failure only, fault

recovery works well. Fault recovery restarts Domino on its current server, and no operating system fail over occurs.

If you configured your operating system cluster to fail over on both hardware and software failures, you don't need fault recovery because the operating system cluster will restart Domino on another server in the cluster. In fact, you should disable fault recovery so you won't have Domino restarting itself while the operating system cluster is also restarting it. This can lead to problems.

By default, fault recovery is disabled. You enable it in the Server document.

1. From the Domino Administrator or the Web Administrator, click the Configuration tab.
2. In the Task pane, expand Server, and click All Server Documents.
3. In the Results pane, select the Server document you want, and click Edit Server.
4. In the Fault Recovery field, choose Enabled.
5. (Optional) Complete any of the following fields that you want.
 - In the Cleanup Script Name field, enter the name of a cleanup script.
 - In the Cleanup Script Maximum Execution Time field, enter the maximum time for a cleanup script to run before being terminated.
 - In the Maximum Crash Limits field, enter the maximum number of restarts allowed during the specified period. If the number of restarts exceeds the limit, the server won't restart.
 - In the "Mail Crash Notification to" field, enter the names of the people to notify each time the server restarts.
6. Make any other changes you want to the Server document, and then click Save & Close.

For more information about fault recovery, see the book *Administering the Domino System*.

Examples of cluster configurations

There are many ways to configure a cluster. Among these, you can:

- Cluster two servers for mail and applications
- Cluster three servers for mail and applications
- Cluster six servers for mail and applications
- Cluster hub servers
- Set up a cluster for disaster preparedness


- Cluster partitioned servers
- Cluster passthru servers
- Use a Domino cluster with an operating system cluster

These examples assume that all servers in a cluster are equal in processing power and resources. If that is not true in your cluster, you may need to make adjustments.

Example of clustering two servers for mail and applications

If you have only two servers in your cluster, you can set them up in one of two ways: You can use one of the servers as the primary server for user access and use the second server as a backup and failover server, or you can equally divide the workload between the two servers and have them fail over to each other. Dividing the workload typically ensures better performance when both servers are running. When one server is not available, performance is the same in both scenarios because one server must process the entire workload of both servers.

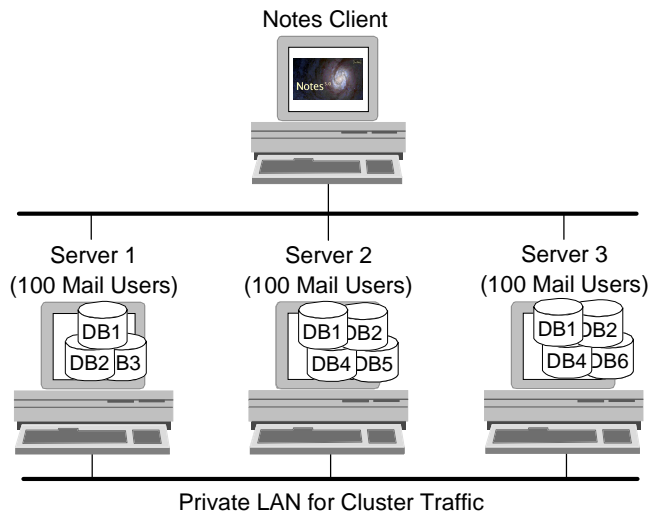
The following figure shows a cluster with two servers with the workload divided between the servers.



Example of clustering three servers for mail and applications

If you have three servers available in the cluster, distribute databases across the three servers. In this scenario, you should replicate critical databases to every server in the cluster. You should replicate non-critical databases to only one other server or to no other server at all. In addition, distribute mail databases equally across the cluster servers. To ensure balance in case of failover, each server should contain replicas of half of the mail databases of each of the other servers. In that way, if a server becomes unavailable, the mail workload of that server is equally divided between the other two servers.

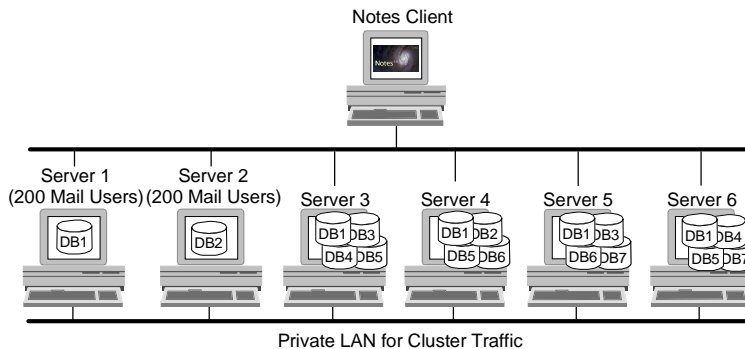
In the following figure, each server contains a replica of the critical databases, DB1 and DB2. There are fewer replicas of the other databases because they are less critical. In addition, each server contains 100 mail databases. Therefore, each server should contain replicas of 50 mail databases from each of the other servers.



Example of clustering six servers for mail and applications

If you have a larger organization and six servers available, you can create a larger cluster. In this configuration, be sure that you do not create more replicas than necessary because of the extra workload created by cluster replication. In a large cluster, you can distribute databases equally across servers or you can dedicate certain servers to particular tasks. For example, Servers 1 and 2 could be mail servers, and Servers 3, 4, 5, and 6 could be application servers with applications divided evenly among them. Critical databases would have replicas on multiple servers. In this configuration, you could even dedicate one of the servers as a backup server for critical databases only. In this case, the server would receive real-time cluster replication events so that the databases are always up-to-date, but users would not be able to access the databases on this server. It would be used strictly to store up-to-date copies of databases in case the primary copies of the databases become damaged or deleted.

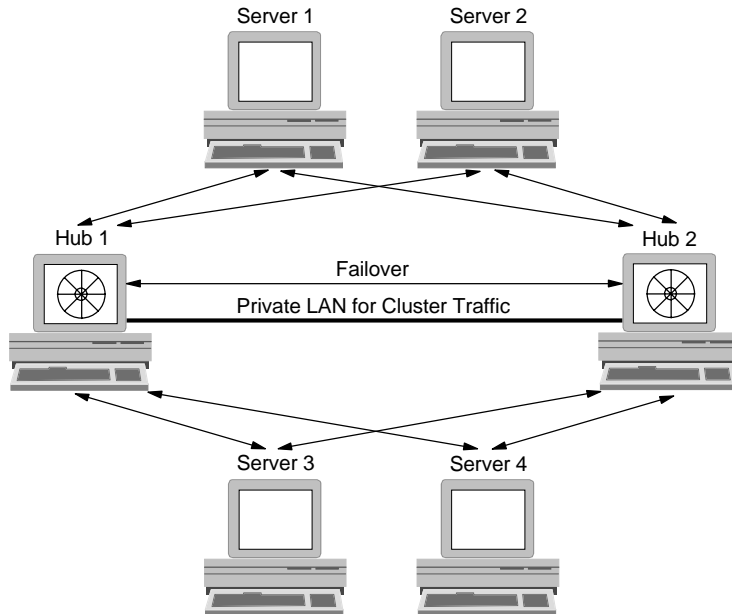
In the following figure, there are six servers — two servers used primarily for mail and four servers used for applications. Some databases, such as DB1, are more important than others, so they have more replicas. In this scenario, you can replicate all of the mail databases from one mail server to the other mail server, if the servers are powerful enough to handle the load if failover occurs. If the two mail servers are not powerful enough, you can distribute some mail replicas to one or more of the other servers in the cluster.



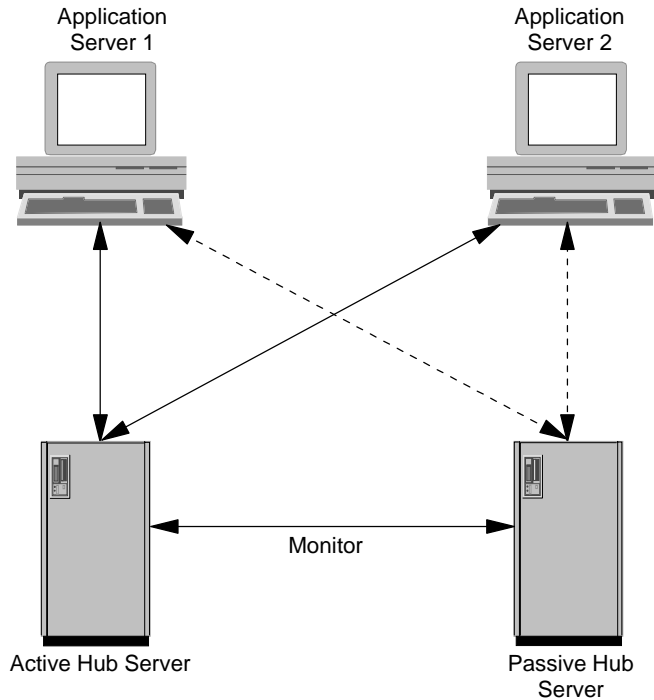
Examples of clustering hub servers

If you use a hub-and-spoke configuration, which is very efficient, you can cluster the hub servers to ensure availability. If one of the hub servers is not available, mail routing or replication fails over to another hub server.

The following figure shows two hub servers in a cluster. The purpose of this cluster is to route mail.




When the primary purpose of the hub servers is replication, the previous configuration may not work well in some enterprises because of the amount of replication that is required. Because both hub servers need to contain the same databases, this configuration causes a lot of replication. Depending on the equipment you use, this configuration could slow down servers and create a lot of network traffic. If that is the case in your enterprise, you might consider using an active-passive operating system cluster for your hub servers, as in the following figure.



Example of clustering for disaster preparedness

You can use a Domino cluster for disaster preparedness by making members of the cluster unavailable to users but still accessible to cluster replication. Although you can use a LAN to do this, it is also a good idea to use a WAN so that your backup servers are at a different location.

The following example shows a simple emergency preparedness setup. The cluster includes three servers, two that are accessible to users and one that is not. The server that is not accessible is at a different location. This remote server is not accessible to users because it has been designated as **RESTRICTED** by using the `Server_Restricted` setting in the Configuration Settings document or in the `NOTES.INI` file. This server contains replicas of databases that are critical. Although users can't access the remote server, cluster replication still keeps replicas on the server up-to-date.




For more information about designating a server as **RESTRICTED**, see the chapter “Managing and Monitoring a Cluster.”

Example of clustering partitioned servers

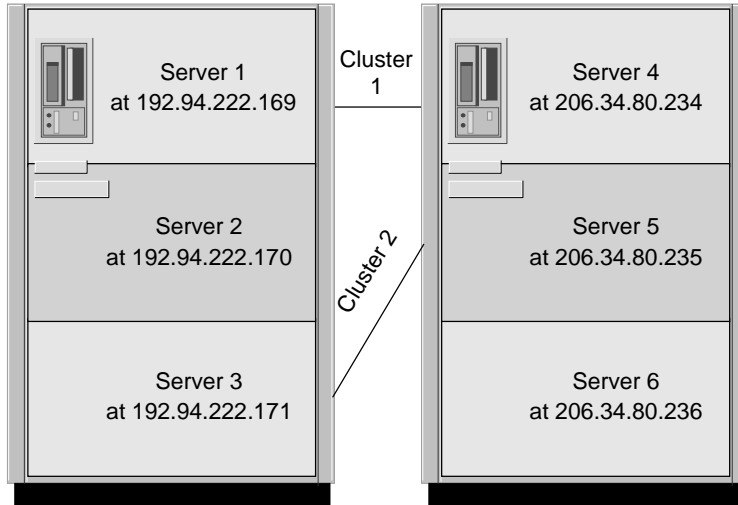
You can include partitioned servers in a cluster. You can cluster partitioned servers with individual servers and with partitioned servers that are on different computers. Do not include partitioned servers that are on the same computer in the same cluster. Doing so can cause significant failover and workload problems in the cluster if that computer fails.

If you use port mapping for your partitioned servers, it is not a good idea to include those partitioned servers in a cluster. Because a failure of the port mapping server makes communication with the other partitioned servers impossible, this can cause failover and workload problems in the cluster.

The following figure shows two computers with three partitioned servers on each. These partitioned servers are configured in three clusters with two partitioned servers in each cluster.



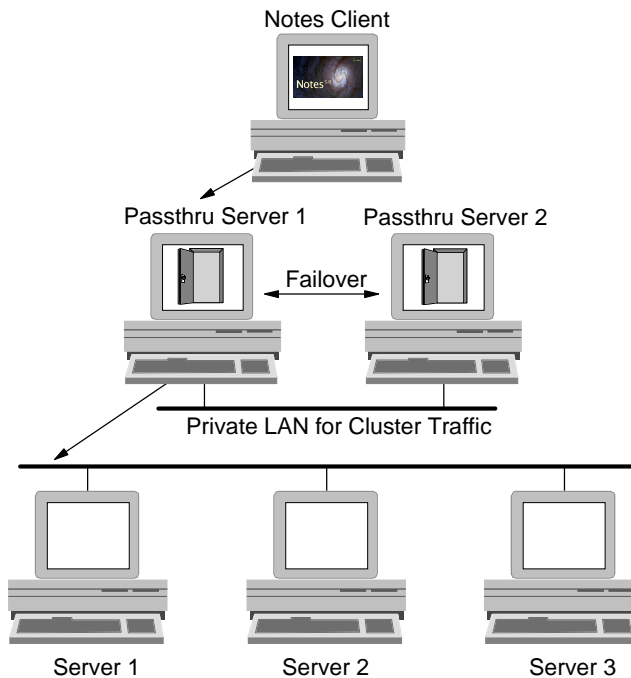
When you include a partitioned server in a cluster, you do not have to include all the partitioned servers on a machine in the cluster. The following figure shows two computers that each have three partitioned servers. Four of the partitioned servers are configured in two clusters, and two of the partitioned servers are not in a cluster.



Example of clustering passthru servers

If you use passthru servers that are critical for providing high availability to your users, you can set up a cluster of passthru servers. To do this, set up two or more passthru servers using the same passthru configuration settings, create a cluster of these passthru servers, and enter the name of one of these servers as the passthru server in users' Location documents. Also, create Passthru Server Connection documents in users' Personal Address Books, and enter the name of one of the cluster servers as the passthru server name.

The following figure shows a client contacting a cluster of two passthru servers in order to connect to another server in the enterprise. The client contacts one of the passthru servers in the cluster, and then one of the passthru servers connects the client to the server the client requested.



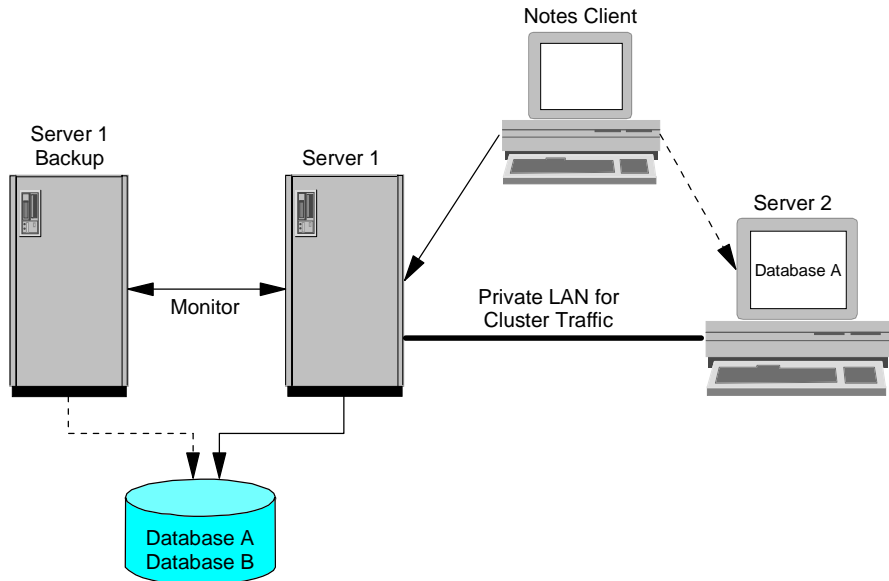
This configuration works only for clients that have a LAN or WAN connection to the passthru servers. For information about setting up mobile clients to use passthru servers, see the chapter "Setting Up a Cluster."

Example of using a Domino cluster with an operating system cluster

If you use operating system clusters as well as Domino clusters, you can use them together to improve availability and ensure that tasks that don't fail over in Domino are still available.

The following figure shows a two-server Domino cluster (Server 1 and Server 2) and a two-server active-passive operating system cluster (Server 1 and Server 1 Backup). In the operating system cluster, Server 1 Backup stands by to take over if Server 1 fails. Because Domino fault recovery is running on the Domino servers, the operating system cluster is configured to fail over only in case of hardware failure. (If Domino fails, fault recovery restarts it.)

The Notes client is using two databases on Server 1. Database A has a replica on Server 2 in the Domino cluster. Database B does not have a replica in the Domino cluster. If there is a hardware failure on Server 1, Domino clustering causes the client to fail over to Server 2 for access to Database A. Then operating system clustering causes Server 1 Backup to start, letting the client resume working in Database B. In addition, there are again two active servers in the Domino cluster so that Domino failover and workload balancing are available again.



Chapter 4

Setting Up a Cluster

This chapter describes how to set up a cluster and make sure it is working correctly.

Before you create a cluster

Before you create the cluster, complete the following tasks:

- Determine which servers will be in the cluster and what they are named.
- Install or upgrade to the Lotus Domino 6 Enterprise server or the Lotus Domino 6 Utility server. These servers enables you to set up a Domino 6 cluster.

Note You can include the following Domino servers in a cluster: the Domino Release 5 Enterprise server, the Domino Release 4.62 Enterprise server, the Domino Release 4.6 Advanced Services server, and the Domino Release 4.5 Advanced Services server. However, these servers will not have access to the features that are new in Domino 6.

- Be sure that you have reviewed the requirements for clustering and that all servers meet these requirements.
- Distribute databases and replicas to the servers you plan to include in the cluster.
- Set up scheduled replication between the cluster servers.

For more information about clustering requirements, see the chapter “Cluster Benefits and Requirements.” For more information about distributing databases, see the chapter “Planning a Cluster.” For more information about setting up scheduled replication, see the book *Administering the Domino System*.

Creating a cluster

To create a cluster, you must have the following access rights:

- Author access and Delete Documents rights and the ServerModifier and ServerCreator roles in the Domino Directory
- Author access with Create documents rights in the Administration Requests database

If possible, use the administration server when creating a cluster. This makes the creation process faster. The administration server does not have to be part of the cluster.

If a server belongs to a different cluster, you do not have to remove the server from that cluster before you add it to the new cluster. The Cluster Administration Process removes the server from the original cluster and then adds it to the new cluster.

Note You cannot use the Web Administrator to create a cluster.

1. From the Domino Administrator, make sure the administration server or another server is current.
2. Click the Configuration tab.
3. In the Tasks pane, expand Server, and click All Server Documents.
4. In the Results pane, select the servers that you want to add to the cluster.
5. Click Add to Cluster.
6. In the Cluster Name dialog box, choose Create New Cluster, and then click OK.
7. Type the name of the new cluster, and then click OK.
8. Choose Yes to add the servers to the cluster immediately, or choose No to submit a request to the Administration Process to add the servers to the cluster.
9. (Optional) If you chose No in Step 8 and you did not add the servers on the administration server, force replication between the server you used and the administration server so that the administration server receives the requested changes sooner.
10. (Optional) If you chose No in Step 8, force replication between the administration server and the cluster servers so the cluster servers receive all the changes sooner.
11. (Optional) If you chose Yes in Step 8, the cluster information is added immediately to the Domino Directory on the server you used to create the cluster. If this server is not part of the new cluster, replicate the changes to one of the servers you added to the cluster.

The cluster creation process

You can choose to add the servers to the cluster immediately or to use the Administration Process to add the servers to the cluster. Although choosing “immediately” adds the servers to the cluster faster, it occasionally leads to replication conflicts.

If you use the Administration Process to add the servers to the cluster, Domino submits a request to the Administration Requests database to add the name of the cluster to the Server document of each server in the cluster. If you are using the administration server to create the cluster, the Administration Process runs immediately and adds the cluster name to the Server documents. If you are not using the administration server, the Administration Requests database on the server you are using must replicate with the administration server before the Administration Process can add the cluster name to the Server documents.

After the Administration Process makes changes to the Server documents in the administration server’s Domino Directory, Domino replicates the changes to the other servers in the cluster, ensuring that every cluster server has an updated list of the servers in the cluster.

If you choose to add the servers to the cluster immediately, Domino immediately makes the changes to the Server documents in the Domino Directory on the server you use to create the cluster. If that server is part of the cluster, the server immediately starts the cluster processes listed below and replicates its Domino Directory with another server in the cluster so that the other cluster servers become informed that they are part of the cluster. If you did not use a cluster member to create the cluster, this process starts when the Domino Directory of the server you used replicates with the Domino Directory of a server in the cluster.

The following occur on each cluster server as part of the cluster creation process:

- The Cluster Administrator and the Cluster Manager start.
- The Cluster Administrator starts the Administration Process, if it is not already running.
- The Cluster Administrator starts the Cluster Database Directory Manager.
- The Cluster Administrator starts the Cluster Replicator.
- The Cluster Database Directory Manager creates the Cluster Database Directory (CLDBDIR.NSF).
- The Cluster Database Directory Manager updates the Cluster Database Directory so it includes a document for each database on the server.

- The Cluster Administrator replicates the Cluster Database Directory and the Domino Directory with the other servers in the cluster so they are all synchronized.
- The Schedule Manager creates the Free Time database (CLUBUSY.NSF).
- The Free Time database replicates with the other cluster servers so they are all synchronized.

When the Domino Directory updates to include the new cluster, each cluster server begins to send messages, known as probes, to the other servers in the cluster. These probes gather information about the status of the other servers in the cluster.

Verifying that the cluster was created correctly

You can do the following to check that the cluster was created correctly.

<i>Action</i>	<i>What you should see</i>
From the Domino Administrator or the Web Administrator, expand Clusters in the Server pane.	The name of the cluster followed by the names of the cluster servers.
<ul style="list-style-type: none"> • From the Domino Administrator or the Web Administrator, click the Configuration tab. In the Task pane, expand Cluster, and then click Clusters. • In the Results pane, open the Server documents of the servers you added to the cluster. 	<ul style="list-style-type: none"> • The name of the cluster followed by the names of the cluster servers displayed in the Results pane. • The name of the cluster in the Cluster name field on the Basics tab.
From the Domino Administrator or the Web Administrator, click a cluster server in the Server pane, and then click the Files tab.	The title “Cluster Directory (6)” in the Tasks pane. This shows that Domino created the Cluster Database Directory.
Compare the replica IDs of the Cluster Database Directories on each cluster server.	The same replica ID on each server.
From the server console, send the following command: show cluster	The name of the cluster, some statistics for the current server, and the names of all the cluster servers.

You can also use Cluster Analysis to generate reports that show if there are any configuration problems in the cluster.

For information about Cluster Analysis, see the topic “Using Cluster Analysis to check the cluster configuration.”

Using Cluster Analysis to check the cluster configuration

You can use Cluster Analysis to verify that a cluster is set up correctly. Cluster Analysis generates reports about the cluster configuration so that you can find existing problems.

