


Lotus Domino Integration Guide for IBM Netfinity Servers

Rufus Credle, Paul Branch, Darryl Miles



International Technical Support Organization

<http://www.redbooks.ibm.com>



International Technical Support Organization

SG24-2102-02

**Lotus Domino Integration Guide for
IBM Netfinity Servers**

December 1998

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix F, "Special Notices" on page 177.

Third Edition (December 1998)

This edition applies to Lotus Domino Release 4.6 Server that is shipped with IBM Netfinity servers for use with Microsoft NT Server, OS/2 Warp Server and Novell NetWare network operating systems.

Comments may be addressed to:
IBM Corporation, International Technical Support Organization
Dept. HZ8 Building 678
P.O. Box 12195
Research Triangle Park, NC 27709-2195

When you send information to IBM, you grant IBM a non-exclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© **Copyright International Business Machines Corporation 1997 1998. All rights reserved.**

Note to U.S. Government Users — Documentation related to restricted rights — Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

Contents

Preface	vii
The Team That Wrote This Redbook	vii
Comments Welcome	ix
Chapter 1. Introduction	1
1.1 How to Use This Integration Guide	1
Chapter 2. Domino Server and Network Design	3
2.1 The Role of Domino and Notes	3
2.2 Domino Services	3
2.2.1 Web Server	4
2.2.2 Electronic Mail	4
2.2.3 Application Development	4
2.2.4 Calendar and Scheduling	5
2.2.5 Replication	5
2.2.6 SMTP MTA	5
2.2.7 POP3 Support - Mail Support for POP3 Clients	5
2.2.8 Domino.Action	5
2.3 Domino Server Packaging	6
2.3.1 Lotus Domino 4.6 Server Mail Access	6
2.3.2 Lotus Domino 4.6 Server for Single Processor	6
2.3.3 Lotus Domino 4.6 Server for Multiprocessor	6
2.3.4 Lotus Domino 4.6 Advanced Services	6
2.4 Starting Your Domino Solution	7
2.4.1 Operating System	7
2.4.2 Network Protocol	7
2.4.3 Naming Conventions	9
2.4.4 Security Considerations	10
2.4.5 Remote Access (Dial-In)	11
2.5 Your Domino Network Gets Bigger	15
2.5.1 Wide Area Networking Issues	16
2.5.2 Replication Schedules	20
2.6 Hub-and-Spoke Network	21
2.7 Other Considerations	25
2.7.1 The Size of Your Domain	25
2.7.2 Systems Management for Your Domino Network	27
Chapter 3. Sizing IBM Netfinity Servers for Domino	29
3.1 Which Platform Is Right	29
3.2 IBM Netfinity for Lotus Domino Servers	31
3.3 IBM Netfinity Family of Servers	33
3.4 About IBM Netfinity	33
3.4.1 The Foundation for Lotus Domino and Domino Applications	34
3.5 IBM Netfinity 3000	36
3.6 IBM Netfinity 5000	38
3.7 IBM Netfinity 5500 and 5500 M10	40
3.8 IBM Netfinity 7000 M10	44
3.9 Enterprise Storage Technology	46
3.9.1 IBM Hard Disk Drives	46
3.9.2 IBM Tape Backup Solutions	46
3.9.3 IBM RAID for Netfinity Servers	47

3.9.4	Fibre Channel: The Future of Storage...Today	47
3.9.5	Tape Backup Expansion Solutions	48
3.9.6	SCSI Hard Disk Expansion Units	48
3.9.7	SSA Storage Expansion Units	49
3.9.8	IBM Rack Solutions	49
3.10	Domino Servers by Function	49
3.10.1	Mail Servers	50
3.10.2	Database (or Application) Servers	50
3.10.3	Web Servers	51
3.10.4	Remote Access Servers (RAS)	51
3.10.5	Hub Servers	52
3.10.6	Passthru Servers	52
3.11	NotesBench Reports	54
3.11.1	Audited Information	54
3.11.2	Interpreting a NotesBench Report	54
3.11.3	Determine Your Client Distribution	56
3.12	Server.Planner	57
3.12.1	Where to Get Server.Planner	58
3.12.2	Getting Started with Server.Planner	58
3.12.3	Query on Data	59
3.12.4	Query Results	60
3.12.5	Graphical Information	61
3.12.6	Saving Queries and Results	62
3.12.7	Standardizing Your Server Hardware	62
3.13	Other Server Requirements and Recommendations	63
3.13.1	Server RAM	63
3.13.2	Server Disk Space	64
3.13.3	RAID and Drive Controllers	66
3.13.4	Processor	67
3.13.5	Network Card	67
3.13.6	Uninterruptable Power Supply	67
3.13.7	Disaster Recovery Software and Virus Protection	68
3.13.8	Recommendations from the Field	68
3.14	Examples of Domino Server Selections	69
3.14.1	Small Sized Domino Environment	69
3.14.2	Medium Sized Domino Environment	70
3.14.3	Large-Sized Domino Environment	72
3.14.4	Company Intranet Environment	75
	Chapter 4. Server Configuration and Optimization	79
4.1	System Installation	80
4.1.1	Disk Partition Configuration	80
4.1.2	Operating System and Service Pack Installation	81
4.1.3	Network Card Driver Installation	82
4.1.4	Domino and Service Pack Installation	82
4.2	Operating System Optimization	82
4.2.1	Windows NT Server	83
4.2.2	Warp Server	87
4.3	Protocol Optimization	88
4.3.1	TCP/IP	89
4.3.2	IPX/SPX	89
4.3.3	NetBIOS	89
4.4	Domino Optimization	90
4.4.1	Ways to Modify NOTES.INI	90
4.4.2	Server Tasks	91

4.4.3 Database Indexing	92
4.4.4 Minimize the Notes Log File	92
4.4.5 Maximum Sessions and Session Timeouts	93
4.4.6 Maximum Users	94
4.4.7 Maximum Mail Threads	95
4.4.8 Number of Databases Cached	96
4.4.9 Multiple Replicators	96
4.4.10 Minimum Mail Poll Time	96
4.4.11 Memory Buffers	97
4.5 Using the Windows NT Performance Monitor	97
4.5.1 Starting Performance Monitor	97
4.5.2 Collecting Data with Performance Monitor	99
4.5.3 Chart View	99
4.5.4 Alert View	103
4.5.5 Log View	105
4.5.6 Report View	108
4.5.7 Installing Domino Statistics As a Performance Monitor Counter	108
4.6 Server Subsystems to Monitor	109
4.6.1 Memory Bottlenecks	112
4.6.2 Processor Bottlenecks	114
4.6.3 Disk Bottleneck	115
4.6.4 Network Bottleneck	116
4.6.5 Other Domino Systems to Monitor	116
4.7 Domino Performance Analysis Example	117
4.7.1 Normal Server Operation	117
4.7.2 Memory Bottleneck Example	120
4.7.3 Monitoring Interrupts, Network Interface and Disk Subsystem	122
4.7.4 CPU Bottleneck Example	122
Chapter 5. Lotus Domino Clustering	125
5.1 Introduction to Clustering	125
5.2 Lotus Domino Server Clustering	126
5.3 Comparing Traditional Notes Replication Versus Cluster Replication	127
5.4 Failover and Availability	130
5.5 Planning and Administration	131
5.6 Network Traffic Considerations	133
5.7 Cluster Setup Strategies	134
5.8 Domino Cluster Planning Guidelines	135
5.9 Domino Cluster Installation and Configuration Steps	136
5.9.1 Installation of the NT Domino Server	137
5.9.2 Configuring the NT Domino Server	139
5.9.3 The First NT Domino Server Startup	141
5.9.4 Registering an Additional Server - the OS/2 Domino Server	142
5.9.5 Configuring a Domino Cluster	143
5.9.6 Verifying the Domino Cluster	144
5.9.7 Restarting the NT Server	144
5.9.8 Installing the OS/2 Domino Server	146
5.9.9 Changes to a Cluster Server's NOTES.INI File	148
5.10 Recovering from Failures	149
5.11 How to Set Up Failover and Workload Balancing to Work Together	149
5.12 Summary	149
Appendix A. Top 10 Ways to Improve Your Server Performance	151
Appendix B. What's New in Domino and Notes Release 5	155

B.1 Lotus Domino R5	155
B.2 Lotus Notes R5	157
B.3 Domino Designer R5	158
Appendix C. Domino and the Internet	161
C.1 The Internet and the World Wide Web	161
C.2 Types of Internet Connections	162
C.2.1 Dial-Up Connection	162
C.2.2 Leased-Line Connection	163
C.2.3 Proxy Connection	164
C.3 Extending Your Notes Network Over the Internet	165
C.3.1 Putting a Domino Server on the Internet	165
C.3.2 Connecting Servers to Servers over the Internet	165
C.3.3 Using NotesNIC to Communicate with Other Domino Servers	167
Appendix D. The Tivoli Enterprise Module for Domino/Notes	169
D.1 Overview of the Tivoli Enterprise Module for Domino/Notes	169
D.1.1 Security Enhancements	170
D.1.2 Task Administration	170
D.1.3 Availability Management	171
D.1.4 Automated Notes Client Installation	172
D.1.5 User Management with the Admin Extension	172
D.2 The Tivoli Enterprise Module for Domino/Notes Module for Domino/Notes Managed Environment	173
Appendix E. Lotus Notes NetBIOS Port Drivers	175
Appendix F. Special Notices	177
Appendix G. Related Publications	179
G.1 International Technical Support Organization Publications	179
G.2 Redbooks on CD-ROMs	179
G.3 Other Publications	179
How to Get ITSO Redbooks	181
How IBM Employees Can Get ITSO Redbooks	181
How Customers Can Get ITSO Redbooks	182
IBM Redbook Order Form	183
Index	185
ITSO Redbook Evaluation	187

Preface

The purpose of the *Lotus Domino Integration Guide for IBM Netfinity Servers* is to provide a resource for describing the roles and capabilities of the IBM Netfinity server platform running within the Lotus Domino environment. The book provides Notes network design and capacity planning recommendations in the following areas:

- Operating system
- Network protocol
- Server placement within your network
- Wide area network issues
- Replication schedules
- Server selection
- Server RAM
- Server disk space
- Server peripherals

Suggested installation and optimization settings for your server are discussed, along with methods to investigate and solve performance bottlenecks on your server.

The Lotus Domino clustering technology is also outlined, which provides high availability and capacity solutions. The simplicity with which these clusters can be configured, thus providing users not only with increased response time, but also constant availability, is a key aspect of the Domino strategy and so is covered here in some depth.

This redbook will help those individuals who need to plan, tune, or size Lotus Domino for the IBM Netfinity servers through detailed examples and scenarios. Particular attention has been paid to the proper size and performance characteristics of various IBM Netfinity servers for disparate environments. Several tools that can aid in the sizing of your servers are discussed, as well as industry-standard "rules of thumb" for configuring small, medium, and large-sized Domino environments.

Some knowledge of Intel-based servers, Lotus Notes, Lotus Domino and network operating systems is assumed.

The Team That Wrote This Redbook

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

Rufus Credle is an Advisory Software Engineer for IBM Netfinity servers at the International Technical Support Organization, Raleigh Center. He conducts residencies and writes redbooks about IBM Netfinity servers and network operating systems. Rufus's various positions during his IBM career included assignments in administration and asset management, system engineering, marketing and services. He holds a BS degree in Business Management from Saint Augustine's College.

Paul Branch is a Staff Technical Marketing Support Specialist for IBM Netfinity servers in the United States. His areas of expertise include IBM Netfinity Servers, IBM PC Servers, Lotus Notes and OS/2 Warp Server. He has developed education courses on OS/2 LAN Server, OS/2 Warp Server and Communications Manager/2.

Darryl Miles is a Systems Specialist at IBM Australia. He has worked at IBM for four years supporting various network operating systems, groupware and electronic mail. He is a Microsoft Certified Systems Engineer, a Certified Novell Engineer (CNE), Warp Server Engineer and Certified Lotus Professional (CLP). He holds an honors degree in Electrical and Computer Systems Engineering from Monash University, Australia. Currently, his main focus is on Lotus Domino, Internet integration and Java.

Thanks to the following people for their invaluable contributions to this project:

Carol Zimmet and James Grigsby, Comms Development Group
Iris Cambridge

Wynndel Burns, CLP, Intranet Manager
Insite Clinical Trials Inc, Cincinnati, OH

Deborah Drakeford, IBM Netfinity Alliance Manager for Lotus
IBM Somers

Sam Emrick, OS/2 External Performance
IBM Austin

Ramsesh Kalkunte, President
Technovations International Incorporated

Dave Laubscher, Manager, IBM Netfinity Technology Center
IBM Raleigh

John Lamb
Advantis

Clement Moy, Lotus Notes Server Performance
IBM Research Triangle Park

Ayudh Nagara, IBM Networking
IBM Australia

Nelson Scott and John McNulty
FCP Technologies

Dana Thompson
KMDS Technical Associates

Andrew Wainwright, Domino Services
IBM United Kingdom

David Watts, IBM Netfinity Specialist
Linda Robinson, Graphic Design
Shawn Walsh, Gail Christensen and Kathryn Casamento, Editors

Debbie Hester, Nancy Lewis
International Technical Support Organization, Raleigh Center.

Martin Obianwu, Roy Bowen and Leisha Penny
IBM Lotus Integration Center - Dallas

Comments Welcome

Your comments are important to us!

We want our redbooks to be as helpful as possible. Please send us your comments about this or other redbooks in one of the following ways:

- Fax the evaluation form found in “ITSO Redbook Evaluation” on page 187 to the fax number shown on the form.
- Use the online evaluation form found at <http://www.redbooks.ibm.com/>
- Send your comments in an Internet note to redbook@us.ibm.com

Chapter 1. Introduction

The information in this redbook is a compilation of experiences gathered from IBM, Lotus, Microsoft, IRIS, KMDS and other business partners and is meant to reflect how an IS professional, manager, technical engineer or sales engineer may approach server selection and implementation for a Domino network environment.

This redbook is meant to provide general direction and guidelines for the tasks of server selection and implementation. Sizing figures are based on audited laboratory tests called *NotesBench* which measures items such as server response time and capacity under quantitative user workload scripts. These user connection figures are meant to approximate total users typically supportable by these average configurations. As is always the case with any laboratory testing, these results should be used as indicators only; actual performance in real end-user installations will vary.

Lotus released a tool called *Server.Planner*, which allows you to query *NotesBench* data based on your requirements, and then will select the appropriate server(s) that will meet this load. *Server.Planner* is available from <http://notes.net> and ships with Lotus Domino Version 4.6. The IBM vendor database is used with the *Server.Planner* tool to select the appropriate IBM Netfinity servers with Domino. The IBM vendor database and *NotesBench* performance reports are available from <http://www.notesbench.org>.

1.1 How to Use This Integration Guide

The flexibility and power of Lotus Domino have led end users to continually push the envelope in how they use Domino and to find new ways to use Domino technology. For that reason, Domino implementations take on many different forms. To start planning your Domino network, sketch out your existing infrastructure on paper. Place Domino servers in your topology as we expand an example solution from a single Domino server to an enterprise design. We discuss the various alternatives for your Domino network; think these over and discuss them with your colleagues and management. Depict where the users are located on your network and what connections they have to the LAN, WAN, Inter/intranet and how those networks are configured. Documenting these factors and reviewing them as you read through this guide will assist you in determining the type of server(s) and platform(s) that may be appropriate for your Domino environment.

Chapter 2. Domino Server and Network Design

Designing a Domino infrastructure can seem like a daunting task. Most Domino networks consist of several distributed Domino servers, each running various services and supporting different workgroups of varying size. There are many questions that you are confronted with, and it may be difficult to know where to start. In this chapter, we first describe the features that Lotus Domino 4.6 offers, and how the software is purchased in one of three different packages. We then start with just a single Domino server and address what issues this presents to a designer and implementer. We expand this example to include remote access, a remote site, and a reasonably sized organizational structure. When the design of your Notes network is complete, and the location and configuration of your servers is decided, you can sufficiently size your Domino servers with the correct hardware. Those details are discussed in Chapter 3, "Sizing IBM Netfinity Servers for Domino" on page 29.

2.1 The Role of Domino and Notes

Lotus Domino 4.6 is an application and messaging server that enables a broad range of secure, interactive business solutions for the Internet and intranet. With Domino, developers can rapidly build, deploy and manage applications that engage co-workers, partners, and customers in online collaboration, cooperation, and coordination of critical business activities.

In 1996, Lotus developed a component program which you could load on your Notes server called HyperText Transfer Protocol (HTTP). This task allowed you to connect to your Notes server using a standard Web browser. This HTTP task was code-named *Domino* during its development.

This new functionality allowed Notes servers to do much more than groupware and electronic mail services had in the past. Lotus then decided to "rebrand" the Notes server and new HTTP task as one complete package called Lotus Domino Server 4.5, Powered by Notes or simply Domino. Clients that connect to this Domino server are Notes clients, Web browsers, and alternate mail clients.

2.2 Domino Services

What is Lotus Domino? The product inherently means different things to different people. For many, Domino is the server that provides them with electronic mail across their organization, or to the Internet. For others, Domino provides them with shared databases that allow their employees to work together better in teams, and discuss new ideas and strategies across disparate locations. More recently, people have come to consider Domino as a powerful dynamic Web server, which allows them to view and modify Notes database information via a simple Web browser. Domino is all of the above services and more.

Domino is best explained as the integration of a Lotus Notes server and a Web server, allowing people using Web browsers to view and change information in Notes applications dynamically. Previously, these applications were only accessible using Notes clients. Now with Lotus Domino, all the groupware applications written for Notes may be delivered to the World Wide Web (WWW).

2.2.1 Web Server

The Web server component of Domino is comprised of two subcomponents. The first is a standard Web server. A standard Web server consists of a server engine which listens for HyperText Transfer Protocol (HTTP) requests from Web browsers. It services these requests by transmitting the requested document to the client in HyperText Markup Language (HTML). A client chooses the document they want to view on a server by using a Universal Resource Locator or URL. An example URL is the IBM World Wide Web site at <http://www.ibm.com/>. Many Web sites are referenced throughout this redbook for you to locate further information.

HTML is a simple tag-based language used for creating documents on the Web. These HTML and graphics files are stored on the Web server in directories. A Web server services Web clients just like a file server services LAN clients (for example, Warp Server, Windows NT or NetWare). This standard Web server is called the Internet Connection Server (ICS) which can display Web documents to any HTML 3.0 compatible browser that supports tables.

It is the second component of Domino that is the most exciting. The Domino engine interfaces between the Notes subsystem and the Web server. This engine translates Notes elements such as documents, forms, views and action buttons into HTML format for a Web browser dynamically. What this means is that data in a Notes database is viewed, modified and created using a standard Web browser on-the-fly, and there is no need for any HTML programming whatsoever. All of Notes strengths such as security, replication, full text searching, agents and application development now apply to the Web site. Creating a Notes document is as easy as using a wordprocessor; it is now that easy to create pages of information on the Web.

2.2.2 Electronic Mail

Notes Mail is the enterprise mail system that is built into Domino. The Notes Mail interface looks like the popular cc:Mail user interface, and offers all of the same base functions as cc:Mail and Microsoft Mail. Notes Mail offers additional functionality such as enterprise calendaring and scheduling, OLE 2.0 support, platform-independent file viewers, the Personal Web Navigator, task management, document libraries, personal journal, phone messages, and unparalleled mobile user support. You can also access your mail on a Domino server using a Web browser, Netscape or Microsoft Mail client.

2.2.3 Application Development

Domino is packaged with a series of application templates to help customers rapidly develop and deploy their Domino-based applications. Application templates include user registration management, threaded discussion and Web-based mail access. Additionally, Lotus announced a line of Net.Apps, which that are packaged applications you can load on your Domino server. You can also write your own Notes databases, but if you don't want to do that you can contact a Lotus Business Partner who can develop an application for you. There are also thousands of Notes and Domino applications already written. You may also wish to subscribe to the free *Lotus Solutions Now* magazine, which is published quarterly. This magazine details new applications that have been developed and are being used by other customers. For more details see <http://www.lotus.com>.

2.2.4 Calendar and Scheduling

Domino now features a rich set of native calendaring and scheduling functionality that is powerful and intuitive for end users, flexible for application developers and scalable to even the largest enterprise. Users gain secure and real-time look-ups of their colleagues' free and busy times. The new calendaring and scheduling feature, modeled on the user interface of the award-winning Lotus Organizer, also leverages services such as replication so users can keep calendars synchronized while they are on the road.

2.2.5 Replication

Domino uses Notes' robust replication facilities to synchronize mirrored sites and distributed intranets. Replication is used to copy and merge database information from server to server, and is also used for mobile access to Notes. This allows Domino to provide services to users regardless of their location. For example, a company can replicate its databases onto an IBM ThinkPad, which can then be disconnected from the network or dial-up connection.

2.2.6 SMTP MTA

The Lotus Notes Message Transfer Agents (MTAs) provide reliable, standards-based messaging backbones for SMTP networks, as well as the Internet, by extending native protocol and directory support to the Domino server. Unlike gateways and connectors in competitive messaging systems such as Microsoft Exchange, the Lotus Notes MTAs are suitable for both simple and sophisticated networks. In more sophisticated networks, the Lotus Notes MTAs are configured to perform fault-tolerant routing so that if one communications path fails another is automatically selected. The MTAs provide this without any degradation of performance or fidelity that result with products that require format translations when switching protocols. In addition, because the MTAs reside on the Notes server, they also leverage the robust multilevel security features of the Notes environment including end-to-end message encryption.

2.2.7 POP3 Support - Mail Support for POP3 Clients

Post Office Protocol (POP3) clients can download mail messages and may be configured to delete or save messages following download. Domino servers' POP3 support includes password level authentication. With this additional capability, Domino servers now support MAPI mail clients, POP3 mail clients and Notes clients. Some examples of these clients are Eudora and Netscape.

2.2.8 Domino.Action

Domino.Action is an architectural framework for rapidly and easily creating and managing an extensible intranet. Domino.Action contains the site creator and eleven template-based modules that allow a webmaster to create a functionally rich site to which others can add content with or without third-party authoring tools. Domino.Action employs Domino's Web application server technology and Lotus Notes to help you deploy your interactive, business process solutions, and bring them to the Internet.

2.3 Domino Server Packaging

Lotus Domino is available in four separate packages to fit different customer needs. Each package is available on multiple operating system platforms.

2.3.1 Lotus Domino 4.6 Server Mail Access

The introduction of the Lotus Domino 4.6 server has given users the ability to use a broader set of clients to use mail services for sending and receiving mail from the Domino Server. Domino Mail Access allows customers using non-Notes mail clients to access Domino 4.6 servers for mail. The Domino Mail Access license fee entitles users to create and use the Domino Server over any of the protocols that allow a client to access that mail file, including, for example; POP3, MAPI and HTTP.

2.3.2 Lotus Domino 4.6 Server for Single Processor

The Domino Server for Single Processor machines provides Domino functionality workgroups and departments. It extends the ability of users to conduct business on the Web through enhancements in key areas such as messaging, scalability, manageability and security, in addition to support of Internet standards such as HTTP, POP3, SMTP, and Secure Sockets Layer (SSL) 2.0. Domino 4.6 also enables developers to rapidly design business applications for both Web browsers and Notes clients, and to seamlessly integrate those applications with existing corporate systems.

Several enhancements have been incorporated into Domino to provide you with rich features. A complete list of enhancements to Domino 4.6 is found at the What's New in the Domino 4.6 Server page, located at <http://www3.lotus.com/home.nsf>.

The enhancements provide organizations with the only server solution that can deliver collaborative groupware applications capable of merging the power of existing corporate systems with the open standards of the Internet and intranets.

A no charge license of Domino -- Single Processor Edition is currently shipping across the IBM Netfinity servers product line.

2.3.3 Lotus Domino 4.6 Server for Multiprocessor

The Domino Server for Multi-processor machines exploit symmetric multiprocessing (SMP). Multiprocessor support allows more simultaneous user sessions per server and overall increased capacity for Domino server functions such as mail routing and replication. Domino accommodates over 2,000 concurrent users per server running on the Intel platform.

2.3.4 Lotus Domino 4.6 Advanced Services

Lotus Domino Advanced Services is an optional capability that can be added to the single or SMP version of the Domino 4.6 server.

The capabilities of the Domino Advanced Services are:

- Clustering, including fail-over and load-balancing - Clustered servers include a minimum of two and up to six interconnected servers, which allow for the gradual increase of capacity by easily adding servers to the cluster as workload and user demand increase. Clusters also ensure high availability by automatically switching to another server in the event the first server

goes down. Clustering allows customers to achieve efficient use of server resources and better user response times. Customers can balance the workload across all servers in the cluster by adjusting the availability threshold of each server. Through its load balancing and fail-over capabilities, the cluster software shifts users attempting to access a database on a loaded (busy) cluster member to a replica on an available cluster member.

- Partitioning - A single physical server can host up to six distinct Domino servers while still providing the same level of security and reliability as if each were hosted on a stand-alone server.
- Billing - The billing feature enables a Domino server to track specific Domino activities. The billing server task collects this information and records the data for billing purposes.

2.4 Starting Your Domino Solution

Implementing a Domino solution on even one or two servers confronts the designer with many of the issues faced in an enterprise Domino network. These initial issues are which operating system, network protocol(s), and security considerations should be used. We discuss each in turn, and provide recommendations when applicable.

2.4.1 Operating System

Domino can run on a variety of operating system platforms such as Windows NT Server, OS/2 Warp Server, several UNIX flavors including AIX, NetWare and Windows 95. This gives you the choice to run Domino on an operating system that you and your colleagues are familiar with. For example, if you have many NT clients, then it is wise to install an NT Domino server. Likewise, if you are using OS/2 Warp, installing an OS/2 Domino server is preferred. If you have a large network of NetWare servers and cannot afford new Domino servers, running Domino on NetWare may be appropriate. This ensures a common environment for you to work in.

It is then important to choose a robust, multithreaded operating system with symmetrical multiprocessor (SMP) capability. For most Domino sites, we discourage the use of Domino running on Windows 95 servers, unless you have a small client base and the cost to upgrade to a more robust platform is prohibitive.

We recommend that for a new, or large-scale Domino infrastructure using the Intel platform, that the Windows NT Server or OS/2 Warp Server operating system be chosen.

2.4.2 Network Protocol

Many networks today, even small LANs, are beginning to use TCP/IP as their primary protocol. The popularity of TCP/IP is due in part to the increased popularity of the Internet, and the many advantages TCP/IP has over existing proprietary protocols. If TCP/IP is available on your network, it should be the first choice for your Domino servers above all other protocols.

For many small networks, the predominant protocol used between servers and clients was not chosen for its technical merit. In most cases the protocol was immaterial, and people used whatever protocol came as standard with the file

and print services operating system they chose. For example, Windows NT Server and OS/2 Warp Server use the NetBIOS protocol as standard. Novell NetWare networks use the IPX/SPX protocol.

With DOS/Windows workstations, it was very difficult to load multiple protocol stacks on one workstation without facing memory problems and other issues. So most clients had to use the same protocol for different applications. Much of this complexity of multiple protocol stacks is now gone with Windows and OS/2 workstations. However, it is still good practice to use existing protocols used on the network for your Domino server. This has two benefits:

- Reduces the amount of different traffic on your network.
- Ensures that any networking problems have been resolved when running file and print services, such as issues with routers and bridges.

However, you must consider that NetBIOS should only be used for small Domino server implementations due to an inherent limitation in the driver on the OS/2 and Windows NT platforms. This limit only allows approximately 250 users per network card on the server.

	TCP/IP	IPX/SPX	NetBIOS	VINES	AppleTalk	X.PC
OS/2 Warp	2000+ ¹	Limited only by server capacity	120/250 ²	50	50	Up to 64 ports
Windows NT	2000+ ¹	Limited only by server capacity	254	Limited only by server capacity	Limited only by server capacity	Up to 64 ports

Because of the popularity of TCP/IP stemming from the Internet, the following are many advantages this protocol has over the existing proprietary protocols:

- TCP/IP is a routeable protocol, whereas NetBIOS is not. This is important if your Domino network expands in the future to include remote sites connected together using routers.
- TCP/IP is built on connectionless technology where information is transferred as a sequence of datagrams, and the network doesn't know that there is any connection between two computers. This makes applications less prone to session timeouts, and provides you a more robust Domino network.
- The dynamic Web server component of Domino only uses the TCP/IP protocol. This is because Web browsers generally require the TCP/IP protocol. If you wish to set up a company-wide intranet, you will need to use the TCP/IP protocol in your network.
- The advanced automatic fail-over and workload balancing features of Domino clustering are only available for TCP/IP clients connected to Domino servers.

¹ 2,000 active sessions have been tested. The actual limit is a function of the server capacity.

² The Notes NetBIOS.dll driver is limited to approximately 120 users. The Notes IBMEENB.dll driver supports approximately 250 users. See Appendix E, "Lotus Notes NetBIOS Port Drivers" on page 175 for a discussion of these two Notes port drivers.

2.4.3 Naming Conventions

All Domino servers and Notes users have a hierarchical name based on the X.500 naming standard. This naming scheme uses a tree structure that reflects the structure of the company using Notes. For example, the following is the hierarchical name of Joe Smith who works in the accounting department of the XYZ organization:

Joe Smith/Accounting/XYZ

Common Name CN = Joe Smith, Organizational Unit = Accounting,
Organization = XYZ

Hierarchical names ensure name uniqueness. For small Notes networks, there is never really a need to have multiple levels in your naming structure. Most companies work with only the Common Name (CN) and Organization (O) level.

In most small Notes networks, you will see names with formats similar to the following:

Jane Wilson/Acme

ServerOne/Acme

Notice how there is no organizational unit for this company Acme, whereas for the XYZ company, Joe Smith was located under the Accounting organizational unit. Organizational units are normally implemented if your Notes network is to be administered at multiple locations, as it provides much tighter security.

If you are implementing an enterprise Notes network, you must do some up-front planning to avoid any work needed to change the naming scheme in the future. These names for servers and people are stored in the Notes Public Name and Address book (NAMES.NSF). This is the heart of Domino security and services, because it lists all people who have access to your Notes domain, and also the configuration of each server. Please refer to pages 20-21 of the Lotus Domino 4.6 *Getting Started with Domino* manual for more details on designing a hierarchical naming scheme.

We recommend that you name users with their full name (for example, Peter Harrison), and not cryptic user IDs or abbreviations (for example, PHARRISO). This is because people's names are automatically searched when you address Notes mail to people. You can distinguish between users of the same name in the same location by middle initials.

Server names should be short (8 - 10 characters), and try to identify their function and location. That way a user knows if a server is located in their own location, or across the globe. An example server naming standard is provided:

RRRCCNN RRR = Server Role
 CCC = City Location of Domino Server
 NN = Number of server

For example:

DEVRAL01 - Development Domino server in Raleigh

GENDAL02 - Second general Domino server in Dallas

2.4.4 Security Considerations

Lotus Domino has three primary layers of security for governing access to the server and the associated databases on the server:

1. Authentication
2. Server access control lists (ACLs)
3. Database access control lists (ACLs)

The authentication and validation process occurs whenever a Notes client, workstation or server attempts to communicate with the Domino server. Requests for communication could be for replication from one server to another, or simply database access by a user. If you wish to know all aspects of Domino security, make sure you read the redbook *The Domino Defense: Security in Lotus Notes and the Internet*, SG24-4848. Some initial security issues you should consider are:

1. First, physically secure your Lotus Notes servers. All IBM servers have a power-on and privileged access password, keyboard password, mechanical lock, selectable boot, tie-down via U-bolt, keyboardless and displayless operation, optional security cover, and an optional tamper-evident cable kit.
2. Ensure your servers are located in a locked, air conditioned computer room.
3. Enable lockup passwords under OS/2, enable a simple Windows NT screen saver with password protection enabled, or monitor lockup passwords under NetWare.
4. Use the SET SECURE command on your Domino server. This password protects the console so you cannot use the LOAD, TELL, EXIT, QUIT or SET CONFIGURATION commands.
5. If your server is providing file and print services as well as Domino, ensure that access to the \NOTES directory and subdirectories has not been given to file server users. This will bypass Notes security completely.
6. Only allow access to the \NOTES directory for the administrator user under Windows NT.
7. Ensure the default access for the Public Names and Address book (NAMES.NSF) on each server is set to Author. But do not allow the user to be able to perform the following:
 - a. Create documents
 - b. Delete documents
 - c. Create personal agents
 - d. Create personal folders/views
 - e. Create shared folders/views
 - f. Create LotusScript agents

You can modify these to suit your needs at a later stage. Set Manager access for Administrators.

8. Take multiple backups of the CERT.ID and SERVER.ID files. Store backups of these files in a very secure location. Also take a copy of these files offsite. Remember that anyone who has access to these files has full access to your Domino network.

9. If you are using Windows NT and not using file and print services on your server, you may wish to turn off the Computer Browser service so your server is not visible to Windows 95 or NT clients using the Network Neighborhood program.
10. Create a group in the Name and Address book called Administrators. Add all administrators to this group. Set the following fields in each server document:

Restrictions			
Server Access	Who can -	Passthru Use	Who can -
Only allow server access to users listed in this Address Book:	No	Access this server:	
Access server:		Route through:	
Not access server:		Cause calling:	
Create new databases:	Administrators	Destinations allowed:	
Create replica databases:	Administrators		

Figure 1. Restrictions Section of Server Document

This sets a tighter restriction on each server so that only administrators can create new and replica copies of databases on your server.

2.4.5 Remote Access (Dial-In)

After your single site is running successfully with Domino, many customers then introduce remote access services for their employees. They often introduce this at a later stage so that they have given their staff sufficient time and training to get familiar with the Notes client. The illustration in Figure 2 on page 12 allows remote access from several different locations.

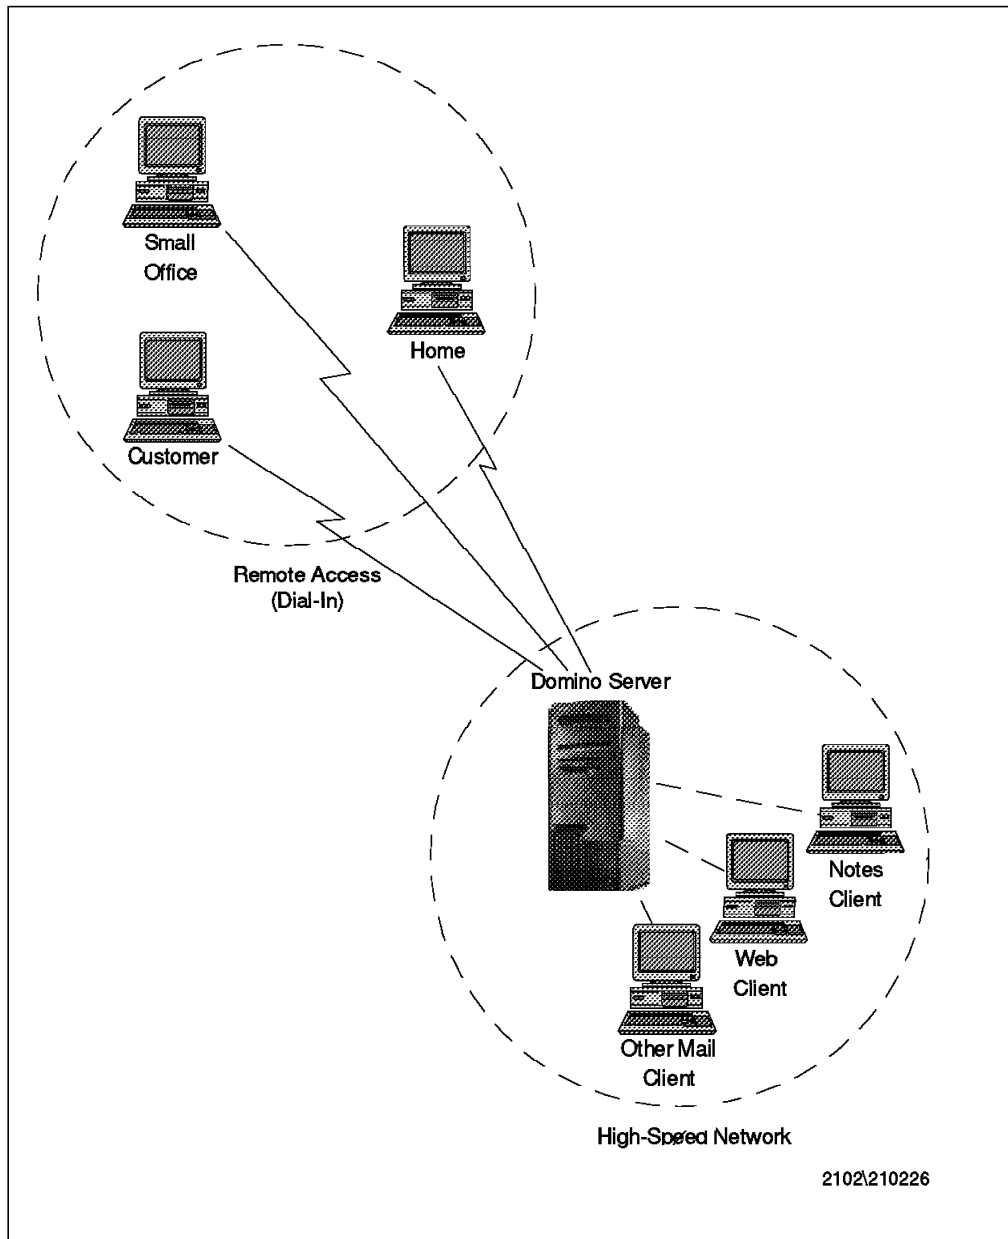


Figure 2. Remote Stations Connected to a Single Domino Site

To allow Notes clients to connect to your server you have two options:

- Notes XPC protocol
- Use other dial-in services such as Warp Server Remote Access Services (RAS), Microsoft RAS or the IBM 8235.

If you only have a few remote clients, we recommend that the Notes XPC solution should be used. Remote access works better if you select a high quality, Hayes-compatible modem with the highest industry recognized speed. We recommend a minimum speed of 33.6 kbps.

Connect this modem to your Domino server and activate the appropriate modem port under Domino. If your server needs to contact other servers to replicate databases and route mail, you should add a second modem to your server. This allows your server to freely contact other servers when a user is connected to

your first modem. Having only two modems connected to your server does not present a significant CPU load.

If your users are complaining that the remote access phone line is always busy, you will probably need to add more modems to support a larger number of remote users. In this case we recommend that you seriously consider a remote access service such as Warp Server RAS, Microsoft RAS or the IBM 8235. This should be installed on dedicated server hardware.

These RAS solutions provide your remote clients with the added benefit that they can connect to other systems on your network, such as AS/400s and RISC servers. These remote access solutions provide callback facilities as well. The remote access services provided by Warp Server RAS, Microsoft RAS and the IBM 8235 all support standard protocols such as TCP/IP, NetBIOS or IPX/SPX. These are also supported by Domino.

Customers using Windows NT or Windows 95 clients should seriously consider Microsoft RAS because of its seamless operation with the Notes client. A connection document can be configured in Notes that uses the underlying RAS software. In Figure 3 on page 14, we show an example location document that has been configured to use a remote RAS server.




Figure 3. Example of Location Document for RAS Connection

The Notes user simply changes their location setting under Notes when they travel offsite. The user will then use Microsoft RAS to connect back to a remote RAS server at the office. Users do not need to additionally start the RAS software on their workstation (32-bit client only). For complete details on setting up a workstation to connect to Microsoft RAS, please refer to pages 146-150 of the Lotus Domino 4.6 *Getting Started with Domino* manual. Windows 95 and NT clients will also be able to connect to Warp Server RAS in the future. Warp Server RAS is currently in development, which will change its existing RAS software to support PPP clients (which means Windows 95 and NT clients also).

Note: The password shown in the server connection document of a user reveals the user's password. This document is saved locally on the user's workstation and can be viewed by anyone. Ensure that physical security steps are taken and such sensitive information is not made available.

2.5 Your Domino Network Gets Bigger

Our example so far has shown a Domino solution located on a high-speed local area network (LAN) including dial-in support. This Domino server, or group of servers, could equally be located on many different LANs in a campus environment. The key point is that clients have sufficient bandwidth to get to the server. This means on most of today's networks 10-Mbps Ethernet or 16-Mbps token-ring. Many networks are starting to implement switching technology, 100-Mbps Ethernet and Asynchronous Transfer Mode (ATM), which provide even greater data rates between the server and the client.

As you are adding more clients to your server you may need to increase its capacity, add more servers, or even implement clustering technology. These are discussed in the following chapters. But, what if you do not have a high-speed link between your server and your clients? What if this link is only 56 kbps? The solution is replication, as discussed in 2.2.5, "Replication" on page 5. Replication is a powerful function of Domino that allows you to replicate databases to additional servers via WAN links or even modems through an analog telephone system. Servers can replicate changes in databases on a scheduled basis to multiple servers. Each server has its own replica copy of a database or multiple databases. It is the frequency of the replication schedule that determines how often the database(s) is kept in sync with its partners. Electronic mail is routed between servers as well. We can see in the diagram in Figure 4 that our Notes network has just expanded one step towards the enterprise.




Figure 4. Single Domino Site Connected to Remote Domino Site

Obviously, if you don't have a WAN infrastructure at your remote site, you will need some sort of dial-up solution. If remote users are located at another LAN,

then another Domino server at that location is the best solution. This is because it allows only one phone call to the remote site to transfer information for all users, rather than each user replicating their changes individually.

2.5.1 Wide Area Networking Issues

So when exactly should you consider having remote Domino servers at other locations? To answer this you must consider a number of factors together, such as the speed of your WAN link, its reliability and whether any other information travels on that link.

Most users accept that the minimum bandwidth needed for acceptable response times for a Notes client is 14.4 kbps. This is the link speed of an analog modem with some form of compression. If the link is any slower than this speed, you find that you are waiting more often for a response than doing productive work. This is what we mean by an *acceptable response time*.

If you have an existing Domino network, obtain a network analyzer and trace all data to and from your server during a busy morning period. Take the busiest period during that trace and determine the largest amount of data that went to or from the server. The reason for this is that most digital WAN links (for example, ISDN, frame relay, and Digital Data Service) are all full-duplex and provide an independent channel in each direction simultaneously. The exception is for dial-up analog (modem) links which operate in half-duplex, where the sum of the in and out traffic should be taken. The collected data should be converted from megabytes (MBs) to bits. From this data rate we can determine the average user traffic load for your company, and use this to determine if a remote Domino server is required.

Our example is a trace of a typical network during the busiest morning period. During that time we traced the traffic to and from our Domino server and we found that 26 MB of data went from our server, and 16 MB of data went to our server. We will use the larger of these two values in our calculation, because we have a frame relay WAN link to our remote site. The trace was over a one-hour period.

1. 26 MB
= 26 x 1024 x 1024 x 8
= 218,103,808 bits of information

We now divide this amount of information into the average bit rate per person when 50 users were connected to our Domino server:

2. Average bit rate per person
= 218,103,808 / (50 x (60 x 60))
= 1211.7 bits/ second (for each user)

Note: 60 x 60 = 3,600 seconds in one hour.

We are expecting to install ten clients on the other side of a 48 kbps link. From our average bit rate per person we expect the average bit rate for ten people to be:

3. Average total bit rate for remote clients
 - = 10 x 1211.7 bits/sec
 - = 12,116.9 bits/sec
 - = 11 kbps

4. Link speed is 48 kbps.

5. Since 11 kbps < 48 kbps = Attach users to Domino server via WAN.

Since our average link rate is less than our link speed we could consider having these users connect back to our central Domino server via the WAN. If you decide to do this, inform the users that they will experience slower response times than local user access. If they do not find this an acceptable solution, then you should consider a remote Domino server. As an alternative, you could have your remote users replicate their most used databases to their local workstation. For example, if a user replicates their mail database to their local machine they can send and receive mail by using the much faster local copy. The database will replicate to the server as a background task on a scheduled basis.

Also consider a remote Domino server if your average total bit rate in (3) was considerably greater than the link speed (4). Remember to take several traces to see the average amount of data to your Domino server for various days of the week.

If you do not have any user data to draw upon, and you are trying to centralize your servers and have Notes users connect to a server via a WAN link, how many Notes users can you realistically support across such a link? There are various guidelines that IBM have developed to help you.

The first is a very conservative rule of thumb is the following:

Remote Domino Server Requirement Formula

If $n \times 14 > \text{link speed in kbps}$, you should consider a remote Domino server.

Where n = number of users at remote site.

For example, if a WAN speed is 48 kbps and there are 100 remote clients at the remote site, because $100 \times 14 = 1400 > 48$, we would recommend a Domino server at the remote site. This rule is used when little bandwidth information is available. The rule of thumb can be used when you want to ensure there is sufficient bandwidth for each remote user.

Since this rule was developed, a study by IBM Global Services and IBM Research on bandwidth requirements for Lotus Notes and Web browsers found the following:

Lotus Notes 3.7 kbps
Web Browsing 5.5 kbps
3270 (PROFS) 1.3 kbps

IBM studies have shown by monitoring usage patterns over a long time with it's PROFS and Lotus Notes network that on average, 1/3 of Notes users will be actively using the system at any one time.

IBM Research also produced similar results by analyzing the Notes server traffic at the IBM Watson Research Center. If you plan to have a number of remote

users access your Domino server across a WAN link at about 60% utilization and assuming that 1/3 of the users will be actively using the Domino network at any one time, you can compute the following:

$$(1 / 3.7) \times 3 \times 0.60 = 0.4864 = 0.5$$

So if your link is x kbps, you can support approximately x/2 number of remote Notes users across the WAN link.

8 kbps	4
16 kbps	8
32 kbps	16
64 kbps	31
128 kbps	62
256 kbps	124
T1	751
E1	994

**Calculated for 16 kbps: $16 / 3.7 \times 3 \times 0.6 = 7.7$ rounded to 8 users.*

These are conservative values and will provide substantial bandwidth for Notes users connected at remote locations connecting to a Domino server across a WAN link.

If the reliability of your link is in question, or you cannot accept any user down time in the event of a WAN link failure, you should definitely consider a remote Domino server. If you also have any other time-critical data which is a higher priority than your Notes data, then also consider a remote Domino server. This is because a large number of remote users may swamp the link with data, which may effect time-critical host communications or other services.

If you look at your own Wide Area Network links you may decide to centralize many of your Domino servers. It is important that you consider operational considerations such as where your support personal are located and the language that they speak. For example, to support users in France it would be logical to have a server located in Paris and have the local support staff manage this server. Having these users connect to a centralized server in the USA simply because sufficient bandwidth is available is probably not the best solution.

A summary of link considerations is presented in Table 2:

<i>Link Consideration</i>	<i>Recommendation</i>
Speed	<p>Determine from network traces the average user bit rate. Multiply this value by the number of users you wish to have at the remote site. Compare this value with the WAN link speed. If greater, consider a remote Domino server.</p> <p>If no user data to go by, use an approximate formula:</p> <p>If the number of users = $n \times 14 > \text{link speed in kbps}$, consider another Domino server at the remote location.</p>
Reliability	<p>Seriously consider a remote Domino server if user downtime is totally unacceptable at the remote site.</p>
Other Data on Link	<p>Seriously consider a remote Domino server if mission-critical data is present on the WAN link, especially if this data is DLC or NetBIOS traffic.</p>

2.5.2 Replication Schedules

As soon as we add another Domino server to a network, whether it be on a LAN, WAN or dial-up link, we must consider the replication schedules that are used. Replication schedules are blocks of time in which you can specify to replicate databases, and route mail with other Domino servers.

Figure 5 is an example of a replication schedule document from ITSORAL01 and ITSODAL01.

SERVER CONNECTION: ITSORAL01/ITSO to ITSODAL01/ITSO			
Basics			
Connection type:	Dialup Modem	Usage priority:	Normal
Source server:	ITSORAL01/ITSO	Destination server:	ITSODAL01/ITSO
Source domain:		Destination domain:	
Use the port(s):	COM1	Destination country code:	
Always use area code:	No	Destination area code:	
		Destination phone number:	919 221 2345
		Login script file name:	
		Login script arguments:	
Scheduled Connection		Routing and Replication	
Schedule:	ENABLED	Tasks:	Replication, Mail Routing
Call at times:	08:00 AM - 05:00 PM each day	Route at once if:	5 messages pending
Repeat interval of:	30 minutes	Routing cost:	5
Days of week:	Sun, Mon, Tue, Wed, Thu, Fri, Sat	Replicate databases of:	Low & Medium & High priority
		Replication Type:	Pull Push
		Files/Directories to Replicate:	(all if none specified)
		Replication Time Limit:	minutes
Comments:	Replication schedule from ITSORAL01 (hub) to ITSODAL01 (spoke)		
Administration			
Owners:	ITSORAL01/ITSO		
Administrators:	ITSORAL01/ITSO		

Figure 5. Example of Server Connection Document

These servers are connected via a dial-up modem. We can see that ITSORAL01 initiates the call each time, starting at 8:00 a.m. and finishing at 5:00 p.m. every day repeating this process every 30 minutes. During those replication periods, databases that are replica copies will be replicated. Mail will be routed during this time. Note that replication will also occur if at any time five normal priority mail messages are pending on ITSORAL01. To view the details on setting up a replication schedule, please refer to pages 186-188 of the Lotus Domino 4.6 *Getting Started with Domino* manual.

2.6 Hub-and-Spoke Network

We recommend that you use a hub-and-spoke replication scheme in your network. You specify which server you would like as the hub, and this server initiates all connections to the spokes as shown in Figure 6.

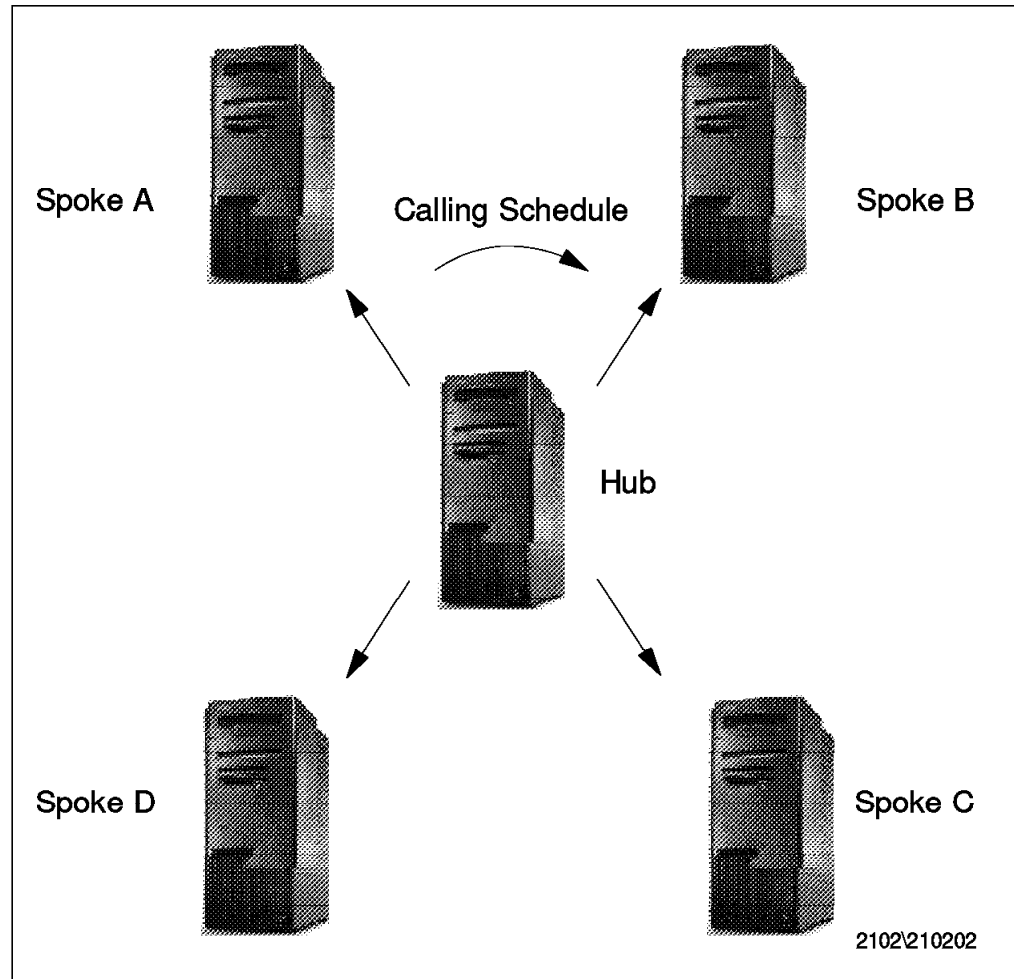


Figure 6. A Graphic Representation of a Small Hub-and-Spoke Network

A hub server is normally located at your main computer site. The many advantages of the hub-and-spoke approach over other topologies are:

- Reduces the number of connection documents required.
- Known flow rate of new information in your network, which allows easier problem determination if a link has failed.
- Reduces the amount of network traffic.
- Scales well as the infrastructure grows. As new servers are added it is easy to integrate this into the existing network (minimal amount of connection documents).
- Allows control and planning of hub load.

As your network grows you can see where new servers fit into your organization, and how and when information will be relayed to them. In Figure 7 on page 22, we can see how our single example Domino site is easily expanded to include three remote Domino sites. These remote sites could be anywhere in the

country or around the world. In addition, each remote site may support additional sites.




Figure 7. Hub-and-Spoke Topology Allows Easy Expansion of a Domino Network

A Notes Hierarchical Architecture includes the following levels:

- Super Hub (level 1) is one that communicates between large sites. These hubs make up a backbone network for replication and mail routing.
- Regional Hub Server (level 2) replicates with both super hubs and application servers (regional servers). Normally, the regional hub server is used to construct the backbone network (an optional solution in some infrastructures).
- Regional Server (Application Server) are those that directly serve users with, for example, mail or shared database services.

Table 3 displays a subset of three connection documents used in the example diagram in Figure 7 on page 22. The connection document examples are two wide area network (WAN) links and one dial-up link.

From Computer	To Call Computer	Call at Times	Tasks	Use Port
Hub	Spoke A	11:00 a.m.	Replication, Mail Routing	LAN0
Hub	Spoke B	12:00 p.m.	Replication, Mail Routing	LAN0
Hub	Spoke C	1:00 p.m.	Replication, Mail Routing	COM1

We can expand this to include more and more remote sites as they come online. Chapter 3, "Sizing IBM Netfinity Servers for Domino" on page 29 discusses how many spokes you can support on your hub server, and whether you should consider installing a dedicated hub server. Many large Domino networks have the hub server contact the spoke servers over a range of time with a specified repeat interval.

Figure 8 shows an example of a hub server, ITSORAL01, which needs to replicate with ten different spoke servers every 30 minutes from 8:00 a.m. to 5:00 p.m. every day. This is configured as 10 connection documents, one for each spoke server.

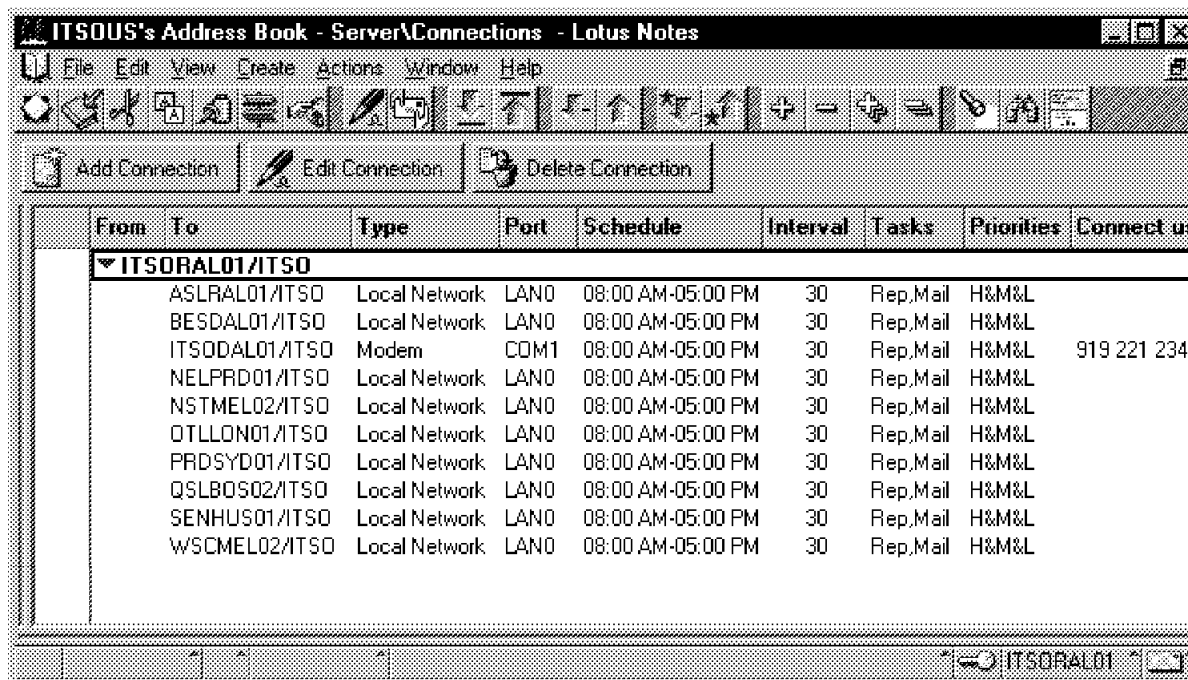


Figure 8. Ten Connection Documents from ITSORAL01 Hub Server

By default your server loads one replicator task to perform all of these replications.

If you create connection documents that schedule a hub for multiple simultaneous or overlapping replications with different destination spoke

servers, you should set up multiple replicators to handle the replication sessions simultaneously. This makes efficient use of server resources, shortens replication cycles and saves replication time. In our example in Figure 8, rather than our hub ITSORAL01 replicating with ASLRAL01 first, then BESDAL01 second and so on until WSCMEL02, it can assign a Replica task to replicate with each spoke simultaneously.

You should add one Replica task for every spoke server your hub will connect too. In our example we would start 10 Replica tasks on our Domino server.

You can set this by adding more than one Replica keyword to the ServerTasks entry in the NOTES.INI file. For example:

```
ServerTasks = Router,Update,Stats,AMgr,  
Adminp,Event,Replica,Replica,Replica,Replica
```

This will start four separate Replica tasks each time your Domino server is started.

Most of our discussion so far has pertained to wide area network (WAN) links. What if the connections to our spokes were via a single modem and not via WAN links? In this case, adding more Replica tasks will not allow us to simultaneously replicate with each spoke, as there can be communication with only a single spoke at a time. We can see from our example list of connection documents in Figure 8 on page 23 that our hub server ITSORAL01 will attempt to replicate with 10 spokes every 30 minutes but now via a modem. Our hub server now proceeds alphabetically through the connection documents, connecting to each spoke server one at a time, replicating and then hanging up before proceeding to the next. At the end of 30 minutes let's say ITSORAL01 still needed to replicate with PRDSYD01, QSLBOS02, SENHUS01 and WSCMEL02. The Replica task will then add the ASLRAL01 server to the list of servers it needs to replicate with, because our connection document for ASLRAL01 specifies a repeat interval of 30 minutes. The Replica task then inspects a list of which servers it needs to replicate with, and chooses which server is alphabetically first. It will then replicate with ASLRAL01 again, even though it has missed replicating with PRDSYD01, QSLBOS02, SENHUS01 and WSCMEL02. Spoke servers that are alphabetically higher in the list will be serviced by a hub server that is overloaded.

An overloaded hub server is one that has been scheduled more work to do than it has time to complete. This can occur even if it has sufficient hardware capacity.

A very useful tool to determine whether your hub server is overloaded is the DYS analyzer. This tool shows the success rate of your replication schedules, average and maximum delays, individual event duration plus other information. This tool can be purchased from DYS Analytics in Newton, Mass.

There are several things you can do if your hub server is overloaded. The quickest solution is to increase the time between replication events. For example, if your average replication takes 12 minutes to complete, and your replication window is only 10 minutes, then increase the time between replications to 20 minutes. This simply means that new and updated information will flow through your organization more slowly. This may or may not be an issue to your users or management. A second solution is to install a second hub server to work in unison with your first server. The final solution is hardware,

software or telecommunication upgrades. Consider which element is worth upgrading in your situation. You may find it prohibitive to upgrade your WAN link speed, or you may already be using the fastest modem you can purchase. If this is the case, then a second hub server may be the solution. You should compare each option carefully.

In many cases, implementing a realistic replication schedule when you first implement Domino will help you avoid many of these replication pitfalls in the future.

2.7 Other Considerations

So far we have given an example of one or more Domino servers located on a central site. Remote access has been provided and we have expanded our network by multiple remote Domino servers connected together via WAN and dial-up links. As our remote network spans several sites, we may install dedicated replication and mail hubs for server capacity reasons. You can scale this design to a large organization. We now outline some additional information you should be aware of that can affect your overall design. These issues also have an effect on what servers you select in Chapter 3, "Sizing IBM Netfinity Servers for Domino" on page 29.

2.7.1 The Size of Your Domain

The heart of your Domino system is the Public Names and Address Book (N&A Book). This is a Notes database stored on each server called NAMES.NSF. The Public N&A Book plays a dual role. On the one hand, it is a directory service that lists all users, groups, and servers and their individual configurations (for example, which server a user's mail file is located in, or what people can access a particular server). On the other hand, it provides information for servers such as replication schedules, certificates, and mail-in databases. It is therefore paramount that security be tightened on this database. Take some time to become familiar with how the database is structured.

The N&A Book is replicated to all servers that are logically grouped together in a Notes domain. Do not confuse a Notes domain with a Windows NT or Warp server domain, which operate independently. Your Notes domain does not have to be restricted to one physical location. For example, your domain may include several Domino servers on one high-speed LAN, and several dial-up connected Domino servers.

As the number of people in your domain increases, so does the size of the N&A book. To give you some examples, at the time of writing the IBMUS domain for the United States contained approximately 86,000 documents and was 831 MB in size. The IBMAU domain for Australia contained 1900 documents and was 19 MB in size. These Domino servers operate on the IBM SP platform. A large domain means a large N&A Book on each server in your domain, with the increased load of frequent searching, indexing and replication between servers. This can dramatically increase the hardware requirements for your servers. If you needed to connect a small remote workgroup to your corporate domain, it too would need to contain the complete N&A Book on its server, and would increase its hardware requirements.

We recommend for typical networks that the N&A Book should be limited to 30,000 entries. If you are expected to use a large amount of lower-cost servers

(for say many remote locations), then you should limit your domain's size even more, to say 10,000 N&A entries.

Breaking a domain into smaller domains will not reduce the functionality of your Domino network. In previous versions of Notes, a single domain would reduce the chance of a user mailing to someone with a nonexistent address. But Domino 4.6 has a new feature called directory assistance which allows you to search for users' names across multiple domains.

If you decide to have multiple domains, we recommend that you break up your domain into geographies. If you break up your domain into country geographies, consider basing the name of your domain on the ISO Country Codes. These codes are listed in Appendix F of Lotus Domino's *Getting Started with the Domino Server* and give every country in the world a two-letter unique name. For example, the code for the United States is US, Denmark DK and Greenland GL.

If your company is called Acme Ltd., consider naming your domains ACME followed by the two letter country codes, for instance, ACMEUS for the United States, ACMEGB for Great Britain. Figure 9 illustrates the company Acme with four Notes domains.