When you run Cluster Analysis, you specify the types of tests you want to run. Domino then creates reports about the areas you selected and places the reports in the Cluster Analysis database (CLUSTA4.NSF) or in a database that you specify. You can generate reports about servers, databases, and the Server Web Navigator. This table describes the types of tests you can run.

<i>Type of analysis</i>	<i>Test</i>	<i>Description</i>
Server	Number of cluster members	Reports the number of servers in the cluster.
	Consistent domain membership	Checks that all servers are members of the same domain, which is required for a cluster to operate correctly.
	Consistent protocols	Checks that servers are running the same protocols. Cluster members cannot communicate if they are running different protocols. (TCP/IP is the only protocol formally supported for cluster communication.)
	Required server tasks	Checks the ServerTasks setting in the NOTES.INI file for required server tasks. There are no required server tasks in NOTES.INI for clustering in Domino Release 6.

continued

<i>Type of analysis</i>	<i>Test</i>	<i>Description</i>
Databases	Consistent ACLs	Compares the ACLs of replicas throughout the cluster to be sure the ACLs are consistent. If they are not, users could fail over to replicas that they can't access or replicas that give them different rights to view and alter database information.
	Disabled Replication (Called "Replication Enabled" in the test report)	Checks to see if cluster replication is enabled for the databases on the server. If users fail over to a database that does not have cluster replication enabled, they may see different information than in the original database.
	Consistent replication formulas	Checks for inconsistent replication formulas among replicas that share the same path. Replicas with the same path should have the same replication formulas.
	Replicas exist within cluster	Checks to see if databases on the current server have replicas in the cluster. Returns "failed" if no replica exists. (Not all databases require replicas.)
Web Navigator	Consistent Server Web Navigators	Checks to see if the Web databases (WEB.NSF) on cluster members are replicas of each other. If they aren't, the Web databases will not fail over to each other.

Running Cluster Analysis

When you run Cluster Analysis, you make decisions about the types of reports you want to generate and the type of detail you want in the reports. While the analysis runs, status messages appear at the bottom of the screen. For large, busy clusters, Cluster Analysis may require several hours or more to complete. If you don't have a dedicated server for analysis tests, you may want to run database analysis during off-peak hours or only for problems that you think exist.

Note Cluster Analysis is not available in the Web Administrator.

1. In the Server pane of the Domino Administrator, choose the server on which you want to run the analysis.
2. Click the Server - Analysis tab.

3. In the Tools pane, expand Analyze, and click Cluster.
4. (Optional) To write the results of the analysis tests to a database other than the Cluster Analysis database, click “Results Database” and specify the server, title, and file name of the database. Then click OK. If the database does not already exist, Domino creates it.
5. (Optional) If a Cluster Analysis database already exists and you want to append the new reports to this database, select “Append to this database.” Otherwise, the cluster analysis overwrites the existing database.
6. Select the types of reports you want Domino to generate: Server, Databases, or Web Navigator.
7. If you selected Databases as a Report type in Step 6, select the type of database tests you want to run: Consistent ACLs, Disabled Replication, Consistent replication formulas, and/or Replicas exist within cluster.
8. Click OK to run the analysis and to open the Results Database.

Viewing the results of a Cluster Analysis

1. Open the Cluster Analysis database if it is not already open.
2. Click one of the following views:
 - By Cluster
 - By Date
 - By Test
3. Open a Cluster Analysis Results document.

Managing user access to databases

Users should have the same access rights in all replicas of a database in the cluster. Otherwise, if users fail over to a replica for which they have fewer access rights, they may be denied access, or they may not see the same documents or be able to perform the same functions as in the original database they were using.

Setting up database ACLs in a cluster

One way to be sure that the access control lists are synchronized across all replicas is to use the following procedure for each database:

1. In the Server pane of the Domino Administrator or the Web Administrator, expand All Servers or expand Clusters.
2. Select a server that contains a replica of the database you want and has Manager access in the ACLs of the other replicas in the cluster.

3. Click the Files tab.
4. Do one of the following:
 - In the Task pane of the Domino Administrator or the Web Administrator, select the folder or view that contains the database you want.
 - In the Task pane of the Domino Administrator only, expand Cluster Directory (6), and then select the view you want.
5. In the Results pane, select the database you want.
6. In the Tools pane, expand Database, and then click Manage ACL.
7. Click the Advanced icon.
8. Choose “Enforce a consistent Access Control List across all replicas of this database,” and then click OK.

This setting ensures that ACLs are consistent across replicas and also enforces the ACL when replicas are accessed locally on either a server or a client.

Another way to keep ACLs consistent across replicas is to give all servers in a cluster Manager access to all databases in the cluster. This ensures that every server can update the ACL of every database.

To give the cluster servers Manager access to all databases, you can create a Group document in the Domino Directory that includes all the servers in the cluster. Then add this group to the ACL of each database, select the user type “Server group,” and give the group Manager access.

It is important that cluster servers have adequate access so they can replicate all data from one replica to another. If there are any restrictions in one replica that are not in another replica, some information will not be available to users when failover occurs. Therefore, be sure that servers not only have Manager access, but that they can all replicate the same data without restrictions.

Private folders replicate differently in a cluster than outside a cluster. When outside a cluster, private folders and their contents do not replicate during server-to-server replication but do replicate during client-to-server replication. In a cluster, however, private folders replicate from server to server so that users are able to access their private folders if they fail over to a different replica. To ensure that private folders replicate between servers in a cluster, be sure to set the user type of the servers in the ACL to “Server” or “Server group.”

Controlling other settings that restrict database access

There are methods of restricting database access in addition to the ACL. It is important that these settings are consistent across databases so that complete replication occurs and failover works transparently to the user. These methods include the following:

- Server restrictions, such as Allow and Deny lists in the Domino Directory
- Access lists in database and directory links
- Readers lists, such as those in documents, views, and folders

For example, if the servers in the cluster contain database or directory links that include access lists, be sure that the cluster servers are in the access lists. Otherwise, they will not have access to those databases or directories and will not be able to replicate with those databases, even if they have Manager access in the ACLs.

If a document in a database includes a Readers field, the cluster servers must be listed in the Readers field or the servers will not have access to that document and will not be able to replicate the document. The same is true if a folder or view includes a Readers field. Because Readers fields are often maintained by a database designer rather than a network administrator, network administrators need to communicate with database designers about this issue.

Creating replicas in a cluster

To create replicas in a cluster, you use the same procedure you would use to create replicas on any server. Domino lets you use one procedure to create multiple replicas on multiple servers simultaneously. This makes it easy to populate the clusters with the replicas you want. While creating the replicas, you can select “Show me only cluster members for cluster: *cluster name*” in the Create Replica dialog box to see which servers are in each cluster. You can then create replicas on every cluster server you want.

For more information about creating replicas, see the book *Administering the Domino System*.

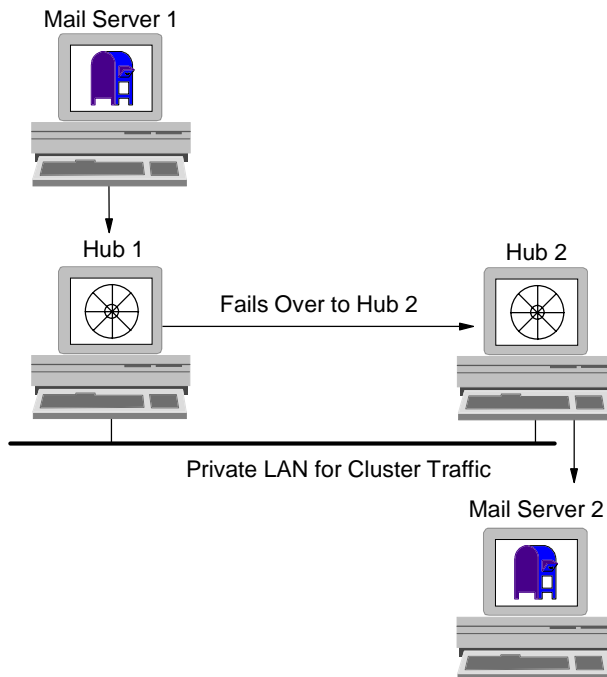
Setting up mail in a cluster

Mail routing fails over if the recipient's mail server is not available when the router attempts to deliver the mail. As long as the mail server is in a cluster, the router delivers the mail to another cluster server that contains a replica of the recipient's mail database. Thus, the recipient continues to receive mail.

By default, mail routing fails over for the last hop of the delivery route only. That is, it fails over for the hop to the user's mail server, if the user's mail server is in a cluster. You can configure mail routing to fail over on any hop in the delivery route, or you can disable mail routing failover altogether.

When you configure mail routing to fail over on any hop, if any server along the route is unavailable but is in a cluster, mail routing fails over to a cluster server and that server continues routing the message. Enabling mail routing failover to occur on any hop is particularly helpful if you are using a hub server to route mail. If that hub server is unavailable but is in a cluster, the router delivers the mail to another hub server in the cluster. That hub server continues to send the message toward its destination.

In the following figure, Mail Server 1 is sending a message to Mail Server 2. It attempts to route through Hub 1, which is the expected route to Mail Server 2. Hub 1 is unavailable, however. Because Hub 1 is in a cluster, Domino routes the message to Hub 2, which routes the message to Mail Server 2.



Changing the mail routing failover setting

To change the default mail routing failover setting, make the following change in the Configuration Settings document for every server in the cluster and every server in the domain that can route mail.

1. From the Domino Administrator or the Web Administrator, click the Configuration tab.
2. In the Task pane, expand Messaging.
3. Click Configurations.
4. Do one of the following:
 - From the Domino Administrator, select the Configuration document for the server or server group you want, and click Edit Configuration.
 - From the Web Administrator, open the Configuration document for the server or server group you want, and click Edit Server Configuration.

If you do not have a Configuration document for the server or server group you want, create one by clicking Add Configuration.

5. Click the Router/SMTP - Advanced - Controls tab.
6. In the Cluster failover field, choose one of the following:
 - Disabled
 - Enabled for last hop only (the default)
 - Enabled for all transfers in this domain
7. Save and close the Configuration document.

Note This setting affects delivery to a client but does not affect sending a message from a client when the mail server is unavailable. If a user sends a message when the mail server is unavailable, the delivery fails over to another server in the cluster, and the router on that server sends the message.

Using shared mail in a cluster

When you send a message to multiple users on a server that is using shared mail, Domino places the header of the message into the recipients' mail databases and places the body of the message into the shared mail database on the server. Before replicating the message to another cluster server, the Cluster Replicator reassembles the message. The server that receives the reassembled message determines if shared mail is being used on that server and if the recipient's mail database is set to always use shared mail. If shared mail is being used, the server deposits the header of the message into the replica of the recipient's mail database and deposits the body of the message

into the shared mail database. If shared mail is not being used, the server deposits the entire message into the replica of the recipient's mail database.

To set up shared mail in a cluster and have replicated messages stored in the shared mail database, you use the same procedure you use for setting up shared mail with replicas that are not in a cluster. This procedure includes the Load Object Set - Always command. You do this on every server that uses shared mail in the cluster.

For more information about setting up shared mail for replica mail databases, see the book *Administering the Domino System*.

Creating mail database replicas in a cluster during user registration

You can use the Domino Administrator or the Web Administrator to create mail database replicas during user registration.

From the Domino Administrator

1. Click the People & Groups tab.
2. In the Tools pane, expand People, and then click Register.
3. In the "Choose a Certifier" dialog box, choose a certifier and click OK.
4. In the Register Person — New Entry dialog box, select Advanced, and then click the Mail tab.
5. In the Mail system field, choose Lotus Notes.
6. Click Mail Server, and choose a cluster server as the Mail server.
7. Click Mail File Replicas.
8. Select "Create mail database replica(s)." A list is displayed of servers in the same cluster as the Mail server.
9. Do one of the following:
 - To create a replica of the mail database on all of the cluster servers, skip this step.
 - To change the list of servers to receive a replica, use the Remove Server(s) button and the Add Server(s) button.
10. (Optional) Select "Create mail replica(s) in background."
11. Click OK, and then complete any other fields you want on the Mail tab.
12. (Optional) If you want to set up the user for roaming in a cluster, follow the procedure in the topic "Setting up roaming in a cluster."
13. Complete the rest of the user registration the way you normally would.

From the Web Administrator

1. Click the People & Groups tab.
2. In the Tools pane, expand People, and then click Register.
3. Choose a CA certifier and, optionally, an explicit policy. Then click OK.
4. In the Register Person dialog box, select Advanced, and then click the Mail tab.
5. In the Mail system field, choose Lotus Notes.
6. In the Mail server field, choose a cluster server as the Mail server.
7. In the Mail template field, choose Mail (6).
8. Complete any other fields you want on the Mail tab, and then click the Replica tab.
9. Select "Create replica(s) of." A list is displayed of servers in the same cluster as the Mail server.
10. Do one of the following:
 - To create a replica of the mail database on all of the cluster servers, skip this step.
 - To change the list of servers to receive a replica, use the Add button and the Remove button.
11. Complete the rest of the user registration the way you normally would.

Using standard replication in a cluster

You should use standard replication in a cluster to be sure that all databases are up-to-date. This is particularly important when you restart a cluster server after it stopped running.

You can:

- Use scheduled replication in a cluster
- Replicate with all servers in a cluster at once

Scheduled replication in a cluster

There are many reasons to use scheduled replication in a cluster:

- To replicate changes that may be lost when a server in a cluster fails
Because cluster replication events are stored in memory only, these events can be lost when a server fails before replicating the events. No data is lost, but the cluster replication events that were pending won't be replicated until you run standard Domino replication.

- To replicate databases for which you have disabled cluster replication
You may have databases that you want to replicate but not every time they are updated. You can disable cluster replication for these databases. To see whether cluster replication is disabled for a database, open the Cluster Database Directory. Databases with the letter “X” in the left column have cluster replication disabled. You can also check this by looking in the Cluster Replication field in the document for each database in the Cluster Database Directory.

For more information about disabling cluster replication and viewing information in the Cluster Database Directory, see the chapter “Managing and Monitoring a Cluster.”

- To replicate based on selective replication formulas
The Cluster Replicator leaves the processing of replication formulas to the standard replicator. Before using replication formulas in a cluster, you should be aware of how this affects cluster replication.

For more information about selective replication in a cluster, see the chapter “How Domino Clustering Works.”

- To replicate replicas that are on the same server

The Cluster Replicator pushes changes to other servers that contain replicas but does not update other replicas on its own server.

Note If there are multiple replicas on a server, the Cluster Manager uses failover by path to select the replica for a user to open during failover. If you put multiple replicas on a server, be sure that all replicas in the cluster that have the same path use the same selective replication formulas. Otherwise, the replica to which users fail over may contain different data than they expect.

You should run standard replication on a regular basis. The number of times per day you run standard replication depends on how important it is for you to keep all replicas synchronized. In most cases, once or twice per day is sufficient. If it is absolutely critical to keep data synchronized at all times, you may want to replicate every hour or two.

In addition, you should replicate whenever you start the server to be sure that all databases are up-to-date. You can create a Program document in the Domino Directory to accomplish this.

Replicating with all servers in a cluster

An additional benefit of creating a cluster is that it can simplify your replication topology and increase the reliability and performance of replication. This is because you can set up a server outside the cluster to replicate with all the servers inside the cluster. You do this by creating a single Connection document that specifies the server outside the cluster as the source of replication and the cluster name as the destination of replication. Then, every database on the source server that has a replica in the cluster replicates. If the cluster contains more than one replica of a database, the source server replicates with only one of those replicas, and cluster replication propagates the changes to the other replicas in the cluster.

Replication with a cluster is more reliable than replication with an individual server because Domino replicates with any server in the cluster that contains a replica of the database it is processing. Therefore, if a server in the cluster is unavailable, replication can still proceed if another replica exists in the cluster. Replication with a cluster can also improve performance because Domino uses workload balancing when choosing a server with which to replicate.

Note When replicating with a cluster name, select Pull Push, Pull Only, or Push Only as the replication type in the Connection document. Pull Pull replication does not work with a cluster name.

The server that initiates replication must be running Domino Release 4.5 or higher. It does not have to be running Lotus Domino 6.

Note Before replicating with a cluster name for the first time, the source server must access one of the clustered servers directly once. When this occurs, Domino populates the source server's cluster cache with the names of the servers in the cluster. This makes the source server aware of the servers in the cluster.

Issuing the Replicate command with a cluster name

From the server console, send the following Replicate commands to replicate databases on a local server with databases in the specified cluster.

<i>Purpose</i>	<i>Command</i>	<i>Explanation of variables</i>
To replicate all the databases that the local server has in common with servers in a specific cluster	replicate <i>cluster_name</i>	<i>cluster_name</i> is the name of the cluster
To replicate a specific database only	replicate <i>cluster_name</i> <i>filename</i>	<i>filename</i> is the file name of a database
To replicate with all the databases in a specific directory	replicate <i>cluster_name</i> <i>local_directory</i>	<i>local_directory</i> is the name of a directory that contains databases

Specifying a cluster in a Connection document

You can create a Connection document to replicate with a cluster. Type the appropriate information in these fields:

<i>Tab</i>	<i>Field name</i>	<i>Information you enter</i>
Basics	Source server	Type the name of a server. The server cannot be a member of the cluster.
	Destination server	Type the cluster name.
Replication/ Routing	Replication task	Choose Enabled.
	Replication Type	Select Pull Push, Pull Only, or Push Only. Pull Pull does not work with a cluster name.
	Files/Directory Paths to Replicate	Leave blank to select all databases; type one or more file names or a directory name to specify databases.

Replicating with a cluster from a Notes client

Replicating with a cluster name is useful when you don't know the location of a database within the cluster.

Note A client must access a server in the cluster directly once before replicating with a cluster name. Doing so makes the client aware of the cluster by adding the names of the servers in the cluster to the CLUSTER.NCF file on the client.

1. Open the database you want, or right-click its bookmark.
2. Do one of the following:
 - If you opened the database in Step 1, choose File - Replication - Replicate.
 - If you right-clicked the bookmark in Step 1, choose Replication - Replicate.
3. If a box appears asking whether to “Replicate via background replicator” or “Replicate with options,” choose “Replicate with options” and click OK.
4. In the “with” field, type the cluster name.
5. Make any other changes you want, and click OK.

Enabling the display of cluster replication status messages

To tell the Cluster Replicator to display status messages on the server console, enter the following in the NOTES.INI file:

```
RTR_logging=n
```

where the possible values for *n* are 1, which enables the display of cluster replication status messages, and 0, which disables the display of cluster replication status messages.

Obeying database size quotas during cluster replication

To make sure that replicas will be identical when users fail over, cluster replication ignores database size quotas by default. This makes cluster replication consistent with the router, which also ignores quotas by default. If you want cluster replication to obey database size quotas, follow these steps.

From the Domino Administrator

1. Click the Configuration tab.
2. In the Task pane, expand Server, and then click Configurations.
3. Do one of the following:
 - If a Configuration Settings document already exists for the server you want, select that document, and then click Edit Configuration.
 - If a Configuration Settings document does not already exist for the server you want, click Add Configuration, and add the name of the server in the “Group or Server name” field on the Basics tab.

4. Click the NOTES.INI Settings tab.
5. Click Set/Modify Parameters.
6. In the Item field, select or enter CLREPL_OBEYS_QUOTAS.
7. In the Value field, enter 1.
8. Click Add, and then click OK.
9. Click Save & Close.

Note To ignore database size quotas again, place 0 (zero) in the Value field in step 7, or delete CLREPL_OBEYS_QUOTAS from the Configuration Settings document.

From the Web Administrator

1. Click the Configuration tab.
2. In the Task pane, expand Server, and then click Configurations.
3. Do one of the following:
 - If a Configuration Settings document already exists for the server you want, open that document, and then click Edit Server Configuration.
 - If a Configuration Settings document does not already exist for the server you want, click Add Configuration, and add the name of the server in the “Group or Server name” field on the Basics tab.
4. Click the NOTES.INI Settings tab.
5. Click Set/Modify Parameters.
6. In the Available Parameters box, click CLREPL_OBEYS_QUOTAS, and then click Add.
7. In the Value field, enter 1, and then click OK
8. Click Save & Close.

Note To ignore database size quotas again, place 0 (zero) in the Value field in step 7, or delete CLREPL_OBEYS_QUOTAS from the Configuration Settings document.

Using directory assistance in a cluster

A server can use directory assistance to look up information in a directory other than its local primary Domino Directory. The directory can be a remote LDAP directory or a local or remote replica of any of the following types of Domino directories: a secondary Domino Directory, an Extended Directory Catalog, or a primary Domino Directory. To configure directory assistance, you create a directory assistance database. This database contains information about directories on other servers and how to connect to them.

If you try to access the directory assistance database on a cluster server that is unavailable, Domino fails over to a replica of the database, if one is available in the cluster. When you use directory assistance in a cluster, be sure you use replicas of the same directory assistance database on all cluster servers that contain a replica. This is important so that Name Lookups fail over correctly.

When the directory assistance database refers to a Domino directory on a cluster server, you can configure Domino to use one of two failover methods for directory assistance: directory assistance failover or cluster failover. When you use directory assistance failover, Domino fails over to one of the servers that you list as having a replica of the directory. Domino does not take into account the current workload of the servers. When you use cluster failover, Domino checks the workload of the servers and then fails over to the least busy cluster server that contains a replica of the directory.

For more information on directory assistance and configuring directory assistance to fail over, see the book *Administering the Domino System*.

Setting up roaming in a cluster

You can set up a cluster of roaming servers. Make sure that replicas of each user's roaming files are on at least two roaming servers in the cluster. When a user tries to access roaming files on a server that is not available, Domino fails over to a different cluster server that has replicas of the user's roaming files.

When you use the Domino Administrator to register new users, you can set them up during registration for roaming in a cluster. You can also upgrade existing users for roaming in a cluster.

When you use the Web Administrator to register new users, you can set them up for roaming but not for roaming in a cluster. To set them up for roaming in a cluster, register them as roaming users, and then manually create replicas of the roaming files on as many cluster servers as you want.

You cannot use the Web Administrator to upgrade existing users for roaming.

Setting up new users for roaming in a cluster

1. From the Domino Administrator, click the People & Groups tab.
2. In the Tools pane, expand People, and then click Register.
3. In the “Choose a Certifier” dialog box, choose a certifier and click OK.
4. In the “Register Person — New Entry” dialog box, select “Enable roaming for this person” and select “Create a Notes ID for this person.”
5. Select Advanced, and then click the Roaming tab.
6. Select the Roaming Server button to make it active, and then click the Roaming Server button.
7. Select a cluster server as the Roaming Server, and then click OK.
8. Click the Roaming Replicas button.
9. Select “Create roaming file replica(s).” A list is displayed of the servers in the same cluster as the Roaming Server.
10. Do one of the following:
 - To create replicas of the roaming files on all of the cluster servers, skip this step.
 - To change the list of servers to receive replicas, use the Remove Server(s) button and the Add Server(s) button.
11. (Optional) Select “Create replica(s) in background.”
12. Click OK, and then complete any other fields you want on the Roaming tab.
13. (Optional) If you want to create replicas of the user’s mail database in a cluster, follow the procedure in the topic “Creating mail database replicas in a cluster during user registration.”
14. Complete the rest of the user registration the way you normally would.