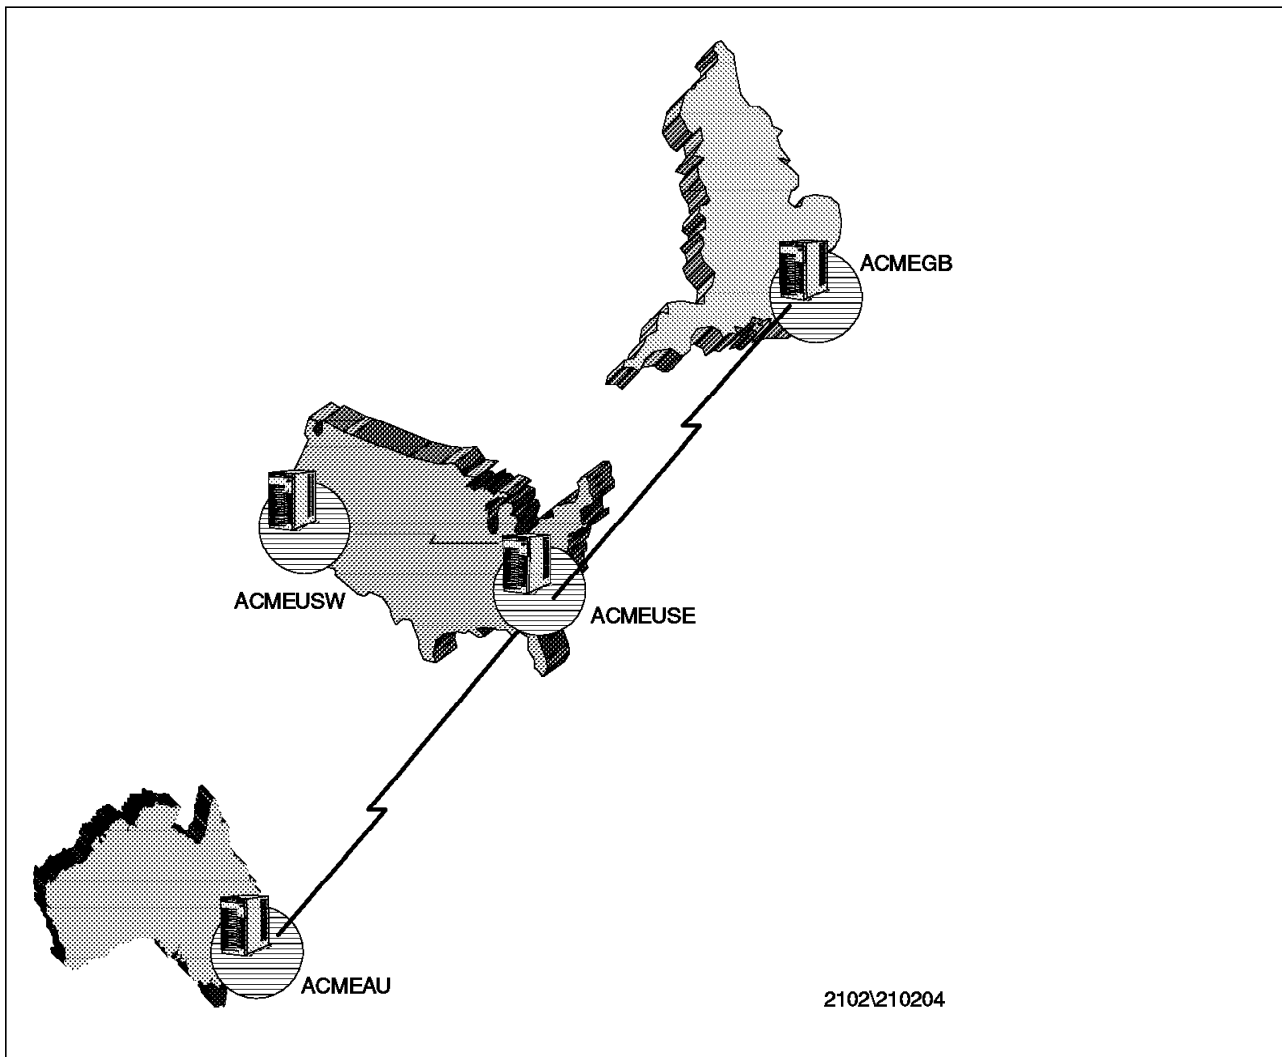


Figure 9. Domain Location for Example Enterprise Network

Note that ACMEUSE is the head office located in Washington, so this is where their central hub server is located as well. This company also has a large amount of remote staff located in different parts of Europe. These users connect to their mail server located in London by the use of modems. Although these users are not physically located in Great Britain, you can still name the domain ACMEGB. The domain name is simply used to route mail to its proper destination. The last two digits can help you identify the physical location of your servers and domains. Whatever domain structure you decide, consider all of the implications before you implement an enterprise Domino network.

2.7.2 Systems Management for Your Domino Network

As your Domino network grows larger and more critical to your business you will want to know what is happening on your Domino network at all times. Effective systems management should therefore be considered seriously.

The first thing you should implement is event monitoring, which is included with Domino. The event server task collects events on your server and reports them to a person, database, or server-management program. You can monitor the following:

- Communication events
- Mail events
- Miscellaneous events
- Replication events
- Resource events
- Security events

You can select which events you would like to monitor, and what to do when they occur. For more details on how to set up event monitoring, please refer to pages 53-68 of Lotus Domino's *Maintaining the Domino System*.

You can also use the Windows NT Performance Monitor if your Domino servers are running on the NT platform to monitor the performance of your servers. The use of the NT Performance Monitor will be discussed in more detail in 4.5, "Using the Windows NT Performance Monitor" on page 97.

You should investigate other Domino management tools, such as the Tivoli Enterprise Module for Domino/Notes. This provides an added server and client agent that provides added functionality:

- Distribute new Notes clients to remote locations.
- Monitor events, and automate responses rather than just reporting them. For example, if a server fails, automatically restart the server. This improves availability of your network.
- Create a policy-based management paradigm, ensuring that management and security are consistent across your organization.
- Make administration tasks quicker than the current interface.

For more details on this Tivoli Enterprise management module, please refer to *Managing Domino/Notes with Tivoli Manager for Domino*, SG24-2104.

Chapter 3. Sizing IBM Netfinity Servers for Domino

The selection of a server platform within the Notes network is not always as easy as selecting a file server for your network. Each Domino server can provide a variety of different functions such as mail, database, Web, passthru or combinations thereof. So depending on the type of functions on your server, and the number of clients, you may need to have more than one server to meet your capacity requirements. In this case, you may define particular roles for your Domino servers, such as a mail server or a replication hub. This can help you better identify the hardware you require.

This chapter outlines the various IBM Netfinity servers line of products and helps you confidently build a reliable foundation for your Domino server. Some of the roles that are given to Domino servers are described, along with the most appropriate IBM Netfinity servers to meet those roles. We then use the Server.Planner tool to define exactly which server we need, and then provide recommendations for disk space and memory, and other server peripherals. We conclude this chapter with four Domino server sizing examples.

3.1 Which Platform Is Right

As you consider what hardware and software you will need for your Notes network, keep in mind your goals for the full Domino deployment. For example, initially you may decide to dedicate one server to the pilot group. Then as you add workgroups, you might dedicate additional servers. This method might not be efficient if workgroups do not fully utilize their server or workgroups are geographically diverse. A better plan might be to deploy a few high-capacity Domino servers that will serve multiple workgroups.




Figure 10. Notes Network Diagram Example Depicting Placement of Servers by Function

Start with a picture of how your entire Domino network environment may be viewed and then balance the requirements for various server functions and workgroup requirements.

When implementing your pilot, be sure to implement all functional aspects of the full Domino deployment. For example, if you decide to use Domino as your mail system, and you need to send mail with additional mail systems, you should investigate and implement the required gateway during the pilot to correct any issues before the production deployment. When selecting hardware for the pilot ensure that it has the capabilities of supporting the production environment. Too often, first time Domino administrators implement server platforms that are incapable of supporting production loads which results in poor performance when the system goes live. Make your system blueprint and stick to it as you develop and roll-out your Notes network.

We also recommend you stress test some of your Domino servers before they go into production to ensure all load issues have been addressed. You can use a

tool such as GroupSizr (formerly known as NoteSizr) and WebSizr to do this testing.

GroupSizr and WebSizr

GroupSizr for Lotus Notes is an application that helps application developers design and implement high-performance and highly scalable Notes and Domino Applications.

WebSizr is a performance Analysis and Sizing framework for HTTP based Servers and applications. Application developers use WebSizr to engineer high-performance Web applications. More information on these products can be reviewed at the following Web site:

<http://www.technovations.com/>

As you plan for additional servers to support Domino, keep in mind that using servers of the same type and configuration will simplify administrative tasks by establishing a common Domino configuration and setup. Include allocating physical space for servers in planning for your hardware needs, and plan for growth when you allocate space. It is much easier to administer multiple servers if they are physically located close to one another. Also consider the physical security of Domino servers. They should be stored in a locked room with access restricted to authorized administrators.

To assist you in determining which IBM server is best suited for your requirements, we begin by providing specifications on each of the IBM Netfinity servers platforms and then examine each function that a Domino server may carry out. To fully understand the capabilities of a particular hardware platform in the Domino environment, you need to consider each of the functions that the server may support in your network.

The following specifications are intended only to be used as a guideline for the capabilities of each server. For specific information when purchasing, refer to your local IBM reseller, the redbook *IBM Netfinity and PC Server Technology Selection Guide*, SG24-4760, or refer to the IBM Netfinity servers home page on the Internet at

<http://www.pc.ibm.com/us/netfinity>.

3.2 IBM Netfinity for Lotus Domino Servers

The IBM Netfinity servers provide their worldwide customers with a reliable foundation for their networked businesses by delivering outstanding power, scalability, control and service. The IBM Netfinity servers family of products helps small, medium, and large companies and enterprises manage their networks, from basic file and print capabilities to the more advanced use of applications, anytime, anywhere. Domino and the IBM Netfinity servers provide their users at any-sized company an integrated e-mail and Web application server and the ability to quickly develop secure, interactive applications for the Internet and intranet.

The following sections outline the features of the IBM Netfinity servers product line, with focus on those specific models that were tested for this document. Although this information was accurate at the time of this writing, server

hardware is updated quite frequently to utilize emerging technologies. Please refer to the IBM Netfinity servers Web page at

<http://www.pc.ibm.com/us/netfinity>

for the latest information.

The information in this section should be used to select the right server platform for your business solution.

3.3 IBM Netfinity Family of Servers

The IBM and Lotus solution for your business starts with a choice of Intel-based servers that can provide exceptional network support. With the IBM server models, you will gain outstanding power, scalability, control and service, a must for today's enterprise to small network environments. Your decision to choose the IBM and Lotus solution will allow you to deploy business-critical applications with confidence. Lotus and IBM are very concerned about the aspects of their customer's business, such as the high availability of network services to access and manage mission-critical information, as well as providing a reliable and secure platform.

This chapter covers the IBM Netfinity 7000, 7000 M10, 5500, 5500 M10, 5000, 3500 and 3000 servers, the first in a new family of enterprise servers. The IBM Netfinity servers leverage IBM's server heritage, combining high-performance application computing based on industry-standard, Intel processor-based technology. The IBM Netfinity servers have been designed to enable customers to run demanding e-business applications, including Web server and online transaction processing (OLTP), as well as data-intensive operations.

3.4 About IBM Netfinity

From the ground up, IBM Netfinity is designed and built to deliver midrange-level performance in critical areas such as e-commerce and e-business. The IBM Netfinity server provides the power and scalability to distribute the most demanding 7 x 24 x 365 applications and data, including your ERP and e-business applications. The IBM Netfinity server includes the following controls for easier management of all your system's resources and the total service that business-critical systems demand:

- **Processing Power** - IBM Netfinity uses the latest Intel Pentium II and Intel Pentium II Xeon processors, integrated into a server system balanced and tuned for high-speed distribution of applications and data throughout your enterprise.
- **Scalability** - IBM Netfinity servers handle the most resource-hungry data and applications with ease. These applications include J.D. Edwards OneWorld, SAP R/3, Lotus Domino, Domino Intranet Starter Pack, Domino.Merchant, Net.Commerce, Microsoft Exchange, Baan, Microsoft SQL Server, Oracle, IBM DB2 and more. The ingenious design of the IBM Netfinity server allows it to be configured as a tower or a rack system, so it's easy to grow your server with your business. Key options and components can be quickly replaced, upgraded or removed - without powering down the server!
- **Control** - IBM Netfinity Manager software and the Advanced Systems Management Adapter card put you in complete control of your server and all your system's resources graphically, from a single point of control. So you can see all your systems at a glance, from virtually anywhere in the world, anytime of the day or night. You'll also be in the position to proactively manage your systems for peak efficiency and output.

- **Service** - IBM's renowned service and support are with you all the way. Listed are the worldwide available Web sites:

Netfinity	www.ibm.com/pc/us/netfinity
Server	www.ibm.com/pc/us/server
ServerProven	www.ibm.com/pc/us/compat
SystemXtra	www.ibm.com/pc/us/systemxtra
TechConnect	www.ibm.com/pc/us/techconnect
Support	www.pc.ibm.com/support
Options	www.ibm.com/pc/us/options
Partners	www.partner.us.pc.ibm.com
Services	www.ibm.com/Services
Year2000	www.pc.ibm.com/year2000/index.html
TCdigest	www.ibm.com/support/tcp

Today, you can configure and order your Netfinity Server directly from the Internet at a very competitive price. To do so, go to the following Web site:

BuyIBM www6.pc.ibm.com/buyibm

If you need assistance to help you in your buying decision, you should use the Sales Assistant located at the BuyIBM Web site.

3.4.1 The Foundation for Lotus Domino and Domino Applications

Lotus Domino gives you an e-business solution designed to meet your needs today and well into the future. The powerful, scalable foundation for Domino is IBM Netfinity.

The IBM Netfinity is ready to provide you with the following features:

1. Latest Pentium II Technology
Up to four-way Pentium II Xeon symmetrical multiprocessing in both tower and rack-mountable IBM Netfinity 7000 M10 models.
2. Ultra Fast Throughput
High-speed intelligent disk subsystems, with Wide Ultra SCSI controllers, Serial Storage Architecture (SSA), and Fibre-Channel adapters and drives, deliver excellent NT application performance.
3. Storage Flexibility
Expend disk arrays while the server is still running using IBM Logical Drive Migration technology with IBM ServeRAID II Ultra SCSI Adapter.
4. Advance Systems Management Adapter
The Advance Systems Management Adapter (ASM) allows you to communicate with your Netfinity server using Netfinity Manager. It can be used to configure and monitor many of your system's features. With the Advanced Systems Management service of Netfinity Manager, you can configure events such as POST, loader, and O/S timeouts, critical temperature, voltage, and tamper alerts and redundant power supply failures. This service also enables you to dial-out and directly access and

control a remote system's ASM adapter or Advanced Remote Management processor.

5. Complete Manageability Software Suite

Includes IBM Netfinity Manager software for proactive server management and control, IBM Capacity Management tool for upgrade planning, and integrated UPS management for safe application shutdown. All are provided as standard to help you reduce your total cost of ownership.

6. Extensive Hardware Management Integration

Disk Drive alerting with S.M.A.R.T. and Predictive Failure Analysis (PFA) provides pre-failure notification for disk drives. Enable full "lights out" remote management including key component status as well as remote power and restart with an IBM system management processor that is standard on some models and optional on all IBM Netfinity systems.

7. High Availability

Minimizes system downtime with ECC memory, RAID subsystems with hot-swappable disk drives, redundant power supply units and cooling fans. Certified for Microsoft Cluster Server to provide a highly available fault-tolerant platform for NT applications to protect against server or application failure.

8. Rack Mountable

IBM Netfinity Rack provides an industry-standard 19" enclosure to support IBM Netfinity 7000 and 5500 rack-mountable servers and other options to save valuable floor space and ease maintenance operations.

9. Compatibility and Performance

Dedicated IBM engineering teams conduct ongoing compatibility testing, certification testing, and performance measurement and optimization to help ensure that IBM Netfinity systems, network operating server, and Lotus Domino run seamlessly together.

10. IBM Start Up Support and IBM Warranty

IBM Start Up Support provides, at no additional charge, telephone assistance for IBM and selected third-party hardware, as well as NT Server, for the first 90 days of ownership. IBM servers are also covered by IBM's three-year limited on-site warranty.

11. TechConnect

A comprehensive program designed exclusively for networking professionals provides unique solution-based technical information, training and certification opportunities, and more.

12. Lotus Domino Included

A Lotus Domino server is included at no additional cost.

3.5 IBM Netfinity 3000

The IBM Netfinity 3000 offers you serious solutions for your file-and-print and application computing needs. This powerful mini-tower machine is designed with the attention to detail you expect from IBM, and is backed by the same world-class service and support found with larger systems.

If your business is at the point where an Intel-based server is a requirement but budget remains a factor, the IBM Netfinity 3000 is the affordable solution that deserves a closer look.

The IBM Netfinity 3000 is the reliable and affordable foundation you need to give your end users ready access to your business-critical applications. The Netfinity 3000 is an excellent entry-level Domino mail or application server.

A potential user of the IBM Netfinity 3000 server may have the following requirements:

- An affordable server that practically installs itself with the IBM ServerGuide CD. (ServerGuide ships with each Netfinity system.)
- A server packed with power and scalability features that support file-and-print serving, workgroup productivity applications such as Lotus Domino Server and Lotus Domino Intranet Starter Pack (both are shipped with the server at no cost), and shared Internet access.
- The scalability of Pentium II processors from 300 MHz to 450 MHz that can be customized with memory from 32 MB up to 384 MB ECC SDRAM.
- Large disk storage requirements for database applications, with a current maximum of 36.4GB internal hard disk. Additional storage can be added externally if required.
- Netfinity Manager software (at no cost), which allows you to apply proactive Systems Management functions across your entire networked business systems. In addition, Netfinity Manager can integrate with key workgroup and enterprise management platforms using industry standard protocols, such as SNMP, MIB and MIF, as well as allowing proprietary systems management software to integrate, by using the critical file monitoring.

Typical Domino functions are:

- Departmental mail server
- Small workgroup applications server
- Remote server
- Passthru server
- Internet/intranet server

The IBM Netfinity 3000 is shown in Figure 11 on page 37.




Figure 11. Exploded View of the IBM Netfinity 3000 Tower Server

The Netfinity 3000 features tower packaging with the following benefits:

- A variety of 300MHz, 350MHz, 400MHz, and 450MHz Pentium II processors.
- A PCI/ISA bus architecture.
- A 512 KB synchronous pipeline bursting L2 write-back cache.
- 32 MB and 64 MB ECC SDRAM memory expandable to 384 MB maximum.

Note: The recommended memory for Domino should be a minimum of 64 MB.

- Preconfigured models with 10/100 Ethernet and 4.51 GB Wide Ultra SCSI disk storage, a maximum storage of 36.4 GB.
- A high-speed Wide Ultra SCSI 32-bit PCI disk controller.
- An 32X CD-ROM IDE drive.
- A total of six bays.
- A total of six PCI card slots, five available.

3.6 IBM Netfinity 5000

The IBM Netfinity 5000 Server delivers the power, scalability and manageability in tower or rack-mount models. The Netfinity 5000 supports e-business, large enterprise applications and small and medium business networked systems. Each value-priced Netfinity 5000 is backed by the same world-class service and support found with our larger systems.

Using the integrated Netfinity Advanced Systems Management processor in conjunction with IBM award-winning Netfinity Manager, it is the total solution for proactive control of your networked business systems, seamlessly integrating with key workgroup and enterprise management platforms.

With flexibility in the models ranging from single 350MHz Pentium II to dual 450MHz Pentium II, up to 45.5 GB of internal Hot-Plug hard disks, up to 1 GB of ECC SDRAM, power supply redundancy, Predictive failure Analysis (PFA) and SMART support, with environmental monitors as standard, this is a very highly manageable enterprise server and e-business server to meet the most demanding small to medium business environments.

The complex business environments of today need the flexibility of growth, and again the Netfinity 5000 is a business server that can grow with your business needs in terms of raw processing power, memory, disk, networking and fault-tolerance solutions, while at the same time providing the balanced power and resilience you have come to expect of a world-class Netfinity server from IBM.

Typical Domino functions:

- Departmental mail server
- Medium-size workgroup applications
- Medium to large database server
- Remote server
- Passthru server
- Internet/intranet server
- Cluster server

The IBM Netfinity 5000 is shown in Figure 12 on page 39.




Figure 12. Exploded View of the IBM Netfinity 5000 Tower Server

The Netfinity 5000 features tower and rack packaging with the following benefits:

- A variety of 350MHz, 400MHz, and 450MHz Pentium II processors.
- SMP Upgrade - two way.
- A PCI/ISA bus architecture.
- A 512 KB synchronous pipeline bursting L2 ECC/write-back cache.
- 32 MB and 64 MB ECC SDRAM memory expandable to 384 MB maximum.
- Preconfigured models with 10/100 Ethernet and one Advanced Systems Management adapter.
- A maximum disk storage of 45.5 GB.
- A high-speed Dual Channel Ultra Fast/Wide disk controller (on planar).
- An 32X - 8X CD-ROM IDE drive.
- A total of eight bays, of which five are hot pluggable.
- A total of five card slots, including three PCI and two PCI/ISA.

3.7 IBM Netfinity 5500 and 5500 M10

The IBM Netfinity 5500 Server has the power, scalability and manageability for the many demands of your networked systems supporting large Domino applications or databases. Furthermore, with the IBM Netfinity 5500, you can maximize your investment through IBM's unrivaled, world-class service and support programs.

In every respect, the IBM Netfinity 5500 system is the affordable, powerful and reliable foundation upon which you can run your business-critical applications with confidence.

IBM Netfinity 5500 systems are priced with business budgets in mind. Tuned for dual symmetrical multiprocessing and fast throughput, Netfinity 5500 servers are driven by the new 350MHz, 400MHz or 450MHz Intel Pentium II and Intel Pentium II Xeon processors with an integrated 100MHz front-end bus (previous generations of Pentium II processors supported only a 66MHz front-end bus), a fast, dual-channel Wide Ultra SCSI RAID controller that handles data up to 40 MBps and 100/10 MBps PCI Ethernet controller as standard.

A potential user of the IBM Netfinity 5500 server may have the following requirements:

- Industry-leading CPU performance capacity of up to two 450MHz Pentium II processors and for the M10 model, two 400MHz Pentium II Xeon processor for compute-intensive applications.
- 100MHz High speed front-end bus for extremely fast memory access.
- Large, fast disk storage requirements, up to 54.6 GB internal storage.
- ECC SDRAM DIMM memory, RAID, disk hot-swap, Hot-Plug PCI and power redundancy capabilities for data security and fault tolerance to support 7 x 24 x 365 business requirements.
- High availability built into the very heart of the server and total control of system resources from remote locations with IBM Netfinity Manager software in conjunction with the integrated Advanced Systems Management Adapter.
- Optional Netfinity High Availability Solution for automatic failover for clustering environments.

The IBM Netfinity 5500 server is intended for customers who need a high-performance, highly available, highly manageable database server or e-business server for business-critical applications in a complex environment. There's support for two-way SMP, integral tape drives and the ultra-fast 10,000 RPM hard disk drives. It is fully in step with Intel's processor technology for today, with the new 100MHz front-side bus for extremely fast memory access and the future, with Pentium II models being fully upgradeable to the new generation of Intel's Pentium II Xeon technology.

Netfinity 5500 is the affordable, powerful and reliable foundation upon which you can run your business-critical applications with supreme confidence.

Typical Domino functions:

- Large to Enterprise database server
- Medium to large workgroup applications

- Large to Enterprise mail server
- Large to Enterprise Internet/intranet Web server
- Large to Enterprise calendar and scheduling
- Cluster server

The IBM Netfinity 5500 is shown in Figure 13. The IBM Netfinity 5500 M10 is shown in Figure 14 on page 42.




Figure 13. Exploded View of the IBM Netfinity 5500 Rack Server




Figure 14. Exploded View of the IBM Netfinity 5500 M10 Rack Server

The Netfinity 5500 and 5500 M10 features tower and rack packaging with the following benefits:

- For the Netfinity 5500, a variety of 350MHz, 400MHz, and 450MHz Pentium II processors. For the Netfinity 5500 M10, the 400 MHz Pentium II Xeon.
- SMP Upgrade - two way (for 5500) and four way (for 5500 M10).
- A PCI (I2O ready)/ISA bus architecture.
- A 512 KB synchronous pipeline bursting L2 ECC/write-back cache.
- 128 MB and 256 MB ECC SDRAM DIMM memory. For the Netfinity 5500, memory expandable to 1 GB maximum. For the Netfinity 5500 M10, memory expandable to 2 GB maximum.

- Integrated with a 10/100 Ethernet card and one Advanced Systems Management processor.
- A maximum hot-swap disk storage of 54.6 GB. Maximum storage is 234.6 GB with optional EXP10 and EXP15 (for M10) external expansion storage enclosures installed in NetBAY3 with ten 18 GB hard disk drives.
- 2-Channel Wide Ultra ServeRAID II PCI disk controller (on planar).
- An 32X - 8X CD-ROM IDE drive.
- A total of ten bays, eight available and six of which are hot pluggable.
- Hot swap power supply.
- A total of five card slots, including three PCI and two PCI/ISA.
- Hot swap cooling devices. Two blowers and one fan supply sufficient cooling even if one device fails. Each device is replaceable without powering down the server so availability to users is not affected.

3.8 IBM Netfinity 7000 M10

The IBM Netfinity 7000 M10 Server is a reliable enterprise server, available in either rack or tower models. This server is able to deliver exceptional power functionality and value to handle complex, business-critical database or application server needs among growing messaging and intranet applications. The Netfinity 7000 M10 system is intended for enterprise customers whose networks depend on a reliable, high-performance platform with excellent scalability for future growth. The Netfinity 7000 M10 can also function as a file server in LAN environments where a significant number of users require a large amount of file server resources.

With the addition of the Netfinity 7000 M10 models to the Netfinity family, even greater intensive CPU requirements and larger databases can be accommodated, meeting the processing power and robustness of even the most demanding and complex business-critical environments in today's highly complex and intensive business markets.

A potential user of the IBM Netfinity 7000 M10 Server may have the following requirements:

- High CPU capacity of up to four Intel Pentium-Pro or four Intel Pentium II Xeon processors for compute-intensive applications (database applications, electronic messaging applications, PC client/server applications, ERP Solutions and environments)
- Large disk storage requirements, currently supporting 109.2 GB of internal hard disks
- Up to 8 GB of ECC memory, RAID, hard disk hot-swap, PCI hot-plug and hot-swap power supplies with PFA support and power redundancy capabilities for data security.
- High availability and total control of system resources from remote locations with IBM Netfinity Manager software, Advanced Systems management Adapter APC Powerchute and IBM Antivirus software.
- Optional Netfinity High Availability Solution for automatic failover for clustering environments

Typical Domino functions:

- Enterprise mail server
- Enterprise database server
- Enterprise Internet/intranet Web server
- Large workgroup applications
- Enterprise calendaring and scheduling
- Enterprise cluster server

The IBM Netfinity 7000 M10 is shown in Figure 15 on page 45




Figure 15. Exploded View of the IBM Netfinity 7000 M10 Server

The Netfinity 7000 M10 features tower and rack packaging with the following benefits:

- Processors - the 400 MHz Pentium II Xeon.
- SMP Upgrade - four way.
- Bus architecture - 3 x 132Mbps PCI (I2O ready) controller.
- A 512 KB and 1 MB L2 ECC/write-back cache.
- 128 MB and 256 MB ECC EDO DIMM memory expandable to 8 GB.
- Integrated Advanced Systems Management processor
- Maximum hot-swap internal disk storage is 36 GB (Netfinity 7000 M10). Maximum storage is 546 GB with optional EXP15 external expansion storage enclosure installed in NetBAY3 or rack with ten 18 GB hard disk drives.
- PCI Disk Controller (on planar) 2-Channel Wide Ultra SCSI
- An 32X - 8X CD-ROM IDE drive.
- Bays - four hot-plug disk bays
- Hot swap power supply.
- Hot plug PCI I/O slots - 12 total (7 x 32Bit plus 5 x 64Bit)
- Hot swap cooling devices.

3.9 Enterprise Storage Technology

This section discusses IBM technology as it relates to ServeRAID, SSA, Fibre Channel, Tape Backup, and Rack Solutions.

The definition of enterprise storage has changed dramatically over the years, from centralized, mainframe-based storage to storage distributed across a client/server network. Your enterprise has changed as well: you've become global, thanks in large measure to the Internet, which has helped introduce your business to the world. The change in your customers is equally dramatic because they demand answers and require information in real time. Today. Right now.

Our storage patent leadership and long-term involvement with every facet of enterprise storage enables us to understand the diverse information technology challenges that face your growing business.

3.9.1 IBM Hard Disk Drives

It's all about your business data. Data your end users need to keep you up and running. To keep you on a par with your competition. To help you surpass them. Your reliance on hard disk drive storage is why we test every one of our drives for system compatibility, component reliability, optimum performance and data integrity. All of which underscores our commitment to deliver the highest performing, most reliable hard disk drives available to you regardless of where in the world you do business.

You can depend on IBM hard disk drives to take your business where you want it to go, because we invented magnetic disk storage technology. IBM has acquired an expansive patent portfolio for such advances as MR (magneto-resistive) Heads, No-ID sector formatting and more. Today, our extensive line of hard disk drives includes industry-leading SCSI and Serial Storage Architecture (SSA) interfaces and hot-swap hard drives, with storage capacities up to 18.2 GB at rotational speeds up to 10,000rpm.

3.9.2 IBM Tape Backup Solutions

It's business truth that system downtime can cost you money. Unproductive end users. Customers who go else-where because you're not up and running. Given the pace at which your business operates, downtime isn't a viable option. That's why reliable, high-capacity tape backup that protects your data is a necessity. You can restore availability to your users across your enterprise in the event of a system failure. You can make copies of your data to test for future challenges such as the Year 2000 (Y2K). Tape backup can help you avoid what a recent study by 3M estimates to be the cost of re-creating lost data: \$3,000 per MB.

For a total backup solution, the value and proven reliability of IBM tape backup technology makes good sense for your business. Tape offers the lowest cost per gigabyte for backup and archival storage and uses the lowest cost media. IBM offers a complete line of tape products in most popular industry-standard tape formats and capacities (4 GB to 1 TB).

Available Tape Formats:

- DLT
- Magstar

- DAT
- 8mm
- Travan NS

All tape drives for Netfinity servers ship with award-winning Seagate Backup Exec software at no extra cost. This upgradable, industry-leading software provides comprehensive protection for your data, as well as your investment in it.

As your backup needs increase, consider the storage management possibilities of IBM tape libraries. With hardware data compression, tape library automation delivers up to one terabyte (1,050 GB) of data backup capability. For large enterprise storage management, look at IBM's ADSTAR Distributed Storage Manager (ADSM). This renowned family of software products can integrate your unattended network backup and archive with storage management and powerful disaster recovery planning functions.

3.9.3 IBM RAID for Netfinity Servers

Your business-critical applications and data are the key resources around which your enterprise revolves, they must be protected. Utilizing a Redundant Array of Independent Disks (RAID) adapter is an excellent storage solution for businesses of all sizes to ensure the high availability and protection of these assets. The recently introduced third-generation of the IBM ServeRAID-3 family offers exciting new possibilities for expanding your RAID capabilities.

When paired with the Netfinity 7000 M10 server, the ServeRAID-3H provides phenomenal performance, with room for growth tomorrow, making it particularly suitable for clustering. The 3H also lets you add an optional battery backup (with its own cache) that can actually save your data even if the system loses power. For less demanding network environments or for departmental use, you should consider the Netfinity 5000 server teamed with the single-channel ServeRAID-3L. Both ServeRAID-3s offer Ultra2 technology capable of incredibly long SCSI cable lengths at impressive data transfer speeds.

The IBM Serial Storage Architecture (SSA) Technology is yet another innovation in storage performance. A full range of SSA RAID options complements our expansive portfolio of storage options, all of which provide industry-leading availability of data stored in your enterprise's networked systems.

3.9.4 Fibre Channel: The Future of Storage...Today

IBM's new Netfinity Fibre Channel Storage products are designed to support the massive data storage requirements of your business-critical applications. These products offer astounding scalability and growth solutions. High-speed data transfers and the wide range of storage configurations available give you the ability to build enterprise class storage systems to meet the growing demands of your business. From basic high-speed configurations offering continuous data access and protection, to long distance, high availability and fully redundant configurations for critical disaster recovery capabilities, Netfinity Fibre Channel solutions meet your enterprise storage requirements.

With data transfer rates of up to 100MBps, access and storage requirements of your high performance business systems are fulfilled. Whether its large database applications, such as enterprise resource planning systems, multimedia applications, decision support systems, or your data warehouse,

you'll benefit from fewer storage bottlenecks. And redundancy support of the Netfinity Fibre Channel Storage solutions and failover capabilities of your operating system will help keep your data and applications ready for business.

The Netfinity Fibre Channel RAID Controller offers six (6) high-speed Ultra-2 LVDS channels, each capable of supporting an IBM EXP 15 Expansion Enclosure with ten (10) Wide Ultra-2 SCSI hard disk drives. This allows for up to sixty (60) disk drives and over one (1) TB of data storage per controller. Add the Netfinity Fibre Channel Hub, with 7 ports supporting both short wave and long wave capabilities, Netfinity Fibre Channel PCI Adapters, and additional Netfinity Fibre Channel RAID Controller Units and your storage capabilities can easily expand to nearly limitless possibilities. All this supports your business-critical systems while protecting your current SCSI hard drive investment. The Netfinity Fibre Channel Hub can also be used to implement multi-node server clusters and storage systems throughout your business, bringing the benefit of reliable high-performance to your users. Flexibility is built-in to Netfinity Fibre Channel Storage solutions. You can configure according to your changing business requirements, from intra-rack, room to room, across the campus or even across town.