Setting up existing users for roaming in a cluster

Note You cannot use the Web Administrator to set up existing users for roaming.

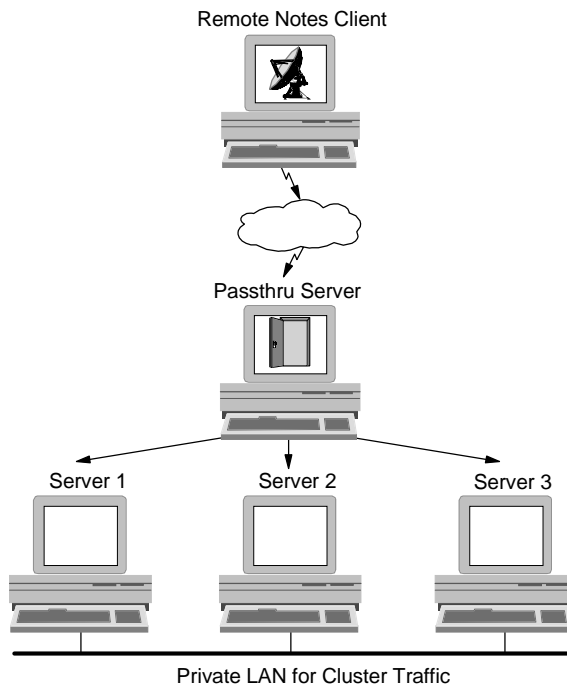
1. From the Domino Administrator, click the People & Groups tab.
2. In the Task pane, click People.
3. In the Results pane, select the people you want to give the ability to roam.
4. In the Tools pane, expand People, and then click Roaming.
5. In the “Where should the users roaming files be stored?” section, select the Roaming Server button to make it active, and then click the Roaming Server button.

6. Select a cluster server as the Roaming server. Domino displays a list of all the servers in the cluster.
7. Select the additional servers you want to receive replicas, and then click OK.
8. Complete the rest of the fields the way you normally would when you upgrade a user for roaming, and then click OK.

Setting up cluster access for mobile users

You can give mobile users the benefits of a cluster by setting them up to access the cluster through a passthru server. The passthru server can redirect the users to a different server in the cluster if the original server is busy or not available. Unless you use a passthru server, users will have to dial each cluster server individually and will not fail over if the server they call is not available.

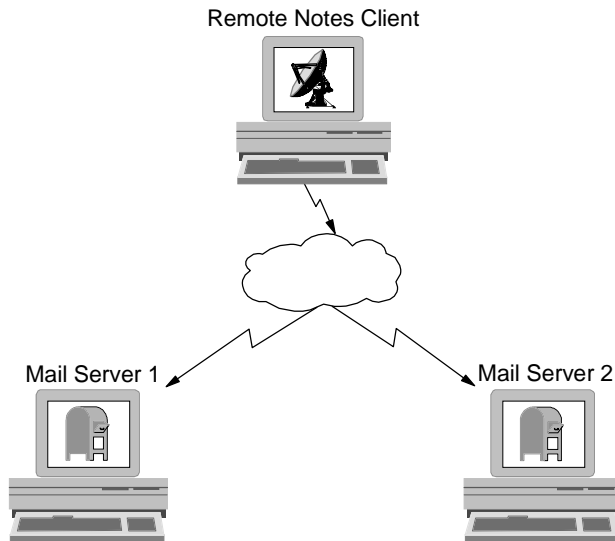
The following figure shows a mobile client calling a passthru server and having access to any server in a cluster. When connecting to the cluster in this way, failover and workload balancing are available to the mobile client.



Because mobile clients can typically call only one server at a time, they cannot take advantage of a cluster of passthru servers. However, you can set up a hunt group of passthru servers to ensure that mobile clients have high availability to passthru servers.

If you do not have access to a passthru server, users can still use their Replicator page to simulate failover if a cluster server is down. To replicate a mail database, for example, set up the Replicator page to call and replicate with the user's mail server, and then call and replicate with the server that contains the user's replica mail database. That way, if either server is available, the user will have access to the mail database. The call to the second server will take very little time if the first call was successful.

The following figure shows a mobile client calling Server 1, which is the user's mail server, and then calling Server 2, which contains a replica of the user's mail database.



Using the Server Web Navigator in a cluster

If you plan to use the Server Web Navigator in a cluster, be sure that the Web databases (WEB.NSF) on the cluster servers are replicas. Since the Web database is generated automatically when you start the Server Web Navigator, it is a common error for the clustered servers to contain Web databases that are not replicas of each other. Therefore, these databases do not fail over to each other.

To be sure that the Web databases are replicas, start the Server Web Navigator on only one of the cluster servers. Then replicate the Web database to the other cluster servers before starting the Server Web Navigator on those servers.

To see if the Web databases are replicas, run Cluster Analysis, as described in the topic “Using Cluster Analysis to check the cluster configuration.”

Setting up a private LAN for a cluster

Setting up a private LAN for your cluster separates the server probes and cluster replication from the rest of your network traffic. This prevents the cluster traffic from slowing down your primary network.

Follow this procedure to set up a private network for your cluster.

1. Install an additional network interface card in each server in the cluster. These cards should be connected through a private hub or switch.
2. Configure the LAN to use TCP/IP.
3. Assign a second IP address to each server. In a cluster using two servers, for example, you could use the addresses 192.168.64.1 and 192.168.64.2.
4. Assign host names for the new IP addresses, and place entries in the HOSTS file or the DNS. In a cluster using two servers, for example, you could assign the host names Acme_clu.acme.com and Acme2_clu.acme.com. The entries in the HOSTS file might then be:
192.168.64.1 Acme_clu.acme.com
192.168.64.2 Acme2_clu.acme.com
5. Verify that the private LAN is operational by making sure that each cluster server can ping both the IP address and the host name of the other cluster servers. Also, be sure that other servers can still ping the cluster servers at their regular LAN addresses.
6. On each server in the cluster, add a new port for the private LAN, and edit the Server document to enable the port.

For example, add a port named CLUSTER, and then add the following information to the Port - Notes Network Ports tab in the Server document to enable the port.

<i>Field name</i>	<i>Example</i>
Port	CLUSTER
Protocol	TCP
Notes Network	Cluster Network
Net Address	Acme_clu.acme.com
Enabled	ENABLED

For more information about adding and enabling a new port, see *Administering the Domino System*.

- Assign each port an IP address from the corresponding subnets, and place this information in the NOTES.INI file in the following form:

```
PORT1_TcpIPAddress=0,a.b.c.d:1352
```

```
PORT2_TcpIPAddress=0,e.f.g.h:1352
```

where PORT1 and PORT2 are the port names and *a.b.c.d* and *e.f.g.h* are the IP addresses for the ports.

If you have ports named TCPIP and CLUSTER, for example, these lines might be:

```
TCPIP_TcpIPAddress=0,192.114.32.5:1352
```

```
CLUSTER_TcpIPAddress=0,192.168.64.1:1352
```

- Do one of the following:
 - Reorder the ports so that the cluster port is first.
This ensures that all cluster traffic uses this port for cluster traffic. Be sure that all other traffic is assigned to use other ports.
 - Add the following line to the NOTES.INI file:
`Server_Cluster_Default_Port=Cluster Port`
where *Cluster Port* is the port you created for the cluster. In this example, this line would be:
`Server_Cluster_Default_Port=CLUSTER`
This ensures that all cluster traffic uses this port for cluster communications no matter what order the ports are in.
Note There is a disadvantage to using the `Server_Cluster_Default_Port` setting to assign a port to the private LAN for cluster traffic. If a cluster server encounters a problem connecting over this port, it will not try another port. Therefore, the server will not be able to communicate or replicate with other cluster

servers. You will have to resolve the network problem or remove this setting from the NOTES.INI file before the server will be able to communicate with the cluster again.

For information about reordering network ports on a server, see *Administering the Domino System*.

9. Restart the server.

Making sure Domino is using the private LAN

To verify that Domino is using the private LAN, you can do the following:

1. At the server console, type **show cluster**.
2. In the results, look for the line “Server cluster default port: *Cluster Port*” where *Cluster Port* is one of the following:
 - If you used the `Server_Cluster_Default_Port` setting to specify a port for the private LAN, *Cluster Port* should be the name of the port you specified, such as CLUSTER.
 - If you did not use the `Server_Cluster_Default_Port` setting to specify a port for the private LAN, *Cluster Port* should be * (asterisk), meaning that the private LAN can use any port.

Making sure that cluster replication is using the private LAN

It is also a good idea to check the cluster to be sure that cluster replication is using the private LAN. To do this, you can look at some of the cluster statistics.

1. From the Domino Administrator or the Web Administrator, click the Servers - Statistics tab.
2. In the list of statistics, expand NET.
3. Expand the port name you gave to the cluster.
4. Look for the following statistics for the private LAN port:
 - BytesReceived
 - BytesSent
5. Expand Replica, and then expand Cluster.
6. Expand SessionBytes, and look for the following replication statistics for the cluster:
 - In
 - Out
7. Compare the `NET.portname.BytesReceived` value with the `Replica.Cluster.SessionBytes.In` value.

These values should be fairly close to each other, although they will not be the same.

8. Compare the NET.*portname*.BytesSent value with the Replica.Cluster.SessionBytes.Out value.

These values should also be close to each other. They won't match exactly because the private network is used for more than just cluster replication.

NOTES.INI settings for setting up a cluster

The following table describes the NOTES.INI settings that pertain to setting up a cluster.

For more information on these settings, see the book *Administering the Domino System*.

<i>Setting</i>	<i>Description</i>
RTR_Logging	Tells the Cluster Replicator to display status messages on the server console
CLREPL_Obeys_Quotas	Tells the Cluster Replicator to obey database size quotas
Server_Cluster_Default_Port	Specifies the port for cluster traffic

Chapter 5

Managing and Monitoring a Cluster

After you've set up a cluster, you can do the following to ensure that it is running smoothly and efficiently:

- Monitor failover, workload balancing, and cluster replication to see how efficiently the cluster is running
- Adjust the failover, workload balancing, and cluster replication settings to improve efficiency or to perform maintenance, such as upgrading a server
- Add or remove servers from a cluster or move servers to different clusters

Monitoring a cluster

Domino provides several ways to find out what is happening in a cluster and make adjustments to keep the cluster running smoothly and efficiently, so that no server is overloaded. When running as part of a cluster, a Domino server constantly monitors its workload, the workload of the other servers in the cluster, and the availability of databases throughout the cluster. In addition, Domino monitors statistics and events that are relevant to a cluster.

There are many ways to view this information. For example, you can view it from the server console or in the log file or in the Statistics pane in the Domino Administrator. In addition, you can collect statistic reports in the Monitoring Results database and then use the Domino Administrator to look at the statistic reports.

Some of the ways to monitor a cluster are:

- Displaying a list of cluster members and their availability
- Enabling statistic reporting in the Monitoring Results database
- Viewing Cluster Manager events and statistics
- Viewing cluster replication events and statistics
- Using Tell commands to display cluster replication information
- Monitoring all the servers in a cluster at the same time

Displaying a list of cluster members and their availability

To display a list of cluster members and their availability, do one of the following.

From the Domino Administrator or the Web Administrator

You can see the availability of each cluster member individually by doing the following:

1. In the Server pane, expand All Servers or expand Clusters.
2. Select the server you want.
3. Click the Server - Statistics tab.
4. In the statistics list, expand Server to see the availability index and the availability threshold of the server you selected.
5. Expand Cluster - Member, and then expand the name of each cluster server to see the availability of the servers in the cluster.

From the Domino Administrator only

You can also see the availability of each cluster member individually by doing the following:

1. In the Server pane, expand All Servers or expand Clusters.
2. Right-click the server you want, and choose Server Properties.
3. In the Server Properties box, click the Cluster tab.

The Cluster tab in the Server Properties box shows you the name of the cluster, the availability threshold and the availability index of the server you selected, the port that is used for cluster traffic, how frequently cluster probes are sent, how many cluster probes the Cluster Manager has sent since the last time the server was started, and the names and availability indexes of all the servers in the cluster.

From the server console

Send the following Domino command from the server console:

```
show cluster
```

This command displays the name of the cluster, several statistics for the current server, and the names and availability indexes of all the servers in the cluster.

Enabling statistic reporting in the Monitoring Results database

One way to monitor what is happening in a cluster is to collect statistics in the Monitoring Results database. You can then view this information from the Domino Administrator or in the Monitoring Results database itself.

In order to enable statistic reporting in the Monitoring Results database, you do the following:

- Enable statistic report generation
- Create a server statistic collection
- Start the Statistic Collector task, if it is not already started

Enabling statistic report generation

Note You cannot use the Web Administrator to enable statistic report generation.

1. From the Domino Administrator, choose File - Preferences - Administration Preferences.
2. Click the Statistics tab.
3. Select “Generate statistic reports while monitoring or charting statistics.”
4. Change any of the other setting you want, and then click OK.

Creating a server statistic collection

A server statistic collection specifies which statistics to collect in the Monitoring Results database.

1. From the Domino Administrator or the Web Administrator, click the Configuration tab.
2. In the Task pane, expand Monitoring Configuration, and then click Server Statistic Collection.
3. Click New Statistics Collection.
4. In the “Collecting server” section of the Basics tab, designate the server that will collect the statistics.
5. In the “Collect from” section of the Basics tab, choose one of the following:
 - All servers in this domain
 - All servers that are not explicitly listed to be collected

- From the following servers.

If you choose “From the following servers,” do the following:

In the Server(s) field, enter the names of the servers from which you want to collect statistics.

To collect statistics from all the servers in a cluster, choose “From the following servers,” and then enter the names of the cluster servers in the Server(s) field.

6. Click the Options tab, and select “Log statistics to a database.”
7. (Optional) Do any of the following:
 - In the “Database to receive reports” field, enter the name of a database to store reports. By default, this is STATREP.NSF (the Monitoring Results database).
 - In the Collection report interval field, enter the number of minutes between reports. The minimum is 15.
 - In the Collection alarm interval field, enter the number of minutes between alarms. The minimum is 15.
 - In the Statistic Filters field, select the types of statistics you do NOT want to collect. By default, the server statistic collection includes all the types of statistics.

To collect cluster replication statistics, do NOT select REPLICATION in this field. The Cluster Manager statistics are always collected in statistic collections.
8. Click Save & Close.

Starting the Statistic Collector task

Use one of the following procedures to start the Statistic Collector task.

Starting the Statistic Collector task automatically at server startup

To start the Statistic Collector task automatically each time you start the server, append the COLLECT parameter to the ServerTasks setting in the NOTES.INI file. For example,

```
ServerTasks=ROUTER,REPLICATION,ADMINP,HTTP,ICM,COLLECT
```

Starting the Statistic Collector task for the current session only

To run the Statistic Collector task for the current session only, do one of the following.

From the Domino Administrator or the Web Administrator

1. In the Server pane, expand All Servers or expand Clusters.
2. Select the server you want.

3. Click the Server - Status tab.
4. In the Task pane, do one of the following:
 - From the Domino Administrator, click Server Tasks.
 - From the Web Administrator, click All Server Tasks.
5. In the Tools pane, expand Task, and then click Start.
6. Select Statistic Collector, and then click Start Task.
7. Click Done.

From the server console

Send the following Domino command from the server console:

```
load collect
```

Viewing Cluster Manager events and statistics

After distributing databases and mail servers to balance the workload, you should track the cluster events and statistics to be sure that the workload is acceptable and that failover occurs as expected. If the statistics show a problem, you may have to make some adjustments.

For information about monitoring cluster replication, see the topic “Viewing cluster replication events and statistics.”

Viewing failover and workload balancing events

Whenever failover occurs, Domino logs the event in the Domino server log file. A failover event may look like this:

```
08/23/2002 11:08:48 AM Failing over from  
Sales/Acme!!Customer.nsf for replica ID 852560C9:007232D,  
directing open to Sales2/Acme
```

When Domino fails over to balance the workload, the event may look like this in the Domino server log file:

```
08/23/2002 11:08:48 AM Load balancing off of  
Sales/Acme!!Customer.nsf for replica ID 852560C9:007232D,  
directing open to Sales2/Acme
```

You can view these events in the log file. Do one of the following.

From the Domino Administrator or the Web Administrator

1. In the Server pane, expand All Servers or expand Clusters.
2. Select the server that stores the log file you want to view.
3. Click the Server - Analysis tab.

4. In the Task pane, expand Notes Log.
5. Click Miscellaneous Events.
6. In the Results pane, open the document you want to view.

From the Domino server log file

1. Open the Domino server log file (LOG.NSF).
2. Open the Miscellaneous Events view.
3. Open the Notes Log document you want to view.

You can also run Log Analysis to gather all of the failover and workload balancing events into a database.

For more information about the Domino server log file and Log Analysis, see the book *Administering the Domino System*.

Viewing Cluster Manager statistics

You can view Cluster Manager statistic reports, or you can view a list of Cluster Manager statistics.

Viewing Cluster Manager statistic reports

Prior to using this procedure, you must enable statistic reporting in the Monitoring Results database.

For information about enabling statistic reporting, see the topic “Enabling statistic reporting in the Monitoring Results database.”

1. In the Server pane of the Domino Administrator or the Web Administrator, expand All Servers or expand Clusters.
2. Select the server you want.
3. Click the Server - Analysis tab.
4. In the Task pane, expand Monitoring Results, and then expand Statistics Reports.
5. Click Clusters.
6. In the Results pane, open the document you want, and look in the “Server cluster statistics” section of the document.

Note If you prefer, you can view these reports directly in the Monitoring Results database (STATREP.NSF). Open the database, expand Statistics Reports, and then click Clusters.

Viewing a list of Cluster Manager statistics

You can view a list of Cluster Manager statistics from the Domino Administrator, the Web Administrator, or the server console.

From the Domino Administrator or the Web Administrator

1. In the Server pane, expand All Servers or expand Clusters.
2. Select the server you want.
3. Click the Server - Statistics tab.
4. In the statistics list, expand Server, and then expand Cluster.

Note To see the availability index, the availability threshold, and the expansion factor of the current server, look in the Server section of the statistics, not the Server - Cluster section.

From the server console

Send the following Domino command from the server console:

```
show stat server.cluster*
```

The Cluster Manager statistics begin with “Server.Cluster.” They give you information about failover, workload balancing, and the state of the servers in the cluster. Among other things, the statistics tell you how often the Cluster Manager attempted failover and workload balancing, and how many of these attempts were successful.

Note To see the availability index, the availability threshold, and the expansion factor of the current server, send the Domino command **show stat server** from the server console.

For an explanation of all the cluster statistics, see the appendix “Cluster Statistics.”

Viewing cluster replication events and statistics

You can view cluster replication events and statistics from the Domino Administrator, from the Web Administrator, or in the server log file.

You can do the following:

- View cluster replication events
- View cluster replication statistic reports
- View a list of cluster replication statistics
- Use cluster replication statistics to find cluster replication backlogs

Viewing cluster replication events

The Cluster Replicator generates Notes Log Replication documents once an hour and records information about all replication performed during that hour. Each Notes Log Replication document provides information about database replication and any errors that currently exist. An error indicates that an item couldn't be replicated. This is often due to the destination server being unavailable. When an error occurs, the Cluster Replicator retries the replication periodically until it is successful. After the replication is successful, the error is no longer recorded. The Cluster Replicator generates one Notes Log Replication document for each server with which it replicates.

You can view cluster replication events that the Cluster Replicator generates. Do one of the following.

From the Domino Administrator or the Web Administrator

1. In the Server pane, expand All Servers or expand Clusters.
2. Select the server that stores the log file you want to view.
3. Click the Server - Analysis tab.
4. In the Task pane, expand Notes Log, and then click Replication Events.
5. In the Results pane, open the replication document you want to view.

From the Domino server log file

1. Open the Domino server log file.
2. Open the Replication Events view.
3. Open the Notes Log Replication document you want to view.

Example of viewing cluster replication events

The following example shows two databases successfully replicated by the Cluster Replicator from the server Sales/Acme to the server Sales2/Acme. A third database, CUSTOMER.NSF, was modified on Sales/Acme but was not replicated to Sales2/Acme because the server became unavailable.

Sales/Acme
09/04 03:50 PM - 09/04 03:50 PM

Remote Server: Sales2/Acme
Initiated By: CLUSTER REPLICATOR ON SALES/ACME
Elapsed Time: 0 minutes

Events

Unable to replicate from customer.nsf to Sales2/Acme customer.nsf:
Remote system no longer responding.

Database	Access	Added	Deleted	Updated	KBrec	KBsent	From
Sales2 cldbdir.nsf	Manager	34	0	1	3	15	cldbdir.nsf
Sales2 names.nsf	Manager	0	2	2	1	13	names.nsf

You can also run Log Analysis to gather all of the replication events into a database.

For more information about the Domino server log file and Log Analysis, see the book *Administering the Domino System*.

Viewing cluster replication statistics

You can view cluster replication statistic reports, or you can view a list of cluster replication statistics.

Viewing cluster replication statistic reports

You can use the Domino Administrator or the Web Administrator to view replication statistic reports in the Monitoring Results database. In addition, you can use the Web Administrator to monitor several important cluster replication statistics.

From the Domino Administrator or the Web Administrator

Prior to using this procedure, you must enable statistic reporting in the Monitoring Results database.

For information about enabling statistic reporting, see the topic “Enabling statistic reporting in the Monitoring Results database.”

1. In the Server pane of the Domino Administrator or the Web Administrator, expand All Servers or expand Clusters.

2. Select the server you want.
3. Click the Server - Analysis tab.
4. In the Task pane, expand Monitoring Results, and then expand Statistics Reports.
5. Click Clusters.
6. In the Results pane, open the document you want, and then look in the “Replica cluster statistics” section of the document.

Note If you prefer, you can view these reports directly in the Monitoring Results database (STATREP.NSF). Open the database, expand Statistics Reports, and then click Clusters.

From the Web Administrator only

You can use the Web Administrator to monitor a predetermined set of cluster replication statistics. These statistics show cluster replication activity, workload, and status. These statistics refresh automatically every minute.

1. In the Server pane, expand All Servers or expand Clusters.
2. Select the server you want.
3. Click the Replication tab.
4. In the Task pane, click Replication Statistics.

Viewing a list of cluster replication statistics

You can view a list of cluster replication statistics from the Domino Administrator, the Web Administrator, or the server console.

From the Domino Administrator or the Web Administrator

1. In the Server pane, expand All Servers or expand Clusters.
2. Select the server you want.
3. Click the Server - Statistics tab.
4. In the statistics list, expand Replica, and then expand Cluster.

From the server console

Send the following Domino command from the server console:

```
show stat replica.cluster*
```

The cluster replication statistics begin with “Replica.Cluster.” They give you information about cluster replication events, such as the number of documents updated, the number of times the Cluster Replicator retried pending replication, and the number of bytes received during cluster replication.

For an explanation of all the cluster statistics, see the appendix “Cluster Statistics.”

Using cluster replication statistics to find replication backlogs

During peak activity periods, servers may show an especially high frequency of replication events. Replication backlogs may occur if the Cluster Replicator is unable to handle all replication requests.

Examine the `Replica.Cluster.WorkQueueDepth` statistic. This statistic shows the number of modified databases awaiting replication. In addition, examine the `Replica.Cluster.SecondsOnQueue` statistic. This statistic shows how long a database waited to be replicated. If the number of databases waiting to be replicated is consistently much greater than zero, or if the amount of time a database waits to be replicated is consistently longer than you would like, consider enabling additional Cluster Replicators to help decrease the replication backlog. You could also decrease the server workload by removing very active databases from the server or by decreasing the number of users who can access the server.

For more information about enabling multiple Cluster Replicators, see the topic “Using multiple Cluster Replicators.” For an explanation of all the Cluster Replicator statistics, see the appendix “Cluster Statistics.”

Using Tell commands to display cluster replication information

You can use Tell server commands to display additional information about cluster replication. This information helps you determine whether to make any changes to the way you configured cluster replication or to the way you balanced the workload in the cluster. The information is displayed on the server console.

From the Domino Administrator

1. In the Server pane, expand All Servers or expand Clusters.
2. Select the server you want.
3. Click the Server - Status tab.
4. In the Task pane, click Server Tasks.
5. In the Results pane, click Cluster Replicator.
6. In the Tools pane, expand Task, and then click Tell.
7. In the Tell Cluster Replicator dialog box, click the Dump tab.
8. Select one of the following, and then click OK.
 - Everything
 - All known cluster mates
 - All databases that have failed to replicate

Note For information about these commands, see “About the Tell commands for cluster replication,” which follows.

From the Web Administrator

1. In the Server pane, expand All Servers or expand Clusters.
2. Select the server you want.
3. Click the Server - Status tab.
4. In the Task pane, click All Server Tasks.
5. In the Results pane, click Cluster Replicator.
6. In the Tools pane, expand Task, and then click Tell.
7. In the Tell Task dialog box, enter one of the following, and then click OK.
 - tell crepl dump
 - tell crepl dump server
 - tell crepl dump retry

Note For information about these commands, see “About the Tell commands for cluster replication,” which follows.

About the Tell commands for cluster replication

The following table describes the cluster replication Tell commands used in the preceding procedures.

<i>Choice in Domino Administrator</i>	<i>Command in Web Administrator</i>	<i>Information displayed</i>	<i>Use</i>
Everything	tell crepl dump	<ul style="list-style-type: none"> • The number of Cluster Replicators running • The work queue depth • The number of cluster replication retries in progress • The last time cluster replication was unsuccessful • The time of the last retry with each of the other cluster servers • Whether retry is currently in progress • The following information for each replication that still must be retried: the name of the database, the time the next retry is due, and the retry interval • Some statistics and other information about every database on the current server that has a replica in the cluster, as long as the replica does not have cluster replication disabled. Also, information about those replicas. 	<ul style="list-style-type: none"> • To view the status of cluster replication • To see which databases did not replicate and which replicas are not up-to-date • To determine if cluster replication is being delayed due to long queues or to replication that needs to be retried • To determine whether to make any changes to the way you have configured cluster replication. • To determine if you want to change the availability threshold on a server or redistribute replicas in the cluster • To determine if you want to force the Cluster Replicator to retry all failed replications immediately rather than waiting for the next scheduled retry

continued

<i>Choice in Domino Administrator</i>	<i>Command in Web Administrator</i>	<i>Information displayed</i>	<i>Use</i>
All known cluster mates	tell clrepl dump server	<ul style="list-style-type: none"> • The number of Cluster Replicators running • The work queue depth • The number of cluster replication retries in progress • The last time cluster replication was unsuccessful with each of the other cluster servers • The last time replication was retried with each of the other cluster servers 	<ul style="list-style-type: none"> • To view the status of cluster replication • To see how long it has been since the Cluster Replicator attempted to replicate each database that needs to be retried
All databases that have failed to replicate	tell clrepl dump retry	<ul style="list-style-type: none"> • The number of Cluster Replicators running • The work queue depth • The number of cluster replication retries in progress • The names of the databases that are waiting to be replicated • The time the next retry is due for each database • The retry interval for each database 	<ul style="list-style-type: none"> • To view the status of cluster replication • To see which databases did not replicate • To see if particular servers are having problems so you can attend to those servers • To determine if you want the Cluster Replicator to retry all failed replications immediately rather than waiting for the next scheduled retry

Viewing the cluster replication information

The cluster replication information is automatically displayed on the server console. You can view the information there. In addition, if you've set up your system to write server console output to a file, you can view this information in the file. You can also view this information from the Domino Administrator or the Web Administrator and copy the information from there to another program.

The following procedure tells how to copy the output information from the Domino Administrator or the Web Administrator to another program.

1. From the Domino Administrator or the Web Administrator, click the Server - Status tab.
2. Do one of the following:
 - In the Task pane of the Domino Administrator, click Server Console. Be sure that the console is not live.
 - In the Task pane of the Web Administrator, click Quick Console.
3. In the Results pane, select the text you want to copy.
4. Choose Edit - Copy.
5. Paste the copied text into the program you want.

Monitoring all the servers in a cluster at the same time

You can use the Domino server monitor to monitor all the servers in a cluster at once. You can decide what information you want to monitor and how to display that information. You can monitor the cluster while you monitor other Domino servers. To tell Domino which servers to monitor and what information to monitor on each server, you create or customize a monitoring profile.

When you start the server monitor, the Statistic Collector task starts, if it is not already running.

Note The Domino server monitor and monitoring profiles are not available in the Web Administrator.

To start the server monitor manually

1. From the Domino Administrator, click the Server - Monitoring tab.
2. In the “Monitoring profiles” field, select the profile for the cluster you want to monitor. By default, Domino creates a profile for each cluster in the domains you are monitoring.
3. Click the Start button.

Once pressed, the Start button becomes the Stop button.

To start the server monitor automatically when the server starts

1. Choose File - Preferences - Administration Preferences.
2. Click Monitoring.
3. Select “Automatically monitor servers at startup.”
4. Make any other changes you want, and then click OK.

Customizing a monitoring profile for a cluster

You can create new profiles and edit existing profiles to customize the tasks and statistics that Domino displays.

1. From the Domino Administrator, click the Server - Monitoring tab.
2. In the “Monitoring profiles” field, select an existing profile.
Selecting a profile initializes the server monitor if it is not already initialized. You cannot make changes to a profile until the server monitor is initialized.
3. To add one or more tasks to monitor, choose Monitoring - Monitor New Task, select the tasks you want to add, and then click OK.
For clustering, it can be useful to monitor the Cluster Database Directory Manager and the Cluster Replicator.
4. To add one or more statistics to monitor, choose Monitoring - Monitor New Statistic, do the following in the “Add Statistic(s) to this profile” dialog box, and then click OK.
 - Expand Replica - Cluster, and then select the statistics you want to monitor for cluster replication.
There are many statistics that are helpful, but SecondsOnQueue and WorkQueueDepth are particularly helpful in determining whether you need to increase the number of Cluster Replicators you are running on the server.
 - Expand Server - Cluster, and then select the other cluster statistics you want to monitor.
If Availability Index and Availability Threshold are not already included in your profile, it is helpful to monitor those. It is also helpful to monitor OpenRedirects - Failover and OpenRedirects - LoadBalance, as well as OpenRequest - LoadBalanced and OpenRequest - ClusterBusy to track how often failover occurs.
5. (Optional) To add a server to the profile, select Monitoring - Monitor New Server, and then select the server from the list; or drag a server from the Server pane to the server monitor.
6. (Optional) To remove a server from the profile, click the name of the server you want to remove, and then select Monitoring - Remove Server.
7. To save your changes to the profile, do one of the following:
 - To save this profile as a new profile while also preserving the original profile, choose Monitoring - Profiles - Save As, and then enter a name for the profile.

- To have this modified profile replace the original profile, you do not have to do anything. The profile is saved automatically when you close the Domino Administrator.

For more information about monitoring Domino servers, see the book *Administering the Domino System*.

Balancing the workload in a cluster

After monitoring your cluster and determining that you want to make an adjustment, you can do the following to better balance the workload across the servers in a cluster:

- Limit the workload on a server (by changing the server availability threshold)
- Change the setting for the maximum number of users who can use a server at the same time
- Move databases to other servers
- Create more replicas of busy databases
- Add one or more servers to the cluster
- Use the IBM Tivoli Analyzer for Lotus Domino tool to analyze and make recommendations about the conditions in a cluster

Keep in mind that workload balancing is not the solution for a general lack of capacity in your enterprise. If your Domino servers are struggling with their current workload, and there are no additional servers to handle the excess load, enabling workload balancing will not solve the problem. To balance the workload, there must be somewhere to send the overflow from the overworked servers. If there is nowhere to send these requests, they will be handled by the original busy servers. However, the process of looking for another available server for each request will increase the workload on the server. If you do not have enough capacity in your enterprise, consider adding more memory or processors or otherwise upgrading your equipment to handle a larger load.

If the workload in your cluster is normally well distributed, consider configuring the cluster for failover only, not for workload balancing.

Limiting the workload of a server

The current workload on a server is measured by the server availability index. To better balance the workload between the servers in your cluster, you can limit the workload of each server by setting the server availability threshold. When the server availability index falls to a point where it is equal to or less than the server availability threshold that you set, the server is designated as BUSY.

When a server is BUSY, access requests fail over to another server in the cluster, if one is available. If no other server is available, the original server takes the access request, even though it is BUSY. No request is denied access because a server is in a BUSY state. Each time an access request is redirected, Domino generates a workload balancing event in the Domino server log file.

The server availability threshold does not affect replication. Replication occurs even when a server is in a BUSY state.

The server availability index

Each server in a cluster periodically determines its own workload based on the response time of the requests the server has processed recently. The workload is expressed as a number from 0 to 100, where 0 indicates a heavily loaded server and 100 indicates a lightly loaded server. This number is called the server availability index. As response times increase, the server availability index decreases.

The server availability index is approximately equal to the percentage of the total server capacity that is still available. For example, if the server availability index is 65, you still have approximately 65% of the capacity of your server available. Although the servers in your enterprise may vary in power and resources, the server availability index represents the same thing on each server — the amount of total availability of that server that is still available.

The server availability index is based on the expansion factor, which is determined by comparing recent response times for specific types of transactions to the minimum time in which the server has ever completed the same types of transactions. For example, if the server is currently averaging 12 microseconds to perform Database Open transactions but the minimum time the server has ever performed a Database Open transaction is 3 microseconds, the expansion factor for Database Open transactions would be 4 (the current time of 12 microseconds divided by the fastest time of 3 microseconds). In other words, the expansion factor determines how many times longer it takes for a transaction to complete currently than it takes under optimal conditions.

Domino stores the minimum time for each type of transaction in memory and in the LOADMON.NCF file, which the server reads each time it starts. When the server shuts down, Domino updates the LOADMON.NCF file with the latest information.

To determine the current expansion factor, Domino tracks the most commonly used types of Domino transactions for specified periods of time. By default Domino tracks these transactions for 5 periods of 15 seconds each. Domino then determines the average time it took to complete each type of transaction and divides that time by the minimum time it ever took to complete that same type of transaction. This determines an expansion factor for each type of transaction. To determine the expansion factor for the entire server, Domino averages the expansion factors for all the types of transactions, giving a heavier weighting to the most frequently used types of transactions.

As the server gets busier, adding more load has an increasingly greater effect on performance and availability. Thus, adding more load to a busy server increases the expansion factor faster than adding more load to a less busy server.

An expansion factor of 64 is considered the maximum load for a server. In other words, if the server is taking 64 times longer to complete transactions than it does under optimal conditions, Domino considers the server to be fully loaded.

How the availability index compares to the expansion factor

To determine the availability index, Domino uses a formula that converts the expansion factor into an approximation of the percentage of the total server capacity that is still available. The following table shows a few examples of expansion factors converted to availability indexes.

<i>Expansion factor</i>	<i>Availability index</i>
1	100
2	83
4	67
8	50
16	33
32	17
64	0

Note The expansion factor and the availability index measure only the response time of the server, which is usually only a small portion of the response time clients experience. For example, the network response time between a client and a server often accounts for a significant portion of the response time the client experiences.

Changing the amount of data used to compute the expansion factor

Although it is not usually necessary, you can use the following NOTES.INI settings to change the amount of data that Domino collects in order to figure the expansion factor.

- To change the number of data collection periods that Domino uses, use the NOTES.INI setting `Server_Transinfo_Max=x` where *x* is the number of collection periods you want Domino to use.
- To change the length of each data collection period, use the NOTES.INI setting `Server_Transinfo_Update_Interval=x` where *x* is the length of each period in seconds.

Choosing the server availability threshold

Setting the server availability threshold on each server is a key factor in balancing the workload in the cluster. Setting the server availability threshold too high can result in failover occurring unnecessarily. Setting it too low can result in poor performance for users who could have received better performance from a different server.

To determine the proper value for the server availability threshold, do the following:

1. During periods of normal to heavy load, use one of the following methods to observe the server availability index:
 - From the Domino Administrator or the Web Administrator, make the server you want current. Then click the Server - Statistics tab. Then, in the statistics list, expand Server.
 - In the Server pane of the Domino Administrator, expand All Servers or expand Clusters, right-click the server you want, choose Server Properties, and then click the Cluster tab.
This method is not available in the Web Administrator.
 - At the server console, type **show cluster**.
 - At the server console, type **show stat server**.

Note Prior to using the next two methods, you must enable statistic reporting.

- From the Domino Administrator or the Web Administrator, click the Server - Analysis tab. In the Task pane, expand Monitoring Results - Statistics Reports - Clusters. In the Results pane, open the Monitoring Results document you want.
- Open the Monitoring Results database (STATREP.NSF), and look in the Statistics Reports - Cluster view.

2. Set an initial availability threshold based on the results of your observation. Consider the following when setting this value:
 - The value should be near the lower end of the values you observed.
 - Add some extra capacity (lower the availability threshold number) to accommodate servers that may fail over to this server.

When a server fails, the workload fails over to other servers in the cluster. If there are only two servers in the cluster, each with the same workload, this would result in approximately a 100% increase in the workload of the remaining server. If there are six servers in the cluster, this would result in approximately a 20% increase in the workload of the remaining servers. Therefore, you should set the availability threshold of each cluster server low enough to allow the server to absorb an adequate portion of the workload if another cluster server fails.

3. Track other cluster statistics to see if the workload is reasonably balanced. The following table lists some of the statistics that are helpful in determining if the workload is balanced.

<i>Statistic name</i>	<i>Description</i>
Server.AvailabilityIndex	The current value of the server availability index. The values range from 0 to 100. A value of 0 indicates that there are no resources available on the server. A value of 100 means that the server is completely available.
Server.ExpansionFactor	The current expansion factor. The values range from 1 to 64. A value of 1 indicates that the server is completing transactions at the minimum time for that server. A value of 64 indicates that it is taking 64 times longer than the minimum time to complete transactions. An expansion factor of 64 is considered fully loaded.
Server.Cluster.OpenRedirects.LoadBalance.Successful	The number of times a BUSY server successfully redirects a client to another cluster member.
Server.Cluster.OpenRedirects.LoadBalance.Unsuccessful	The number of times a BUSY server is unsuccessful in redirecting a client to another cluster member. A server will be unsuccessful if the target server is also in a BUSY state or otherwise unavailable.

continued

<i>Statistic name</i>	<i>Description</i>
Server.Cluster.OpenRequest.ClusterBusy	The number of times a BUSY server tries to redirect a client request when all other cluster servers were also BUSY.
Server.Cluster.OpenRequest.LoadBalanced	The number of times a user tried to open a database on this server when the server was BUSY.

These statistics are cumulative from when the server was started.

4. Compare these statistics for all the servers in the cluster.
5. Adjust the server availability threshold on any servers that do not seem well balanced.

Note Workload balancing is not a substitute for having adequate computer resources for your enterprise. If your servers are already overloaded, workload balancing merely increases the problem because there is no place for a BUSY server to send client requests. Looking for an available server, however, decreases the performance on an already busy server.

Setting the server availability threshold

The server availability threshold specifies the lowest acceptable server availability index. Approximately once each minute, Domino computes the server availability index and compares it to the server availability threshold that you set. If the availability index is less than or equal to the availability threshold, the server is marked as BUSY. When a server is marked as BUSY, requests to open databases are redirected to another server, if one is available. When the availability index becomes higher than the availability threshold again, the BUSY condition is withdrawn.

From the Domino Administrator

1. Click the Configuration tab.
2. In the Task pane, expand Server, and then click Configurations.
3. Do one of the following:
 - If a Configuration Settings document already exists for the server you want, select that document, and then click Edit Configuration.
 - If a Configuration Settings document does not already exist for the server you want, click Add Configuration, and add the name of the server in the “Group or Server name” field on the Basics tab.
4. Click the NOTES.INI Settings tab.
5. Click Set/Modify Parameters.
6. In the Item field, select or enter SERVER_AVAILABILITY_THRESHOLD.

7. In the Value field, enter the number you want for the server availability threshold.

The higher the number you enter, the less workload the server can carry before going into the BUSY state. Entering the number 100 automatically puts the server into the BUSY state, regardless of its actual availability. Entering the number 0 disables workload balancing for that server. The default value is 0.

8. Click Add, and then click OK.
9. Click Save & Close.

From the Web Administrator

1. Click the Configuration tab.
2. In the Task pane, expand Server, and then click Configurations.
3. Do one of the following:
 - If a Configuration Settings document already exists for the server you want, open that document, and then click Edit Server Configuration.
 - If a Configuration Settings document does not already exist for the server you want, click Add Configuration, and add the name of the server in the “Group or Server name” field on the Basics tab.
4. Click the NOTES.INI Settings tab.
5. Click Set/Modify Parameters.
6. In the Available Parameters box, click SERVER_AVAILABILITY_THRESHOLD, and then click Add.
7. In the Value field, enter the number you want for the server availability threshold, and then click OK.

The higher the number you enter, the less workload the server can carry before going into the BUSY state. Entering the number 100 automatically puts the server into the BUSY state, regardless of its actual availability. Entering the number 0 disables workload balancing for that server. The default value is 0.

8. Click Save & Close.

Using the availability threshold when you restart a server in a cluster

When you restart a server in a cluster, it is a good idea to make the server BUSY until all replication to the server is complete. This ensures that users access up-to-date information in the databases on the server. You can make a server BUSY by setting the availability threshold to 100. When replication is complete, make the server available to users.

Using the server availability threshold to control failover to specific servers

In some cases, you may want to limit failover to a server. For example, if you set up a cluster over a WAN and one of the cluster servers is more distant than the other servers, you may want to limit failover to the distant server. You can limit failover to this server by setting its availability threshold very high.

For example, if you have three servers — one in Boston, one in New York, and one in Hong Kong — the Boston server would fail over to the Hong Kong server if it is more available than the New York server. However, if you set the availability threshold on the Hong Kong server to 100, the other cluster servers will not fail over to the Hong Kong server unless no other available cluster server contains a replica of the requested database.

When you control failover in this manner, be sure that the other cluster servers (the servers in Boston and New York in the example) have enough resources to handle most of the failover in the cluster.

Setting the maximum number of users on a server

You can also balance the workload in a cluster by using the `Server_MaxUsers` setting. This setting specifies the maximum number of active users allowed on a server at one time. When the server reaches this limit, the server goes into the `MAXUSERS` state and rejects any additional requests until the number of active users falls below the `Server_MaxUsers` limit. When Domino rejects an access request because of a `MAXUSERS` state, the Cluster Manager attempts to redirect the request to another cluster server that contains the appropriate replica. If no other server is available, Domino rejects the access request and displays an explanatory message.

Note The `Server_MaxUsers` setting does not affect replication. Replication occurs even when a server is in a `MAXUSERS` state.

From the Domino Administrator

1. Click the Configuration tab.
2. In the Task pane, expand Server, and then click Configurations.
3. Do one of the following:
 - If a Configuration Settings document already exists for the server you want, select that document, and then click Edit Configuration.
 - If a Configuration Settings document does not already exist for the server you want, click Add Configuration, and add the name of the server in the “Group or Server name” field on the Basics tab.
4. Click the NOTES.INI Settings tab.

5. Click Set/Modify Parameters.
6. In the Item field, select or enter SERVER_MAXUSERS.
7. In the Value field, enter the maximum number of users you want to access the server at the same time.
8. Click Add, and then click OK.
9. Click Save & Close.

From the Web Administrator

1. Click the Configuration tab.
2. In the Task pane, expand Server, and then click Configurations.
3. Do one of the following:
 - If a Configuration Settings document already exists for the server you want, open that document, and then click Edit Server Configuration.
 - If a Configuration Settings document does not already exist for the server you want, click Add Configuration, and add the name of the server in the “Group or Server name” field on the Basics tab.
4. Click the NOTES.INI Settings tab.
5. Click Set/Modify Parameters.
6. In the Available Parameters box, click SERVER_MAXUSERS, and then click Add.
7. In the Value field, enter the maximum number of users you want to access the server at the same time, and then click OK.
8. Click Save & Close.

To see how often requests are redirected, check the log file for workload balancing events or check the Cluster Manager statistics.

For information about viewing the log file and the Cluster Manager statistics, see the topic “Monitoring Cluster Manager events and statistics.”

Note You can use the Server_MaxUsers setting with any Domino server. However, only the servers in a cluster redirect access requests to another server when a server is in a MAXUSERS state. Servers that are not in a cluster reject the access requests.

Redistributing replicas

Often you can redistribute replicas to other servers in the cluster in order to better balance the workload. For example, if one of the servers in the cluster is significantly more busy than the other servers, consider moving one or more replicas from the busy server to the less busy servers. You can also create more replicas of a busy database so that the workload is shared by more servers, thus reducing the workload on a busy server.

To move or create replicas in a cluster, you use the same procedures you use to move or create replicas on any Domino server. While moving or creating the replicas, you can select “Show me only cluster members for cluster: *cluster name*” in the “Move Database” or “Create Replica” dialog box. This causes the Domino Administrator or the Web Administrator to display only the names of the servers in the cluster. You can then easily place the replicas on every cluster server you want.

For more information about moving replicas and creating replicas, see the book *Administering the Domino System*.

Using the IBM Tivoli Analyzer for Lotus Domino in a cluster

The IBM Tivoli® Analyzer for Lotus Domino consists of two parts — Activity Trends and Server Health Monitoring. This topic describes how these tools work in a cluster.

Note The IBM Tivoli Analyzer for Lotus Domino is a separate product offering from Tivoli software from IBM.

Using Activity Trends in a cluster

If you are using Activity Trends, you should be aware that this tool does not recognize when servers are in a cluster. If you create an Activity Trends server group that includes servers from different clusters or some servers that are in a cluster and some servers that are not in a cluster, Activity Trends may suggest moving a database out of a cluster in order to balance the resources in the server group. To prevent this, you can create a separate server group for each cluster and a separate group for nonclustered servers. If you decide to include cluster servers and nonclustered servers in the same server group, you can use the Activity Trends pin features to exclude specific databases from resource balancing.