3.9.5 Tape Backup Expansion Solutions

To keep pace with the growth of your business, you will need to expand your tape backup systems. The flexible Netfinity NetMEDIA Storage Expansion Unit EL can help. It can house two DLT or four 8mm, DAT or Travan tape drives in your Netfinity rack or stack solution. It features dual independent power supplies and high-efficiency cooling. When used with the Netfinity NetMEDIA Systems Management Adapter and appropriate software, the NetMEDIA EL can help you implement tape mirroring or arrays for even greater data security.

Additionally, the flexible NetMEDIA Systems Management Adapter provides the added benefit of an Ultra-2 SCSI channel (also known as low voltage differential SCSI or LVDS) that handles both Ultra2 and earlier generation SCSI devices. Plus, the onboard electronic self-termination function lets you power off the NetMEDIA unit without disrupting the controlling server.

3.9.6 SCSI Hard Disk Expansion Units

For increased hard disk drive storage expansion, look to the 10-bay Netfinity EXP15. The 3U rack-mountable EXP15 offers high data availability thanks to redundant components that include powerful hot-pluggable cooling and power supplies. Further, by allowing increased cable distances, the EXP15 can be located away from the server, promoting the growth of DASD farms and other flexible storage configurations. The Netfinity EXP15 also offers investment protection by providing a high-performance platform for mounting future generations of hard disk drives required by your growing business.

Added systems management functionalities of the EXP15 include the provision for remote monitoring and diagnostics capability. For example, if the system self-diagnoses a potential overheating problem, it can automatically send a high temperature alert to your systems management console for immediate action.

3.9.7 SSA Storage Expansion Units

The IBM SSA disk subsystem is another high-availability enterprise storage management solution. Like the Netfinity EXP15, its rackable design has hot-pluggable cooling and power systems and can operate at greater distances from the host server for increased floor space flexibility. If SSA is your storage technology of choice, you should know that IBM offers the widest range of SSA components in the world.

3.9.8 IBM Rack Solutions

Let's assume your business and the technology that supports it grow in concert, but your floor space remains the same. Even if your IT budget can't support a new building this year, you can still effectively manage your floor space while maintaining easy access to crucial components with an IBM rack solution. An IBM rack solution can conveniently house multiple servers, SCSI or Netfinity Fiber Channel Solutions, an uninterruptible power supply (UPS) or NetMEDIA and EXP15 Storage Expansion Units.

For scalability and investment protection, all IBM NetBAY products promote easy configuration flexibility, tidy, efficient cabling and easy access to key system components within the enclosure. Plus, you can connect additional NetBAY enclosures to gain "intra-rack" solutions and capabilities. It's the perfect building block for protecting your equipment investment because it lets you easily implement rack and/or stacking expansion options as your business expands.

With the innovative NetBAY3 stackable enclosure you can add up to three pedestals for a total 9U of expansion space in your Netfinity 7000 or 5500 server system. If you need more, consider the NetBAY22 half-height rack with 22U capacity. Or move to the full-size IBM Netfinity Rack with 42U capacity, capable of holding up to eight servers (or five when configured with a console and UPS).

Take some time. Carefully consider your customer's storage needs. And then choose IBM storage solutions.

We highly recommend that you order the latest redbook, *Implementing ServeRAID SCSI and SSA RAID Disk Subsystems*, SG24-2098-01 to obtain a complete update and examples on how to fully utilize these adapters. If this book is not available, you should use the user's guide shipped with the adapters.

For more information on IBM Storage solutions, as well as other solutions (memory, networking, monitors and video cards, and processor upgrades) we recommend that you visit the following IBM Web site:

<http://www.pc.ibm.com/us/accessories/netfinity/index.html>

3.10 Domino Servers by Function

If your Notes network grows to a very large size, it is more manageable to dedicate certain Domino tasks per server. Limiting a few Domino tasks to a server is like giving it a role. For example, if you specify your Domino server should primarily support the HTTP or Web task, your Domino server takes on the role of a Web server. These Domino tasks or server roles are explained below. Remember that these are not hard and fast specifications, but just common server roles that many customers implement in their Domino environments.

3.10.1 Mail Servers

Mail servers are used in the Notes network to store and route mail messages between users, servers, and Notes applications. Dedicating one or more servers to these responsibilities provides the following:

- Easier administration of mailboxes and users. This provides for improved auditing and simplified installation processes.
- Minimize the processor load on the server by reducing the number of databases that need to be replicated with other servers, and by reducing the requirement to index databases on a regular basis.
- Strategic organization of users per mail server will significantly reduce the impact of traffic on your network. Mail traffic delivered to users on the same server will not be routed across the network. This will also allow for immediate mail delivery by eliminating the need to create mail delivery schedules.
- Provide for more specific requirement goals, such as response time, thereby establishing a better scalable model for expansion as usage and storage needs increase.

3.10.1.1 Server Issues

For mail routing we have found that the first place to look for a potential bottleneck is in the disk subsystem. Please refer to 4.6.3, “Disk Bottleneck” on page 115 where we discuss how to identify server bottlenecks in your Domino server. After reviewing the disk subsystem, if you determine that the disk resources are adequate and the disk subsystem is properly tuned the following is a suggested order to look at other potential bottlenecks: memory, CPU, then network adapters.

3.10.2 Database (or Application) Servers

Database servers are used to store Notes application databases, such as discussion, tracking, news, and online documentation databases. Customized applications such as help desk databases, support tracking systems, and so on, should typically be dedicated to specific servers. Notes application databases are becoming very robust and sophisticated systems. Consequently, interaction between numerous databases is quite common. Placement of interdependent databases across multiple servers could produce increased stresses on the servers and the network depending on the amount of interaction required by applications. Dedicated servers to database tasks will provide the following benefits:

- Enable you to group applications by their usage. This will typically reduce replication needs, and simplify administration of both server and database security.
- Allow you to optimize server performance for application database usage without regard to mail usage.
- Provide for more specific requirement goals, such as response time. Thereby establishing a better scalable model for expansion as usage and storage needs increase.

3.10.2.1 Server Issues

Application server usage has a large impact on the CPU and disk subsystem, more than any other Domino server function. Specifically, replication and indexing are two of the most resource consuming tasks on a Domino server. Both of these tasks are directly associated with application type databases. This is because application databases are larger in size (such as the N&A book for large domains) and are updated more frequently than other databases.

3.10.3 Web Servers

Web servers are used to display information to Web browsers. This information may be stored as text files on the server stored in HTML format, or from Notes databases. Any information that is seen from a Web browser should be placed on a single server, to reduce the stresses on the server and network as discussed in 3.10.2, "Database (or Application) Servers" on page 50. Dedicating servers to Web tasks will provide the following benefits:

- Allow you to tighten security on one server for both intranet use or Internet connection. Only place databases that should be viewed by Web browsers on this server, especially if this server is to be browsed by Internet clients.
- Allow you to optimize server performance for Web usage without regard to mail usage. Provide for more specific requirement goals, such as response time, thereby establishing a better scalable model for expansion as usage and storage needs increase.

3.10.3.1 Server Issues

Web server usage is similar to application server usage in that it has a large impact on the CPU and disk subsystem. This is due to the dynamic HTTP task that converts information in Notes databases into HTML format and accesses HTML documents stored on disk.

3.10.4 Remote Access Servers (RAS)

Remote Access Servers are used to provide dial-in access to Notes using a modem. As previously discussed, we recommend you consider RAS solutions such as Warp Server RAS or Microsoft RAS if you have more than a few users at a time who need to dial in to your Notes network. If you decide to use a Domino remote server, consider the limitation that Domino can only support a combined total of 64 LAN and COM ports. These limitations are dependent upon the internal design of Lotus Domino. Therefore, designing remote servers with high-end IBM Netfinity 7000 server platforms will not provide additional remote user access.

3.10.4.1 Server Issues

Utilization of serial ports produces a significant impact on CPU utilization. Due to the slow throughput rates of serial ports, transactions being processed by the CPU are queued for transmission. This queuing process results in prolonged wait states for the CPU, at which times no other CPU processes can occur. The use of intelligent multiport communications adapters significantly reduces the impact on the CPU and is therefore highly recommended.

3.10.5 Hub Servers

Hub servers are designed to be used as dedicated servers to route mail and replicate databases to other servers within the network. This typically results in hub servers only being used in larger Domino environments to minimize the performance impact on production servers.

Since the typical use of hub servers within the concept of the Domino topology is only to facilitate communications between other servers, users normally don't have access to these servers (or know they exist) which eliminates the requirements for large numbers of simultaneous connections. Therefore, you might assume that a hub server would not necessitate the need for a high-end server platform to fulfill the requirements. However, this is typically not the case. You must consider the fact that each hub server is required to support the storage of all common databases and be capable of replicating these databases to each additional server in a timely fashion. Each hub server must also be capable of routing mail between servers ensuring that messages are delivered within an acceptable time frame. To fulfill these requirements, each hub server should be designed with extensive CPU and disk subsystem capabilities.

3.10.5.1 Server Issues

Hub servers generally provide database replication and mail routing to several other servers within the network. These servers produce heavy loads on all components of the hardware. Generally, hub servers are designed with high-capacity scalable platforms. Other suggestions:

- The system should have a fast disk subsystem for replicating and archiving.
- Full-text indexing should not be allowed on these systems since index updates degrade performance and users will not be accessing these databases. Please refer to Chapter 4, "Server Configuration and Optimization" on page 79 for more details on how to disable full-text indexing on your server.
- Restrict access to the server for all users so they do not degrade performance.

3.10.6 Passthru Servers

Passthru servers are used to provide a "stepping-stone" to allow users to connect to other Domino servers without knowing all the routing steps required to make the connection. Implementing passthru servers will provide your remote Notes users and Domino servers access to servers within your network, by way of dial-in or X.25 connections.

Passthru servers are also used to provide connectivity between servers or workstations on the LAN that do not share a common protocol. For example, if a workstation is only configured to communicate via IPX and the destination servers only communicate via TCP/IP, a passthru server using both protocols on one or more adapters could allow communication between these two systems.

Since passthru servers are only employed to provide connections from one Domino host to another, the requirements for the disk subsystem are confined to the support of the operating system and the Domino server application. This leaves the primary concerns for passthru server designs with the CPU and network interface. Depending on the combined number of simultaneous connections that your remote server must support, each of the high-end IBM Netfinity platforms is appropriate.

Each passthru server is configured to specify which destination servers the passthru server can access, and which users and servers can use that server for passthru capabilities. This feature allows you to control the maximum number of Notes clients and Domino servers that can utilize each passthru server at any given time, as well as a secure environment.

3.10.6.1 Server Issues

Passthru servers are only used to provide access to other servers and Notes services within the network. Therefore, the impact to server hardware is confined to the CPU and the network interface in that respective order. Selection of the appropriate hardware is dependent only on the number of sessions that need to gain access through the server.

3.11 NotesBench Reports

Before we explain how to determine what server(s) you will need for your Notes network, we explain a performance testing tool called NotesBench developed by Lotus. NotesBench enables hardware vendors to directly provide Lotus Domino customers with Domino server capacity and performance information on various platforms and configurations. This basically means that all server hardware from various manufacturers is tested using a common benchmark so that you can compare one server to another on a level playing field. These NotesBench reports provide a great deal of information on the capacity of a Domino server. We use a published IBM Netfinity 7000 M10 server NotesBench report running Lotus Domino Server 4.62a for Windows NT 4.0 as an example to explain what information can be obtained from these reports.

3.11.1 Audited Information

Always make sure that you are interpreting an audited NotesBench report. This is a report that was performed to Lotus's standards, and so all vendors products were tested equally.

The latest official NotesBench reports are available at:

- <http://www.notesbench.org>
- <http://www.pc.ibm.com/techconnect/tech/benchmark.html>

3.11.2 Interpreting a NotesBench Report

NotesBench provides a workload to a test system by way of a driver. A driver can execute one of nine NotesBench workloads. These workloads are listed in Table 4:

Workload Type	Description
Replication Hub	This workload models a replication hub server that exists to propagate changes among a collection of other servers.
Mail Routing Hub	This workload models a mail routing hub server that exists to route messages to other servers (a "pure" router) and possibly to also deliver messages to local users.
Idle Usage	This workload models users who are connected to a Domino server, but who are not doing anything. It establishes the upper bound on the number of sessions that a particular server configuration can support.
Mail Only	This workload models new Notes users who use the Domino server only for mail communication. This type of user may only exist in an organization for three or four months before moving on to behave more like a Mail and Shared Database Notes user.
Mail and Shared Database	This workload models active Notes users who use the Domino server for performing mail and simple shared database operations.
Groupware_A	This workload models "power users" who use the Domino server for sending large mail messages, adding documents with attachments to shared databases, performing full-text searches, and replicating changes from their local workstation to the server. Currently, Groupware_A is the most strenuous workload within the NotesBench suite of workloads.
Web Walker	This workload models users who execute simple Web browsing strategies such as viewing pages.

<i>Table 4 (Page 2 of 2). NotesBench Workload Types</i>	
Workload Type	Description
Web Buyer	This workload models users who perform more complex activities such as entering data into a form, uploading and downloading data. This is more of a commercial-transaction.
Calender and Scheduling	This workload models various user operations within the Calendar and Scheduling environment, such as looking up the free time of a person or booking a room.
Shared Discussion	This workload models a server of active users, who are only performing heavy shared database operations. The operations occur in a shared database, navigation of unread documents, additions, and updates to documents in a shared database.
Cluster Mail and Shared Database	This workload models a server in a cluster supporting several power users who are sending large mail messages, adding, updating documents to shared databases with replica copies throughout the cluster.

Although we mention the several workload types available, our example of Notesbench data obtained from the IBM Netfinity 7000 M10 illustrates Mail-Only. The Mail-Only workload was run on a four-processor configuration.

If we look at part of the IBM Netfinity 7000 M10 NotesBench report for Lotus Domino 4.62a we can see the following:

Test Script	Maximum Users	NotesMark (tpm)	Response Time (sec)	\$/User	\$/NotesMark
Mail-Only	6,100	8,816	0.419	\$11.25	\$7.78

We can see key values such as our server can support 6100 active Mail-Only users. The response times are worth noting for the these particular users. This is equivalent to the time a user would have to wait for a request. NotesBench auditing rules allow a maximum response time of up to 5 seconds, although most vendors and customers prefer sub-second response times.

Note

Be aware that these results are from a server that has been operating in a test environment which is generally running at a very high CPU utilization (a much higher CPU utilization than you would prefer for a production server). So do not use these values directly in your capacity planning calculations. Simply be aware that these results are upper limits only, and that in practice the number of supported users and hubs for your Domino server will be lower.

From the NotesBench Report

The IBM Netfinity 7000 M10 server with RAID-1 demonstrated leadership performance running Lotus Domino Server Release 4.62a on Microsoft Windows NT Server Version 4.0 with Service Pack 3. The IBM Netfinity 7000 M10 was configured as a partitioned server with four 400MHz Intel Pentium II Xeon processors, 2.5 GB of memory, and seventeen 4.51 GB 10K rpm hard disk drives. The server under test supported a Mail-only workload of 6,100 active mail users. Three Domino 4.62a Partitioned Servers ran on the IBM Netfinity 7000 M10 system, using identical copies of the Public Name and Address Book. Mail messages generated by users connected to the three partitioned servers were all routed to six destination servers.

In addition to the server under test, the benchmarked client/server network used six destination servers, 15 client driver systems, and three controller client systems. All systems were connected via six 100 Mbps Ethernet LAN segments, using the TCP/IP network protocol.

From this configuration, decide if there are any settings that may be different in your server configuration. Compare the speed of the processors from the server you are considering to the one measured in the report. Maybe the operating system you will use will not be Windows NT or maybe you wish to have more RAM in your server. Also, note what operating system service packs have been applied, and what network card driver levels are used. You do not have to mimic the server that was tested, but you may find valuable configuration tips that you can apply to your Domino server. A well-written NotesBench report will describe why the tester arrived at the optimum server configuration. This can aid you in your server tuning as well.

You may notice that RAID-1 was used for this NotesBench test. Although RAID-0 would have given even better performance values, you will have a higher fault tolerance level with RAID-1 as well as RAID-5 in your environment. You should consider this in your capacity calculations.

A more detailed and complete NotesBench report of IBM Netfinity 7000 M10 can be obtained from the Web site <http://www.notesbench.org> as well as updated reports on the IBM Netfinity 5500 M10.

These results demonstrate that the IBM Netfinity 7000 M10 system provides leadership performance and price/performance as a hardware platform for the Domino Server environment.

3.11.3 Determine Your Client Distribution

We have determined that the type of NotesBench workload used to test a Domino server will have a dramatic effect on how many users you can support on your server; therefore, you need to determine what distribution of client types you have in your company. For example, what distribution of your users fit the Mail-Only, MailDB and Groupware_A client types. If you do not know the distribution of your client types, a common percentage distribution is given in Table 5 on page 57:

<i>Table 5. Common Client Type Distribution</i>	
Client Type	Percentage Distribution
Mail-Only	20
MailDB	50
Groupware_A	30

We find that after a few weeks of using Notes, most users have advanced from simply using a few Notes mail functions to also using several shared databases on your server. We also find that many Notes users become very familiar with advanced Notes functions. This explains the high percentage distribution of Groupware_A clients (30%).

In most Domino environments a typical user profile will change from a Mail-Only user profile to a MailDB user profile. If all users are very familiar and capable of using the features of Domino, they are classified as Notes "power-users". This group of users are represented as Groupware_A users in the table.

Note: The more skilled your users are, the more likely they will send large mail messages, add attachments, and perform full-text searches. This will cause a higher load on the server, and reduce the number of users you can support.

Many vendors currently provide their own rule-of-thumb guidelines to measure the number of users and different client types that can be supported on a server. Most vendors use the CPU value measured in the report and scale the user values based on a maximum CPU utilization value of 50%. This provides extra capacity for tasks that are not running in the test environment such as agents, replication and indexing of a larger Public N&A Book.

These guidelines vary from vendor to vendor but all provide a degree of comfort in their calculations so that you would feel confident that your server would support the load that has been quoted. The solution now is to replace these subjective guidelines with a tool developed by Lotus called Server.Planner. This tool allows you to enter the exact numbers of users and other requirements, and Server.Planner will select which server will meet your requirements from a list of servers in its database. This tool is discussed in detail in the next section.

3.12 Server.Planner

Before Lotus developed Server.Planner (pronounced Server dot Planner) there were primarily two ways in which you could determine if your choice of server would have sufficient capacity before it went into production. The first was to set up a simulated environment and test the capacity of your server using your known client distribution. For example, if you determined that 90% of your clients are Notes power users, you would test the server with a 90% Groupware_A NotesBench client script. These scripts could be modified to send more mail or open more databases than the standard script. You could also tune the server by adding more memory, for example, to better meet your requirements, and test this new configuration again. This would give you a great level of confidence that the server you have selected will support your user load when it goes into production.

3.12.1 Where to Get Server.Planner

Server.Planner is available from the NotesBench Consortium Web site at <http://www.notesbench.org>. However, you must be a registered member of the NotesBench Consortium. Currently, IBM along with Advantis, Compaq, Digital, Lotus, HP, NCR and KMDS are members of the NotesBench consortium. All members of the Consortium will work with Lotus to ensure data quality, as well as ensure Server.Planner is enhanced in the future. Server.Planner is also shipping with Lotus Domino 4.6. The IBM vendor database will list all of the IBM servers that have been tested using the NotesBench testing process.

3.12.2 Getting Started with Server.Planner

The first step is to download the Server.Planner data set into the \NOTES\DATA directory on your computer. Also download the IBM vendor database and store this in your \NOTES\DATA directory as well. Vendor databases all start with the prefix CP. Therefore, the IBM vendor database will be named CPIBM.NSF. This name may change when the database is finally posted to the IBM Web site.

To start Domino Server.Planner, open the Domino Server.Planner Analyst database (DSP.NSF). You should see the following Notes Navigator screen.



Figure 16. Domino Server.Planner Navigator Screen

We will now query the IBM vendor database to select the appropriate IBM Netfinity system.

3.12.3 Query on Data

To query the information in your IBM vendor database:

1. Click on the **Query on Data** graphic in the Navigator.
2. Choose **Create** and **Query on Data** on the menu bar.
3. Select the tasks your users can perform. Select Replication Hub if your server does database replication. Select Mail Hub if your server handles mail. Select User Tasks if your server performs end-user database activities. You may also select a combination of any of these activities.
4. Enter the specific information needed for each task.
5. Select an acceptable end-user response time for a new person connecting to your server.
6. Select the percentage of responses that fall within range. Specify how often you want the server to provide the given response time.
7. Select the Query Tolerance % range.
8. Select the operating system platforms. Click the **Select All** button to query all operating systems listed.
9. Click the **Vendor Selection** button to select vendor databases. Downloaded vendor databases appear. You can query all your vendor databases or choose a smaller set.
10. Click the **Query** button on the action bar to get results.

Figure 17 is a sample Query on Data screen for the server tasks, response time and operating system platform.

The screenshot shows a graphical user interface for configuring a query. At the top is a menu bar with icons and labels for Exit, Save Query, Save Query and Results, Query, Graph Responses, Graph Costs, and Help. Below the menu bar, the interface is organized into several sections:

- Tasks:** A box containing three checkboxes: Replication Hub (unchecked), Mail Hub (checked), and User Tasks (checked).
- Response Time:** A box containing three radio buttons: Fast (selected), Medium (unchecked), and Slow (unchecked).
- % Responses fall within Range:** A box containing four radio buttons: 50 (unchecked), 75 (selected), 90 (unchecked), and 95 (unchecked).
- Query Tolerance % Range:** A box containing four radio buttons: 0 (unchecked), 5 (unchecked), 10 (selected), and 25 (unchecked).
- Specify OS Platform(s) to Include:** A box containing two columns of checkboxes. The first column includes AIX/RS6000, AIX/SP2, HP-UX, NetWare 3.12, NetWare 4.1, and NT 3.51. The second column includes NT 4.0, OS/2 2.11 SMP, OS/2 Warp Connect, OS/2 Warp Server, Solaris 2.5, and Win95 1.0. Several of these are checked.

At the bottom of the form area, there are three buttons: "Vendor Selection", "Select All", and "Deselect All".

Figure 17. Server Requirements Query on Data Screen

Figure 18 on page 60 is a sample Query on Data screen for the Mail Hub and number of users and their distribution.

Mail Hub:

5 Spokes Messages/Hour MB Routed/Hour

Number of Active Users: 700

Mail (Low): 20 %

Mail & Database (Medium): 50 %

Groupware (High): 30 %

Figure 18. Mail Hub and User Requirements Query on Data Screen

We can see from our example in Figure 17 on page 59 and Figure 18 that we require a server that will support 700 active Notes clients and also perform the function of a mail hub to five spokes. We require a fast response time (less than one second) at least 75% of the time, and would like to choose a server that will meet this requirement within a 10% variance. The major Intel-based operating systems have been selected to run on this server such as NetWare, Windows NT and OS/2. The standard user distribution of 20% Mail, 50% MailDB and 30% Groupware has been entered. Once we have entered our requirements we can proceed to query the vendor database.

3.12.4 Query Results

Domino Server.Planner presents the results of a query at the bottom of the Query on Data form after you click the **Query** button.

Your results are summarized in a table containing Vendor Name, Machine ID, Machine Description, Response Time, CPU Type Number and Platform. Given our previous example requirements, Figure 19 shows that the IBM Netfinity 5500 will meet this user and mail hub requirement. Please note that these results are from a beta version of the IBM vendor database. Results may vary with the officially released database.

Database: Summary: Results of query for 700 users, 5 spokes and fast response time. Beta version of IBM Vendor

Results:

Platform:	Vendor:	Machine ID:	Description:	Response Time:	CPU/#:
Win NT 4.0	IBM	8660-XXXX1	IBM Netfinity 5500	Fast	Pentium II/400

Figure 19. Results from Server.Planner

3.12.5 Graphical Information

Server.Planner creates graphs of results when you click the **Graph Responses** or **Graph Costs** button after you get your query results. Graphs are sorted by Machine Name and Response Time, or Machine Name and Cost. Click the **Sort** button to toggle between ascending and descending order. If you installed Lotus Components, you can choose from many graphical formats. At the time of writing, the Lotus Components software was only available for Windows 95 and NT.




Figure 20. Average Response Time Output of Results




Figure 21. Servers Sorted by Cost

3.12.6 Saving Queries and Results

Click the **Save Query and Results** button to store both query and results in the Server.Planner Decision Maker Database. Enter the name of your server and a file name for the Decision Maker Database to create. Store the results of many different queries in the same query database. Server.Planner provides specific machine information to support your results.

3.12.7 Standardizing Your Server Hardware

Using the Server.Planner tool you can choose the appropriate server for each location in your organization. We recommend that you limit your choice of servers to only one or two models. For example, you may have to size all the Domino servers for a large organization using the Server.Planner tool. You will determine a list of perfectly sized servers, and this list may include one or more of the following models: IBM Netfinity 3000, 5000, 5500, 5500 M10, and the 7000 M10. However, many organizations standardize on only one or two models for several reasons:

- Increased staff knowledge in their servers by only needing education in one or two models, rather than the complete server range.
- Allows increased on-site part inventory for those server models.
- Easier to order and install for a customer who is implementing a large network in a short period of time.
- Easier to diagnose hardware and software-related problems by implementing a standard operating environment of particular models, network cards, BIOS levels, RAM and disk space levels.

You can then select one or two server models which will meet the minimum load requirement for each location. Therefore, let us suppose that you need to install 15 Domino servers and the Server.Planner tool has chosen the following number of IBM servers:

- 4 x IBM Netfinity 7000 M10
- 8 x IBM Netfinity 5500
- 3 x IBM Netfinity 3000

You may decide that the minimum server platform will be the IBM Netfinity 5500, and the high-end server model will be the IBM Netfinity 7000 M10. This would reduce your server selection to two server models. The added advantage is that the three sites that have the Netfinity 3000 model have extra capacity for the future.

Also try to standardize on the amount of RAM and disk space you have in your servers. This may not be possible in all cases, but try to reduce the variation in the systems you are purchasing as this has many long-term advantages for an enterprise Domino environment. (See 3.13, "Other Server Requirements and Recommendations" on page 63 for additional information.)

3.13 Other Server Requirements and Recommendations

Using the Server.Planner tool, we can easily identify which IBM Netfinity system will satisfy your Domino server requirements. There are still some important sizing issues which we still need to quantify, and these are the amount of RAM and the amount of disk space storage you will need. We provide some other server recommendations for your server to ensure high performance and smooth operation.

3.13.1 Server RAM

The amount of random access memory (RAM) you need in your Domino server is based on two factors. The first is a base amount of RAM to meet your operating system requirements and general Domino operation. The second variable component depends on how many Notes and Web clients have sessions with your server, and other tasks your server may perform.

Use the following formula to calculate how much RAM you need in your server:

Minimum Memory Requirement

Platform RAM component + (Number of concurrent users /3) MB

The Platform RAM component is dependent on the operating system Domino is running on. A list of the most popular Intel-based operating systems is shown in Table 6:

Operating System	Platform RAM Component
Microsoft Windows NT 3.51 and 4.0 (Intel, Alpha)	64 MB or more recommended
IBM OS/2 Warp Server 4 and Warp Connect	64 MB or more recommended
Microsoft Windows 95	24 MB or more recommended
Novell NetWare 3.12 and 4.1	96 MB or more recommended

Remember that this RAM requirement only includes the operating system and Domino running on your server. If you have any file and print services, backup and recovery or anti-virus software running on your server, you will also need to add additional RAM to your server.

Do not install small increments of RAM in your server. If you have a requirement for 150 users on a Domino server, this calculates to approximately 114 MB (64 MB + 150/3 MB). However, you would not try to install exactly 114 MB of memory in your server. Choose a convenient memory arrangement which uses the minimum amount of memory DIMMs, and is above your calculated minimum RAM amount. This would be two DIMMs of 64 MB each, giving a total server memory of 128 MB.

This algorithm is appropriate for mail and application servers, and mail hubs. However, it is not appropriate for replication hubs. Usually, replication hubs are heavily used but have very few, if any, sessions open to active users.

Domino provides the ability to have multiple replicators on one server to be able to handle more concurrent replication requests. However this requires more memory, in fact about another 4 MB per additional replicator.

We recommend that you always use Error Checking and Correcting (ECC) RAM. All IBM Netfinity systems currently support this memory. This memory is less prone to errors than normal parity memory.

3.13.2 Server Disk Space

The amount of disk space your Domino server requires can differ greatly depending on what tasks the server needs to perform. The Domino services that require disk space are outlined below so you can easily determine how much disk space your server requires.

3.13.2.1 Operating System and Domino

First determine how much disk space you need for your operating system and Domino software. Lotus recommends a value of 500 MB. This is a good initial value for several reasons:

- It covers the minimum operating system requirements for Microsoft NT Server, OS/2 Warp Server, and Novell NetWare. These are the most popular Intel server operating systems.
- It gives you enough disk space in case you decide to install all of the Domino features, now or in the future.
- Provides sufficient disk space for future enhancements such as the SMTP MTA or Domino.Fax gateway.

3.13.2.2 Operating System Paging File

Both Windows NT Server and OS/2 Warp Server use a paging file stored on disk which allows the operating system to provide a virtual address space larger than the physical RAM size. Therefore, as you run more and more applications on your server, you will reach a point when the amount of physical RAM is exhausted. This will not cause your programs to fail. The operating system will simply page memory blocks between physical RAM and disk to meet the application's needs. For OS/2 Warp this paging file is called SWAPPER.DAT; for Windows NT this file is called PAGEFILE.SYS. Because accessing a disk is slower than RAM, a server's performance will dramatically decrease the more it has to page memory blocks between RAM and this paging file. Therefore the more memory in your server the better.

By default, Windows NT sets the PAGEFILE.SYS to the size of physical memory plus 12 MB, but making it larger often resolves virtual memory shortages.

It is recommended that you make the initial size of this paging file for Windows NT equal to the amount of memory you have in your server plus 10%. This paging file can be smaller for Warp Server. This is discussed in more detail in Chapter 4, "Server Configuration and Optimization" on page 79. If you had a smaller paging file that grows and reduces in size, the operating system has to vary this paging file size depending upon application load. This requires more CPU utilization than is necessary. Having a large paging file has other advantages. It has been shown that a server will gain performance improvements if this file is a contiguous set of disk blocks, because it reduces the chance of the paging file becoming fragmented.

We recommend additional disk space equal to the amount of RAM in the server plus 10%. Therefore a server fitted with 100 MB of RAM will need 110 MB of free disk space to store the paging file on disk.

3.13.2.3 Mail Users

You should add approximately 30 MB of disk space per mail user. This may seem a very high value. However, most users when they start using Notes have the tendency to keep a copy of all their mail correspondence. This happens with most mail systems. Also, mail users from time to time store mail documents with large file attachments, such as executables (EXEs) or compressed files (ZIPs). The disk space required for a user's mail file will also increase as they use features such as calendar and scheduling. Providing a large amount of disk space for mail users initially gives you plenty of time to monitor the size of user's mail databases, and prepare you to add more disk space if necessary.

3.13.2.4 Application Servers

The types of application databases that can run on Domino can vary widely. Some databases are internal ordering systems that do not grow quickly in size, while others are discussion databases that are updated by many users each day. A standard Notes database is one that grows slowly in size, while a discussion type is more dynamic and has the tendency to grow quickly. We recommend that you initially provide 30 MB per standard Notes database and 100 MB per Notes discussion database.

3.13.2.5 Mail Hubs

Mail hubs do not need as much disk space as other servers because they only route mail. However, if a link to a remote server is unavailable and users have sent many e-mails with large attachments, then your mail hub can quickly run out of disk space. These mail items are temporarily stored in the MAIL.BOX database before they are delivered to another server or user. You should initially provide 100 MB of additional disk space for your mail hub. If you have a very large number of users dependent on this mail hub, then add additional disk space as a precaution.

3.13.2.6 Replication Hubs

These servers need large amounts of disk space to store complete replica copies of databases. These databases are replicated from other servers to remote spoke servers. Determine how many application servers replicate to this hub. If all the databases are different on each application server, then the total amount of disk space you will need would be the total of the disk space on each of the application servers added together. If some databases are common across some of the application servers, then your total disk space requirement may be lower.

3.13.2.7 Summary of Disk Space Requirements

In summary, you can use the following formulas to calculate the minimum disk space you will need on your Domino server:

- 500 MB for operating system and Domino software.
- An additional page file disk space requirement: amount of RAM + 10%.
- If application or mail server, add the following disk space:

(Number mail users + Number standard databases) x 30 + Number of discussion databases x 100 MB

- If mail hub, add 100 MB of additional disk space for MAIL.BOX.
- If replication hub, add up total amount of disk space on application servers that this hub replicates with.

For example, we have a server that needs to support 100 mail users. The server is fitted with 128 MB of RAM. There are only two discussion databases planned for this server, and the server acts as a mail hub to a remote Domino server.

The server would require the following disk space:

500 MB for operating system and Domino

128 MB + 128/10 MB for paging file

30 x 100 = 3000 MB for 100 mail users

2 x 100 MB for the two discussion databases

100 MB for the MAIL.BOX mail routing database

A total of 3940 MB of disk space is recommended as a minimum.

3.13.3 RAID and Drive Controllers

Several of the IBM Netfinity servers have Redundant Array of Inexpensive Disks (RAID) technology as standard. This allows data to be striped across independent hard drives, but is seen by the operating system as one large disk drive. The three most popular levels of RAID that are used are RAID 0, RAID 1 and RAID 5. RAID 0 generally gives the highest disk performance on a server, but there is no fault tolerance if one of the individual drives fails. RAID 1 provides excellent fault tolerance, and is essentially disk mirroring. However, half of your disk drives are used for fault tolerance, so if you have 20 GB of disk drives in your server, you can only use 10 GB of storage. RAID 5 technology uses parity blocks so that if one drive fails, the other working disk drives can rebuild the information on a replacement hard drive. This provides you the fault tolerance of RAID 1 but is more efficient in use of disk space.

We recommend the following for your Domino server:

- Implement RAID 5 on all of your servers if possible. As your Domino servers become more critical to your business, you will want the highest disk tolerance available.
- Only use hardware-controlled RAID arrays. Software-controlled RAID arrays cause additional strain on your server in terms of CPU utilization. The CPU should be saved for Domino functions, not providing a RAID subsystem.
- If possible, install an extra hot spare hard drive in addition to your other RAID drives. Hot spares mounted in the array cabinet can replace failed disks automatically, or be managed by the administrator.
- Choose the fastest hard drives. Hard drives rate the internal platter speed in rpm (revolutions per minute). At the time of writing the most popular high speed drives operate at 10K rpm. The disk subsystem has larger latency times than other components of your server, so choosing the fastest drives will give considerable performance benefits.
- Do not mix SCSI Fast with SCSI Fast/Wide drives.

- Do not mix drives that spin at different speeds.
- Do not mix drives of different sizes.
- Enable write-back cache on the RAID controller, this setting caches information before it is written back to disk. This provides extra performance for high-demand Domino Web servers. Some controllers do not have battery backups, so this data may be lost in the event of a power failure. You should seriously consider a server UPS (Uninterruptable Power Supply) if this setting is used.
- From the NotesBench report for your server, determine the RAID stripe size that was used. Use this value for best results. The stripe size used is usually 16K versus the default 8K.

3.13.4 Processor

Many of the Domino functions are processor intensive. These tasks are routing messages, indexing databases, searching databases and dynamically creating HTML in memory to send to Web clients. Therefore, we recommend that you purchase an IBM Netfinity server that is capable of symmetric multiprocessor (SMP) capability. You may not need or be able to afford the extra CPU today, but you are providing plenty of room for growth in the future.

Domino takes advantage of extra CPUs in many ways. Some examples are allocating multiple replicator and router tasks to different CPUs.

Domino is a highly scalable application that can take advantage of the benefits of SMP. However, nothing within Notes needs to be tuned or configured to take advantage of this functionality.

3.13.5 Network Card

We recommend that you purchase a high-performance network card in your server, especially if your Domino server is used by a large number of users. We recommend the following network card and configuration for your server:

- PCI Ethernet or token-ring card with full-duplex capability.
- A brand name network card which will provide you with excellent service and support, and robust network drivers for the operating system you are using.
- For servers that will support a large number of users, consider connecting one or two network cards into an Ethernet or token-ring switch. This will provide extra throughput into the server.

3.13.6 Uninterruptable Power Supply

We recommend that you protect your server from data loss by connecting it to the power supply via an Uninterruptable Power Supply (UPS). This will allow adequate time for your server to save data from the SCSI controller and operating system cache to disk before shutting down properly. This will reduce the possibility of database corruption, which can result if a Domino server is not stopped correctly. Make sure that the UPS you choose has compatible software that will allow commands to be executed on your server in the event of a power failure. These commands can stop the Domino server, and also perform a backup of important databases such as the Public Name and Address Book, before the server is shut down.

3.13.7 Disaster Recovery Software and Virus Protection

Hardly any network today exists without some sort of backup software system. It is important that you purchase a high-quality backup and recovery system which allows you to easily and quickly restore your system in the event of a complete operating system failure. Many companies only install software that allows them to recover files from tape. Unfortunately, if your server crashed it could take you quite some time to install your operating system again, and restore your Domino server from tape. Therefore, purchase a backup and recovery system that allows you to completely restore the operating system and applications by a simple boot diskette process. This will enable you to get your Domino server operational as quickly as possible.

Ensure that your backup system can back up the open files, such as the Public Name and Address Book database. This database is always open whenever Domino is running on a server.

Domino is the center of your electronic mail and company collaboration. Users will then use Domino to e-mail others with software that may be infected with a computer virus. We recommend that you purchase anti-virus software that can scan files at your workstations. Also consider anti-virus software that can scan attachments in Notes documents for viruses.

3.13.8 Recommendations from the Field

In this section, we've provided information pertaining to several customer engagements that have taken place out in the field. Understand that each customer engagement was based on a different environment, however, the rules of thumb listed are based on a collection of actual experiences:

- For mail systems only with Notes Clients (not Web clients).
 1. Treat the NotesBench number as "registered" users and not active or concurrent.
 2. Due to the inflated nature of NotesBench numbers, reduce them to one-third of their published state.
 3. For clustering subtract an additional 25%.
 4. For calendaring and scheduling subtract an additional 15%.

Note: This should give you a relative registered user base from which an average of 50% of the users can be concurrent.

- Server memory requirements per user for workload
 - HTTP user: 1 MB/per user
 - Notes Mail: .5 MB/per user
 - Notes Mail, C&S: .7 MB/per user
- Web sizing

This is a difficult topic in Domino and must be addressed with caution on an individual case by case basis. Due to extremely high I/O for a mail server and increased utilization from the HTTP task, usually you're able to see 200-250 concurrent HTTP mail users on a 4-way Netfinity 7000.

Note: These numbers are used by IBM engagement teams in customer sizings and they have provided this rough algorithm from installations experienced at different customer locations. Due to limitations in Notes 4.6x,

the team has generally seen performance degradation once there are more than 750 concurrent users on a Domino server. The team's sizing for new models such as the Netfinity 7000 M10 are a bit skewed. The only bottleneck hardware wise would be disk I/O, but with current disk structures we can support the higher user load above the 750 mark. Since the issue lies within Notes, there will still be a performance degradation past the 750 user mark. Until Domino and Notes 5.0 is released, the team is unable to verify the performance gains with the new Pentium II Xeon processors.

3.14 Examples of Domino Server Selections

We now discuss four examples with several different operating systems using the Server.Planner tool and recommendations we have just outlined. Hopefully, these will explain more clearly how you can size your IBM Netfinity servers for your Domino environment. These are examples only. You should consider your own environmental specifics when sizing your Domino servers. Always consult with your colleagues and management before you make your server and configuration decisions.

3.14.1 Small Sized Domino Environment

A customer wants to install Domino for 50 users at a single location for electronic mail and a single discussion database. There will be approximately 10 remote users who will dial in to get their mail from time to time. All of the users are connected to an Ethernet LAN running at 10 Mbps. The users are using OS/2 Warp clients and connect to file and print services on a Warp server using NetBIOS.

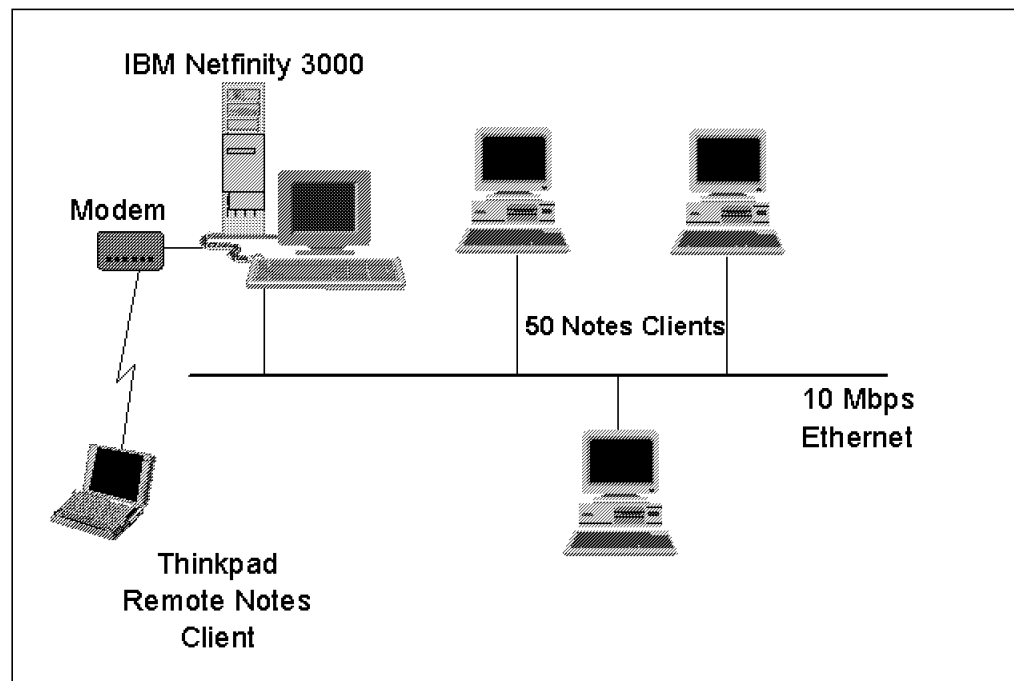


Figure 22. Small-Sized Domino Environment

We decided that for this location our new Domino server will run on a Warp server and use the NetBIOS protocol. This leverages the customer's existing experience with OS/2 and will not require additional protocols to be loaded at the client workstations. The customer was not sure what types of users would

be using the Domino server, so we decided to use the standard distribution user profile of 20% Mail users, 50% MailDB and 30% Groupware_A. Entering these values into the Server.Planner tool, we found that the IBM Netfinity 3000 would meet their requirement. The customer did not expect large growth in the short term and cost constraints prohibited a more expensive server.

Using our memory formula we calculated that we would need a minimum of 80 MB of RAM. The customer was going to run backup and restore and virus protection software, so another 32 MB of RAM was required, for a total of 112 MB of RAM.

The minimum disk space requirement was determined by the following components:

Base operating system and Domino:	500 MB
Page file requirement:	105 MB
50 user mail files:	1500 MB
One discussion database:	100 MB
Total:	2205 MB

To meet this minimum requirement, one 4.51 GB disk drive was installed in the server. This would meet the minimum requirement, and give sufficient disk space growth for the medium term.

A single 56 Kbps Hayes modem was installed on the server to allow remote users to dial in to check their mail and access the discussion databases. Remote users were educated to replicate their mail and discussion databases to their computer's hard drive so they could work offline. This was a popular choice, because users found this much faster than viewing their mail on the server at modem speeds. It also meant the modem was more available for other remote users as well.

3.14.2 Medium Sized Domino Environment

Another customer wanted to use Domino to support 200 users for a new location. They wanted their users to have electronic mail within the organization and to the Internet as well. Their users were grouped into four organizational units that would each require their own discussion database. There was approximately 50 remote users who would dial into the office to get their mail. This customer intended to use a token-ring network running at 16 Mbps. The customer had previously used Windows NT for servers and Windows 95 for the workstations at other locations.




Figure 23. Medium-Sized Domino Environment

It was decided that for this new location the new Domino server would run on Windows NT Server 4.0 and Windows 95 clients because the customer was familiar with these operating systems. As this was a new site, the customer was happy to use the TCP/IP protocol because it was the best choice technically, and would fit with their future expansion to more Internet/intranet tools. The customer was not sure what types of users would be using the Domino server, so we decided to use the standard distribution user profile of 20% Mail users, 50% MailDB and 30% Groupware_A. The customer felt that the server might have to support an extra 50 users in the medium term so this was added to the requirement of 200 initial clients along with the 50 remote users. Entering these values into the Server.Planner tool we found that the IBM Netfinity 5500 would meet their requirements. The extra CPU load for the SMTP MTA (Message Transfer Agent) was not considered enough to increase the server's CPU load significantly enough to warrant a larger server. An extra processor would be added in the future if the server was overloaded due to a large increase in Internet mail through the MTA, or the users were using the server for more demanding work. The IBM Netfinity 5500 has the capability for an extra processor if required.

Using our memory formula we calculated that we would need a minimum of 147 MB of RAM. The customer was going to run backup and restore and virus protection software so another 32 MB of RAM was required. This gave us a total of 179 MB of RAM. The customer would install one 256 MB SDRAM ECC DIMM.

The minimum disk space requirement was determined by the following components:

Base operating system and Domino:	500 MB
Page file requirement:	211 MB
250 User Mail Files:	7500 MB

Four discussion databases:	400 MB
Total:	8611 MB

To meet this minimum requirement, four 4.51 GB drives were installed with a RAID 5 configuration. This gave an effective disk space of 13.5 GB. This would meet the minimum requirement, and give sufficient disk space growth for the medium term.

Because of the large number of remote users that would need to connect to the server each day a single modem would only allow each user 28 minutes of average connect time (24 hours x 60 minutes / 50 users). It was expected that most users would prefer to connect to the server at the start or end of the day at their remote locations.

The customer decided to install an IBM 8235 device with 10 modems attached. This would give a better chance for each remote user to connect to the network to access Domino for their electronic mail, and access other devices on the network. More modems will be added in the future if there is sufficient demand.

3.14.3 Large-Sized Domino Environment

A large customer has the requirement for a reliable and expandable electronic mail and groupware infrastructure using Domino. They want to support approximately 900 users at one site and another 200 users at a remote location connected together by a 48-kbps frame relay link. The users at both sites need to use a variety of Notes databases, two discussion databases, a supply ordering database, customer tracking database, LAN help desk database, and travel and expenditure database, or a total of six Notes databases initially. There was expected to be no remote users at this early stage. As users are supplied with IBM ThinkPads to work from home and other locations, a dedicated RAS solution was going to be installed by the customer.

The customer felt that the users at both sites would quickly use much of the functionality of Notes because of the intensive training program that each user was required to attend. The customer had a 10Base-T Ethernet network at both sites, connecting back to IBM 8271 Ethernet switches. The 8271 Ethernet switches support a 100-Mbps connection and 200-Mbps full-duplex connection.

The customer wished to use Windows NT for both their servers and client workstations and use the TCP/IP protocol for all services. Because of the expectation that users would be more skilled with Notes than an average Notes user, the user profile was broken into 80% MailDB and 20% Groupware_A. Entering these values into the Server.Planner tool they found that the IBM Netfinity 7000 using two Pentium Pro processors would meet their requirements. However, the customer felt that the number of users could grow substantially within the next six months, possibly to 1400-1500 users for the larger site. Therefore, it was decided to use one Netfinity 7000 server as a mail server, and another Netfinity 7000 server for an application and replication hub server. The customer felt that if the number of users expanded in the future at the central site that a Domino clustering solution would be implemented to ensure high availability.

The customer did not have any data as to the amount of traffic generated by their average user. So the general WAN link speed rule was used to determine whether the 200 users at the remote site should connect back to the main Netfinity 7000 servers. The calculation determined:

If $n \times 14 >$ link speed in kbps, you should consider a remote Domino server (here n = number of users at remote site).

$200 \times 14 > 48$, so they will install a remote Domino server.

Entering the value of 200 users into the Server.Planner tool they found that the IBM Netfinity 5500 would meet their requirement for the remote Domino server.




Figure 24. Large-Sized Domino Environment

3.14.3.1 Central Site - Mail Server

The customer was expecting to grow to 1500 users in the medium term. Using our memory formula we calculated that we would need a minimum of 564 MB of RAM. The customer was going to run backup and restore and virus protection software, so another 32 MB of RAM was required. This gave us a total of 596 MB of RAM. The customer would install four 128 MB and four 64 MB ECC DIMMs to give a total server memory of 768 MB of RAM.

The minimum disk space requirement was determined by the following components:

Base operating system and Domino:	500 MB
Page file requirement:	704 MB
1500 User Mail Files:	45000 MB
Disk space for MAIL.BOX	100 MB
Total:	46304 MB (approx. 45.2 GB)

Note: The discussion and other databases are stored on the application Domino server.

To meet this minimum requirement, two banks of six 9.1 GB Wide Ultra SCSI 7200 rpm drives were installed with a RAID 5 configuration, a total of 12 drives in all. This gave an effective disk space of 45 GB. This would meet the minimum requirement, and give sufficient disk space growth for the medium term.

Two PCI 100Base-T Ethernet network cards were configured at full duplex and connected back to an IBM 8271 Ethernet switch. This would ensure the Ethernet network was not a bottleneck for the server.

3.14.3.2 Central Site - Application and Replication Hub Server

To minimize the variation of server systems, the customer decided to install the application and router hub server with 640 MB of RAM.

The minimum disk space requirement was determined by the following components:

Base operating system and Domino:	500 MB
Page file requirement:	704 MB
Two discussion databases:	400 MB
Four standard databases:	120 MB
Disk space for MAIL.BOX:	100 MB
Total:	1824 MB

To meet this minimum requirement, three 4.51 GB Wide Ultra SCSI drives were installed with a RAID 5 configuration. This gave an effective disk space of 9 GB. This was a lot more than the minimum requirement, but the customer felt that there was no budget within the next six months to spend further capital on extra computer hardware so the extra disk space was purchased now.

The application and replication hub server was fitted with two PCI Ethernet cards because of the large number of users that can connect to this server.

3.14.3.3 Remote Site

The remote site was expected to support 200 users. Using our memory formula we calculated that we would need a minimum of 130 MB of RAM. The customer was going to run backup and restore and virus protection software so another 32 MB of RAM was required. This gave us a total of 162 MB of RAM. The customer would install one 128 MB and 64 MB SDRAM ECC DIMMs. This would give us a total server memory of 192 MB of RAM.

The minimum disk space requirement was determined by the following components:

Base operating system and Domino:	500 MB
Page file requirement:	211 MB
200 User Mail Files:	6000 MB
Two discussion databases:	400 MB
Four standard databases:	120 MB
Disk space for MAIL.BOX:	100 MB
Total:	7331 MB

To meet this minimum requirement, four 4.51 GB drives were installed with a RAID 5 configuration. This gave an effective disk space of 13.5 GB. This would meet the minimum requirement, and give sufficient disk space growth for the medium term.

3.14.3.4 Other Issues

The customer decided that the replication hub server at the main site would provide mail and be the replication hub for the organization. Therefore, as more remote sites are installed with Domino they will replicate back to this hub server. This would also allow maintenance to be performed on the central site mail server after hours, and mail from the remote sites would be stored at the mail hub until the central mail server was operational.

The replication schedule for the hub server was to contact the remote site every hour for normal-priority databases such as the discussion and supply ordering. The hub would contact the spokes every 30 minutes for high-priority databases such as customer tracking and help desk. This was the standard for extra spoke servers as they came online. This gives the hub server plenty of time to contact many spokes if and when they are installed in the future. It also gives the Notes administrators time to determine whether the replication schedule needs to be changed or link speeds increased.

3.14.4 Company Intranet Environment

A large customer had the requirement for a company-wide intranet for their organization. The company had approximately 5000 employees around the world that had access to company resources through a worldwide TCP/IP WAN. The intranet would be located at the company's central computer site, which is running on an Ethernet network with 10 Mbps connections for workstations and 100 Mbps for servers.

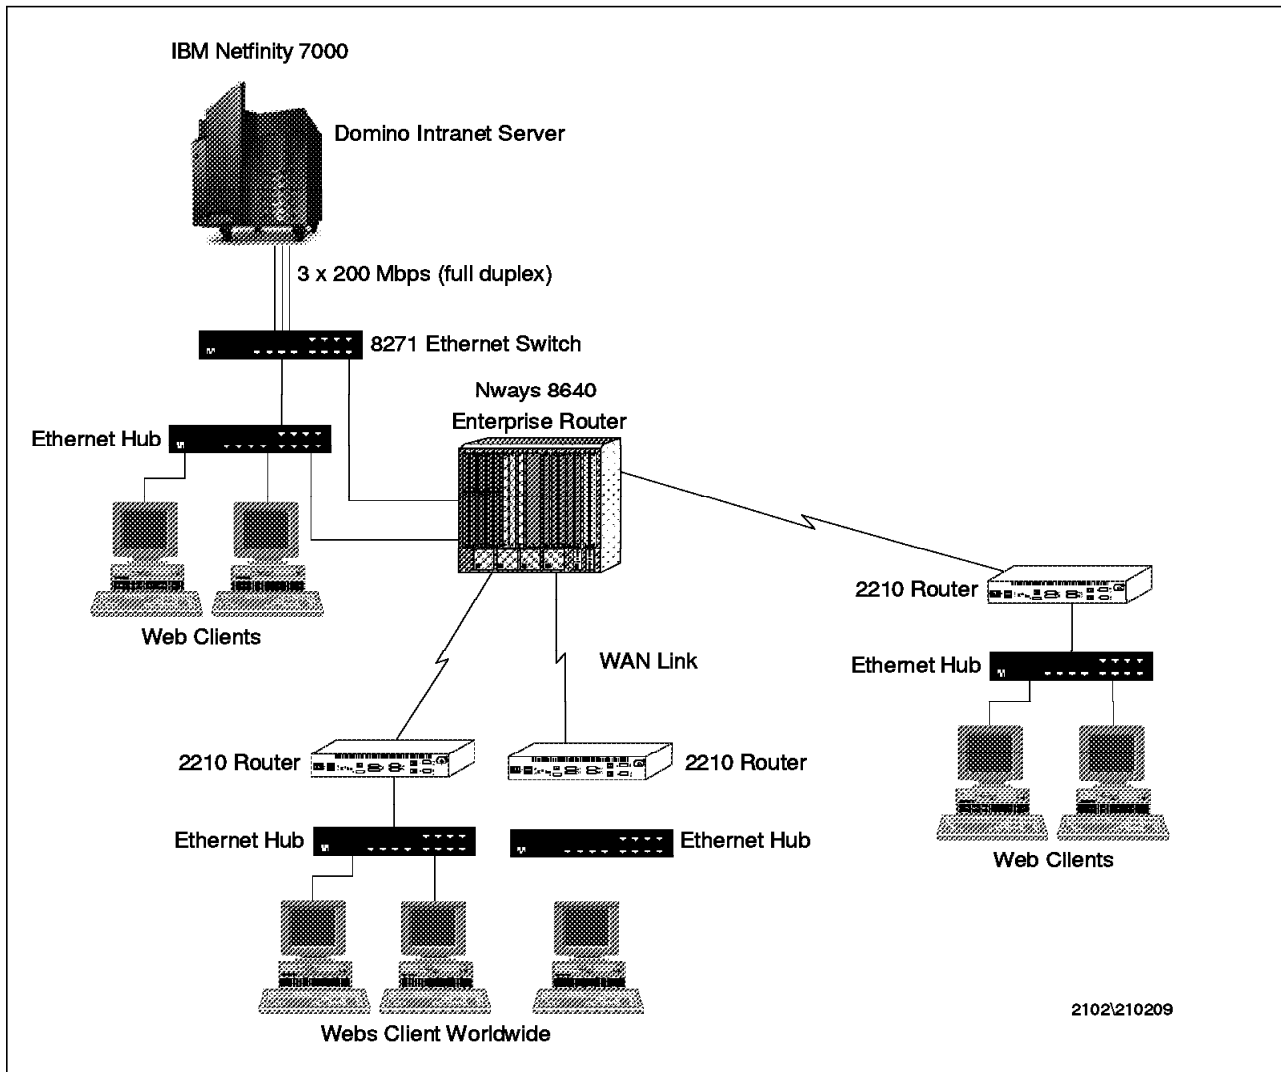


Figure 25. Company Intranet Environment

It was decided that the high-performance IBM Netfinity 7000 using two Pentium Pro processors would meet their requirement, given previous NotesBench results and that the customer only expected one third of their staff at a time to be accessing the intranet for information. This is based on distribution of staff around the world in different time zones and the fact that information on the intranet did not need to be accessed continually by the company's staff.

The customer was running Windows NT for other server functions, so we decided for this new intranet that our new Domino server would run on Windows NT Server 4.0.

Given that their server was expecting no more than a third of the 5000 staff members to access it at one time, we can specify the number of active users to be $5000/3 = 1666$. Using our memory formula we calculated that we would need a minimum of 619 MB of RAM. The customer was going to run backup and restore and virus protection software so another 32 MB of RAM was required. This gave us a total of 651 MB of RAM. The customer would install four 128 MB and four 64 MB ECC DIMMs to give us a total server memory of 768 MB of RAM.

Note: The Netfinity 7000 systems have two memory banks each with eight DIMM sockets. All memory DIMMs installed in each bank must be of the same speed and size and must be installed in matched sets of 4, 8, or 16.

This is more than sufficient memory for the intranet site to perform excellently.

The customer was expecting to supply information from about eight different Notes databases. These databases were currently using 5 GB of disk space. These databases are replicated to the Domino intranet each hour from another Domino server.

The minimum disk space requirement was determined by the following components:

Base operating system and Domino:	500 MB
Page file requirement:	704 MB
Homepage database:	300 MB
Eight Notes databases:	5120 MB
Total:	6624 MB

Note: We have planned for a very large home page database of 300 MB. In practice your home page database would probably never get this large.

To meet this minimum requirement, four 4.51 GB drives were installed with a RAID 5 configuration. This gave an effective disk space of 13.5 GB. This would meet the minimum requirement, and give sufficient disk space growth for the medium term.

Three PCI 100Base-T Ethernet network cards were configured at full duplex and connected back to an IBM 8271 Ethernet switch. This would ensure the Ethernet network was not a bottleneck for the server.

The customer configured this intranet as its own Notes domain. This ensured that security could be tightened, and that the intranet would not need to support a large N&A Book. The server's security was such that no Notes clients could connect to this server, and accessing Notes databases such as the N&A Book from Web clients was not allowed.

Chapter 4. Server Configuration and Optimization

How you install and configure your operating system and Domino server software are equally important in having a high performance and reliable server. We describe in this chapter how to configure Windows NT or OS/2 Warp for a Domino environment based on our experience with both operating systems.

If you have followed our server selection method and chosen the appropriate server, along with adequate memory, CPU and disk subsystem, for the majority of cases, your Domino server will be optimized and ready to face the work you have planned for it. However user loads can change, and demands may grow beyond anyone's anticipation and you may find that you need to investigate server performance in the future. This chapter discusses what Domino server parameters you might need to tune, and concludes with using the Windows NT Performance Monitor to help you investigate bottlenecks in your Domino server.

Note: For the latest performance information on Lotus Domino, visit the Performance Zone at: <http://www.lotus.com/performance>

The following list is a summary of what tasks you should complete to install Domino on an IBM Netfinity server. This list will hopefully remind you to complete all tasks so that the installation is a smooth process:

1. Install the latest BIOS for your server.

(Located at <http://www.us.pc.ibm.com/listfiles.html>.)

Note: You should use the HardwareGuide CD from ServeGuide to ensure your system is ready for software installation.

2. Install the latest SCSI card BIOS if applicable.

(Located at <http://www.us.pc.ibm.com/listfiles.html>.)

3. Configure ServeRAID adapter if applicable.

4. Partition drives (for example C: and D:).

5. Install the operating system on the C: drive with or without the aid of the IBM ServerGuide CD.

Note: The ServerGuide CD loads the necessary drivers for a quick and easy networking operating system installation.

6. Install the operating system service packs.

7. Install the latest driver version of the network card.

8. Test the network subsystem using the TCP/IP ping command or other connectivity tests.

9. Install the correct video drivers for the IBM Netfinity system.

10. Install Domino on the D: drive of the server.

11. Install the Domino service pack (Incremental installers).

12. Optimize the operating system for Domino by setting application responsiveness, file system cache and size/location of the operating system paging file.

13. Optimize the protocols if necessary.

14. Optimize Domino using the NOTES.INI if necessary.

15. Implement Domino security to NAMES.NSF and other databases.
16. Implement the Administrator group, and set which users can create new databases and replica copies on server.
17. Implement operating system and physical server security.

We discuss the majority of these steps in this chapter in more detail.

4.1 System Installation

The following section outlines how to configure your operating system, and provides a checklist of things you should do such as applying service packs, network card drivers and Domino service packs.

4.1.1 Disk Partition Configuration

Your operating system provides the convenience and flexibility to have various disk partitions of different sizes and file systems. Some computer engineers prefer to have a single partition, while others seem to prefer several partitions of various sizes. Then there is the question of what file system you should use. There are benefits to each disk partition decision; however, we recommend for your Domino server the following disk configuration and file systems:

<i>Table 7. Recommended Domino Disk Configuration</i>			
Volume Name	Size	File System	Software
C:	500 MB	FAT	Operating System
D:	Remaining Space	NTFS (Windows NT) HPFS386 (Warp Server)	Domino Software and Databases

Having this partition and file system configuration has several advantages:

- Domino software is separated from operating system software. This logically separates the operating system from applications and data.
- If a server fails you can boot the server with a floppy diskette to try and repair the operating system partition. You can feel reassured that your Notes databases are on a separate partition.

Note: DOS doesn't understand HPFS or NTFS, so files stored on an HPFS or NTFS volume are invisible to DOS. In the event your server fails and you would like to have access to your HPFS or NTFS volumes as well, you should have bootable disks available for your OS/2 or NT operating systems.
- If you cannot repair the error, you can reinstall the C: partition with the operating system quickly and easily.
- The Domino server software is located on a high-performance file system such as the NT File System (NTFS) available on Windows NT, or High Performance File System (HPFS) 386 available on Warp Server. These file systems are the most advanced file systems for large disk subsystems.

If you have a RAID configuration, we recommend that you still use a single C: and D: partition. The RAID system can protect against a single drive failure. If this happens, then your server will still operate due to the data striping across all drives at the hardware level. Some RAID controllers do not allow you to have

a very large D: partition. You may need to break the D: partition into multiple smaller partitions such as a D:, E:, F: for example.

If your server does not support RAID, we recommend that you allocate one complete disk drive for the C: partition, and allocate each other drive as a separate partition. Some operating systems allow you to have a logical drive spanning multiple physical hard drives. For example, the C: drive might span several physical drives. We advise against this because if one drive fails, you will lose all of your data. Therefore, for each non-RAID drive you have in your server, specify a single drive volume for each physical drive.




Figure 26. RAID 5 and Non-RAID Configuration and Disk Partitioning

4.1.2 Operating System and Service Pack Installation

Follow the instructions that come with your operating system to install the software on your server. You may decide to install the operating system using the CD-ROM that comes with your operating system, or use an operating system installation aid called ServerGuide. IBM ServerGuide provides a semi-automated installation of OS/2, NetWare, Windows NT and SCO UNIX as well as a wealth of additional information and utilities to assist in the initial installation tasks. We highly recommend that you consult the following redbooks when you are installing Windows NT or Warp Server with IBM Netfinity servers:

- *OS/2 Warp Server Integration Guide for IBM Netfinity and IBM PC Servers*, SG24-2125
- *IBM PC Server and Windows NT Integration Guide*, SG24-4763

In most cases the operating system you install will have service packs that fix known problems. At the time of writing the following Web sites gave the latest service pack information:

Warp Server <http://www.software.ibm.com/os/warp/support>

Windows NT Server <http://www.microsoft.com/downloads>

We recommend that you install the latest service pack for your operating system when you first install your system. You may wish to install service packs if you encounter problems in the future which the service packs fix, but only apply these if they solve a specific problem. We do not recommend applying service packs as soon as they are released simply because they are available.