Using Server Health Monitoring in a cluster

Server Health Monitoring in a cluster works the same way it works outside of a cluster. Unlike Activity Trends, Server Health Monitoring assessments are based on each individual server rather than on a group of servers.

Managing failover in a cluster

When you set up a cluster, you create replicas of databases so that users fail over automatically to a different replica if the database or server they try to access is unavailable. In addition, you may want to cause failover to occur intentionally to:

- Upgrade the server software or the operating system
- Perform maintenance on the server
- Replace a server with a different server

Causing failover to occur

To cause failover to occur, you can use the `Server_Restricted` setting. This setting tells a server to deny new open database requests and places the server into a `RESTRICTED` state. This prevents new users from accessing a server, although users who have active connections to databases on the server retain their connections. This setting is useful when you want to do maintenance on a server, upgrade a server, or take a server out of service for another reason. This is also useful when users have failed over to a server and you want them to fail back to the original server.

When a server is in a `RESTRICTED` state, the Cluster Manager redirects new open requests to other servers in the cluster. When an attempt to redirect is unsuccessful, the user receives an explanatory message and cannot access the server. For each redirection attempt, Domino generates a failover event in the log file.

Note The `Server_Restricted` setting does not affect replication. Replication occurs even when a server is in a `RESTRICTED` state.

From the Domino Administrator

1. Click the Configuration tab.
2. In the Task pane, expand Server, and then click Configurations.
3. Do one of the following:
 - If a Configuration Settings document already exists for the server you want, select that document, and then click Edit Configuration.
 - If a Configuration Settings document does not already exist for the server you want, click Add Configuration, and add the name of the server in the “Group or Server name” field on the Basics tab.
4. Click the NOTES.INI Settings tab.
5. Click Set/Modify Parameters.
6. In the Item field, select or enter `SERVER_RESTRICTED`.

7. In the Value field, enter 0, 1, or 2 where these numbers mean the following:
 - 0 - The server is unrestricted
 - 1 - The server is RESTRICTED for the current session only. Restarting the server clears the setting.
 - 2 - The server is RESTRICTED persistently, even after the server restarts.
8. Click Add, and then click OK.
9. Click Save & Close.

From the Web Administrator

1. Click the Configuration tab.
2. In the Task pane, expand Server, and then click Configurations.
3. Do one of the following:
 - If a Configuration Settings document already exists for the server you want, open that document, and then click Edit Server Configuration.
 - If a Configuration Settings document does not already exist for the server you want, click Add Configuration, and add the name of the server in the “Group or Server name” field on the Basics tab.
4. Click the NOTES.INI Settings tab.
5. Click Set/Modify Parameters.
6. In the Available Parameters box, click SERVER_RESTRICTED, and then click Add.
7. In the Value field, enter 0, 1, or 2 where these numbers mean the following:
 - 0 - The server is unrestricted
 - 1 - The server is RESTRICTED for the current session only. Restarting the server clears the setting.
 - 2 - The server is RESTRICTED persistently, even after the server restarts.
8. Click OK, and then click Save & Close.

From the server console

Send the following Domino command from the server console:

```
set config server_restricted=n
```

where *n* can be 0, 1, or 2. These numbers mean the following:

0 - The server is unrestricted

1 - The server is RESTRICTED for the current session only. Restarting the server clears the setting.

2 - The server is RESTRICTED persistently, even after the server restarts.

If you want to restrict a server and do not want to wait for all users to close their existing sessions, enter the Drop All command at the console after you put the server into the RESTRICTED state. The Drop All command closes all existing sessions on the server. When users try to reopen the databases they were using, they fail over to a different server, if one is available.

Replacing a cluster server with a different server

If you want to replace a cluster server with a different server, do the following:

1. Set up Domino on the new server.
2. Set the server as RESTRICTED so users won't be able to access it until it is active in the cluster.
3. Create the replicas you want on the new server.
4. Add the new server to the cluster.
5. After the new server becomes active in the cluster, remove the restriction from the new server, and set the old server as RESTRICTED.

This causes user requests to fail over to the new server, and lets you remove the old server without disrupting users.

6. Remove the old server from the cluster.

Note You can use the Server_Restricted setting for any Domino server. This setting is not limited to clusters.

Other ways to manage failover

If you have set a server to be RESTRICTED or BUSY but have now brought the server back online and want users to fail back to this server, you can make the server that users had failed over to RESTRICTED or BUSY.

You can also set up a server as a backup to another server. You can set the availability threshold to 100 on the backup server so the server is in a BUSY state at all times. This is similar to a RESTRICTED state, except a BUSY server accepts new open requests if no other replica is available, while a RESTRICTED server does not. Therefore, the backup server accepts open database requests only when the primary server is unavailable.

Managing replication in a cluster

You should use both cluster replication and standard replication in a cluster and run them both at the same time. Cluster replication is the primary means of replication in a cluster, while standard replication updates any databases that the Cluster Replicator couldn't update.

You should schedule replication in a cluster on a regular basis to be sure that databases are properly updated even when they aren't replicated by the Cluster Replicator.

For more information about using standard replication in a cluster, see the chapter "Setting Up a Cluster."

As you move or copy databases into a cluster, be sure to check the access control list and other restrictions that are copied with the database. All servers must be able to replicate all replicas of a database for a cluster to work properly. In addition, users' access rights should be the same in every replica of a database so that they see the same thing and can perform the same tasks if they fail over. Make sure that any access restrictions to a database in a cluster are uniform across all replicas.

Configuring and controlling cluster replication

When you add a server to a cluster, Domino starts the Cluster Replicator automatically. Each time you start the server, the Cluster Replicator starts. You can also do the following:

- Use multiple Cluster Replicators
- Disable cluster replication for an entire server
- Disable cluster replication for specific databases
- Pause cluster replication
- Resume cluster replication
- Force the Cluster Replicator to retry failed replications
- Force the Cluster Replicator to update the Cluster Database Directory information it stores in memory
- Force the Cluster Replicator to generate a log document

Using multiple Cluster Replicators

If you have heavily used databases that consistently overload the Cluster Replicator with update requests, consider running multiple Cluster Replicators. When you run multiple Cluster Replicators on a server, they work simultaneously to replicate changes to other servers. If one Cluster Replicator is busy replicating changes to one database, a second Cluster Replicator can begin replicating changes to another database. By sharing the replication workload, multiple Cluster Replicators ensure that data updates are made quickly and databases stay tightly synchronized.

Determining how many Cluster Replicators to run

To determine how many cluster replicators to run, monitor the cluster replication statistics. For example, `Replica.Cluster.WorkQueueDepth` shows the number of modified databases that are waiting to be replicated. If the number of databases waiting is frequently greater than zero, consider adding one or more Cluster Replicators. However, the number of databases waiting to be replicated may also be greater than zero if you don't have enough network bandwidth to process the transactions fast enough. If this is the case, you should consider setting up a private LAN for your cluster. Also keep in mind that adding processors or memory to your server can enhance replication performance.

`Replica.Cluster.SecondsOnQueue` shows the number of seconds that the last database spent waiting in the replication queue before being replicated. Since the cluster replicator checks its queue every 15 seconds, this number should be under 15 during periods of light load. If this number is frequently higher than 30, you should consider adding one or more cluster replicators.

You should add one Cluster Replicator at a time until the cluster replication statistics are where you want them to be.

Starting multiple Cluster Replicators

You can configure Domino to start multiple Cluster Replicators automatically each time the server starts, or you can start multiple Cluster Replicators for the current session only.

Starting multiple Clusters Replicators automatically at server startup

You can use the Domino Administrator or the Web Administrator to configure Domino to start multiple Cluster Replicators at server startup.

From the Domino Administrator

1. Click the Configuration tab.
2. In the Task pane, expand Server, and then click Configurations.
3. Do one of the following:
 - If a Configuration Settings document already exists for the server you want, select that document and then click Edit Configuration.
 - If a Configuration Settings document does not already exist for the server you want, click Add Configuration, and add the name of the server in the “Group or Server name” field on the Basics tab.
4. Click the NOTES.INI Settings tab.
5. Click Set/Modify Parameters.
6. In the Item field, select or enter CLUSTER_REPLICATORS.
7. In the Value field, enter the number of Cluster Replicators you want to run on this server.

Note Entering 0 (zero) in the Value field does not stop all Cluster Replicators. One Cluster Replicator will still run. To turn off all Cluster Replicators, see the topic “Disabling cluster replication for an entire server.”
8. Click Add, and then click OK.
9. Click Save & Close.
10. Restart the server so the setting takes effect.

From the Web Administrator

1. Click the Configuration tab.
2. In the Task pane, expand Server, and then click Configurations.
3. Do one of the following:
 - If a Configuration Settings document already exists for the server you want, open that document, and then click Edit Server Configuration.
 - If a Configuration Settings document does not already exist for the server you want, click Add Configuration, and add the name of the server in the “Group or Server name” field on the Basics tab.
4. Click the NOTES.INI Settings tab.
5. Click Set/Modify Parameters.
6. In the Available Parameters box, click CLUSTER_REPLICATORS, and then click Add.
7. In the Value field, enter the number of Cluster Replicators you want to run on this server, and then click OK.

Note Entering 0 (zero) in the Value field does not stop all Cluster Replicators. One Cluster Replicator will still run. To turn off all Cluster Replicators, see the topic “Disabling cluster replication for an entire server.”

8. Click Save & Close.
9. Restart the server so the setting takes effect.

Starting multiple Cluster Replicators for the current session only

To run multiple Cluster Replicators for the current session only, do one of the following.

From the Domino Administrator or the Web Administrator

1. In the Server pane, expand All Servers or expand Clusters.
2. Select the server you want.
3. Click the Server - Status tab.
4. In the Task pane, do one of the following:
 - From the Domino Administrator, click Server Tasks.
 - From the Web Administrator, click All Server Tasks.
5. In the Tools pane, expand Task, and then click Start.
6. Select Cluster Replicator.
7. Click “Start Task” once for each Cluster Replicator you want to start, and then click Done.

From the server console

Send the following command from the server console once for each Cluster Replicator you want to start.

```
load clrepl
```

Each time you send this command, the server starts another Cluster Replicator.

Disabling cluster replication for an entire server

You can disable cluster replication to investigate a problem or to improve performance. You can stop cluster replication for the current server session only, or you can prevent cluster replication from starting when the server starts. If you stop cluster replication for the current session only, cluster replication starts again the next time the server starts. Disabling cluster replication has no effect on scheduled replication, which still occurs.

Use one of the following procedures to disable cluster replication.

Note Disabling the Cluster Replicator prevents only replication from that server to other cluster servers. It does not prevent replication to the server from other cluster servers.

Disabling cluster replication automatically at server startup

You can use the Domino Administrator or the Web Administrator to configure Domino to disable cluster replication at server startup.

From the Domino Administrator

1. Click the Configuration tab.
2. In the Task pane, expand Server, and then click Configurations.
3. Do one of the following:
 - If a Configuration Settings document already exists for the server you want, select that document, and then click Edit Configuration.
 - If a Configuration Settings document does not already exist for the server you want, click Add Configuration, and add the name of the server in the “Group or Server name” field on the Basics tab.
4. Click the NOTES.INI Settings tab.
5. Click Set/Modify Parameters.
6. In the Item field, select or enter `DISABLE_CLUSTER_REPLICATOR`.
7. In the Value field, enter 1 (one).
8. Click Add, and then click OK.
9. Click Save & Close.
10. Restart the server so the setting takes effect.

Note To restart cluster replication, set `Disable_Cluster_Replicator` to 0 (zero) or remove this line from the Configuration Settings document. Then restart the server.

From the Web Administrator

1. Click the Configuration tab.
2. In the Task pane, expand Server, and then click Configurations.
3. Do one of the following:
 - If a Configuration Settings document already exists for the server you want, open that document, and then click Edit Server Configuration.
 - If a Configuration Settings document does not already exist for the server you want, click Add Configuration, and add the name of the server in the “Group or Server name” field on the Basics tab.
4. Click the NOTES.INI Settings tab.

5. Click Set/Modify Parameters.
6. In the Available Parameters box, click DISABLE_CLUSTER_REPLICATOR, and then click Add.
7. In the Value field, enter 1 (one), and then click OK.
8. Click Save & Close.
9. Restart the server so the setting takes effect.

Disabling cluster replication for the current session only

To disable cluster replication for the current session only, do one of the following.

From the Domino Administrator or the Web Administrator

1. Click the Server - Status tab.
2. In the Task pane, do one of the following:
 - From the Domino Administrator, click Server Tasks.
 - From the Web Administrator, click All Server Tasks.
3. In the Results pane, select a Cluster Replicator.
4. In the Tools pane, expand Task, and then click Stop.

From the server console

Send the following Domino command from the server console:

```
tell clrepl quit
```

Using these procedures shuts down all Cluster Replicators, even if multiple Cluster Replicators are running.

Disabling cluster replication for specific databases

In a cluster, all databases are automatically set to replicate in real time so that replicas are always synchronized. However, real-time replication is needed only for databases that require constant synchronization of data. For other databases, you can disable cluster replication.

1. In the Server pane of the Domino Administrator or the Web Administrator, expand All Servers or expand Clusters.
2. Choose a server in the cluster.
3. Click the Files tab.
4. In the Task pane, expand Cluster Directory (6).
5. Click one of the four views:
 - Databases by Filename
 - Databases by Pathname
 - Databases by Replica ID

- Databases by Server

The Databases by Filename view is particularly useful for disabling cluster replication on specific databases.

6. In the Results pane, select the databases for which you want to disable cluster replication.
7. Do one of the following:
 - From the Domino Administrator, click Tools - Disable Cluster Replication on Selected Databases.
 - From the Web Administrator, click Tools - Disable Replication.

Tip Databases with the letter “X” in the left column in the Cluster Database Directory have cluster replication disabled.

Disabling cluster replication of a database prevents only replication of changes from that database to other servers in the cluster. It does not prevent replication to the database from other cluster servers. Disabling cluster replication has no effect on standard replication.

Note To reenables cluster replication for specific databases, follow the procedure above with the following exception: In Step 7, click Tools - Enable Cluster Replication on Selected Databases from the Domino Administrator, or click Tools - Enable Replication from the Web Administrator.

Creating an agent to disable cluster replication for specific databases

As an alternate method of disabling cluster replication for specific databases, you can create an agent.

1. Open the Cluster Database Directory (CLDBDIR.NSF) on any server in the cluster.
2. Choose Create - Agent.
3. In the Name field of the Agent properties box, type a name for the agent.
4. In the Options section, choose Shared or Private, depending on who will use this agent.
5. In the Trigger field of the Runtime section, choose On event.
6. Choose “Action menu selection” as the trigger event.
7. In the Target field, choose “All selected documents.”
8. Click the Security tab of the Agent properties box, and specify any security options you want.
9. Close the Agent properties box.
10. In the Programmer’s pane, click Add Action.

11. In the Action field, choose Modify Field.
12. In the “Modify by” field, choose Replacing.
13. In the “The value in field” field, choose ClusterReplicate.
14. In the “With the new value” field, type 0 (zero).
15. Click Add.
16. Choose File - Save to save the agent.

Creating an agent to reenable cluster replication for specific databases

If you create an agent to disable cluster replication on specific databases, you might want to create another agent to reenable cluster replication on specific databases. To create this agent, follow the previous procedure with the following exceptions:

- In step 3, type a different name for the agent.
- In step 14, type 1 (one) in the “With the new value” field.

Running the agents

1. Open the Cluster Database Directory.
2. Select the databases for which you want to disable cluster replication or enable cluster replication.
3. Choose the appropriate agent name from the Actions menu.

Pausing cluster replication

When you want to stop the Cluster Replicator temporarily, it is better to pause it than to stop it. When you pause the Cluster Replicator, it continues to store pending cluster replication events in memory and continues to store the Cluster Database Directory in memory. If you stop the Cluster Replicator, these things are deleted from memory. Therefore, you should replicate manually when you restart the Cluster Replicator to replicate any changes that were lost when the Cluster Replicator stopped. You do not have to replicate manually when you pause the Cluster Replicator.

When you pause the Cluster Replicator, it finishes processing the current task and then stops taking new requests until you resume cluster replication.

From the Domino Administrator

1. In the Server pane, expand All Servers or expand Clusters.
2. Select the server you want.
3. Click the Server - Status tab.
4. In the Task pane, click Server Tasks.

5. In the Results pane, select the Cluster Replicator.
6. In the Tools pane, expand Task, and then click Tell.
7. Click Pause, and then click OK.

From the Web Administrator

1. In the Server pane, expand All Servers or expand Clusters.
2. Select the server you want.
3. Click the Server - Status tab.
4. In the Task pane, click All Server Tasks.
5. In the Results pane, select the Cluster Replicator.
6. In the Tools pane, expand Task, and then click Tell.
7. Type **tell clrepl pause** and then click OK.

From the server console

Send the following Domino command from the server console:

```
tell clrepl pause
```

Resuming cluster replication

After pausing cluster replication, you can resume cluster replication by following this procedure.

From the Domino Administrator

1. In the Server pane, expand All Servers or expand Clusters.
2. Select the server you want.
3. Click the Server - Status tab.
4. In the Task pane, click Server Tasks.
5. In the Results pane, select the Cluster Replicator.
6. In the Tools pane, expand Task, and then click Tell.
7. Click Resume, and then click OK.

From the Web Administrator

1. In the Server pane, expand All Servers or expand Clusters.
2. Select the server you want.
3. Click the Server - Status tab.
4. In the Task pane, click All Server Tasks.
5. In the Results pane, select the Cluster Replicator.
6. In the Tools pane, expand Task, and then click Tell.
7. Type **tell clrepl resume** and then click OK.

From the server console

Send the following Domino command from the server console:

```
tell clrepl resume
```

Forcing the Cluster Replicator to retry failed replications immediately

Occasionally, cluster replication is unsuccessful. This is usually due to the destination server being unavailable. The Cluster Replicator stores unsuccessful replication events in memory and retries these replications periodically. The interval between retries starts at one hour and increases over time to a maximum of one day. You can see if any cluster replication events are waiting to be retried by looking in the server log file.

If you do not want the Cluster Replicator to wait to retry these replications, you can force it to retry the replications immediately. This is especially useful when you restart a different cluster server that was unavailable.

From the Domino Administrator

1. In the Server pane, expand All Servers or expand Clusters.
2. Select the server you want.
3. Click the Server - Status tab.
4. In the Task pane, click Server Tasks.
5. In the Results pane, select the Cluster Replicator.
6. In the Tools pane, expand Task, and then click Tell.
7. Click "Retry all failed replications immediately," and then click OK.

From the Web Administrator

1. In the Server pane, expand All Servers or expand Clusters.
2. Select the server you want.
3. Click the Server - Status tab.
4. In the Task pane, click All Server Tasks.
5. In the Results pane, select the Cluster Replicator.
6. In the Tools pane, expand Task, and then click Tell.
7. Type **tell clrepl retry** and then click OK.

From the server console

Send the following Domino command from the server console:

```
tell clrepl retry
```

Forcing the Cluster Replicator to update the Cluster Database Directory information immediately

The Cluster Replicator stores information from the Cluster Database Directory in memory and uses this information to replicate changes to the other cluster servers. Every 15 seconds, the Cluster Replicator checks the Cluster Database Directory for changes, such as databases that were added or deleted or databases that have a different cluster replication status. If the Cluster Replicator detects changes, it updates the information it has stored in memory.

You can force the Cluster Replicator to check immediately for changes to the Cluster Database Directory.

From the Domino Administrator

1. In the Server pane, expand All Servers or expand Clusters.
2. Select the server you want.
3. Click the Server - Status tab.
4. In the Task pane, click Server Tasks.
5. In the Results pane, select the Cluster Replicator.
6. In the Tools pane, expand Task, and then click Tell.
7. Click “Refresh cluster configuration changes,” and then click OK.

From the Web Administrator

1. In the Server pane, expand All Servers or expand Clusters.
2. Select the server you want.
3. Click the Server - Status tab.
4. In the Task pane, click All Server Tasks.
5. In the Results pane, select the Cluster Replicator.
6. In the Tools pane, expand Task, and then click Tell.
7. Type **tell clrepl refresh** and then click OK.

From the server console

Send the following Domino command from the server console:

```
tell clrepl refresh
```

Forcing the Cluster Replicator to log immediately

If you do not want to wait for the Cluster Replicator to generate a log document, you can force the Cluster Replicator to generate a log document immediately. The Cluster Replicator logs all replication events that occurred since the last time it logged. This includes replication errors and replications that are waiting to be retried.

From the Domino Administrator

1. In the Server pane, expand All Servers or expand Clusters.
2. Select the server you want.
3. Click the Server - Status tab.
4. In the Task pane, click Server Tasks.
5. In the Results pane, select the Cluster Replicator.
6. In the Tools pane, expand Task, and then click Tell.
7. Select “Write a Replication Event Log record immediately,” and then click OK.

From the Web Administrator

1. In the Server pane, expand All Servers or expand Clusters.
2. Select the server you want.
3. Click the Server - Status tab.
4. In the Task pane, click All Server Tasks.
5. In the Results pane, select the Cluster Replicator.
6. In the Tools pane, expand Task, and then click Tell.
7. Type **tell clrepl log** and then click OK.

From the server console

Send the following Domino command from the server console:

```
tell clrepl log
```

Creating configuration settings for all servers in a cluster

If there are configuration settings that you want to apply to some or all servers in a cluster, you can create a Configuration Settings document that applies to the servers you want.

1. From the Domino Administrator or the Web Administrator, click the People & Groups tab.
2. In the Task pane, click Groups.
3. In the Results pane, click Add Group.
4. Create a group that includes the servers you want as members, and then click Save & Close.
5. Click the Configuration tab.
6. In the Task pane, expand Server and then click Configurations.
7. In the Results pane, click Add Configuration.

8. in the “Group or Server name” field on the Basics tab, type the name of the group you created.
9. Choose the settings you want.
10. Save and close the document.

Note You can include a server name in more than one Configuration Settings document, but there are specific rules for which document takes precedence.

For more information about the rules of precedence in Configuration Settings documents, see the book *Administering the Domino System*.

Managing database availability in a cluster

There are three database attributes that let you specify whether a database is available for user access. These are:

- Out of Service
- In Service
- Pending Delete

Making a database unavailable for user access

On occasion, you may want to make a database unavailable for user access. You would do this if you want to do some maintenance on the database or if you want users to fail over to a different replica of the database because the server is reaching a high level of use. To make a database unavailable to users, you mark the database “out of service.”

When you mark a database “out of service,” users cannot open the database. Open database requests fail over to a replica, if one is available. If no replica is available, Domino denies users access to the database and displays an explanatory message.

Users who are using the database when you mark it “out of service” continue to have access until they close the database. If users close an out-of-service database and then try to reopen it, they fail over to another replica on an available server, if there is one. This means that the database is gradually brought to an OUT OF SERVICE state without disruption to users who are actively using it. In addition, replication updates from other replicas continue to occur even when a database is marked “out of service.”