4.1.3 Network Card Driver Installation

We also recommend that you download the latest driver for your network card, whether that be Ethernet, token-ring or ATM. Many people use the drivers that are supplied on the diskette that comes with the network card. This is fine in the majority of cases, but since you are installing a critical server for your business you should take the time now to determine whether there is a later version of the driver for your network card.

The majority of IBM network card drivers are located at:

<http://www.networking.ibm.com/netprod.html>

You will then see a list of IBM hardware networking products. Select LAN adapters to quickly get to the list of IBM networking adapters. You can then select your IBM network card and download the latest driver. For non-IBM adapters, check that vendor's Web site.

4.1.4 Domino and Service Pack Installation

After you have installed your operating system and network drivers, but before you install Domino, ensure your network protocols are working well. If you are using the TCP/IP protocol, make sure you can ping other devices and they can ping you. Ensure you have your TCP/IP host addressing correct so that you can ping other devices on your network by a host name rather than a TCP/IP address.

When you have ensured your network transport is running correctly you can install Domino from the CD-ROM. Remember to install Domino to the D: partition as we discussed in 4.1.1, "Disk Partition Configuration" on page 80.

When you have completed the installation of Domino you should install the latest service packs for Domino. These are called Incremental Installers and are downloaded from the Iris Web site at:

<http://notes.net/down.nsf/welcome>.

You will need to determine what version of Domino software you are currently running. This is found at the server console when the server first starts or you can issue the Show Server command at the Domino command prompt. You will also need to choose the appropriate Incremental Installer based on the operating system your server operates on.

4.2 Operating System Optimization

You can aid the performance of your Domino server by changing various operating system settings for Windows NT and Warp Server. We recommend you implement these changes just after you have installed your Domino server. Some of the following sections are taken from the Lotus Notes Knowledgebase, which provides technical information on Lotus products. The Knowledgebase

database is purchased from Lotus as a CD-ROM subscription or purchase a license to connect to Lotus Knowledge Base via the Lotus Notes Network.

Note: If you purchase the license, you can download as often as you like, paying a per-minute access charge. For more information, go to the following Web site:

<http://www.support.lotus.com/css.nsf>

4.2.1 Windows NT Server

Domino is supported on Windows NT Server 3.51 and 4.0. Windows NT is designed as a server platform that primarily services file storage and retrieval functions. This poses unique demands on an operating system. To utilize Windows NT as a Domino server platform, the operating system must be configured to service the unique demands of an application server.

Specifically, there are two parameters that must be changed in order for NT to be an effective application server. If these changes are not made, whenever the server is invoked it will indicate that the Registry values need to be changed. Both parameters are configurable via the Windows NT Graphical User Interface (GUI):

4.2.1.1 Application Responsiveness

The first parameter that must be changed is Foreground/Background Responsiveness. This parameter is configured by doing the following:

On a Windows NT 3.51 Server:

1. Open **Control Panel** and select the **System** icon.
2. Click the **Tasking** button.
3. Select **Foreground and Background Applications Equally Responsive** for the Foreground/Background Responsiveness parameter (instead of the default of Best Foreground Application Response Time).

On a Windows NT 4.0 Server:

1. Open **Control Panel** and select the **System** icon.
2. Move to the Performance Panel.
3. In the Application Performance section, slide the Boost indicator from Maximum to None.

This setting is shown in Figure 27.




Figure 27. Equal Response for Foreground and Background Applications

By default, the Foreground/Background Responsiveness is set to Best Foreground Application Response Time. With this setting, applications running on the Windows NT server would receive a priority boost if they are in the foreground. However, it is more than likely that other windows will be in the

foreground with the Domino server task in the background (for instance, when the administrative client is being utilized). By selecting **Foreground and Background Applications Equally Responsive**, the operating system does not give priority to any one task. This will allow the Domino server to have the same priority as any other process on the machine.

The following Registry parameter is changed from a value of two to zero:
HKEY_LOCAL_MACHINE\System\CurrentControlSet\
Control\PriorityControl\Win32PrioritySeparation.

4.2.1.2 File System Cache Size

The second parameter that needs to be changed is the Server setting in the list of Installed Network Software.

This parameter is configured by doing the following:

On a Windows NT 3.51 Server:

1. Open **Control Panel** and select the **Network** icon.
2. In the scrolling list of Installed Network Software, locate Server.
3. Open the **Server** option.
4. From the four choices, select **Maximize Throughput for Network Applications** (instead of the default choice of Maximize Throughput for File Sharing).

On a Windows NT 4.0 Server:

1. Open **Control Panel** and select the **Network** icon.
2. Select the **Services** tab.
3. Highlight the **Server** icon in the list and select the **Properties** button.
4. From the four choices, select **Maximize Throughput for Network Applications** (instead of the default choice of Maximize Throughput for File Sharing).

This setting is shown in Figure 28.

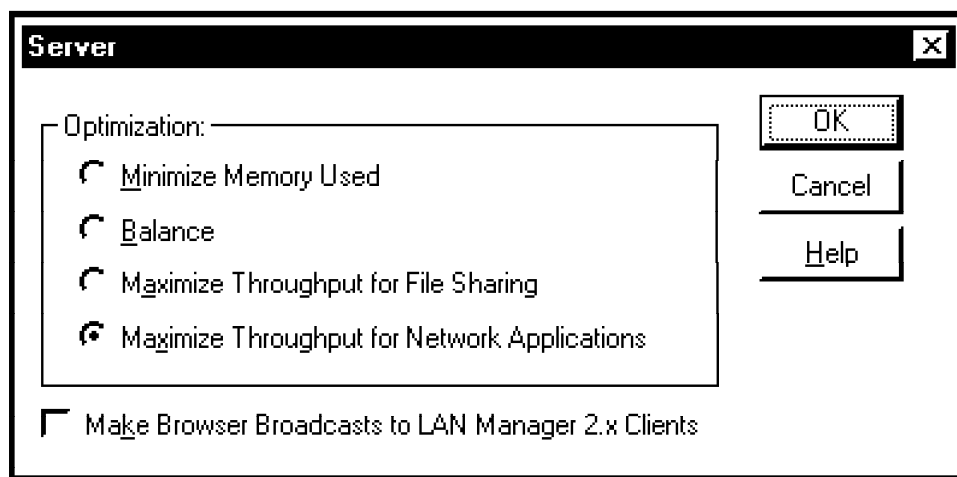


Figure 28. Set Windows NT to Application Server Role

By default, the Server parameter is set to Maximize Throughput for File Sharing. This allocates memory for an unlimited number of connections to the server. By selecting **Maximize Throughput for Network Applications**, memory is still

allocated for an unlimited number of connections to the server, and less memory is set aside for cache. This results in more memory available for servicing network throughput.

The following Registry parameter is changed from a value of one to zero:
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\Session
Manager\Memory Management \LargeSystemCache.

4.2.1.3 Paging File

In the previous chapter we said that Windows NT will choose a paging file equal to the amount of memory in your server plus 12 MB. To view the paging file settings, on a Windows NT 4.0 Server:

1. Open **Control Panel** and select the **System** icon.
2. Select the **Performance** tab.
3. Click the **Change** button.

You should see a window similar to the one shown in Figure 29.

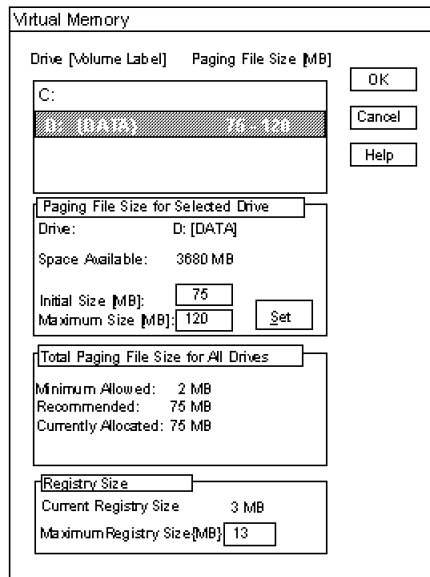


Figure 29. Virtual Memory Settings for Windows NT

We recommend you set the minimum size of your paging slightly larger than the Windows NT recommended value equal to the amount of memory in your server plus 10%. For example, if the amount of RAM in your server is 64 MB, set the initial size of the paging file to 71 MB. We also recommend that you move your paging file (PAGEFILE.SYS) to the D: partition, because this will locate the file on a higher performing file system such as NTFS rather than FAT.

If you have different drives in your server of different performance, always store the PAGEFILE.SYS on the highest performing disk drive. The maximum value for this file is your own preference. A value of 50% beyond the recommended value is a good start. For example, if the recommended paging file is 100 MB then set the maximum value to 150 MB. This is large enough to accommodate any short term expansion and will not be too large to use up large amounts of disk space. Windows NT will warn you if it needs to make the paging file larger. You can increase the maximum value of the paging file or add more memory to your server.

4.2.2 Warp Server

Domino is supported on Warp Server 4.0. You can purchase Warp Server in two packages, Warp Server or Warp Server Advanced. Warp Server Advanced is recommended for supporting more than 100 active users. The key difference between the two packages is that Warp Server Advanced includes a special, optimized file system called HPFS386 and has a free add-on feature that supports symmetric multiprocessing (SMP).

4.2.2.1 Application Responsiveness

As with Windows NT, OS/2 Warp Server will give the foreground application a priority boost over background applications, which on a server could be other applications apart from Domino. To set an equal priority for foreground and background applications, edit the CONFIG.SYS file on the server and add the following setting:

```
PRIORITY_DISK_IO=NO
```

You will need to save the CONFIG.SYS file, and shut down, and restart the server to make this change effective.

4.2.2.2 File System Cache Size

If you are running Warp Server Advanced with the HPFS386 file system on one partition, Warp Server detects how much memory is installed in your server, and if the value is greater than 20 MB of RAM, it will allocate 60% of your memory to the HPFS386 file system cache. Therefore, if your server is fitted with 256 MB of RAM, you will only have 102 MB for your operating system and Domino requirements. Many customers are unaware of this and Domino will run poorly as a result. As with Windows NT, you need to reduce the size of the file system cache so Warp Server can meet the demands of an application server rather than a file server.

You need to edit the HPFS386.INI file and change the following setting:

```
CacheSize = nnnn
```

Where nnnn is the number of kilobytes of memory HPFS386 will allocate for cache.

If you are using Warp Server solely for Domino, then choose a cache value at 10% of system memory. If your server is also performing some file serving services, you may wish to choose a cache size closer to 20% of system RAM.

4.2.2.3 Swap File

The swap file for Warp Server is called SWAPPER.DAT. By default this file is located in the \OS2\SYSTEM directory and is initially set to 2 MB in size. If the SWAPPER.DAT file is left at this default setting, the file system will dynamically increase the size of this file as more virtual memory is needed. The problem with leaving the SWAPPER.DAT file at the 2 MB default size occurs when the file grows over non-contiguous areas on the hard drive. This will degrade the performance of the server by interrupting data transfer when the drive heads move to each location. The solution is to set the initial size of the SWAPPER.DAT file to a larger value and move the file to the root directory of a high-performance file system. These are now discussed.

Warp Server does not provide a recommended value for the paging file. You can choose an initial value equal to the size of memory plus 10%; however, if your server has a large amount of RAM it may never need to page that much memory to disk. We suggest the initial values of the SWAPPER.DAT as shown in Table 8.

Memory in Server (RAM)	SWAPPER.DAT Size
less than 128 MB	64 MB
greater than 128 MB	128 MB

We also recommend that you move the location of the SWAPPER.DAT file to the root of the D: partition, because this will locate the file on a higher performance file system such as HPFS386.

To set the initial value and location of the SWAPPER.DAT file, edit the CONFIG.SYS file and change the following line: `SWAPPATH=D:\ 40960 <initial size>`

Therefore a SWAPPER.DAT size set to an initial size of 128 MB would have the following line in the CONFIG.SYS file: `SWAPPATH=D:\ 40960 131072`

This setting will also give a warning message at the server's console when the disk space on the D: drive falls below 40 MB. You will need to save the CONFIG.SYS file and shut down and restart the server to make the change effective.

Monitor the size of your SWAPPER.DAT file and if it increases in size when your Domino server is running, increase its initial size by 10%.

4.2.2.4 Thread Tuning

On powerful servers with plenty of RAM you may wish to consider changing the MAXWAIT parameter in the CONFIG.SYS from a value of 3 to 2. This parameter sets the maximum amount of time that a thread can sit idle before OS/2 raises its priority. Reducing MAXWAIT means that OS/2 will do more switching between threads, which improves the smoothness of the multitasking. Determine the response time of your server before and after this change to determine whether the server's performance was improved.

4.3 Protocol Optimization

Optimization of protocols within the Notes environment really has very little to do with the Domino server application. Instead, the issues we discuss here are common guidelines that should be used for any network environment regardless of whether Notes is installed.³

³ For more information on each protocol, refer to the *Notes R4.5 Network Configuration Guide*.

4.3.1 TCP/IP

Using the TCP/IP protocol in a Notes environment has historically proven to provide the best performance and most efficient use of network adapter hardware. This fact is attributed directly to the nature of how TCP/IP is routed and transmitted on the network. The use of TCP/IP on a LAN is highly recommended, especially when implementing a geographically disperse network using WAN links. More often than not, TCP/IP is the protocol used by routers for communication across WAN links. By implementing your Notes environment with the TCP/IP protocol on the LAN, the overhead associated with each router examining and forwarding packets is reduced.

With respect to Domino, a common problem with Domino servers is when the packet buffer size is increased by the administrator for performance or by the installation of a new client/server application. For proper operation within Domino, the NOTES TCP BUFFER PACKET SIZE parameter within the NOTES.INI file must match for all servers and clients on the LAN.

4.3.2 IPX/SPX

Using the IPX/SPX protocol in local area networks has also proven to produce very high performance results for Notes environments. However, when communicating across wide area links, the performance is severely degraded due to packet segmentation and server-to-server SPX traffic.

There are two types of NetWare services supported by Domino: Bindery Service for NetWare 3.x servers and NetWare Directory Services (NDS) for NetWare 4.x servers. The default service support when installed on an OS/2 server is NDS. In a mixed NetWare 3.x and 4.x environment, the Domino server will first perform a Directory Services request. If no 4.x server responds, it will then attempt a Get Nearest Server (GNS) request. Only NetWare 4.x servers will respond to the initial Directory Services request. Therefore, if you want to have the Domino server connect to a 3.x server, add the following to the NET.CFG file:

```
DIRECTORY SERVICES OFF
```

For Windows NT 3.51, Directory Services is not supported by the Windows NT Microsoft NetWare Client. To utilize NetWare 4.x Directory Services, you will need to disable the Microsoft NetWare Client and install the Novell NetWare Client. Furthermore, if the Microsoft NetWare Client is being used on the Domino NT server, both the NwLink IPX/SPX and Gateway Services for NetWare options need to be installed. Both of these options are located on the Windows NT CD-ROM.

4.3.3 NetBIOS

Use of the NetBIOS protocol for Domino servers is hampered by a relatively low limit on the number of active users that are supported through the Domino NetBIOS port drivers. Both OS/2 and NT-based servers can accommodate at most about 250 active sessions simultaneously.

For this reason, using this as the only protocol on a robust Domino server would result in the server being underutilized. The NetBIOS limitations would be reached long before the server would otherwise reach capacity.

Having said that, NetBIOS is still a very important protocol and many installations will require it to be supported on the server, although usually in conjunction with TCP/IP or sometimes even IPX/SPX.

In an OS/2 environment, if you do require NetBIOS, Domino supplies two port drivers that are to be used. These are named NETBIOS.dll and IBMEENB.dll. You can select one or the other, but not both. From a Domino application standpoint, the two drivers are equivalent.

The NETBIOS.dll driver uses an interface supplied by the OS/2 LAN Requester and therefore requires the LAN Requester to be loaded in order to use it. This driver also limits the number of users it can support to approximately 120.

If you require more NetBIOS users, you can use the IBMEENB.dll driver. This driver will support approximately 250 users by spreading the network load across multiple LAN adapters. However, setup of the IBMEENB.dll driver is not as easy because it requires that the resources in the NOTES.INI, PROTOCOL.INI, and IBMLAN.INI files be coordinated. The Warp Server Tuning Assistant can provide help with this task.

If NetBIOS is required in your Notes environment, we recommend that the NETBIOS.dll driver be used if possible. If this driver does not allow enough users for your environment, then consider the IBMEENB.dll driver as a way to get some more Notes NetBIOS client connectivity, although be prepared for some additional time and effort to get it configured.

4.4 Domino Optimization

Domino installed straight out of the box is optimized for most configurations. However, there are some tuning parameters in specialized situations that can give your Domino server some added performance and also protect the server from overloading. The Domino server is controlled by the values stored in the NOTES.INI file. We recommend you take a backup of your NOTES.INI file before you make any modifications in case your settings are not accepted by the system. A full listing of the NOTES.INI parameters are listed in Appendix A of Lotus Domino's *Getting Started with Domino Server*.

4.4.1 Ways to Modify NOTES.INI

Before making changes to the NOTES.INI file, use the Show Configurations server command at the console to check the existing settings. To edit the settings in the NOTES.INI, use one of the following methods:

Server Configuration document This document is a form that exists in the Domino server N&A Book. Use a Server Configuration document to specify settings for a single server, a group of servers, or all servers in a domain. You cannot specify all NOTES.INI settings in a Server Configuration document. The Set/Modify Parameter list in the Server Configuration document lists a set of default settings that you can specify.

Set Configuration command	The Set Configuration server command (issued at the Notes server console) writes the new setting to the NOTES.INI file. If you previously specified the setting in any Server Configuration document that affects this server, Domino writes the new setting to the Server Configuration document specific to the server being configured or creates a new document if necessary.
Editing the NOTES.INI file	You can edit the NOTES.INI file using a text editor, however, we do not recommend using this method because of the possibility of introducing errors into the file and impairing the operation of the Domino server.

4.4.2 Server Tasks

Minimizing the number of server tasks that are run on the server, the frequency in which they run, and the time in which they run will allow you to increase the performance of the server. The Domino server tasks to be automatically launched by the server are set in the NOTES.INI variables:

```
ServerTasks=Replica,Router,Update,Stats,AMgr,Adminp
ServerTasksAt1=Catalog,Design
ServerTasksAt2=UpdAll,Object Collect mailobj.nsf
ServerTasksAt22=Statlog
```

Each of these variables control the schedules for automatic server and database maintenance tasks. The time is entered in 24-hour format, where 0 is 12:00 a.m. and 23 is 11:00 p.m. In the example above, Catalog and Design tasks would initiate at 1:00 a.m., and the Statlog task would initiate at 10:00 p.m.

Following are some suggestions on how to increase performance related to Domino server tasks:

- Remove the Stats server task (installed by default). This task is installed only for backward compatibility with Notes Release 2.x servers that have been set up for statistics collecting and can be removed (although it takes up very few resources). This is the server task that produces the messages Stats agent started and Stats agent shutdown on the server console once each hour.
- Both the Replicator and Router tasks can be removed if they are not being used on the server, as each of these tasks takes up a fair amount of server resources when loaded. For example, if you only have one Domino server in your organization which is used for both applications and mail routing, you may not need the Replicator task as you don't have any other servers to replicate from (and because clients will be replicating from the Domino server, not vice versa). Another example might be if you have a hub replication server that never routes mail. In this case, you can remove the Router task. Be careful with this example because, although the server might not be used for routing mail between users, you may have some mail-enabled Domino database applications, in which case you will need the Router task and should not remove it.
- Daily server tasks should be run when other server tasks are not running and at times when few or no users are using the Domino server such as before users begin using the server in the morning, during lunch time or

after normal working hours. This allows the maximum amount of server resources to be available for each server task that is currently executing and for user sessions. Examples of such server tasks are Design, Catalog, Statlog, and customized Notes API programs that only need to be run once a day.

4.4.3 Database Indexing

Use the indexer server task to maximize disk space used for database view indexes and minimize the effort required to update view indexes. In general, the fewer view indexes that the Indexer server task must update, the fewer server resources are used by this task. You can use the NOTES.INI variable `Default_Index_Lifetime_Days` to minimize the amount of effort required by the Indexer task when updating view indexes. The `Default_Index_Lifetime_Days` variable controls how long a view index is kept in a database before it is deleted due to non-use.

For example:

```
Default_Index_Lifetime_Days= number of days
```

The default value for this variable is 45 days (unless a database designer has specified a different lifetime in the Design - View - Attributes Index Options dialog box). Setting this number to a lower value will reduce the number of view indexes on a server and therefore reduce the amount of time and effort that the Indexer server task must spend in updating view indexes.

Disable the updating of full-text indexes on a server if you do not have any full-text indexed databases on your server (and do not intend to). The NOTES.INI variable `Update_No_Fulltext` can be used to disable all full-text index updating on the server. You may want to use this variable to disable the updating of full-text indexes if, for example, you do not want users to create full-text indexes on their mail files on a mail server in order to save disk space on that server (the space which would be used for the full-text indexes) and to save the Indexer task the time and resources of updating these full-text indexes. This is a very good setting for mail and replication hub servers, which in most circumstances do not have any user connections.

For example:

```
Update_No_Fulltext=0 or 1
```

Setting this value to 0 causes full-text indexes to be updated each time the `Updall` (a server command task) is executed and setting the value to 1 disables all full-text indexing.

4.4.4 Minimize the Notes Log File

Keep the Notes log file (LOG.NSF) as small as possible and minimize the amount of database activity that is logged. Minimizing the activity of the Notes log will help to increase server performance by minimizing the updates to the log file.

In general, you should only log Domino server activities when you are keeping regular statistics on certain activities, trying to diagnose a problem, or track the use of a specific resource. The NOTES.INI variables below all apply to logging Domino server activities and can be disabled if the results of their logging are not being used.

For example:

```
Log=logfilename, log_option, not_used, days, size
Log_AgentManager=0 or 1
Log_MailRouting=0, 10, 20, 30, or 40
Log_Replication=0 or 1
Log_Sessions=0 or 1
Log_Tasks=Yes or No
Log_View_Events=0 or 1
PhoneLog=0, 1 or 2
```

For specific information on each of these variables refer to Appendix A in the *Notes R4.6 Administrator's Guide*.

4.4.5 Maximum Sessions and Session Timeouts

You may wish to set the maximum number of sessions that can run concurrently on your server. This is configured to prevent the server overloading by a large amount of users connecting to your server simultaneously. When a new user attempts to log on, if the current number of sessions is greater than the value of `SERVER_MAXSESSIONS=` (in the `NOTES.INI`), the Domino server closes the least recently used session. In order for a session to be considered for closing, it must have been inactive for at least one minute. For example, if this parameter is set to 100, and the 101st person tries to access the Domino server, the Domino server drops the least-used session from the server in favor of this new session.

Note: Reducing the `Server_MaxSessions` to a specific number will not prevent the server from allowing more than that number of concurrent active users on the server, but will drop the sessions soon after they become inactive. This frees up resources. Conversely, Domino will not close any session that has been idle for less than one minute regardless of the demand on the server.

Session details are found by using the following console command:

```
show stat net
```

The following is the output of the `show stat net` command from the console of a Domino server:

```
> show stat net

NET.TCPIP.BytesReceived = 6,677,642
NET.TCPIP.BytesSent = 9,724,217
NET.TCPIP.Sessions.Established.Incoming = 336
NET.TCPIP.Sessions.Established.Outgoing = 20
NET.TCPIP.Sessions.Limit = 65535
NET.TCPIP.Sessions.LimitMax = 65535
NET.TCPIP.Sessions.LimitMin = 10
NET.TCPIP.Sessions.Peak = 7
NET.TCPIP.Sessions.Recycled = 0
NET.TCPIP.Sessions.Recycling = 0

>
```

This server had a single network port named `TCPIP`. If your server has a different port name, then this is shown after the word `NET`. The key session details you will want to investigate are `Sessions.Established.Incoming`, `Sessions.Established.Outgoing` and `Sessions.Peak`. Note these session values

are per network port. So, if you have more than one network card in your server you will need to add up all session values. If your server is overloaded, use these values to determine the value of the Server_MaxSessions value. The format of the Server_MaxSessions parameter in the NOTES.INI file is:

```
Server_MaxSessions = number
```

You can also specify this setting in a Server Configuration document in the Public Name and Address Book.

The default server session timeouts are operating system dependent for network connections:

Operating System	Session Timeout
OS/2	4 hours
Windows NT	30 minutes
NetWare	30 minutes

The default timeout for a COM port connections is 2 hours. After this time, it will drop the database server connection. Inactive database server sessions do not take up many resources on a Domino server, but it is a good idea to minimize them anyway in order to regain the resources that they do take up. To do this, use the following NOTES.INI variable:

```
Server_Session_Timeout=<minutes>
```

For example:

```
Server_Session_Timeout=45
```

Note: 45 minutes is the minimum recommended setting for this variable for OS/2. Setting a shorter time may cause the server to have to reopen database server sessions too often, which can negatively affect the server's performance as it must reevaluate the database server sessions as users or other servers recommence their use of the server.

4.4.6 Maximum Users

To protect your server from overloading, you can also set the maximum number of users that are allowed to access a server. When you set this value and the maximum user limit is exceeded, Notes users are prompted with the message at their workstation as shown in Figure 30 on page 95 when they try to access a database on the server.



Figure 30. Message Box to Indicate that Maximum Users Exceeded

If the Domino server is a part of a cluster, the server state of MAXUSERS is reached and the request is passed to another server in the cluster without the user being aware of this happening.

To set the maximum number of users, set the following NOTES.INI value:

```
Server_MaxUsers = number
```

4.4.7 Maximum Mail Threads

The MailMaxThreads setting determines the maximum number of concurrent processes that the mail router can create to perform its mail transfers efficiently. The default setting is one thread per server port. Using more router threads allows for mail to be routed more quickly, because there is more than one router thread looking in the MAIL.BOX file and distributing mail documents at the same time, instead of only one router which routes the MAIL.BOX messages one by one. This is especially important on mail hub servers.

While it is often desirable to route mail as quickly as possible, it is also possible to have too many router tasks running. If you are using multiple router tasks, you should monitor the performance and use of each of the tasks to be sure that none of them remains idle the majority of the time, as extra router tasks can take up significant system resources. We recommend you also temporarily change the NOTES.INI Log_MailRouting setting to show detailed thread information so you can gain a better understanding of task usage. The following is the output at the server console when Log_MailRouting was set to 40. The output has been truncated for readability.

```
05/16/97 09:55:30 AM Router: Transfer to server <$LocalDelivery> is ready with 1 messages
05/16/97 09:55:31 AM Router: Starting delivery thread 000000E2
05/16/97 09:55:31 AM Router: Delivery thread 000000E2 searching for work
05/16/97 09:55:31 AM Router: Message 004C7330 delivered to Rufus Credle/ITSO from
    ITSORAL01/ITSO UNID - OFB3FBB11A:4DAED790 ON85256499:004C7330 Size: 1K Time: 00:00:01 Hop Count: 1
05/16/97 09:55:31 AM Router: Idling delivery thread 000000E2
05/16/97 09:56:30 AM Router: Transfer to server <$LocalDelivery>
05/16/97 09:56:30 AM Router: Delivery thread 000000E2 searching for work
05/16/97 09:56:30 AM Router: Message 004C901D delivered to Darryl Miles/ITSO from
    ITSORAL01/ITSO UNID - OFA365580C:0BE62ACE ON85256499:004C901D Size: 1K Time: 00:00:01 Hop Count: 1
05/16/97 09:56:30 AM Router: Idling delivery thread 000000E2
```

You will see from this information the usage of various threads and whether adding more will be more effective for your server.

To change the maximum number of mail threads set the following NOTES.INI value:

```
MailMaxThreads = number
```

You can also specify this setting in a Server Configuration document in the Public Name and Address Book.

4.4.8 Number of Databases Cached

If your server has sufficient memory, you can improve the performance of the server by increasing the number of databases that Domino can cache in memory at the one time. The default is 25 or the NSF Buffer Pool Size divided by 300K, whichever value is greater. The maximum number of databases that can be cached in memory is approximately 2000. Monitor the Database.DbCache.Hits statistic on your server. This indicates the number of times a database open request was satisfied by finding the database in cache. A high value indicates database cache is working effectively. If the ratio of Database.DbCache.Hits to InitialDbOpen is low, you might consider increasing NSF_DbCache_Maxentries.

To set the number of databases that a server can hold in its database cache at one time, set the NOTES.INI value as follows:

```
NSF_Database_Maxentries = number
```

4.4.9 Multiple Replicators

In 2.5.2, “Replication Schedules” on page 20 we noted that you can load multiple replicator tasks on a hub server to save replication time. In order to have multiple replicators running on your server you need to increase the replicator’s NOTES.INI value.

To set the number of Replicator tasks that can run concurrently on the server, edit the NOTES.INI value:

```
Replicators = value
```

4.4.10 Minimum Mail Poll Time

You can control the minimum allowable time in which Notes clients may poll for new mail. It is possible that your Domino server resources are being overtaxed by being constantly bombarded by requests for new mail from the Notes client machines if users have changed their default new mail notification check time from 15 minutes to a smaller number such as 2 or 5. You can control the minimum frequency of these requests from the server by using the MinNewMailPoll NOTES.INI variable. This variable determines the minimum allowable checking time from clients regardless of the value that is specified on the client machines. The syntax of this variable is as follows:

```
MinNewMailPoll= minutes
```

For example:

```
MinNewMailPoll=15
```

4.4.11 Memory Buffers

One of the most important memory pools in a Domino server is the NSF buffer pool. This memory pool is used for buffering I/O transfers between the NIF indexing functions and disk storage. The maximum size is 256 MB at the time of writing. This memory pool size is automatically calculated by the server. To increase the default size of the NSF buffer pool, the following value must be changed within the NOTES.INI file:

```
Nsf_Buffer_Pool_Size = value in bytes
```

We recommend, however, that the NSF buffer size be allowed to take the default value. Changing this value from the default will stop the automatic calculation of this memory pool by the server and your server may actually perform worse.

4.5 Using the Windows NT Performance Monitor

If you are running Domino on a Windows NT server you can use a powerful tool called the Performance Monitor. Performance Monitor is a graphical tool for measuring the performance of your own Windows NT-based computer or other Windows NT-based computers on a network. It is located in the Administrative Tools group of both the Windows NT Workstation and Windows NT Server products. On each computer, you can view the behavior of objects such as processors, memory, cache, threads, and processes. Each of these objects has an associated set of counters that provide information on such things as device usage, queue lengths, and delays, as well as information used for throughput and internal congestion measurements. It provides charting, alerting, and reporting capabilities that reflect current activity along with ongoing logging. You can also open log files at a later time for browsing and charting as though they were reflecting current activity.

We discuss how to use the Performance Monitor tool to help you investigate the performance of your Domino server and look for bottlenecks if you feel your Domino server is not performing as well as you would prefer. If you are running Domino on an OS/2 Warp Server, you can use equivalent tools such as System Performance Monitor (SPM/2) or Netfinity.

Note: In our example of the Windows NT Performance Monitor tool, you will see items such as Lotus Notes Objects and Notes Statistics. Both of these items represent the Domino server.

4.5.1 Starting Performance Monitor

You can activate Performance Monitor by clicking on **Start** → **Programs** → **Administrative Tools** → **Performance Monitor**, or by typing PERFMON on the command line. In either case, you will see the window shown in Figure 31 on page 98.




Figure 31. Windows NT Performance Monitor

Note

No disk activity is monitored by Performance Monitor unless you first issue DISKPERF -Y on the command line and reboot the server.


By enabling the disk monitoring you are placing additional loads on the processor. We recommend that you disable this feature once you have completed your tests. This is done by executing DISKPERF -N from the command line.

If you do not have a Network Interface object when you start Performance Monitor, you will need to install the SNMP Service for Windows NT. You can do this by starting the Network program located in the Control Panel. Select the **Services** folder and then click on **Add**. Select **SNMP Service** from the list of services to be installed.


4.5.2 Collecting Data with Performance Monitor

To collect and display Performance Monitor data, you can either choose from the View pull-down menu one of the following options, or you can click on one of the smart icons located at the top of the Performance Monitor main panel.

The options are:

1. Chart (smart icon )


It graphically displays (in real-time) how the values of counters change in response to changes in system load conditions. Two display modes are available: line graph or histogram. Refer to 4.5.3, “Chart View” for more details.

2. Alert (smart icon )

This view, discussed in 4.5.4, “Alert View” on page 103, is particularly useful for monitoring servers in a distributed environment where centralized administration and alerts are tracked by the logging mechanism. For each alert that you are monitoring, you have to supply a threshold value. When the specified alert threshold value is reached, the alert is written to the application log. You also have the option to start a corrective procedure to perform tasks that are aimed at containing the problem that triggered the alert.

3. Log (smart icon )

In order to perform time-lapse analysis, real-time data is stored to disks for retrieval later. This view is described in 4.5.5, “Log View” on page 105.

4. Report (smart icon )

This is useful for observing numeric values of multiple counters simultaneously and helps in deciding which counters to place in a chart. This view, detailed in 4.5.6, “Report View” on page 108, is also useful for the exportation of specific sets of numeric data from Performance Monitor to applications such as Excel or Lotus 1-2-3.

4.5.3 Chart View

The chart view is displayed by selecting **Chart View** from the View pull-down menu, or clicking on the smart icon  .

When a system object is selected for display, the values of the specified counter are charted in a graphical format as shown in Figure 32 on page 100. This graph is displayed in real-time and reflects any changes in system load. The example in Figure 32 on page 100 shows only one object being measured, but multiple objects can be selected for simultaneous charting.





Figure 32. Performance Monitor - Chart View

4.5.3.1 Adding Objects to Be Monitored

From the Edit pull-down menu of the Performance Monitor, select **Add to chart**, or click on the smart icon . To select an object to monitor, first choose an object and a related counter (see Figure 33).





Figure 33. Adding Objects to Be Monitored

The Select Computer button () to the right of the Computer: field lets you select a remote server running Windows NT. You can then monitor components on that machine. Figure 34 on page 101 shows an example of this window.

Note: It may take a while to display this window while the network is checked for other Windows NT systems.

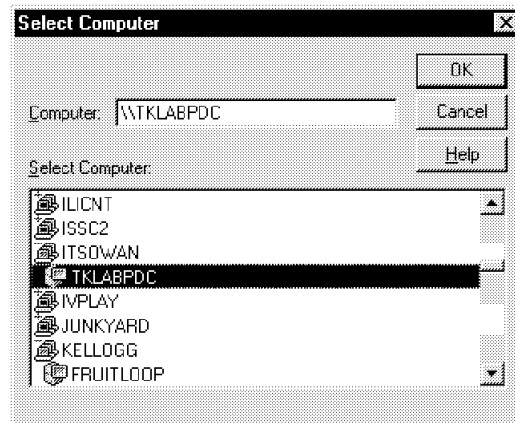


Figure 34. Adding Objects from Remote Servers

4.5.3.2 Knowing What the Objects Are Measuring

If you click on the **Explain >>** button, you will see a brief description of the selected object (see Figure 35). This option is extremely useful and can usually provide valuable information regarding the selected object.

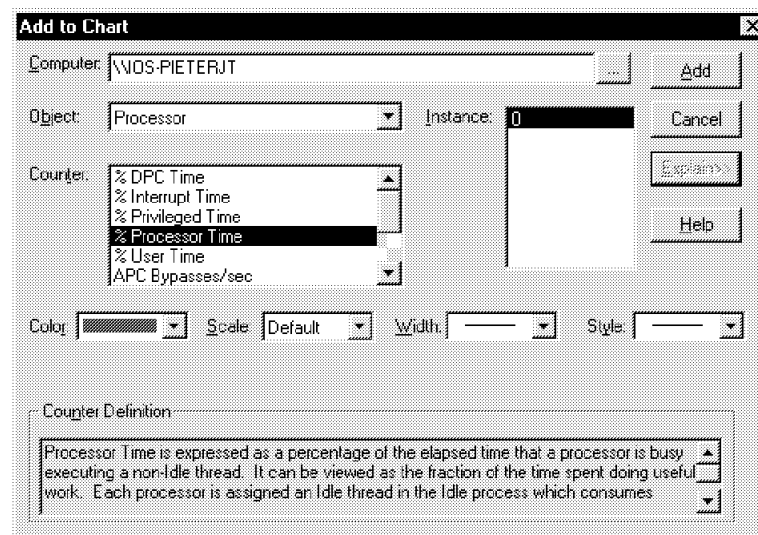


Figure 35. Using the Explain Button to Get Details

4.5.3.3 Determining Which Objects Can Be Measured

In Figure 36, you can see a list of objects that can be selected. Notice that the only objects displayed are components that have been installed in your server. For example, if you do not have the IPX protocol installed on your system, you will not see any objects relating to this protocol.




Figure 36. A List of Installed Objects

4.5.3.4 Choosing an Object Instance

Some objects have multiple instances. For example, a counter can have multiple instances relating to the same measurement for different but equivalent devices, such as multiple CPUs. This is shown in Figure 37.





Figure 37. Several Instances of the Same Object

4.5.3.5 Changing the Chart Appearance

There are various options relating to the appearance of the charts. To change them, select **Chart** from the Options pull-down menu, or click on the Options

smart icon (). You will see a window similar to the one in Figure 38 on page 103.

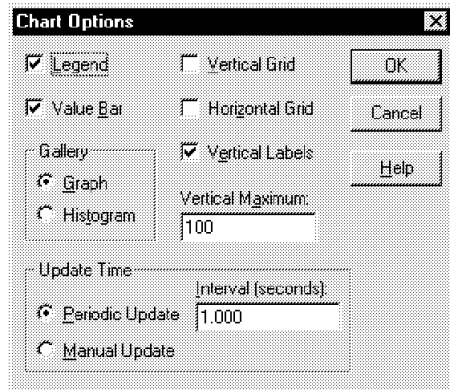


Figure 38. Chart Options Window

Selecting **Histogram** from the Gallery box is useful for simultaneously monitoring many instances of a given counter. With the Histogram view, you cannot see the history of the measured objects but you will gain a feel for how they relate to each other.

4.5.4 Alert View

Alert view is useful when you want to determine how many times a certain occurrence takes place in relation to a certain object. Alternatively, use Alert view if you want to take corrective action when the measured object reaches a predetermined threshold.


To select the Alert view, click on the **View** pull-down menu and then select the **Alert** option. Alternatively, you can click on the smart icon  .

Figure 39 on page 104 shows the Alert view with some alerts generated for the processor object that has exceeded the predetermined 10% utilization threshold. Times and dates are shown in the panel, as well as the server from which the measurement was taken.

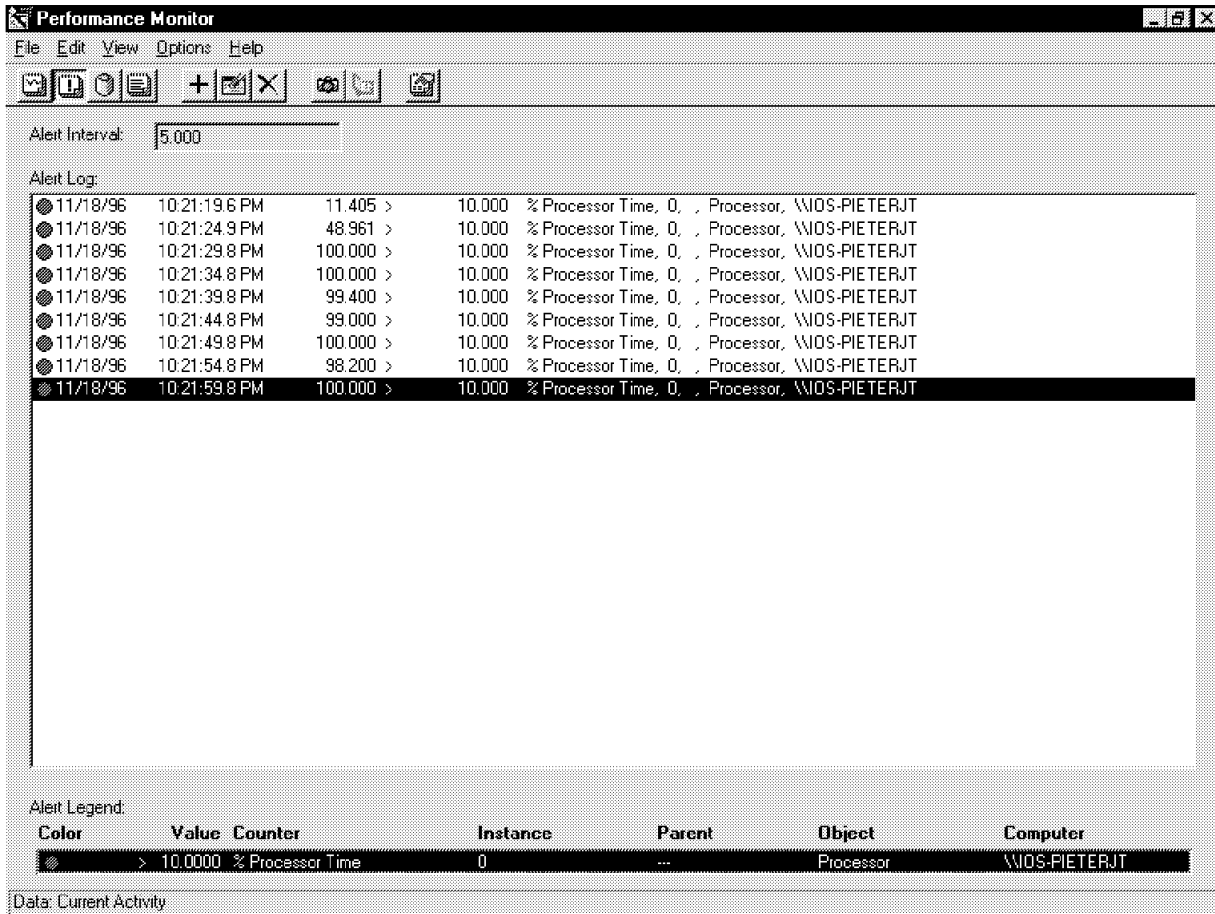



Figure 39. Performance Monitor - Alert View

4.5.4.1 How to Measure Objects Using Alert View

To add an object for monitoring in Alert view, click on **Add to Alert** from the Edit pull-down menu, or click on the smart icon . You will see a window similar to the one in Figure 40.

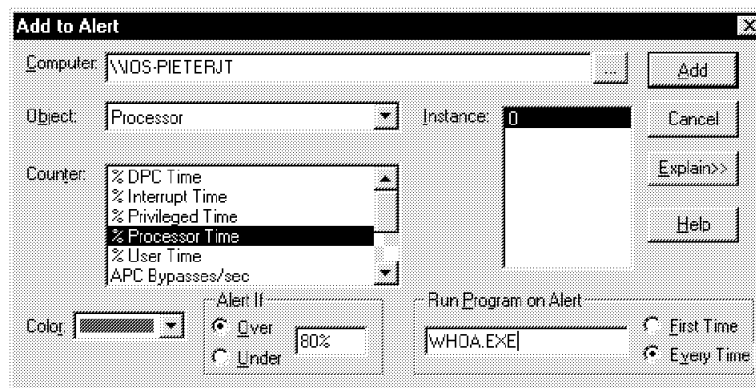


Figure 40. Adding an Alert


4.5.4.2 How to Set Up Alerts

The procedure for adding an object to be measured in Alert view is similar to the Chart view function, but you also have to complete the Alert If field. The steps are as follows:

1. Choose an object to measure.
2. Choose the counter to measure.
3. Choose the instance of the object to measure.
4. Supply a numeric value in the Alert If field (see Figure 40 on page 104).
This value is the threshold upon which the alert is generated when the object measured reaches it.
5. You can supply an optional program or batch file name in the Run Program on Alert field (see Figure 40 on page 104). This program or batch file is triggered once the threshold value is reached.

4.5.5 Log View

Log view is useful for capturing and storing data to disk for analysis at a later stage.

To select the Log view, click on the **View** pull-down menu and then select the **Log** option. Alternatively, you can click on the smart icon . You will see a panel that resembles the one in Figure 41.





Figure 41. Performance Monitor - Log View

4.5.5.1 Selecting Objects to Be Added to the Log File

From the pull-down menu, click on **Edit** then **Add to Log** or click on the smart icon . Figure 42 appears. Select each object you wish to log and click on the **Add** button. When you are finished selecting objects to log, click on the **Done** button. (The Done button appears after clicking Add.)

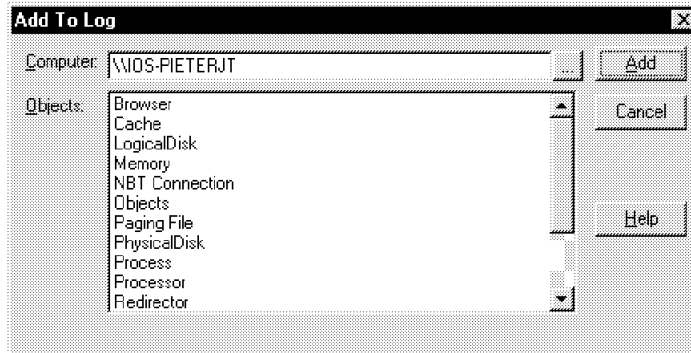



Figure 42. Adding Objects to the Log File

Other Servers

You will notice that you do not have the option to select individual counters. However, you can select objects from different servers by changing the Computer field.

4.5.5.2 Starting or Stopping the Log File

Once you have selected the objects to log, you can choose the **Options** → **Log** or click on the smart icon . You will see a window similar to the one in Figure 43.

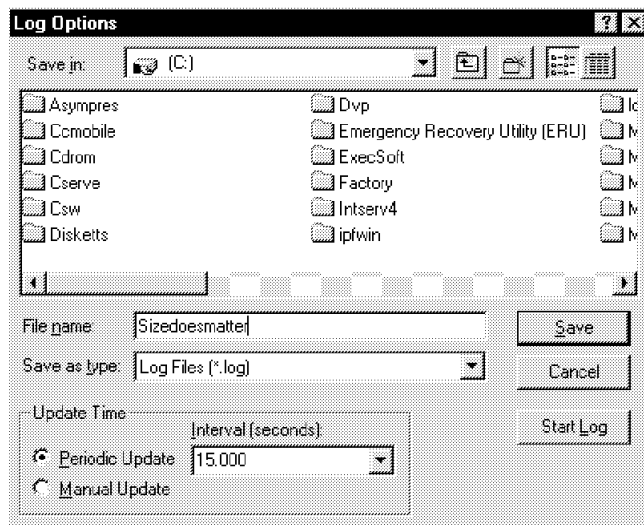


Figure 43. Log Options Window

Type in a file name to save the log to. Modify the logging interval as required.

To start the logging, click on the **Start Log** button.

To stop the logging once it starts, click on **Stop Log** from the Log Options window.

Note: To create a log file, first select the objects that you want to capture data from. If you have not selected any objects, the Start Log button shown in Figure 43 on page 106 is grayed out.

4.5.5.3 Retrieving Data from the Log File


By default, Performance Monitor displays current data. In order to view logged data instead you need to *switch* the source of data from Current to Log File.

To use the data collected in the log file, click on **Options** → **Data From...** from the main window. Figure 44 appears.



Figure 44. Selecting the Source of Data

Select **Log File** and type in the file name of the log you specified in 4.5.5.2, “Starting or Stopping the Log File” on page 106.

Now, you have the ability to use the other views (Chart, Report and Alert) to view the data you have captured. You add the objects and counters to the selected view as before, using the  smart icon. See 4.5.3, “Chart View” on page 99, 4.5.6, “Report View” on page 108 and 4.5.4, “Alert View” on page 103 for details.

4.5.5.4 Selecting a Time Slice from the Log File to Analyze

To select a particular time frame from the log file that you want to analyze, select **Time Window** from the Edit menu. The Input Log File Timeframe window is displayed (Figure 45).




Figure 45. Time Slice Option Window

You can use the slider bar in Figure 45 to change the time window of the data that you are analyzing.


4.5.5 Exporting the Log

Once you have displayed the log information using either the Chart, Report or Alert views, you can export that view in text format for analysis by other applications. The data is exported in either:

- .TSV format (tab separated variable)
- .CSV format (comma separated variable)

To export data, select **File** → **Export...** from the main pull-down menu bar.

4.5.6 Report View

The report view is displayed by selecting **Report View** from the View pull-down menu, or clicking on the smart icon  .

Report view is useful for displaying measurements from the individual objects in a numerical form. You can export this data for trend analysis to an application such as Excel or Lotus 1-2-3 by choosing the **Export** function from the File menu.

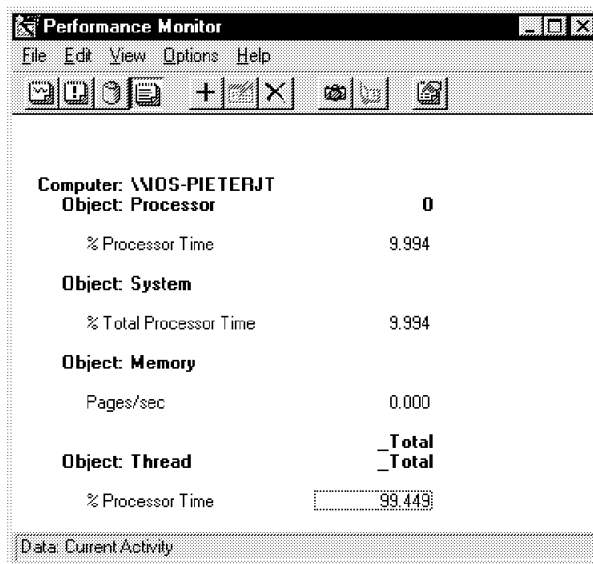


Figure 46. Chart Options Window

You can select multiple objects to be added in the Report view, as shown in Figure 46. Unlike charts, there is no graphical display, but the real-time numeric values of the objects are shown.

4.5.7 Installing Domino Statistics As a Performance Monitor Counter

You can install Domino server statistics as a counter within the Windows NT Performance Monitor. The Performance Monitor lists all numerical Domino server statistics, including statistics generated by add-in programs.

Type the following command from the Notes program directory on the server:

```
notesreg.bat <directory>
```

Where directory is the full path to the Notes program directory.

For example, type:

notesreg.bat D:\notes

This will add a Lotus Notes object to the Performance Monitor program as shown in Figure 47.

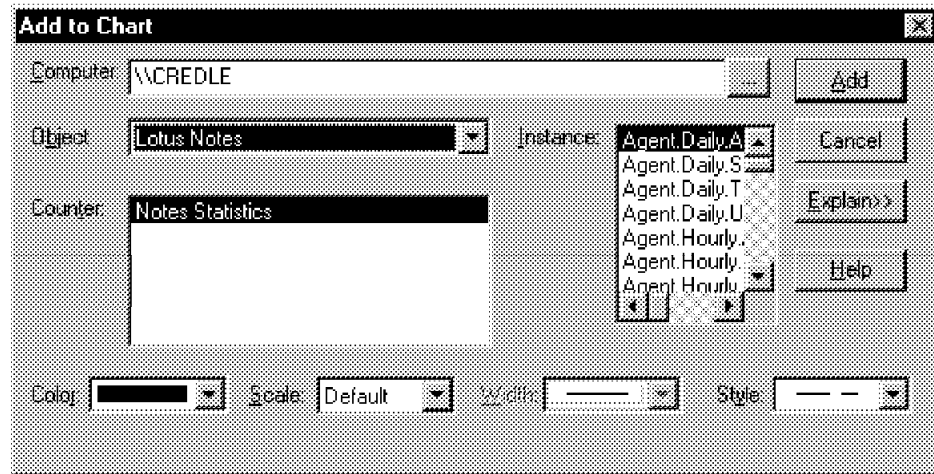


Figure 47. Lotus Notes Performance Monitor Object

When viewing the Notes statistics using the NT Performance Monitor be aware of the following:

- Notes statistics don't appear as instances in the Performance Monitor until Domino or an add-in program assigns or updates a statistic. To force this to happen, initialize Notes statistics on the server, for example, by typing Show Stat at the server console.
- If the Domino server or add-in program running on the server terminates abnormally, you must stop the Performance Monitor before restarting the Domino server or add-in program.
- You can use the Performance Monitor to view the statistics of a remote Domino server. If the remote Domino server or add-in program running on the server terminates abnormally and you stop the Performance Monitor but can't restart the Domino server or add-in program, you may need to reboot the remote Domino server machine.

4.6 Server Subsystems to Monitor

The critical subsystems of Domino servers are:

- CPU
- System memory
- Network subsystem
- Disk subsystem

The single resource that consumes the most time during a task's execution is that task's bottleneck. Bottlenecks can occur because resources are not being used efficiently, resources are not being used fairly, or a resource is too slow or too small. For example, a high-performance server has plenty of RAM and disk space but the bottleneck might be a slow network speed to the server. This bottleneck can make the server seem slow from a user's perspective even if the

server has sufficient capacity. Another example is a server that has a very high CPU utilization. The bottleneck is actually an insufficient amount of memory that is causing high amounts of paging, which therefore affects the CPU. Be aware that a bottleneck in your server in one area may show up in other server systems.

Table 10 lists the key subsystems you should monitor on your Domino server along with a brief description, and when each item can cause a server bottleneck. These bottlenecks are discussed later in this chapter in more detail.

<i>Table 10 (Page 1 of 3). Windows NT and Domino Subsystem Objects to Monitor</i>				
Object	Counter	Instance	Description	Bottleneck
CPU				
Processor	% Processor Time	CPU number	The % Total Processor Time is the average percentage of time that all the processors on the system are busy executing non-idle threads.	Average > 80 % then possibly add another CPU, or increase speed.
System	Processor Queue Length	CPU number	Processor Queue Length is the instantaneous length of the processor queue in units of threads. All processors use a single queue in which threads wait for processor cycles.	If number is large, processes need to be run, but can't get enough CPU time.
Processor	Interrupts/sec	CPU number	Interrupts/sec is the number of device interrupts the processor is experiencing. A device interrupts the processor when a task has been completed or when it otherwise requires attention. Normal thread execution is suspended during interrupts.	> 1000 look for inefficiency of hardware I/O devices.
System Memory				
Memory	Available bytes		Available Bytes displays the size of the virtual memory currently on the Zeroed, Free, and Standby lists.	< 4 MB consistently then add more memory

Table 10 (Page 2 of 3). Windows NT and Domino Subsystem Objects to Monitor

Object	Counter	Instance	Description	Bottleneck
Memory	Pages/sec		Pages/sec is the number of pages read from the disk or written to the disk to resolve memory references to pages that were not in memory at the time of the reference. This is the sum of Pages Input/sec and Pages Output/sec.	Average: > 10 warning > 20 urgent Add more memory
Paging File	% Usage Peak	\\?.\D:\pagefile.sys	The peak usage of the Page File instance in percent.	See also the Process Page File Bytes Peak; use to determine size of memory paged to disk.
Disk subsystem				
Physical Disk	% Disk Time	_Total	Disk Time is the percentage of elapsed time that the selected disk drive is busy servicing read or write requests.	> 70%, server swapping too much or faster disk system needed
Network subsystem				
Network Interface	Bytes Total/sec		The number of bytes the server has sent to and received from the network. The value provides an overall indication of how busy the server is.	> 2 million bytes/sec for 16 Mbps TR. > 1.25 million bytes/sec for 10 Mbps Ethernet.
Domino Server Workload				
Lotus Notes		Mem.Allocated	Useful for assessing the load on a Domino server. This is dynamic memory allocated as users connect to the server.	A consistently high value may indicate another server.
Lotus Notes		Server.Trans. PerMinute	Number of transactions that took place in the last minute.	When continually high, redistribute the load to other servers if this load is high relative to other servers.

<i>Table 10 (Page 3 of 3). Windows NT and Domino Subsystem Objects to Monitor</i>				
Object	Counter	Instance	Description	Bottleneck
Lotus Notes		Server.Trans. PerMinute.Peak	Peak number of transactions that took place in any given minute since the server started.	Good for comparing to other Domino servers.
Lotus Notes		Server.Trans.Total	Total number of transactions the server has serviced since it was started.	Good for comparing to other Domino servers.

There are quite a large number of things you may need to monitor on your server to detect a performance problem. However, if you break a server into its components the list of things you may need to monitor shown in Table 10 on page 110 will seem less daunting. We now explain each subsystem individually, and also show some Performance Monitor results from a test Domino server to aid you with your performance analysis.

4.6.1 Memory Bottlenecks

Lack of memory is by far the most common cause of serious performance problems in computer systems. Our guideline discussed in 3.13.1, "Server RAM" on page 63 is a good starting point for the minimum amount of RAM you should use in your server. However, if more users simultaneously connect to your server than you have planned, then the amount of RAM in your server may not be enough. To handle this lack of memory the system starts paging. The performance of your server will decrease, because physical disk access takes much longer than memory access. Your Notes and Web clients will see this bottleneck at the server as slower response times.

The Memory - Available Bytes counter displays the amount of free virtual memory. If this counter stays consistently below 4 MB on your Domino server, paging is occurring and performance is less than optimal.

The Performance Monitor Memory - Pages/sec counter is the number of pages read from the disk or written to the disk to resolve memory references to pages that were not in memory at the time of the reference. As a rule, you can assume that if the average of this counter is consistently greater than 10, then memory is probably becoming a bottleneck in the system. Once this counter starts to average consistently at 20 or above, performance is significantly degraded and disk thrashing is probably occurring.

If Memory - Pages/sec is increasing, yet Memory - Available Bytes is not decreasing, then you may actually not have a memory bottleneck. The Memory - Pages/sec counter also increases for each non-cached read from disk, so access to some databases may make this counter less useful.

With Windows NT 4.0 you can use the task manager to get the amount of memory being used by each task. Simply right click on the task bar and choose **Task Manager**. This can give information as to whether there are any tasks using more memory than we expect. We can possibly stop these tasks or add more memory.




Figure 48. Windows NT Task Manager

If you decide you need to add more memory to your Domino server, you can use the following guide to determine how much memory to add:

Approximate Memory to Add Formula

$$\text{Paging File \% Usage Peak} \times \text{Paging File size (on disk)} = \text{amount of RAM needed}$$

For example, if your page file is 100 MB and the Paging File % Usage Peak is 20%, then you would need approximately 20 MB of additional RAM to have a system that does minimal paging. We recommend that if you add additional memory you also investigate if you can add more secondary cache memory as well.

At the server console enter the command:

```
show stat mem.availability
```

This will give Domino's general assessment of free memory available: Painful, Normal or Plentiful. Add more memory if the assessment is anything other than Plentiful.

4.6.2 Processor Bottlenecks

The CPU of a server is the most demanded resource on a Domino server. This is quite different from file servers that are more disk intensive. So if your server has enough memory then the CPU is released from the great burden of memory paging (discussed in the previous section).

If the Processor - % Processor Time counter consistently registers (averages) above 80%, the processor may be the bottleneck. If this occurs, you should determine what task is using the CPU. You can use the Performance Monitor or Task Manager to do this. You can tell if the CPU activity is due to applications or to servicing hardware interrupts by monitoring Processor - Interrupts/sec. This is the number of device interrupts the processor is experiencing. A value over 1000 should cause you to look at the efficiency of hardware I/O devices such as the disk controllers and network cards. For example, if an inappropriate ISA LAN adapter is used in a server, this could result in a large interrupt value and would therefore be the bottleneck in the system.

Another important metric is System Processor Queue Length. This is a measure of how many threads are in the queue waiting to be executed. If this number is small, this indicates the processor is able to keep up with demands. If this number is large, this indicates that you have processes that are waiting to run but can't get enough CPU time. This is a good indication that you have a processor bottleneck.

Monitor the Domino task `Server.Trans.PerMinute`. This value is only calculated once per minute. If this value rises when the CPU rises then you can determine whether Domino is the primary CPU load. Also monitor the `Server.Trans.PerMinute` peak value as well. You can then determine whether your server is consistently overloaded or this is merely a short-term peak.

You can also set `Server_Show_Performance=1` in the `NOTES.INI` file. This will display the Transactions/minute value and the number of users connected on the server's console every minute. Change this value back to zero to turn off console performance monitoring.

What can you do if your CPU seems to be the bottleneck in your server? The most efficient way to overcome this bottleneck is to add another CPU if your server has SMP capability. If you read through our server sizing chapters, then you will have noticed that we recommended you purchase an SMP-capable server, even if you did not require it initially. If your CPU is the bottleneck in your server, you can now reap the benefits of your initial decision and add another CPU in your server.

Another choice you have to reduce CPU load is to move certain Domino tasks to another server. If your server is both a mail and application server, you may need to consider having two servers providing these Domino services, or consider implementing a clustering solution as discussed in Chapter 5, "Lotus Domino Clustering" on page 125.

4.6.3 Disk Bottleneck

Since hard disks are mechanical devices and therefore one of the slowest devices in the system, a bottleneck here can be very costly in terms of system performance. The metric to track in the disk subsystem is the percent of time that the disk(s) is being utilized. The timer, Physical Disk - % Disk Time shows you the overall disk utilization for a specific instance of a hard disk (or one array). This can be further analyzed in terms of the read percentage (% Read time), and the write percentage (%Write time).

Remember that one bottleneck can show up in other areas. If you discover a disk bottleneck, the first thing you need to determine is whether it's really more memory that you need. If you are short on memory, you will see the lost performance reflected as a disk bottleneck.

If the Physical Disk object's % Disk Time counter consistently registers above 70%, the physical disk is the bottleneck. This counter is the percentage of elapsed time that the selected disk drive is busy servicing read or write requests, including time waiting in the disk driver queue. You can check that paging is not causing a disk bottleneck by the following formula:

Excessive Paging Calculation Formula

$$\% \text{ disk time used for paging} = 100 \times (\text{Memory Pages/sec} \times \text{PhysicalDisk Avg DiskSec/Transfer})$$

If this is more than 10% of the total disk activity, then paging is excessive.

The first step to alleviate a disk bottleneck is add more RAM, as this adds extra file system cache. The next step is to install a faster disk and controller and this makes the biggest difference to disk subsystem performance. Adding more physical drives in a RAID 5 configuration can result in significant performance improvements when the disk subsystem is the bottleneck. However, adding more controllers usually does not significantly improve performance. When using high-performance disk controllers, the physical drive access times are usually the performance limiting factor for the disk subsystem.

A disk subsystem bottleneck is one of the hardest to overcome when a Domino server is running in production. This is because adding extra hard drives to a RAID 5 configuration usually involves rebuilding the entire disk system. This means that production data needs to be backed up to tape, the operating system and Domino re-installed, and your data restored. This can take several hours to complete, which would normally have to be done out of business hours. You can possibly add a new disk controller and disk drives, and configure a new RAID logical drive as an alternative to adding more drives to your first. This may remove the load with your original disk subsystem.

From our previous discussions we hope you can see that implementing the best disk subsystem you can buy when you purchase your server will reduce the need to upgrade your disk subsystem for the life of your server.

4.6.4 Network Bottleneck

The network subsystem is monitored by observing the network bytes/sec. Network bytes/second shows you the throughput of your LAN subsystem. If you know the LAN bandwidth of your LAN adapter, then you can see whether the load is approaching the limits of the adapter. For example, the throughput of an IBM Auto LANStreamer MC32 approaches the speed of the media. At 16 Mbps, this equates to a throughput of roughly two million bytes/sec. If you monitor the network bytes/second and see that your server is passing one million bytes/second through this adapter on a sustained basis, then it is time to add a second network adapter.

Make sure you have a high-performance network card that does not need the CPU to transfer data between disk and the network such as a PCI full-duplex adapter.

If you have your server network adapters connected on the same segment as your users, you may see a large number of collisions on an Ethernet network, or generally high network utilization. Network Segment % Network Utilization should be considered when things start slowing down to the point they are no longer acceptable. Some say that this point is around 40 to 50%. Then the network is the bottleneck. Therefore connecting your server into a high-performance switch can reduce the collision domain and increase bandwidth into the server. Full-duplex network cards are a wise investment for your Domino server because they allow the server to transmit and receive simultaneously.

4.6.5 Other Domino Systems to Monitor

You should make sure that the number of users that connect to your server is not greater than the capacity limits you initially set. For example, the NotesBench report might have stated that the maximum number of users is 400 Groupware_A users for your server. You planned for 300 Groupware_A users and the Server.Planner tool chose a particular server to meet this requirement. You should monitor your server to make sure that there are no more than 300 clients currently using your server each day. If there are, then you will most likely need to increase the capacity of your server if possible, redistribute users to other servers, or choose a more powerful server.

One of the key statistics in Table 11 on page 117 is the Server.Sessions.Dropped. This indicates the number of users who gave up waiting for a server response, usually by pressing the Ctrl+Break key sequence. This can indicate a server is too slow to satisfy a user's request. In which case you need to look for a server bottleneck. It can also indicate a slow link speed. For example, if you have a high-performance server but the link speed is very slow, users can give up waiting for information to be sent down the slow speed link. This may occur when opening their mail database on a remote server. In this case the link speed is the bottleneck and not the server.