To mark a database “out of service,” follow this procedure:

1. In the Server pane of the Domino Administrator or the Web Administrator, expand All Servers or expand Clusters.
2. Select the server that contains the database you want to mark “out of service.”
3. Click the Files tab.
4. Do one of the following:
 - In the Task pane in the Domino Administrator or the Web Administrator, select the folder or view that contains the database you want.
 - In the Task pane in the Domino Administrator only, expand Cluster Directory (6), and then select the view you want.
5. In the Results pane, select the database or databases you want.
6. In the Tools pane, expand Database, and then click Cluster.
7. Click “Out of service,” and then click OK.

To mark all databases on a server “out of service,” use the Server_Restricted setting. A restricted server does not accept new database open requests.

For information about the Server_Restricted setting, see the topic “Managing failover in a cluster.”

Making an “out of service” database available again

If you’ve marked a database “out of service” and now want to restore access to the database, you must mark the database “in service.” When you put an out-of-service database back into service, it becomes fully available to users.

1. In the Server pane of the Domino Administrator or the Web Administrator, expand All Servers or expand Clusters.
2. Select the server that contains the database you want to mark “in service.”
3. Click the Files tab.
4. Do one of the following:
 - In the Task pane in the Domino Administrator or the Web Administrator, select the folder or view that contains the database you want.
 - In the Task pane in the Domino Administrator only, expand Cluster Directory (6), and then select the view you want.
5. In the Results pane, select the database or databases you want.

6. In the Tools pane, expand Database, and then click Cluster.
7. Click “In service,” and then click OK.

Deleting databases from a cluster

To minimize disruption to users, you can set a database to be deleted after every active user has finished using the database. To do this, you mark the database as “pending delete.” When a database is pending delete, the database does not accept any new database open requests. After all users have closed the database, Domino pushes changes to another replica and then deletes the database. You should mark a database “pending delete” if you plan to remove a database that is obsolete or if you are copying the database from one server to another and want to delete the database from the original server.

1. In the Server pane of the Domino Administrator or the Web Administrator, expand All Servers or expand Clusters.
2. Select the server that contains the database you want to delete.
3. Click the Files tab.
4. Do one of the following:
 - In the Task pane of the Domino Administrator or the Web Administrator, select the folder or view that contains the database you want.
 - In the Task pane of the Domino Administrator only, expand Cluster Directory (6), and then select the view you want.

The Databases by Filename view is a convenient view from which to delete replicas, especially if you want to delete multiple replicas of a database from the cluster.

The Databases by Server view is also convenient, especially if you want to delete multiple databases from a particular server.

5. In the Results pane, select the database or databases you want.
6. In the Tools pane, expand Database, and then click Cluster.
7. Click “Pending delete,” and then click OK.

Adding a server to a cluster

To add a server to a cluster, you must have the following access rights:

- Author access and Delete Documents rights and the ServerModifier and ServerCreator roles in the Domino Directory
- Author access with Create documents rights in the Administration Requests database

If possible, use the administration server when adding a server to a cluster. The administration server does not have to be part of the cluster.

If a server is part of a different cluster, you do not have to remove the server from that cluster to add it to the new cluster. The Cluster Administration Process removes the server from its original cluster and then adds it to the new cluster.

Note You cannot use the Web Administrator to add a server to a cluster.

1. From the Domino Administrator, make sure the administration server or another server is current.
2. Click the Configuration tab.
3. In the Task pane, expand Server, and click All Server Documents.
4. In the Results pane, select the server(s) you want to add to the cluster.
5. Click Add to Cluster.
6. Choose the name of the cluster you want to add the server(s) to, and then click OK.
7. Choose Yes to add the server to the cluster immediately, or choose No to submit a request to the Administration Process to add the server to the cluster.
8. (Optional) If you chose No in Step 7 and you did not use the administration server to add the server to the cluster, force replication between the server you used and the administration server so that the administration server receives the requested changes sooner.
9. (Optional) If you chose No in Step 7, force replication between the administration server and the cluster servers so the cluster servers receive the changes sooner.
10. (Optional) If you chose Yes in Step 7, the cluster information is added immediately to the Domino Directory on the server you used to add the server to the cluster. If the server you used is not part of the cluster, replicate the changes to one of the cluster servers.

When you add a server to a cluster, there can be a performance impact because of the amount of replication that must take place initially. Depending on the types of databases on the server and the number of

replicas you create, adding a server can affect CPU performance, Input/Output, and network traffic. Therefore, it is a good idea to add only one server at a time to a cluster, depending on the ability of the equipment in the cluster. If you set up a private LAN for cluster traffic, adding servers is less of a concern because it does not affect your primary network.

What occurs when you add a server to a cluster

When the added server detects a change to its Domino Directory that indicates the server was added to a cluster, the Cluster Manager and the Cluster Administrator start automatically. In addition, the following changes occur:

- The Cluster Administrator starts the Administration Process, if it is not already running.
- The Cluster Administrator starts the Cluster Database Directory Manager.
- The Cluster Administrator starts the Cluster Replicator.
- The Cluster Database Directory Manager creates the Cluster Database Directory (CLDBDIR.NSF).
- The Cluster Database Directory Manager adds a document to the Cluster Database Directory for each database on the new cluster member.
- The Cluster Administrator replicates the Cluster Database Directory and the Domino Directory on the new cluster member with the most available member of the cluster.
- The Schedule Manager creates the Free Time database (CLUBUSY.NSF).
- The Free Time database replicates with the most available member of the cluster.

When the added server's Domino Directory updates, the added server begins to send probes to other servers in the cluster. In return, cluster members begin to probe the new server when their Domino Directories update. This is the way that all the servers in the cluster keep track of the availability and status of the other servers.

To verify that the server was added correctly to the cluster, follow the same procedure you follow to check that a cluster was created correctly.

For more information, see the chapter "Setting Up a Cluster."

The role of the Administration Process in adding a server to a cluster

You can choose to add a server to a cluster immediately, or you can submit the request to the Administration Process. If you submit the request to the Administration Process, Domino creates an Add Server to Cluster request to the Administration Requests database on the server from which you initiated the Add to Cluster request. If you make the request on the administration server, the Administration Process acts immediately on the request. If you make the request from another server, that server's Administration Requests database must replicate with the administration server before the Administration Process can process your request. The Administration Process then adds the cluster information to the Server document in the administration server's Domino Directory to show that the server is part of the cluster. The next time the administrative server replicates with the added server, Domino replicates these changes to the added server's Domino Directory. Cluster membership changes do not take effect until the added server receives the changes to the Server document.

If you choose to add a server to a cluster immediately, Domino immediately makes the changes to the Server document on the server from which you initiated the Add to Cluster command. If you initiate the request on the server you are adding, Domino updates the cluster information immediately on the server you're adding. You do not have to wait for the Administration Process to update the Domino Directories on the cluster servers. Although this adds the server to the cluster faster, it can also lead to replication conflicts.

Removing a server from a cluster

To remove a server from a cluster, you must have the following access rights:

- Author access and Delete Documents rights and the ServerModifier and ServerCreator roles in the Domino Directory
- Author access with Create documents rights in the Administration Requests database

If possible, use the administration server when removing a server from a cluster. The administration server does not have to be part of the cluster.

Note You cannot use the Web Administrator to remove a server from a cluster.

1. From the Domino Administrator, make sure the administration server or another server is current.
2. Click the Configuration tab.

3. In the Task pane, expand Cluster and click Clusters.
4. In the Results pane, select the server you want to remove from the cluster.
5. Click Remove from Cluster.
6. Choose Yes to remove the server from the cluster immediately, or choose No to wait until the Administration Process removes the server from the cluster.

What occurs when you remove a server from a cluster

When the removed server detects the change to its Domino Directory that indicates the server was removed from a cluster, the Cluster Administrator on that server does the following:

- Stops the Cluster Replicator.
- Stops the Cluster Database Directory Manager.
- Replicates the Cluster Database Directory, the Domino Directory, and the Free Time database (CLUBUSY.NSF) to the most available cluster server. This ensures that the other cluster servers receive all changes from the server being removed.
- Removes the Cluster Database Directory (CLDBDIR.NSF) from the server.
- Removes the server Database documents from the Cluster Database Directory of the most available cluster server. These changes then replicate to other members of the cluster. You can view the Cluster Database Directory to verify that the removal occurred correctly.

Note It is possible to remove a server from a cluster even if the server is shut down. However, the documents for this server's databases remain in the Cluster Database Directory. To remove these documents, open the Cluster Database Directory, select the Databases by Server view, and manually remove the documents for the server.

The role of the Administration Process in removing a server from a cluster

You can choose to remove a server from a cluster immediately, or you can submit the request to the Administration Process. If you submit the request to the Administration Process, Domino creates a Remove Server from Cluster request to the Administration Requests database on the server from which you initiated the Remove from Cluster request. If you make the request on the administration server, the Administration Process acts immediately on the request. If you make the request from another server, that server's Administration Requests database must replicate with the administration server before the Administration Process can process your

request. The Administration Process then removes the cluster information from the Server document in the administration server's Domino Directory. The next time the administrative server replicates with the removed server, Domino replicates these changes to the removed server's Domino Directory. Cluster membership changes do not take effect until the removed server receives the changes to the Server document.

If you choose to remove a server from a cluster immediately, Domino immediately makes the changes to the Server document on the server from which you initiated the Remove from Cluster command. If you initiate the request on the server you are removing, Domino updates the cluster information immediately on the server you're removing. You do not have to wait for the Administration Process to update the Domino Directories on the cluster servers. Although this removes the server from the cluster faster, it can also lead to replication conflicts.

Moving a server from one cluster to another

To move a server from one cluster to another, you simply add the server to the new cluster. Domino automatically removes the server from the old cluster.

For more information, see the topic "Adding a server to a cluster."

Viewing information in the Cluster Database Directory

The Cluster Database Directory contains one document for each database in the cluster. The document tells you such information as the server, path, and replica ID of the database; whether cluster replication is enabled for the database; and whether the database is out of service or pending delete. The Cluster Manager and the Cluster Replicator use this information to determine if the database is available and how to access it during failover and cluster replication, respectively.

You can view the Cluster Database Directory to see what databases are in the cluster, where the databases are located, and the status of the databases.

1. In the Server pane of the Domino Administrator or the Web Administrator, expand All Servers or expand Clusters.
2. Select a server in the cluster.
3. Click the Files tab.
4. In the Task pane, expand Cluster Directory (6).
5. Click one of the four views:

- Databases by Filename
- Databases by Pathname
- Databases by Replica ID
- Databases by Server

6. In the Results pane, open the database document you want.

The following table describes the views in the Cluster Database Directory.

<i>View</i>	<i>Description</i>	<i>Use</i>
Databases by Filename	Shows databases in a cluster sorted by database name	To find which databases are in a cluster and on which servers replicas reside. Using this view is a convenient way to see which databases have cluster replication enabled or disabled and to alter that setting. This is also a convenient view for deleting replicas and creating new replicas
Databases by Pathname	Shows databases in a cluster sorted by the path name relative to the Domino Data directory	To find information about a database when you know the database file name or when there is more than one replica of a database on the same server. You can also use this view to see which databases have cluster replication enabled or disabled.
Databases by Replica ID	Shows databases in a cluster sorted by Replica ID	To find which databases are in the cluster, where the databases are located, and how many replicas of each database exist in the cluster.
Databases by Server	Shows databases in a cluster sorted by the name of the server on which they reside	To find a database when you know where it is located or to find which databases are on each server. You can also use this to determine if the databases on a server are added to the Cluster Database Directory when the server is added to the cluster, and to determine if the databases on a server are removed from the Cluster Database Directory when a server is removed from a cluster.

Creating a new Cluster Database Directory

If the Cluster Database Directory on a server gets deleted, you can create a new one. To do this, follow these steps.

1. Shut down the Cluster Database Directory Manager by sending the command **tell CLDBDIR quit** from the server console.
2. Replicate CLDBDIR.NSF from another cluster server to the server that does not have a replica of the database.

3. Restart the Cluster Database Directory Manager by sending the command **load CLDBDIR** from the server console.

It is important that you create the replica before starting the Cluster Database Directory Manager. Otherwise, the Cluster Database Directory Manager will create a new Cluster Database Directory and add to it a document for each database on the server. However, the existing Cluster Database Directories on the other cluster servers already contain a document for each database on this server. Therefore, after these Cluster Database Directories replicate, they would contain two documents for every database on this server. Although these duplicate documents are deleted in subsequent replications, it can be confusing to view the Cluster Database Directory while it contains so many duplicate documents.

NOTES.INI settings for managing a cluster

The following table describes the NOTES.INI settings that pertain to managing a cluster.

For more information on these settings, see the book *Administering the Domino System*.

<i>Setting</i>	<i>Description</i>
Server_Availability_Threshold	Specifies the lowest acceptable availability index
Server_MaxUsers	Specifies the maximum number of active users allowed on a server at one time
Server_Restricted	Prevents users from accessing a server
Cluster_Replicators	Specifies the number of Cluster Replicators to run
Disable_Cluster_Replicator	Stops all Cluster Replicators

Chapter 6

Clustering Domino Servers That Run Internet Protocols

This chapter describes how to configure and use the Internet Cluster Manager (ICM) to extend the benefits of Domino clusters to HTTP clients and Domino Web servers. It also describes an alternative method of using failover with HTTP servers, as well as methods to use failover on POP3, IMAP, and LDAP servers.

How the Internet Cluster Manager works

The Internet Cluster Manager (ICM) lets you use Domino clusters to provide failover and workload balancing to HTTP clients (Internet browsers) when they access Domino Web servers. This makes your Web servers and databases highly available to clients. You can run the ICM on a Lotus Domino 6 Enterprise server, a Lotus Domino 6 Utility server, or a Domino Release 5 Enterprise Server. You install and configure Domino clusters as you normally would, and then you configure the ICM. The ICM supports the HTTP and HTTPS protocols.

The ICM acts as an intermediary between HTTP clients and the Domino Web servers in a cluster. When Domino Web servers are running in a cluster, they generate URLs that direct HTTP client requests to the ICM. The ICM maintains information about the availability of servers and databases in the cluster. When the ICM receives a client request, it redirects the client to the most available server that contains a replica of the requested database.

The ICM sends periodic probes to the Web servers in the cluster to determine their status and availability. When the ICM receives a client request, it looks at the information in the Cluster Database Directory to find a server that contains the requested database. The ICM determines the most available server that contains the requested database and then redirects the client to that server. This results in the client closing the session with the ICM and opening a new session with the selected server. The user may see this as a change in the host name in the URL. The user may also see the path to the database change in the URL because the database may have a different path on the target server.

If the page that a Web server displays to a client includes links to other databases, the Web server includes the host name of the ICM in the URLs to those databases in the following instances:

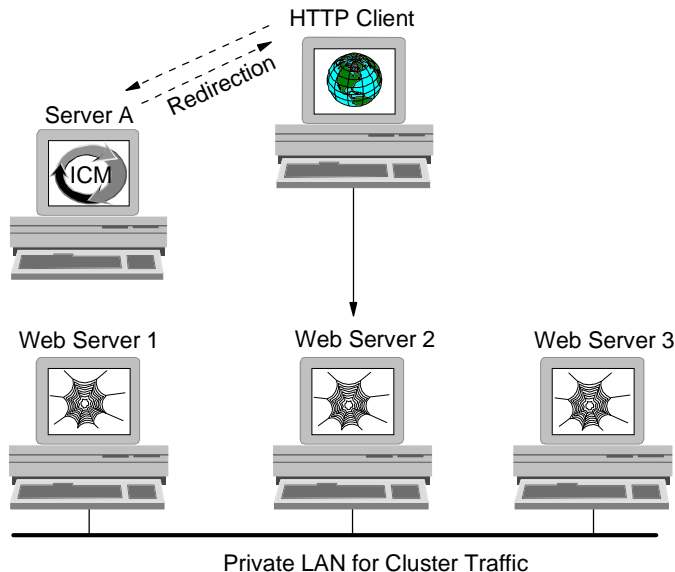
- When generating URLs to databases on the same server as the original database
- When generating URLs to databases on different servers if there are replicas of those databases on the server that contains the original database

This ensures that users accessing those links go through the ICM.

Note In cases not mentioned above, you can use the Redirect URL command to create links to other servers.

For information about the Redirect URL command, see the book *Administering the Domino System*.

The following figure shows an HTTP client asking the ICM to open a database, and the ICM redirecting the client to the best server that contains the requested database, Server 2. The client then connects directly to Server 2.



The ICM can run on a server in the cluster or outside the cluster. When the ICM runs on a server in the cluster, it accesses the local copy of the Cluster Database Directory. When the ICM runs on a server outside the cluster, it selects a server in the cluster and accesses the Cluster Database Directory on that server. If the server that the ICM selects becomes unavailable, this connection fails over to another server in the cluster.

The ICM always uses its local copy of the Domino Directory. Therefore, the ICM must be in the same Domino domain as the cluster.

Performance considerations

In most cases, users will experience better performance when you use the ICM. The overhead of using the ICM is very small, but the benefit to performance from workload balancing can be significant. In cases where the workload was already balanced, there will not be a significant increase or decrease in performance.

Generating URLs that refer to the ICM

The Domino Web server in a cluster reads the Server document to find the host name of the ICM. The Web server then generates URLs that refer to the ICM. The ICM accepts and processes all URLs supported by the Domino Web server. These include URLs that do the following:

- Open servers, databases, and views
- Open forms, navigators, and agents
- Open, edit, and delete documents
- Open documents by name from a view
- Open image files, attachments, and OLE objects
- Create search queries

The following conditions can affect the way the ICM generates URLs:

- If a URL includes a path, the ICM may not be able to resolve this URL to a single database because multiple servers could contain a database with the same path and file name but with different replica IDs. When this occurs, the ICM displays a list of the possible databases from which the user can choose the correct database.
- If a URL includes a replica ID, the ICM may not be able to resolve this URL to a single database if there are multiple replicas of the database on a server, as could be the case with selective replicas. When this occurs, the ICM redirects the client request to a server that contains at least one replica of the requested database, and the Web server selects the replica to present to the user. To access a specific replica, users must specify the path of the replica.
- If a URL includes a NoteID, the ICM may give incorrect results when processing the URL. This is because an object's NoteID may not be the same in all replicas. Unlike a NoteID, a Notes object name and universal ID are identical in all replicas of a database.

Planning to use the ICM

You should plan the cluster that the ICM will service and then plan where to run the ICM itself and how many ICMs to run.

Planning the cluster

You can add Web servers to an existing cluster, use existing servers in a cluster as Web servers, or create a cluster for Web traffic only. You plan the cluster by considering the same factors you consider when you create any Domino cluster. For example, you should consider the processing power of the servers available and how much traffic there will be in the cluster. If there is a lot of Web traffic on a server, you should consider dedicating the server to Web traffic only. You should also distribute databases in a way that balances the workload.

In general, you should consider the following factors:

- The number of servers to include in the cluster
- The number and placement of replicas in a cluster
- How to distribute databases across servers
- Whether to create a private LAN for cluster traffic
- Where to locate the ICM

For more information about planning a cluster, see the chapter “Planning a Cluster.”

Planning the location of the ICM

You can run the ICM on a server in the cluster or on a server outside the cluster. In addition, you can run multiple ICMs to be sure that one ICM is available at all times.

You can configure the ICM in several ways. For example you can:


- Configure a single ICM and place it outside the cluster
- Configure multiple ICMs outside the cluster
- Configure a single ICM as part of the cluster
- Configure multiple ICMs as part of the cluster
- Configure one ICM outside the cluster and one inside the cluster

The following examples describe the benefits of each of these configurations.

Example of a single ICM outside the cluster

When you set up a single ICM outside the cluster, it is a good idea to dedicate the server to the ICM alone. The server should not contain any databases or server tasks other than those required to run the server. Configuring the ICM in this way makes it more reliable because there are fewer activities performed on the server that could interfere with performance and lead to server failure.

The following figure shows an HTTP client accessing a single ICM that resides outside the cluster. The ICM can redirect the client to any of the four Web servers in the cluster.

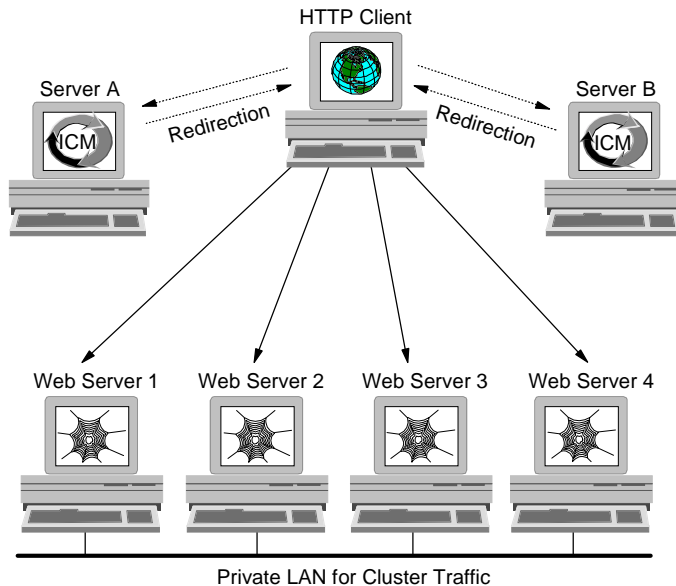


Example of multiple ICMs outside the cluster

You can improve the availability of the ICM by configuring more than one ICM to handle user requests. In this configuration, if one ICM becomes unavailable, the other remains available so that client requests are still handled. You configure both ICMs with the same host name so that if one of the ICMs fails, the other ICM takes over without affecting users.

Although two ICMs are running and handling user requests simultaneously, the real benefit of multiple ICMs is availability, not increased performance. This configuration increases performance significantly only in cases where a single ICM would be overwhelmed by client requests.


The following figure shows an HTTP client with access to two ICMs outside the cluster. Each ICM can redirect the client to any of the four Web servers in the cluster.



Example of a single ICM inside the cluster

You can include the ICM in the cluster by running the ICM on any server in the cluster. Be sure that the server can handle the added traffic that the ICM generates. You should run the ICM on the most powerful server in the cluster or the server that receives the least traffic.


The following figure shows an HTTP client accessing a single ICM that resides on a server in the cluster. The ICM can redirect the client to any of the four Web servers in the cluster.



Example of multiple ICMs inside the cluster

To improve the availability of the ICM, you can install more than one ICM in the cluster. Be sure that the servers on which you install the ICMs can handle the additional traffic.


The following figure shows an HTTP client with access to two ICMs that reside on servers in the cluster. Each ICM can redirect the client to any of the four Web servers in the cluster.



Example of one ICM outside the cluster and one ICM inside the cluster

You can configure one ICM to run outside the cluster and one to run inside the cluster. If the dedicated server outside the cluster becomes unavailable, you have a backup ICM available without having to dedicate a server to the additional ICM.

The following figure shows an HTTP client with access to two ICMs, one outside the cluster and one inside the cluster. Each ICM can redirect the client to any of the four Web servers in the cluster.



Configuring the ICM

You configure the ICM by making entries in the Internet Cluster Manager section of the Server document. You can also set up a separate IP address for the ICM. You can then start the ICM.

You can configure the ICM settings on one server and have more than one ICM access these settings. This lets ICMs on different servers share a common configuration. You should include the ICM configuration information on every Web server in the cluster, not just the server on which you run the ICM. You do this because each Web server uses its own Server document to determine how to generate URLs that refer to the ICM. The Web server obtains the host name of the ICM from the Server document. The Web server then uses that host name to generate URLs that reference the ICM.

To configure the ICM

1. From the Domino Administrator or the Web Administrator, click the Configuration tab.
2. In the Tasks pane, expand Server and click All Server Documents.
3. Do one of the following:
 - In the Results pane of the Domino Administrator, select the Server document for the server on which you want to run the ICM. Then click Edit Server.
 - In the Results pane of the Web Administrator, open the Server document for the server on which you want to run the ICM. Then click Edit Server.
4. Click the Server Tasks - Internet Cluster Manager tab.
5. Complete the following fields:

<i>Section</i>	<i>Field name</i>	<i>Description</i>
Basics	Cluster name	The name of the cluster the ICM will service. If this field is blank, which is the default, Domino uses the name of the cluster that contains this server.
	ICM Notes port	The name of the Domino port the ICM will use to communicate with HTTP clients. If this field is blank, which is the default, the ICM can use any Domino TCP/IP port to communicate with HTTP clients. Enter a port name only if you want to restrict ICM communication to one specific port.
	ICM SSL keyfile	The name of the SSL key file that contains certificates to identify the ICM when communicating with HTTP clients. If this field is blank, the ICM uses the key file specified on the Ports - Internet Ports tab of the Server document.
	Allow users to browse databases in the cluster over HTTP	Lets HTTP clients view a list of the databases in a cluster. When you enable this field, users can enter <code>http://icmhostname/?OpenServer</code> as the URL to access. Entering this URL displays a list of databases on the servers in the cluster associated with the ICM named in <i>icmhostname</i> .

continued

<i>Section</i>	<i>Field name</i>	<i>Description</i>
Configuration	Get configuration from	Lets you specify a different Server document from which to get configuration information. This field lets multiple ICMs share the same configuration.
	Obtain ICM configuration from	This field appears when you select "another server document" in the field "Get configuration from." Enter the name of the server whose Server document contains the configuration you want to use.
	ICM hostname	The fully qualified name of the host that clients should use to communicate with the ICM. This can be the registered DNS name or the IP address. The Domino Web server uses this field to create URLs that reference the ICM. If this field is blank, the Web server will not be able to generate URLs that refer to the ICM.
ICM HTTP Port Settings	TCP/IP port number	Enter the port number for the ICM to use. If you are running the ICM on the same server as the Web server, you must avoid address and port conflicts. If you do not give the ICM its own IP address, be sure the port number the ICM uses is different from any of the other port numbers you use on the server.
	TCP/IP port status	To enable HTTP communication with the ICM, choose enabled. To disable HTTP communication with the ICM, choose Disabled.
	SSL port number	Enter the port number to use for SSL. If you are running the ICM on the same server as the Web server and you do not give the ICM its own IP address, be sure the SSL port number is different from any of the other port numbers you use on the server.
	SSL port status	To enable HTTPS communication with the ICM, choose enabled. To disable HTTPS communication with the ICM, choose Disabled.

6. Click Save & Close to save the Server document.

When the ICM starts, it looks at the Server document on the server on which it is running to find the ICM cluster name and its network address. It then obtains the host name and port settings from the same Server document or from the Server document specified in the field “Obtain ICM configuration from.”

If you run the ICM on the same system as a Domino Web server, you must avoid IP address or port number conflicts. The best approach is to assign the ICM its own IP address. You can also have the ICM share an IP address with the Web server if you specify different port numbers for the ICM and the other protocols on the Web server.

Setting up a separate IP address for the ICM

When you run the ICM on a Web server, you can give the ICM its own IP address to avoid conflicts.

1. Use your operating system to make the IP address available.
2. In the Server pane of the Domino Administrator or the Web Administrator, choose the server that contains the ICM.
3. Set up a new port by doing the following:
 - Click the Server - Status tab.
 - In the Tools pane, expand Server, and then click Setup Ports.
 - Click New.
 - Specify a name for the new port, such as ICMPORT.
 - Choose TCP as the driver.
 - Click OK twice.

Note As an alternative method of setting up the port, you can do the following on the server that contains the ICM.

- Add this line to the NOTES.INI file:

Portname=TCP,adapter number or network number,number of sessions,data buffer size

For example: ICMPORT=TCP,0,15,0

- Add the name of the port, such as ICMPORT, to the Ports setting in the NOTES.INI file.

4. Add the following to the NOTES.INI file:

Portname_TcpIpAddress=0,IPAddress

where *portname* is the name of the new port you configured, such as ICMPORT, and *IPAddress* is the IP address you are using for the ICM.

Example: ICMPORT_TcpIpAddress=0,192.94.222.169

5. From the Domino Administrator or the Web Administrator, click the Configuration tab.
6. In the Tasks pane, expand Server, and click All Server Documents.
7. Do one of the following:
 - In the Results pane of the Domino Administrator, select the Server document for the server on which you want to run the ICM. Then click Edit Server.
 - In the Results pane of the Web Administrator, open the Server document for the server on which you want to run the ICM. Then click Edit Server.
8. Click the Server Tasks - Internet Cluster Manager tab.
9. In the field “ICM Notes port,” enter the name of the port you configured, such as ICMPORT.
10. If you want to use port 80 for both the ICM and the Web server, you must do the following:
 - In the Server document, click the Internet Protocols - HTTP tab.
 - In the “Host name(s)” field, enter the IP address or host name of the Web server.
 - In the “Bind to host name” field, select Enabled.
11. Click Save & Close.

Starting the ICM

You can run the ICM on a single server or on multiple servers to ensure that an ICM is running at all times, even when one ICM is out of service. When you run more than one ICM, you can choose to use the same configuration settings for each ICM. You set this in the “Get configuration from” field on the Server Tasks - Internet Cluster Manager tab in the Server document when you configure the ICM.

For more information about configuring the ICM, see the topic “Configuring the ICM.”

Use one of the following procedures to start the ICM.

Starting the ICM automatically at server startup

To start the ICM automatically each time you start the server, append the ICM keyword to the ServerTasks setting in the NOTES.INI file. For example,

```
ServerTasks=ROUTER,REPLICA,ADMINP,HTTP,ICM
```

Starting the ICM for the current session only

To start the ICM while the server is running, use one of the following methods. The ICM will remain running until the server stops but will not restart automatically the next time the server starts.

From the Domino Administrator or the Web Administrator

1. In the Server pane, expand All Servers or expand Clusters.
2. Select the server you want.
3. Click the Server - Status tab.
4. In the Task pane, do one of the following:
 - From the Domino Administrator, click Server Tasks.
 - From the Web Administrator, click All Server Tasks.
5. In the Tools pane, expand Task, and then click Start.
6. In the Start New Task dialog box, select Internet Cluster Manager (ICM).
7. Click Start Task, and then click Done.

From the server console

Send the following command from the server console.:

```
load icm
```

Failover and workload balancing

When using the ICM, failover and workload balancing work the same as in standard Domino clusters. Domino computes the server availability index based on all open sessions, whether they are from Notes clients, HTTP clients, or other Domino services. To configure workload balancing and failover, you use the same settings, such as `Server_Restricted`, `Server_Availability_Threshold`, and `Server_MaxUsers`. For database availability, you also use the same settings, such as marking a database out of service or pending delete.

Note Unlike in a standard Domino cluster, the ICM can direct a client to a server that is in the `MAXUSERS` state, if no other server is available.

For information about how failover and workload balancing work, see the chapter “How Domino Clustering Works.” For information about configuring failover and workload balancing and database availability, see the chapter “Managing and Monitoring a Cluster.”

The ICM maintains the following information so that it can find a replica when a client asks for one:

- Information about which databases are available in the cluster and where they are stored. The ICM obtains this information from the Cluster Database Directory.
- Information about the availability of each server. The ICM obtains this information each time it probes the servers in the cluster.
- Information about which Web servers are configured for HTTP and which are configured for HTTPS. The ICM obtains this information from the Server documents of each server in the cluster.

To determine which replica of a database to open, the ICM does the following:

- Determines where replicas reside and whether they are marked out of service or pending delete.
- Checks the server availability index of each server that contains a replica.
- Checks the availability of the server by pinging the HTTP port or the HTTPS port, depending on the client request.
- Eliminates any servers that are not reachable or are RESTRICTED.
- Eliminates any servers that are BUSY or in the MAXUSERS state.
- Selects a server from those remaining. If there are no servers remaining, the ICM chooses a server that is BUSY or in the MAXUSERS state, if one is available. If there are multiple servers remaining, the ICM chooses the server with the lightest current workload.

After choosing the server to access, the ICM looks at the Server document to determine which port to use to access the server.

When a server fails

When an HTTP client is connected to a server that fails, the client receives a message saying that the server is not responding. To fail over to a different replica, if one is available, the user must contact the ICM again. The user can do this in the following ways:

- Click the Back button in the browser one or more times to connect to a page through the ICM
- Use a bookmark
- Type the URL

The user may or may not have to reauthenticate with the new server. This is determined by the following factors:

- If the user already authenticated with the new server during this session, no authentication is necessary
- If the HTTP client and the server both support SSL3, reauthentication occurs automatically

Security

The ICM supports SSL. The ICM can use the same SSL certificates that the Domino Web server uses, or you can specify a different set of SSL certificates for the ICM. You configure this on the Server Tasks - Internet Cluster Manager tab of the Server document. The ICM uses settings on the Ports - Internet Ports tab of the Server document to determine the SSL protocol version and whether to accept expired certificates.

In addition, normal Domino server and database security are in effect when using the ICM. The ICM, however, does not participate in the security process. When an HTTP client wants to access a database, it sends an anonymous request to the ICM. The ICM responds by telling the client which server to access. The client then redirects its request to the appropriate server. The server then establishes a dialog with the client and uses whatever security measures are in effect on that server to authenticate the user. If you want to protect the ICM itself from unauthorized access, you can use a firewall or another hardware security system.

For more information about SSL, firewalls, and other network security methods, see the book *Administering the Domino System*.

Managing and monitoring the ICM

The ICM records significant events, such as not being able to find a database, in the Miscellaneous Events view in the log file. You can also view ICM statistics from the Domino Administrator, the Web Administrator, or the server console. You can use these events and statistics, in conjunction with the Cluster Manager and Cluster Replicator events and statistics, to determine which servers in the cluster are the busiest so you can better balance the workload in the cluster.

Viewing the log file

To view ICM events in the log file, do one of the following.

From the Domino Administrator or the Web Administrator

1. In the Server pane, select the server that stores the log file you want to view.
2. Click the Server - Analysis tab.
3. In the Task pane, expand Notes Log, and then click Miscellaneous Events.
4. In the Results pane, open the document you want to view.

From the Domino server log file

1. Open the Domino server log file (LOG.NSF).
2. Open the Miscellaneous Events view.
3. Open the document you want to view.

Viewing ICM statistics

The ICM generates statistics that indicate how the cluster and the ICM are performing. To view these statistics, you can do one of the following.

From the Domino Administrator or the Web Administrator

1. In the Server pane, select the server you want.
2. Click the Server - Statistics tab.
3. In the Statistics list, expand ICM.

From the server console

Type **show stat icm**

To see descriptions of the ICM statistics, as well as the other cluster statistics, see the appendix “Cluster Statistics.”

Monitoring all the ICM servers at the same time

You can use the Domino server monitor to monitor all the ICM servers at once. You can decide what information you want to monitor and how to display that information. You can monitor the ICM servers and the entire cluster while you monitor other Domino servers. To tell Domino which servers to monitor and what information to monitor on each server, you create or customize a monitoring profile.

Note The Domino server monitor and monitoring profiles are not available in the Web Administrator.

Starting the server monitor

You can start the server monitor manually while the server is running, or you can start the server monitor automatically when the server starts.

To start the server monitor manually

1. From the Domino Administrator, click the Server - Monitoring tab.
2. In the “Monitoring profiles” field, select the profile you want.
3. Click the Start button.

Once pressed, the Start button becomes the Stop button.

To start the server monitor automatically when the server starts

1. Choose File - Preferences - Administration Preferences.
2. Click Monitoring.
3. Select “Automatically monitor servers at startup.”
4. Make any other changes you want, and then click OK.

For more information about monitoring Domino servers, see the book *Administering the Domino System*.

Creating or customizing a monitoring profile for the ICM

You can create new profiles and edit existing profiles to customize the tasks and statistics that Domino displays.

1. From the Domino Administrator, click the Server - Monitoring tab.
2. In the “Monitoring profiles” field, select an existing profile.
Selecting a profile initializes the server monitor if it is not already initialized. You cannot make changes to a profile until the server monitor is initialized.
3. To add one or more tasks to monitor, choose Monitoring - Monitor New Task, select the tasks you want to add, and then click OK.
To monitor the ICM task, select “Internet Cluster Manager (ICM).”
4. To add one or more statistics to monitor, choose Monitoring - Monitor New Statistic.
5. In the “Add Statistic(s) to this profile” dialog box, expand ICM, select the ICM statistics you want to monitor, and then click OK. For example, the Command - Redirects statistics and the Requests statistics tell you how busy each ICM is. This can help you to balance the workload.
6. (Optional) To add a server to the profile, select Monitoring - Monitor New Server, and then select the server from the list; or drag a server from the Server pane to the server monitor.
7. (Optional) To remove a server from the list, click the name of the server you want to remove, and then select Monitoring - Remove Server.

8. To save your changes to the profile, do one of the following:
 - To save this profile as a new profile while also preserving the original profile, choose Monitoring - Profiles - Save As, and enter a name for the profile.
 - To have this modified profile replace the original profile, you do not have to do anything. The profile is saved automatically when you close the Domino Administrator.

Compatibility with previous releases of Domino

The ICM directs HTTP client requests to any Domino Web server in a cluster, whether it is a Lotus Domino 6 server or a server using an earlier release of Domino. Therefore, using the ICM in a mixed cluster has no negative effect on clients. Only Lotus Domino 6 and Domino Release 5 Web servers, however, generate URLs that include references to the ICM. When a user selects URLs that contain these references, the ICM is able to balance the workload and redirect the user to an appropriate server.

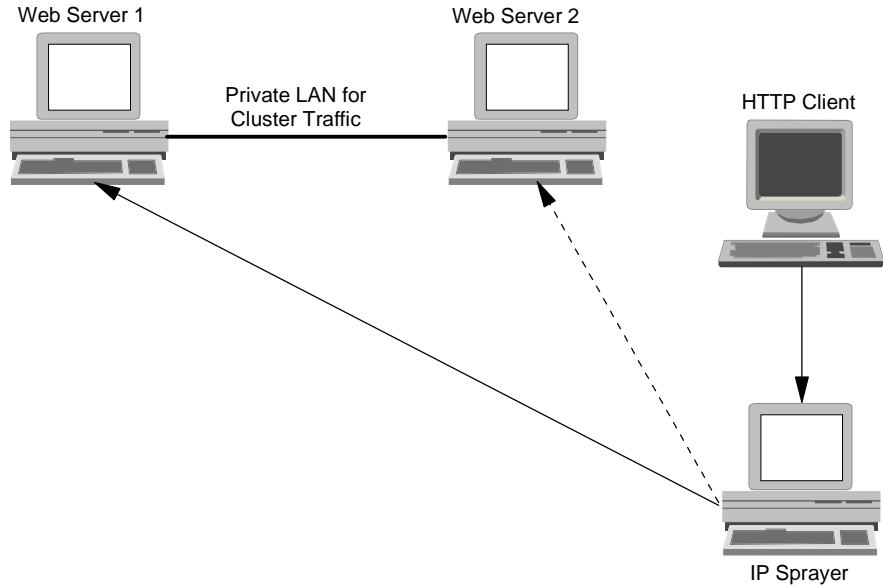
If the ICM redirects the client request to a server running a previous release of Domino, the URLs that server generates do not contain references to the ICM. As the client moves between Web pages, the client continues to access pages on that server only. To take advantage of the ICM at that point, the user must select or enter a URL that contains a reference to the ICM.

Using an IP sprayer with Domino for HTTP and POP3 failover

Instead of using the ICM, you can use an IP sprayer to let HTTP clients access Domino Web servers in a cluster. The client requests go through the IP sprayer, which distributes the requests to available cluster servers that have replicas of the requested databases. If a cluster server is unavailable, the IP sprayer directs the client request to a different cluster server that contains a replica of the requested database.

Using an IP sprayer gives you some distribution of the workload, but it does not use intelligent failover. In other words, the IP sprayer does not know what the load is on the Domino servers so it does not redirect requests to the server with the smallest load. In addition, the IP sprayer cannot obtain information from the Cluster Database Directory. Therefore, you must configure the sprayer so that each server it looks on for a particular database contains a replica of that database.

In the following figure, an IP sprayer directs an HTTP client to Web Server 1. If Web Server 1 is not available, the IP sprayer directs the client to Web Server 2, which contains a replica of the database the client requested. You can configure the IP sprayer to alternate between Web Server 1 and Web Server 2 when directing HTTP client requests. This helps to balance the workload. If either server becomes unavailable, the IP sprayer directs all requests to the server that is still available.

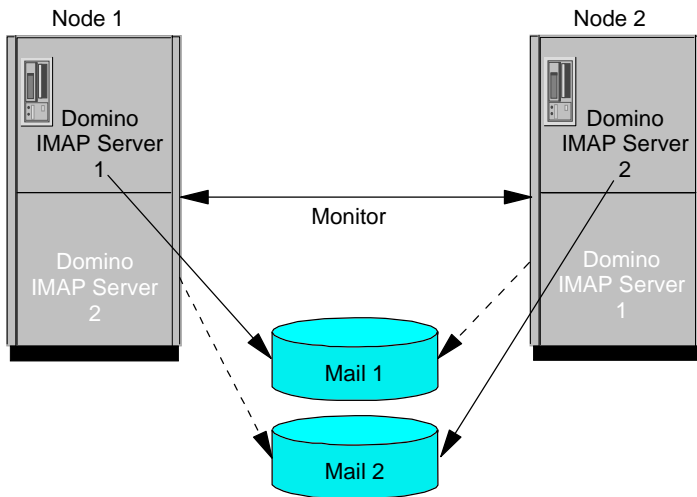


You can also use an IP sprayer for failover with POP3 clients.

Setting up failover for IMAP

Because IMAP requires unique identifiers for mail messages, you cannot use replicas of a database with IMAP. You have to access the same database each time. Therefore, you should use operating system clusters if you want to set up failover for IMAP.


The following figure shows an active-active operating system cluster. Each node has an IMAP server running in one Domino partition. The other partition contains a copy of the resources of the active partition on the other server. If either node fails, the other node takes over and runs both IMAP servers in different partitions.



Using an IP sprayer with Domino for LDAP failover

You can cluster LDAP servers to increase the availability of data. One way to do that is to create a Domino cluster of LDAP servers, and then use an IP sprayer to direct client requests.

In the following figure, there are two LDAP servers in a Domino cluster. Both servers contain a replica of the LDAP directory. An IP sprayer directs an LDAP client to LDAP Server 1. If LDAP Server 1 is not available, the IP sprayer directs the client to LDAP Server 2.



You can also use an operating system cluster for LDAP failover.

If you are using directory assistance, you can use either directory assistance failover or Domino failover.

For more information on configuring directory assistance to fail over, see the book *Administering the Domino System*.

Appendix A

Cluster Statistics

There are three categories of cluster statistics that help you analyze clusters:

- Cluster Manager statistics
- Cluster Replicator statistics
- Internet Cluster Manager statistics

This appendix describes these statistics.

For more information about viewing statistics, see the book *Administering the Domino System*.

Cluster Manager statistics

Other than AvailabilityIndex, AvailabilityThreshold, and ExpansionFactor, each statistic name includes the prefix Server.Cluster, which is not included in the names in this table. AvailabilityIndex, AvailabilityThreshold, and ExpansionFactor include the prefix Server.

<i>Statistic name</i>	<i>Description</i>
Member	Shows the names and the availability indexes of the servers in the cluster
AvailabilityIndex	A measure of a server's availability. 0 (zero) indicates no available resources; 100 indicates complete server availability
AvailabilityThreshold	The current availability threshold of the server. When the availability index is less than or equal to the availability threshold, the server is put into the BUSY state.

<i>Statistic name</i>	<i>Description</i>
ExpansionFactor	The current expansion factor. This value is used to compute the availability index. The values range from 1 to 64. A value of 1 indicates that the server is completing transactions at the minimum time for that server. A value of 64 indicates that it is taking 64 times longer than the minimum time to complete transactions. An expansion factor of 64 is considered fully loaded, and results in an availability index of 0 (zero).
OpenRedirects.Failover.Successful	Total times that server successfully redirects a client to another cluster member after the client fails to open a database by replica ID
OpenRedirects.Failover.Unsuccessful	Total times that server is unable to redirect a client to another cluster member after the client fails to open a database by replica ID
OpenRedirects.FailoverByPath.Successful	Total times server successfully redirects a client to another cluster member after the client fails to open a database by path name
OpenRedirects.FailoverByPath.Unsuccessful	Total times server is unable to redirect a client to another cluster member after the client fails to open a database by path name
OpenRedirects.LoadBalance.Successful	Total times server successfully redirects a client to another cluster member after the client tries to open a database by replica ID when the server is BUSY
OpenRedirects.LoadBalance.Unsuccessful	Total times server is unable to redirect a client to another cluster member after the client tries to open a database by replica ID when the server is BUSY
OpenRedirects.LoadBalanceByPath.Unsuccessful	Total times that server is unable to redirect a client to another cluster member after the client tries to open a database by path name when the server is BUSY
OpenRedirects.LoadBalanceByPath.Successful	Total times that server successfully redirects a client to another cluster member after the client tries to open a database by path name when the server is BUSY
OpenRequest.ClusterBusy	Total client requests when all servers are BUSY

continued

<i>Statistic name</i>	<i>Description</i>
OpenRequest.DatabaseOutOfService	Total times a client tries to open a database that is marked out-of-service on the server
OpenRequest.LoadBalanced	Total times a client tries to open a database on the server when the server is BUSY
PortName	Default port used for intra-cluster network traffic or an asterisk, which indicates there is no default port and any available active port can be used
ProbeCount	Total times that a server completes a probe of the other cluster members
ProbeError	Total times that a server receives an error when probing another server
ProbeTimeout(mins) *	Shows the interval at which an intracuster probe occurs

* This statistic does not appear in the Cluster statistics report, but you can use the Show Stat command to view it or you can add it to the Cluster statistics report form.

Cluster Replicator statistics

Each statistic name includes the prefix Replica.Cluster, which is not included in the names in this table.

<i>Statistic name</i>	<i>Description</i>
Servers	Total servers in the cluster that are receiving replications from this server
Successful	Total successful replications since the server started.
Failed	Total failed replications since the server started
Docs.Added	Total documents added by the Cluster Replicator
Docs.Updated	Total documents updated by the Cluster Replicator
Docs.Deleted	Total documents deleted by the Cluster Replicator
Files.Local	Number of databases on the current server for which there are replicas on other servers in the cluster
Files.Remote	Number of databases on other servers to which the Cluster Replicator pushes changes

continued

<i>Statistic name</i>	<i>Description</i>
Retry.Skipped	Total times the Cluster Replicator did not attempt to replicate a database. The retry is skipped when the destination server is known to be unreachable or the database is waiting for another retry attempt.
Retry.Waiting	Total number of replicas that are waiting for retry attempts
SecondsOnQueue	Total time, in seconds, that the last database replicated spent on the work queue
SecondsOnQueue.Avg	Average time, in seconds, that a database spent on the work queue
SecondsOnQueue.Max	Maximum time, in seconds, that a database spent on the work queue
SessionBytes.In	Total bytes received during cluster replication
SessionBytes.Out	Total bytes sent during cluster replication
WorkQueueDepth	Current number of databases awaiting replication by the Cluster Replicator
WorkQueueDepth.Avg	Average work queue depth since the server started
WorkQueueDepth.Max	Maximum work queue depth since the server started

Internet Cluster Manager statistics

The following statistics help you track the performance of the ICM.

<i>Statistic name</i>	<i>Description</i>
ICM.AvailabilityIndex	Shows the names and the availability indexes of the servers in the cluster
ICM.Command.Total	The number of URL commands the ICM received
ICM.Command.Unknown	The number of URL commands the ICM did not recognize
ICM.Command.Redirects.Successful	The number of times the ICM successfully redirected a client URL request to a cluster member
ICM.Command.Redirects.Unsuccessful	The number of times the ICM could not redirect a client URL request to a cluster member
ICM.Command.Redirects.ClusterBusy	The number of times the ICM received a client request when all servers were BUSY
ICM.Receive.Error	The number of times the ICM could not process a client request because of a communication problem between the client and the ICM

continued

<i>Statistic name</i>	<i>Description</i>
ICM.Requests.Per1Hour.Total	The number of HTTP requests the ICM received in the past hour
ICM.Requests.Per1Minute.Total	The number of HTTP requests the ICM received in the past minute
ICM.Requests.Per5Minutes.Total	The number of HTTP requests the ICM received in the past 5 minutes
ICM.Server.Running	Tells whether the ICM task is running
ICM.Sessions.Inbound.Accept.Queue	The number of new connections that have been detected and are waiting to be serviced by a server thread
ICM.Sessions.Inbound.Active	Current number of inbound connections
ICM.Sessions.Inbound.Active.SSL	Current number inbound connections that are SSL connections
ICM.Sessions.Inbound.BytesReceived	Total number of bytes received by all inbound connections since the server started
ICM.Sessions.Inbound.BytesSent	Total number of bytes sent by all inbound connections since the server started
ICM.Sessions.Inbound.Peak	The maximum number of concurrent inbound connections since the server started
ICM.Sessions.Inbound.SSL	The maximum number of concurrent inbound SSL connections since the server started
ICM.Sessions.Inbound.Total	Total number of inbound connections since the server started
ICM.Sessions.Inbound.Total.SSL	Total number of inbound SSL connections since the server started
ICM.Sessions.Outbound.Active	Current number of outbound connections
ICM.Sessions.Outbound.Active.SSL	Current number outbound connections that are SSL connections
ICM.Sessions.Outbound.BytesReceived	Total number of bytes received by all outbound connections since the server started
ICM.Sessions.Outbound.BytesSent	Total number of bytes sent by all outbound connections since the server started
ICM.Sessions.Outbound.Peak	The maximum number of concurrent outbound connections since the server started
ICM.Sessions.Outbound.Peak.SSL	The maximum number of concurrent outbound SSL connections since the server started

continued

<i>Statistic name</i>	<i>Description</i>
ICM.Sessions.Outbound.Total	Total number of outbound connections since the server started
ICM.Sessions.Outbound.Total.SSL	Total number of outbound SSL connections since the server started
ICM.Sessions.Threads.Busy	Number of threads that are currently busy
ICM.Sessions.Threads.Idle	Number of threads that are currently idle
ICM.Sessions.Threads.InThreadPool	Current number of threads created
ICM.Sessions.Threads.Peak	Maximum number of threads since the server started

Appendix B

Troubleshooting Cluster Problems

This appendix includes information about error messages and other problems that may arise when using a cluster.

Some database changes are not replicating quickly to other servers.

Use the following checklist to troubleshoot problems related to slow replication between servers in a cluster. Check that the following conditions exist:

- The Cluster Replicator is started on the server where the modified database is located. You can check this on the Servers - Status tab of the Domino Administrator or by sending the following command from the server console:

```
show tasks
```
- The modified database and its replicas on other servers are listed in all the Cluster Database Directories.
- All replicas of the modified databases have the same Replica ID. To check this, open the Databases by Replica ID view in the Cluster Database Directory.
- The Cluster Replicator is not encountering errors when it attempts to replicate to other servers in the cluster. Check the `Replica.Cluster.Failed` and `Replica.Cluster.Retry.Waiting` statistics to see if error conditions exist. Also, examine the Replication Events log documents generated by the Cluster Replicator.
- The Cluster Replicator is able to keep up with the current server replication workload. Check the `Replica.Cluster.WorkQueueDepth` and `Replica.Cluster.SecondsOnQueue` statistics to determine if there is a backlog of replication requests. If so, consider starting an additional Cluster Replicator.
- Cluster replication is enabled for all replicas of the database. Open the Cluster Database Directory, and check the left column for the letter "X." Databases with the letter "X" in the left column have cluster replication disabled.

- CLREPL_OBEYS_QUOTAS is set to 0 (zero) if you used this setting in the Configurations Settings document or in the NOTES.INI file. When this setting is set to 1, cluster replication obeys the database size quotas that were set by the database administrator. The Cluster Replicator will not push changes to a replica if the changes would result in the replica exceeding its size quota. If CLREPL_OBEYS_QUOTAS is set to 0 or is not present at all, the Cluster Replicator ignores database size quotas.

Client requests do not fail over for certain databases even though the replicas are listed in the Cluster Database Directory

When there are two or more replicas of a database on a server, the Cluster Manager uses failover by path, not failover by replica ID. To ensure that client requests fail over correctly, do not include multiple replicas of a database on the same server; or if you do, create the replicas using the same names and paths as the replicas to which you want to fail over.

Although I marked a database Out of Service, users can still open it

The database that users can open may be a replica on a non-cluster server. Marking databases “out of service,” “in service,” or “pending delete” works only for databases that are in a cluster. In addition, new replicas of the databases do not inherit these attributes, even in a cluster.

The Cluster Database Directory includes two copies of the database documents for all the databases on a particular server

If the Cluster Database Directory on a server is deleted, the Cluster Database Directory Manager recreates it and then populates it with a document for each database on the server. These documents then replicate to the other servers in the cluster. Since each server in the cluster already has documents for this server’s databases, their Cluster Database Directories will then contain two documents for each database on this server. This is a temporary condition and causes no system errors. The next time the server’s Cluster Database Directory Manager starts, it detects the problem and removes the extra documents.

To avoid creating duplicate documents, replicate the Cluster Database Directory from another server to the server on which the Cluster Database Directory was deleted before you restart the server.

The value of the Replica.Cluster.Retry.Waiting statistic is greater than zero

This statistic indicates that the Cluster Replicator could not complete some replications, and these replications are waiting to be retried. To see why the replications were not successful, force the Cluster Replicator to generate a Replication Event Log document, which includes information about all the cluster replications waiting to be retried. To force the Cluster Replicator to log this information, send the following command from the server console:

```
tell clrepl log
```

To view the error conditions, examine each of the Log documents generated by this command (one for each server being replicated to), and then correct the errors. You can sometimes correct the problem by restarting a server that is currently unreachable. When the errors are corrected, cluster replication succeeds, and the Replica.Cluster.Retry.Waiting statistic becomes zero.

“Cluster Replicator was unable to configure using Cluster Database Directory cldbdir.nsf: File does not exist”

This message can occur for the following reasons:

- The Cluster Replicator cannot find the Cluster Database Directory. This often occurs when you first add a server to a cluster and the Cluster Replicator starts before the Cluster Database Directory Manager has created the server’s Cluster Database Directory. If this is the cause of the problem, it will resolve itself.
- The Cluster Database Directory has been deleted. Replicate the Cluster Database Directory from a different cluster server.

“Cluster Replicator was unable to configure using Cluster Database Directory cldbdir.nsf: Invalid replica ID for cluster database directory. If cluster name changed, delete cluster database directory and restart cldbdir task.”

The ClReplID field in the Server document in the Domino Directory does not match the replica ID of the Cluster Database Directory. To fix this, you can delete the Cluster Database Directory from this server and then replicate it from a different cluster server. If this doesn’t correct the problem, remove the server from the cluster and add it to the cluster again.

Private folders do not replicate from one clustered database to another

Check the access control list (ACL) of the databases to be sure that the User type of the servers is “Server” or “Server group.”

“HTTP Server Initialization error. Could not bind port 80. Port may be in use.”

This message can occur if you have conflicting IP addresses or port numbers when you attempt to start the Domino Web server on a server that is running the ICM. The most likely cause is that the ICM and HTTP task are both attempting to use the same IP address and TCP/IP port.

Check the Server document to ensure that the ICM and the HTTP server have been assigned different TCP/IP port numbers or that the ICM is configured to use a different IP address than the HTTP server. If the ICM and HTTP server are both using port 80, but on different IP addresses, make sure that you have chosen “Enabled” in the “Bind to host name” field on the Internet Protocols - HTTP tab in the Server document.

The Server Web Navigator does not fail over

Check the replica IDs of the Web databases (WEB.NSF) on the servers in the cluster. If the replica IDs are not identical, the databases will not fail over to each other. To fix this problem, replicate the Web database from one server to all the other servers in the cluster.

This problem often occurs because the Web task creates the Web database when it first starts, if this database doesn't already exist. Therefore, it is a good idea to start the Web task on one cluster server only, and then replicate the Web database to the other cluster servers before you start the Web task on those servers.

Clients receive the message “Server Not Responding” instead of failing over

If a server becomes unavailable while a database is open, failover does not occur. The user should reopen the database. Reopening the database causes failover to occur if a replica exists on an available server in the cluster. If the user was editing a document when the server became unavailable, the user can copy the document to the replica.

“Server Not Responding” can also appear when a user tries to send and save a message when the user’s mail server is unavailable. The message is sent successfully because the mail router fails over. (The user can see that the message was sent successfully by clicking the status history list in the status bar.) However, saving a message or document does not cause failover. To save the message, the user can reopen the database, which causes failover if a replica mail database exists on an available server. The user can then copy the sent message to the replica.

Index

A

- ACL
 - enforcing on replicas, 4-7
- Activity Trends
 - clusters and, 5-26
- Administration Process
 - adding servers, 5-46
 - creating clusters with, 4-3
 - removing servers with, 5-47 to 5-48
- Availability index
 - described, 5-18
 - expansion factor and, 5-19
 - viewing, 5-7
- Availability threshold
 - described, 5-20
 - server, 5-18
 - setting, 5-20, 5-22, 5-24
 - viewing, 5-7
 - workload balancing and, 2-5

B

- BUSY state
 - servers, 2-5, 5-18, 5-22
- BUSYTIME.NSF
 - free time database, 2-13

C

- Calendar and scheduling
 - clusters and, 2-13
- CLUBUSY.NSF
 - free time database, 2-13
- Cluster Administrator
 - described, 2-9
- Cluster Analysis
 - described, 4-5
 - results, 4-7
 - running, 4-6
- Cluster Analysis Database
 - described, 4-5
- Cluster Database Directory
 - creating, 5-50
 - described, 2-8, 5-49

- troubleshooting, B-2
- updating information stored,
 - 5-40
- viewing, 5-49
- Cluster Database Directory Manager
 - described, 2-9
- Cluster Manager
 - described, 2-7
 - failover and, 2-2
 - locating replicas, 2-5
 - monitoring events, 5-5
 - removing server documents, 5-48
 - statistics, A-1
 - workload balancing, 2-5
- Cluster Name Cache
 - described, 2-8
- Cluster names
 - replication and, 4-16
- Cluster replication
 - configuring, 5-30
 - described, 2-10, 5-30
 - disabling, 5-33, 5-35
 - pausing, 5-37
 - resuming, 5-38
- Cluster Replicator
 - backlog, 5-11
 - described, 2-9, 5-30
 - log file, 5-8, 5-40
 - replication history, 2-11
 - retrying failed replications, 5-39
 - statistics, A-1, A-3
 - Tell commands, 5-11, 5-13
 - troubleshooting, B-3
 - updating Cluster Database
 - Directory information, 5-40
 - using multiple, 5-31, 5-33
- Clusters
 - adding servers, 5-45 to 5-46
 - analyzing, 4-5 to 4-6
 - backup servers, 5-29
 - balancing workload, 5-17
 - benefits of, 1-1
 - calendar management in, 2-13
 - components of, 2-7
 - configuring, 3-1, 3-12 to 3-22, 5-41

- Connection documents, 4-16
- controlling access to, 4-7, 5-24
- controlling database
 - availability, 5-42
- creating, 4-1 to 4-4
- creating mail databases in, 4-12
- database distribution in, 3-7, 3-9, 5-26
- database information, 2-9
- database size quotas, 4-17
- deleting databases from, 5-44
- described, 1-1
- Directory assistance
 - database, 4-19
- disabling replication, 5-33, 5-35
- disaster preparedness, 3-18
- displaying members, 5-2
- failover, 2-2 to 2-5, 5-27
- fault recovery in, 3-11
- hardware for, 3-2
- HTTP servers in, 6-1, 6-19
- IMAP servers in, 6-21
- LANs and, 3-10, 4-23, 4-25
- LDAP servers in, 6-22
- log file, 5-5, 5-7
- logging replication events, 4-17
- mail databases in, 3-7, 2-12
- mail routing failover, 4-10
- memory requirements, 3-2
- mobile user access to, 4-21
- monitoring, 5-1, 5-15 to 5-16
- moving servers, 5-49
- NOTES.INI settings, 4-26, 5-51
- operating system clusters, 2-14, 3-22, 6-21
- partitioned servers in, 3-19 to 3-20
- pausing replication, 5-37
- planning, 6-4
- removing servers, 5-47 to 5-49
- replacing servers, 5-29
- replica distribution in, 2-1, 3-3
- replicating folders, 2-11, 4-8
- replication, 2-9, 4-15
- replication history, 2-11
- replication schedule, 4-13

- replication types, 5-30
- requirements, 1-3
- resuming replication, 5-38
- roaming users, 4-19
- server availability, 5-18
- shared mail, 2-12, 4-11
- statistics, 5-3, 5-15 to 5-16, A-1, A-4
- WANs and, 3-11, 3-18, 5-24
- Web databases in, 4-23
- Web servers in, 6-1, 6-19
- workload balancing, 2-5, 5-11, 5-17

Configuring

- clusters, 3-1, 3-12 to 3-22, 5-30, 5-41
- Internet Cluster Manager, 6-5 to 6-13

Connection documents

- replicating, 4-16

D

Data synchronization

- described, 1-2

Database access

- controlling, 4-9
- managing, 4-7

Database quotas

- obeying in a cluster, 4-17

Databases

- controlling access to, 4-7
- deleting from a cluster, 5-44
- disabling replication, 4-14
- distributing, 3-3, 3-7, 3-9, 3-13 to 3-15
- failover causes, 2-3
- information about, 2-8
- making available for access, 5-43
- making unavailable for access, 5-42
- marking in service, 5-43
- marking out of service, 5-42

Directories

- finding, 4-19

Directory assistance

- clusters and, 4-19

Disaster preparedness

- clusters and, 3-18

Drop All command

- restricting servers with, 5-29

Duplicate documents

- avoiding creation of, B-2

E

Examples

- active-active operating system cluster, 2-16
- active-passive operating system cluster, 2-15
- cluster configuration, 3-12 to 3-22
- cluster for disaster preparedness, 3-18
- Cluster Manager events, 5-5
- Cluster Replicator events, 5-9
- Domino and operating system clusters together, 3-22
- failover, 2-2
- hub server cluster, 3-16
- IMAP server cluster, 6-21
- Internet Cluster Manager configuration, 6-5 to 6-13
- IP sprayer and Web servers, 6-19
- LDAP server cluster, 6-22
- partitioned server cluster, 3-19 to 3-20
- passthru server cluster, 3-21
- workload balancing, 2-6

Expansion factor

- described, 5-18

F

Failover

- causes of, 2-3
- causing, 5-24, 5-27
- Cluster Manager and, 2-8
- described, 1-1 to 2-2
- HTTP servers, 6-14, 6-19
- IMAP servers, 6-21
- Internet Cluster Manager, 6-14
- IP sprayer and, 6-19, 6-22
- LDAP servers, 6-22
- log file and, 5-5
- mail databases, 2-12
- operating system clusters, 2-14
- troubleshooting, B-2
- Web servers, 6-14, 6-19

Fault recovery

- clusters and, 3-11
- enabling, 3-12

Folders

- replicating, 2-11, 4-8
- troubleshooting, B-4

Free Time database

- described, 2-13

H

Hardware

- cluster, 3-2

HTTP clients

- Internet Cluster Manager and, 6-1
- security, 6-16

HTTP servers

- failover, 6-14, 6-19
- troubleshooting, B-4

Hub servers

- configuring, 3-16

Hub-and-spoke topology

- clusters and, 3-16

I

IBM Tivoli Analyzer for Lotus

- Domino tool
- clusters and, 5-26

IMAP

- failover, 6-21

Internet

- security, 6-16

Internet Cluster Manager

- compatibility with Domino releases, 6-19

- configuring, 6-5 to 6-13

- described, 6-1

- failover, 6-14

- IP address, 6-9, 6-12

- location of, 6-4

- log file, 6-16

- performance, 6-3, 6-17

- planning, 6-4

- security, 6-16

- starting, 6-13

- statistics, 6-17, A-1, A-4

- troubleshooting, B-4

- URLs, 6-3, 6-9, 6-19

- workload balancing, 6-14

IP sprayer

- failover, 6-19, 6-22

L

LANs

- setting up in a cluster, 3-10, 4-23, 4-25

LDAP service

- failover, 6-22

M

Mail databases

- distributing, 3-7, 3-13, 3-15, 4-12
- failover, 2-12

Mail files

- clusters and, 4-12

Mail routing

- failover, 4-10

Memory requirements

- for clusters, 3-2

Mobile users

- cluster access for, 4-21
- replication and, 4-22

Monitoring

- cluster workload, 5-1, 5-15 to 5-16
- events and statistics, 5-5
- statistics, 5-15 to 5-16

N

Notes client

- requirements for cluster, 1-4

NOTES.INI file

- Internet Cluster Manager variable, 6-13
- workload balancing variables, 2-5

NOTES.INI settings

- clusters, 4-26, 5-51

O

Operating system clusters

- benefits of, 2-17
- described, 2-14
- failover, 2-14
- IMAP and, 6-21
- with Domino clusters, 3-22

P

Partitioned servers

- clustering, 3-19 to 3-20

Passthru servers

- configuring, 3-21
- mobile user access with, 4-21

R

Replica IDs

- troubleshooting, B-3 to B-4

REPLICA task

- using, 2-10

Replicas

- access rights to, 4-7

ACLs and, 4-7

- creating, 3-4 to 3-6, 4-9

- distributing, 2-1, 3-3

- moving, 5-26

- troubleshooting, B-2

Replication

- backlog, 5-11

- cluster names and, 4-16

- clusters and, 2-9, 4-15, 5-30

- disabling, 4-14, 5-33, 5-35

- displaying status messages, 4-17

- error information, 2-10

- folders, 2-11

- formulas, 4-14

- history, 2-11

- LANs and, 4-25

- logging events, 4-17

- mobile users and, 4-22

- pausing, 5-37

- resuming, 5-38

- retrying failed, 5-39

- statistics, 5-7, 5-10

- troubleshooting, B-1

- types of, 5-30

- Web databases, 4-23

Replication Log

- error information, 2-10

Response time

- server, 5-18

Roaming users

- change from nonroaming, 4-20

- clusters and, 4-19

- registering, 4-20

S

Scalability

- clusters and, 1-1

Schedule Manager

- described, 2-13

Scheduling information

- clusters and, 2-13

Secure Sockets Layer. See SSL, 6-16

Security

- Internet, 6-16

Selective replication formulas

- processing, 2-11, 4-14

Server availability index

- described, 5-18

- expansion factor and, 5-19

- viewing, 5-7

Server availability threshold

- described, 5-20

- setting, 5-22 to 5-24

- viewing, 5-7

Server documents

- Internet Cluster Manager settings, 6-9

Server Health Monitoring

- clusters and, 5-26

Server statistic collection

- creating, 5-3

Server Web Navigator

- starting in a cluster, 4-23

Servers

- adding to clusters, 5-45 to 5-46

- availability index, 5-7, 5-18 to 5-20, 5-22

- availability threshold, 5-7, 5-18, 5-20 to 5-24

- backing up, 5-29

- BUSY state, 2-5, 5-18, 5-22

- cluster requirements, 1-3

- clustering, 3-1, 3-12 to 3-22

- failover causes, 2-3, 5-27

- log file, 5-5, 5-7

- removing from cluster, 5-47 to 5-49

- restricting access, 5-27, 5-29

Shared mail

- clusters and, 2-12, 4-11

Size quotas

- cluster replication and, 4-17

SSL

- Internet security and, 6-16

Standard replication

- described, 5-30

- server failure and, 4-13

Statistic Collector

- Cluster Manager statistics, 5-6, 5-10

- starting, 5-4

Statistics

- cluster, 5-3, A-1, A-4

- Cluster Manager, 5-6, 5-10

- Internet Cluster Manager, 6-17

Statistics Reporting database

- Cluster Manager statistics, 5-6, 5-10

- replication statistics, 5-10

T

Tell commands

- Cluster Replicator, 5-11 to 5-15, 5-37 to 5-41

Tivoli Analyzer. See IBM Tivoli
Analyzer for Lotus
Domino, 5-26

Troubleshooting

- Cluster Database Directory, B-2
- Cluster Replicator, B-3
- failover, B-2
- folder replication, B-4
- HTTP servers, B-4
- Internet Cluster Manager, B-4
- IP addresses, B-4
- replica IDs, B-3 to B-4
- replication, B-1, B-2
- Web Navigator, B-4

U

URLs

- Internet Cluster Manager and,
6-3, 6-9, 6-19

W

WANs

- using in a cluster, 3-11, 3-18, 5-24

Web databases

- replicating, 4-23

Web Navigator

- troubleshooting, B-4

Web servers

- failover, 6-14, 6-19
- security, 6-16

Workload balancing

- Activity Trends and, 5-26
- Cluster Manager and, 2-8
- clusters and, 2-5, 5-11, 5-17
- controlling cluster access, 5-24
- database distribution for, 3-7, 3-9
- Internet Cluster Manager, 6-14
- log file and, 5-5
- moving replicas, 5-26
- multiple Cluster Replicators,
5-31, 5-33
- NOTES.INI file variables, 2-5
- Server Health Monitoring
and, 5-26
- servers and, 5-17 to 5-26