We strongly advise you to monitor the following Notes statistics along with the server objects outlined in Table 11 on page 117.

<i>Table 11. Important Domino Statistics to Monitor</i>		
Statistic Label	Statistic Name	Description
Number of current users	Server.Users	Number of users with sessions open on the server. Useful to monitor overall server use. If this number is consistently higher than that of other servers and performance is a problem, redistribute the server load to other servers. Ensure this value is not greater than your initial sizing limits.
Peak number of users	Server.Users.Peak	The maximum number of concurrent users with open sessions since the server was started.
Number of dropped sessions in mid-transaction	Server.Sessions.Dropped	Number of users who gave up trying to connect to server (usually by restarting or pressing Ctrl+Break). Useful to monitor server response time. A high number of dropped sessions indicates a slow server, or a slow link.

4.7 Domino Performance Analysis Example

We now use the NT Performance Monitor tool to look for server bottlenecks on a test Domino server. In this example, an IBM PC Server 720 Pentium 90 MHz fitted with 64 MB of RAM and running Windows NT 4.0 and Domino 4.51, and connected to a single token-ring LAN operating at 16 Mbps.

To create a simulated user load we used a tool called NoteSizr (today, NoteSizr has been renamed to GroupSizr for Lotus Notes) developed by Technovations. This tool allows us to program a series of tasks that a user would perform on a Domino server such as opening databases, searching, and creating documents. It can simulate up to several thousand users. Using the NoteSizr tool and the NT Performance Monitor, we highlight Domino server bottlenecks, which will prove useful in identifying performance problems on your servers.

4.7.1 Normal Server Operation

Our test server was loaded with 50 clients and the results of some of the key server performance indicators are shown in Figure 49 on page 118. You can see that the CPU utilization is fluctuating between zero and at one point 80%. This is normal for a server; the key point is that the CPU is not averaging above 80%. We look later in this chapter at a CPU bottleneck in a server. The Available Bytes counter is larger than 10 MB which is why the line is graphed at 100%. The servers Pages/sec counter is also very low, which is very good because this means that there is adequate memory in this server to meet the user load.




Figure 49. Normal Server Operation Performance

The screen shown in Figure 50 on page 119 is the NoteSizr Run-Time window. This screen is displayed on the client workstation that is performing the load on the server. The window lists a series of Notes operations and the number of times that each operation has been tested.




Figure 50. NoteSizr Run-Time Monitor Screen

In Figure 51 on page 120 you can see the Test Results screen that is presented after the test is completed. NoteSizr provides information on the time it took to complete various Notes tasks.

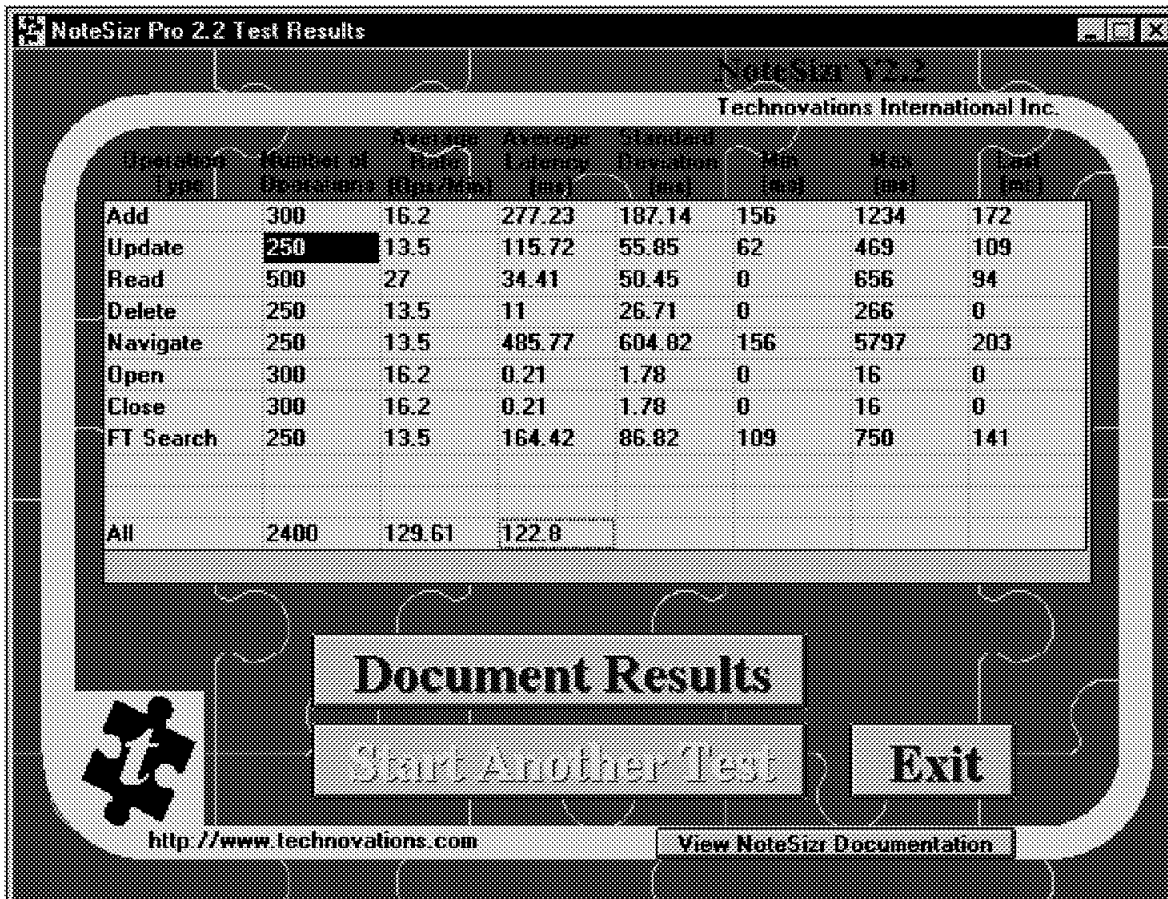


Figure 51. NoteSizr Test Results Screen

4.7.2 Memory Bottleneck Example

With 50 simulated clients connected to our test server, the minimum memory formula recommends the following amount of RAM:

$$48 + 50/3 = 64.3 \text{ MB.}$$

Therefore, our server has approximately enough RAM to meet this number of clients. As we load this server with more clients, Domino will require more RAM per active client. If we do not add the required RAM our server will perform more paging.

Figure 52 on page 121 shows our test server loaded with 1000 clients using NoteSizr. You can see that the Pages/sec value increases as Windows NT pages memory between real RAM and the paging file on disk. The Available Bytes value falls in value as real RAM is used by Domino, and then increases when Windows NT pages memory to disk. We know that if this value falls below 4 MB, then memory needs to be added to the server. On this server, the Available Bytes value is approximately 4 MB or lower.

The Pages/sec value (represented by the black hyphenated line) is very high, and is averaging above 20 Pages/sec. This is another indicator that memory must be added to this server. You can see that insufficient memory is causing excessive paging. This causes the CPU utilization to be much higher than it would otherwise be.




Figure 52. Performance of Server with Memory Bottleneck

As the test continues, the user load script changes to adding documents, and performing searches, which is very CPU intensive. The lack of sufficient memory is exacerbating the CPU utilization. The server eventually spends most of its time paging memory to and from disk, and the response time for users falls dramatically.




Figure 53. Server with Disk Thrashing Due to Insufficient RAM

4.7.3 Monitoring Interrupts, Network Interface and Disk Subsystem

We previously discussed that you should monitor Interrupts/sec, Network Interface and Disk subsystem. Figure 54 shows a graph of these three items when our test server was loaded with 1000 clients.

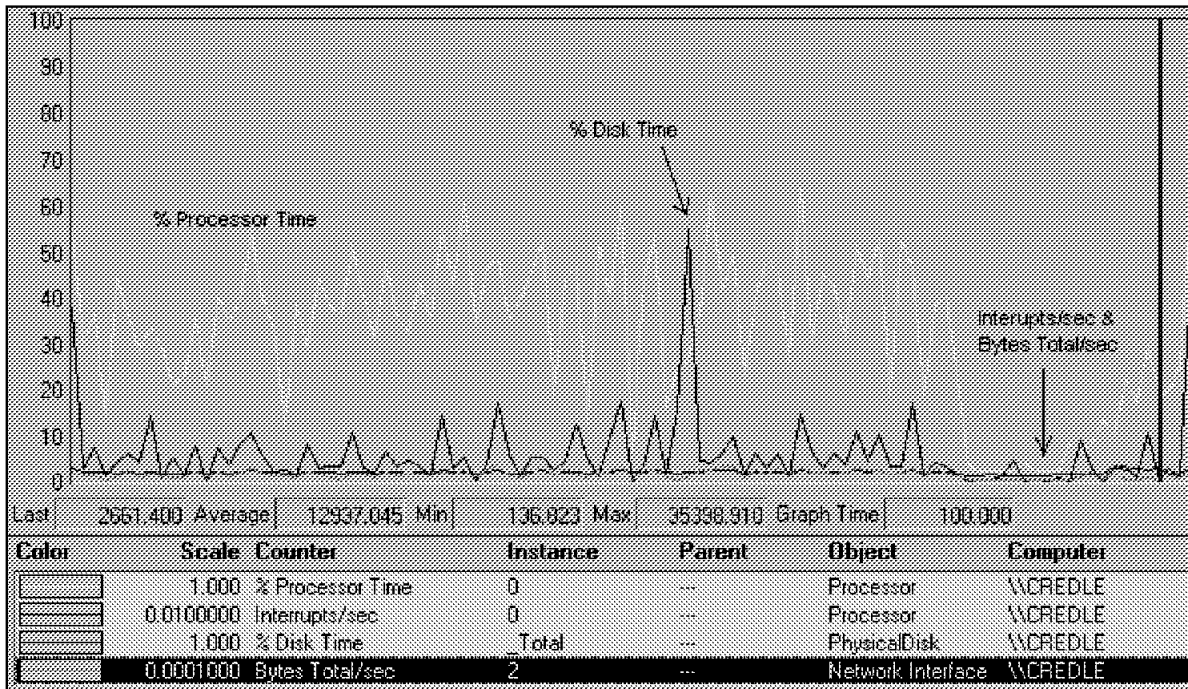


Figure 54. Other Important Variables to Monitor

We can see that although our server is under significant load, the % Disk Time value is below the warning level of 70% so we do not have a disk subsystem bottleneck. The Interrupts/sec value is also very low, so our network cards are not a bottleneck for the CPU. The network card was a Micro Channel busmaster network card which does not require the processor to transfer data between memory and the network. The Bytes Total/sec is very low as well so our network is not a bottleneck either. These performance indicators should be measured on your server as they may be the bottleneck in your server.

4.7.4 CPU Bottleneck Example

If you have added sufficient memory to your server, then the Pages/sec counter should be on average a small value. However, your CPU might be a bottleneck. In Figure 55 on page 123 we can see that the CPU (represented by the white line) is averaging at about 63%. This is a high enough value for you to want to monitor your server more closely. If the CPU average rises above 80%, you will need to add additional processors to your server or transfer server load to other Domino servers. You can see that as the server becomes increasingly stressed the number of sessions dropped to Notes clients also increases.




Figure 55. Server with CPU Bottleneck

We added an additional two processors to our test server, bringing the total number of CPUs to three. We then ran our test again and you can see the results (represented by the white line) in Figure 56.




Figure 56. Two Additional CPUs Gives Significant Performance Benefits

We have only plotted the CPU of processor one for clarity, but you can see that the average utilization is much lower. Another benefit of multiple processors is

that the processor queue length is almost always zero. This means that there are hardly ever any threads waiting to be serviced by a processor. The number of server transactions per second is also much larger with multiple processors.

Chapter 5. Lotus Domino Clustering

Lotus Domino Server 4.6, Powered by Notes (formerly known as Lotus Notes Server) introduces a new functional option called Advanced Services. The Lotus Domino Advanced Services license provides a set of progressive features that address availability concerns of large Lotus Domino enterprises and is intended for use by demanding, mission-critical Domino applications. There are three fundamental components provided by Advanced Services:

Domino Server Clusters	Domino Server Clusters allows up to six Domino servers to be grouped into a cluster. Domino Clusters provides high availability, workload balancing, and scalability.
Domino Partitioned Servers	Domino Partitioned Servers allows up to six Domino servers to run on a single hardware platform. Each server is isolated from the other through unique Name and Address Books (NABs).
Domino Billing	Domino Billing collects server usage statistics and tracks server activities. This information is used in monitoring, billing and system management.

Although Advanced Services provides all of the above functionality, this chapter focuses on the planning and installation of the Domino Server Clustering feature. In addition, we assume the reader has Lotus Notes administration experience.

5.1 Introduction to Clustering

In simple terms, computer *clustering* is the linking of a group of computers together to appear as one computer to the client. Similar resources on servers that are clustered are viewed by clients as a single, highly available resource. In other words, when a client requests the use of a resource, for example a database, the client should not be concerned where the database resource resides or which server is handling their request. In addition, if a server is unavailable or too busy to handle the request, the request is transparently passed to another member in the cluster, and that server handles the client's request for the resource. From the client's point of view, they are simply gaining access to the resource requested in a timely reliable manner.

Today computer clustering is regarded as an important emerging technology, but in reality it is not all that new. The clustering of computers has been implemented for years, on different types of computer platforms, from mainframes to powerful RISC workstations. It has been implemented in research organizations and commercial enterprises. If clustering is not new, why is it getting so much attention today? Recent cluster announcements are focusing on clustering of Intel-based processor servers. These servers typically get the most attention primarily because of sheer numbers and widespread use. Intel-based servers are becoming increasingly more powerful and can contain quantities of data that would have been associated only with large mainframe computers. More importantly many companies are using these servers for their mission-critical business applications, they are often "betting their businesses" on them. In addition, they are being implemented as Web servers, a relatively new area for many companies. If their Web server is unavailable, they lose

money. Given these circumstances, the following clustering characteristics are extremely important:

- High availability - If one member of the cluster fails, other member(s) of the cluster seamlessly assume the failing member's workload. The client's work should not be interrupted by a server failure.
- Scalability - It should be easy to add (or remove) a member from the cluster as needed. Adding servers takes on new importance when addressing unanticipated workloads (Web servers or new applications). Easy removal of servers also provides administrative flexibility for scheduled or unscheduled server maintenance.
- Workload balancing - Workload is distributed to other members of the cluster. This balancing should increase overall system performance.

Today's Intel-based clustering solutions have implemented some of these functions, but many have limited function (for example, support for only two processors). But cluster implementations are still evolving and under development.

5.2 Lotus Domino Server Clustering

The Lotus Domino Server Advanced Services license provides an "application" clustering solution for the Lotus Notes environment. Many of the non-Domino cluster solutions are based on specific operating systems (such as MVS, NT, or UNIX). In other words, all the cluster members must be running the same operating system. Domino clusters are different. They are supported on any operating system supported by Lotus Domino Advanced Services. A Domino cluster could consist of an OS/2-based Domino server, an NT-based Domino server, an RS/6000 AIX-based Domino server, and a NetWare-based Domino server (or any combination). Depending on the operating system mix of your Domino environment, this could be a significant advantage. For example, clustering could be used in operating system or hardware migration efforts.

In addition to operating system independence, Domino clusters provide other advantages:

- Failover protection for business-critical databases and Domino servers. When a clustered Domino server fails, the Notes clients accessing the server can continue to have access to the database. When the user attempts to access the database on the malfunctioning server the Domino cluster servers "failover" users by redirecting database requests to other servers in the cluster. This function addresses a cluster's high availability characteristic.
- Workload balancing is incorporated into Domino clusters. When a Notes client requests services from a heavily used Domino server, the client's request is passed automatically to other cluster servers. This ensures that workload is evenly distributed across the cluster. Workload balancing helps you achieve optimum performance for your Domino servers.
- Cluster replication is unique to Domino clusters. It is important that replicas of critical Domino databases are maintained on two or more servers. Cluster replication ensures that all changes, whether to a database or to the cluster membership itself, are immediately passed to other databases or servers in the cluster. The databases are tightly synchronized to provide

high availability of information. The basic difference in the two types of replication are:

- Traditional Domino replication is scheduled or performed manually.
- Cluster replication is event driven.

Cluster replication is discussed in more detail later.

- **Scalability.** Domino clusters give you an easy way to grow. They let you expand your processing resources effortlessly as the number of users or the workload grows. You can spread your user accounts across groups of servers. You can create multiple database replicas to maximize their data availability or move users to other Domino servers or clusters as you plan future growth.

The need to easily and quickly add a server to a cluster is becoming increasingly important. You might be thinking, "If you use proper capacity planning techniques, you can predict, plan and schedule an upgrade. Why is this quick upgrade capability in Domino clusters so important?"

Quite simply, in today's emerging Web application server environment, the rules and disciplines of capacity planning may not always apply. We read an astute observation by a networking analyst regarding the debut of a new Web site: "The Web is like a ship. Everybody runs over to one side, and it tilts in that direction". Many a new Web site has almost capsized because of Web surfers. We have read horror stories about new Web sites or Web applications that have been brought to their knees because of unanticipated workload. How do we accurately predict the workload demand of an unidentified audience? The ability to easily and quickly add more resources (such as a server) to your Web site is imperative in these situations. Domino clusters provide you with the ability to easily add one or more Domino servers to the cluster or replicate heavily used Domino databases onto additional servers in the cluster.

In summary, Lotus Domino server clusters have the following characteristics:

- Up to six Advanced Services servers can be configured into one cluster.
- Any database on any of the clustered servers can be clustered 2-6 times. Databases can be clustered to reliably achieve a service level commitment or to support large-scale database demands.
- Cluster servers can be in close proximity to one another or they can be geographically dispersed.
- Cluster servers have the option of all running the same operating system or running on different supported Domino Server Advanced Services operating systems to form a cluster. In addition to *operating system independence* this also translates into some hardware platform independence, for example, Intel or RISC systems.

5.3 Comparing Traditional Notes Replication Versus Cluster Replication

A major difference between a traditional group of Lotus Notes servers and a cluster of Domino servers is the method used to replicate databases.

First, what exactly is Lotus Notes replication? Quite simply Lotus Notes gives you the ability to keep multiple replicas of the same database on multiple servers. Replication is an important and powerful Lotus Notes process that

provides synchronization of documents in a database that may reside on different Lotus Notes servers. This process updates databases on the different servers so that changes made to a database eventually propagate to all the other replicas of the database. This ensures that documents in a database residing on more than one server are synchronized with the other servers that contain replica copies.

How does traditional Notes replication happen? It occurs one of two ways:

- It is triggered on a timed basis or a set schedule (as shown in Figure 57).
- It can be started manually from the server console or another server.

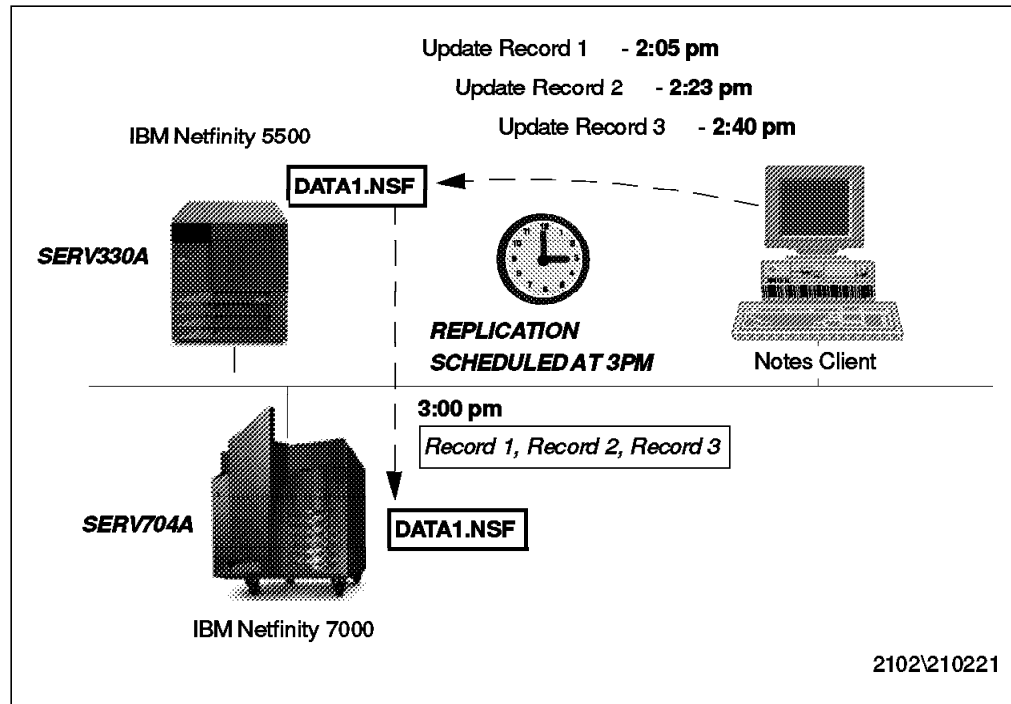


Figure 57. Replication Schedule

In the diagram, the Notes client is updating records in the Notes database DATA1.NSF (on a Domino server named SERV330A). There is a scheduled replication every hour, in this case, 3 p.m. When the user saves, updates or adds a record to the DATA1.NSF database, the DATA1.NSF on SERV330A is updated, but its replication copy on SERV704A will not be updated until the next scheduled replication at 3 p.m. During this time period the databases on the two servers are not synchronized.

How does Lotus Domino cluster replication differ? Cluster replication is automatic and event driven. This means that as database changes occur, Notes immediately replicates the changes to other replicas of the database located in the cluster. These changes are replicated immediately and continuously to all the other servers in the cluster that are maintaining a replica copy of the database. Using this replication methodology, no matter which replica a user tries to access after a server failure the information is identical.

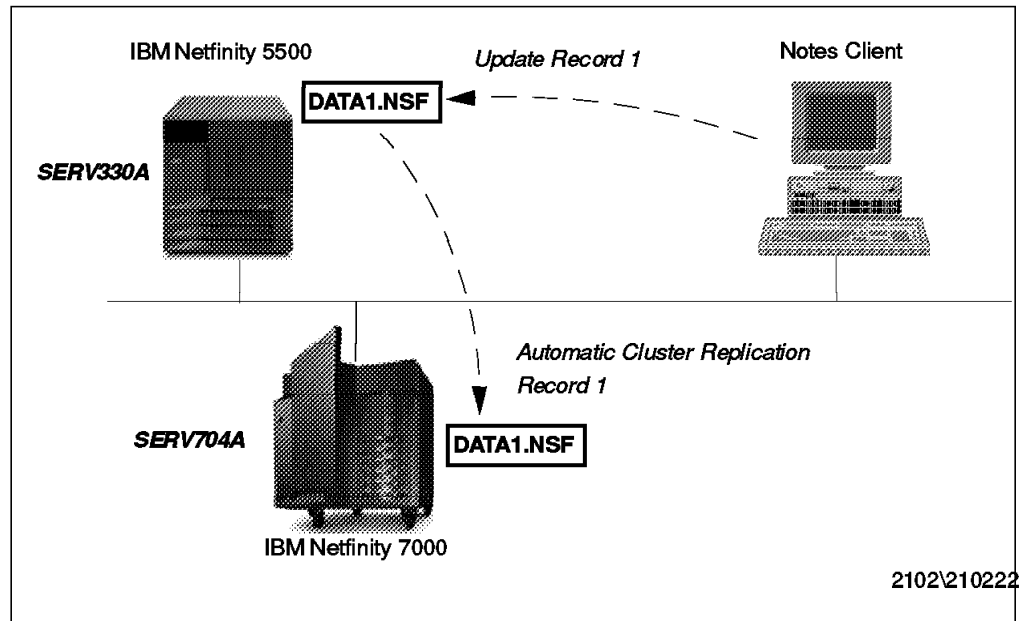


Figure 58. Automatic Cluster Replication

In Figure 58, the Notes client is updating a record in the Notes database DATA1.NSF (on a Domino server named SERV330A). When the user saves, updates or adds a record to the DATA1.NSF database, the cluster replication will take place instantly and automatically replicate this change to any server in the cluster that has a replica copy of the database. In our example, there is a replica of DATA 1.NSF on SERV704A. Cluster replication is automatic and immediate and does not wait for a timed schedule.

You can use both cluster replication and standard Notes replication tasks within a cluster and run them both at the same time. The better strategy is to primarily run the cluster replication task within the cluster and occasionally use standard Notes replication between cluster members to update databases that the cluster replication task may not have been configured to replicate. Regularly scheduled standard Notes replication is useful within a cluster to ensure that databases not replicated by the Cluster Replicator are properly updated.

Note

There are occasions when the cluster replication cannot occur. Domino stores modification events in memory only, therefore, both the source and destination servers must be available in order for the Cluster Replicator to push changes to the destination. This is not a problem if the destination server is not available when the modification event occurs. The event continues to reside in memory of the source server so that the Cluster Replicator can retry the replication at periodic intervals.

It can be a problem if the source server is shut down before the replication to the replicas completes; the modification event may be lost. When the source server restarts, the Cluster Replicator has no information about changes made to databases on the source server before the shutdown and cannot complete replication to other replicas. This is the reason you will want to perform a scheduled replication using the REPLICIA task with all members of the cluster if a server unexpectedly shuts down during a heavy workload period.

5.4 Failover and Availability

A cluster's ability to redirect database 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, the user connects to a replica of the database on another server in the cluster. When users open a database, although they have been switched to a different server, the switch is transparent.

When does failover happen? In general, a failover occurs when a user cannot access a database in a cluster. This happens either because they cannot access the database server or they cannot access the database itself.




Figure 59. Failover Process

In Figure 59, the clients want to open and work with the DATA1.NSF database on a Domino server called SERV330A. However, let's assume that SERV330A has

been taken down for an OS upgrade and is unavailable. After the open request is issued, the clients are redirected to an available server that has a replica copy of DATA1.NSF, in this case, SERV704A. Transparently, the end user has the same access and database information from a different server in the cluster.

In a non-clustered Domino environment our client would have received a Server not responding message and they could not work on the database until the server became available again.

What happens if the client is already working in a database and its server fails? This is a good news/bad news situation. Basically, the failure scenario is the same as a non-clustered Domino environment. The client receives the Server not responding message and all the unsaved work is lost. The good news, though, is the user needs to simply exit the failing database, then try to open it again. When the client re-opens the failed database, failover occurs and our client can start working again on the database, although this time the client is working on another server in the cluster. Remember, failover occurs upon opening a database.

Besides losing access to the server or database there are other reasons for a failover to occur. When the server is too busy, the failover also occurs automatically for the user. This type of failover provides Domino clusters with its workload balancing functionality. There are two algorithms you can choose when setting up workload balancing: setting a performance index threshold or setting a maximum number of users allowed on the server at one time.

What does the client see when this type of failover occurs? Nothing. The failover is transparent; the client is not aware of the failover (except the icons on the client's Notes workspace might be stacked). Do we need to install special code on my Notes client workstations to take advantage of the cluster failover feature? No, the client workstations only need Notes R4 (or higher) client code to handle failover. There is no special license required on the clients to use clusters. At this time, however, the client failover features are not available to Web browser clients.

How do Domino administrators know a failover occurred? Besides seeing the event in the console log, the administrator can look at the Notes log to see if a failover or workload balance event occurred.

5.5 Planning and Administration

As we have been discussing, you can use the cluster capabilities of failover and workload balancing together to achieve greater database availability. Although both capabilities are desirable, there is additional overhead involved when incorporating them; therefore, proper planning of the cluster is important. Let's discuss some different strategies you can use to optimize your Domino database and server access.

Do we need to replicate all databases on every server in the cluster? The answer is no. Determining how many replicas to create and where to place them requires some planning. Before you begin to create replicas in a cluster, consider how frequently users access a database and their need for data redundancy. For heavily used databases, you may want to create more than one replica and locate these on your most reliable computers.

In general, the more replicas of a database, the more accessible the data. If you create too many replicas, however, you can add unnecessarily to the overhead of maintaining a Domino system and impact performance. In Figure 60 on page 132, we have five servers belonging to a Domino cluster. We distributed different databases throughout the cluster.

Practically speaking, having more than three replicas of a database may not provide you with significant additional availability. Check the databases and if users can adequately access them from one or two servers only, do not increase the number of replicas in the cluster. The probability of losing both servers at the same time is minimal. There are situations where users require the constant availability of a specific database. In these cases consider placing replicas on every member in the cluster provided you have adequate disk space and resources to do this. Not all databases require multiple replicas in a cluster. For example, a server log does not need to have a replica on another server.

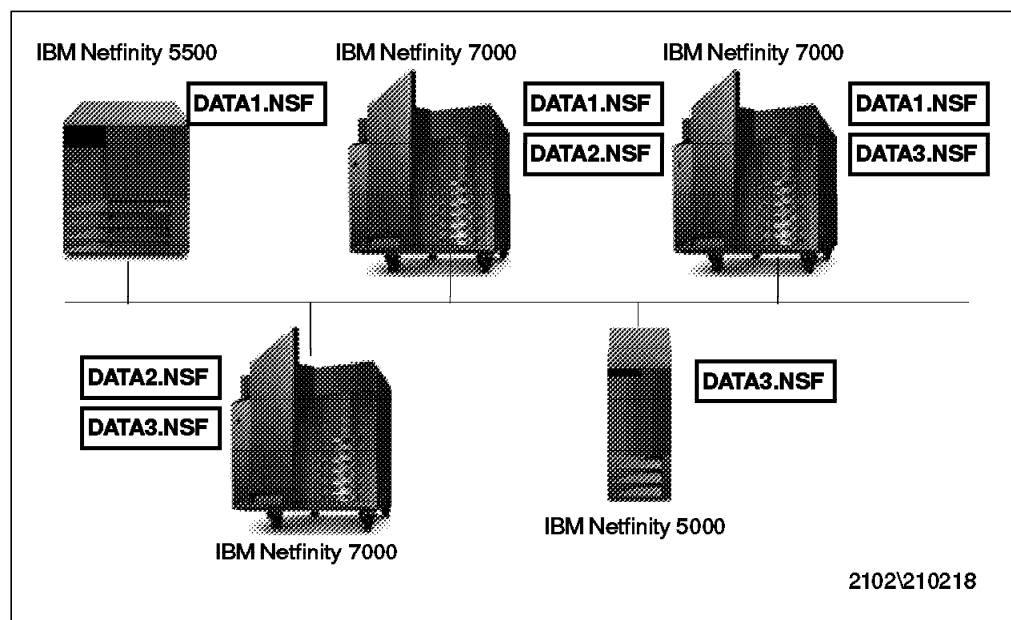


Figure 60. Full Cluster Environment

It may occur to you that having multiple replicas of databases throughout the cluster is an inefficient use of your disk resources. You may read about some other cluster solutions that permit multiple clustered servers to access and share a group of disks (see Figure 61 on page 133).




Figure 61. Shared Disk Environment

Is there a single point of failure in the diagram? Remember, an important objective of clusters is to provide a high availability solution. Computer professionals (for example, the experienced pessimists) are aware that if a single point of failure exists in a configuration, the computer gremlins will find it and cause havoc, and always when you need the system the most. Some might argue that Domino clusters contain multiple copies of selected databases; therefore, it is wasting disk space. Although additional space is required, this high availability solution certainly makes it harder for the gremlins to pick a single point of failure.

In addition, no specialized high availability hardware is required. Many specialized shared disk hardware solutions also have geographic limitations - for example, the cable lengths can only be a particular distance from the servers. Domino clusters can be located on different floors of a building, different locations, or different countries. Is there any potential performance bottleneck in the diagram? These configurations exist today in other clustered environments and provide adequate performance with the aid of specialized hardware and connections.

Experienced administrators realize that providing adequate disk throughput on a single server can be challenging at times. Imagine the administrative task of analyzing and balancing the disk workload with multiple clustered servers trying to access the same group of disks and data. As the number of servers participating in a cluster increases this administrative concern becomes more and more complex. Domino clusters do not require specialized disk hardware and can use traditional techniques for analyzing disk performance.

5.6 Network Traffic Considerations

Cluster replication is constantly updating replicated databases over your network. Although these transactions are primarily small records, potential congestion on the network may be a performance concern for a Domino cluster administrator. If the network traffic caused by cluster replication is a concern, you can isolate the cluster traffic by optionally setting up a private LAN

dedicated to intra-cluster communication between servers (as shown in Figure 62 on page 134).

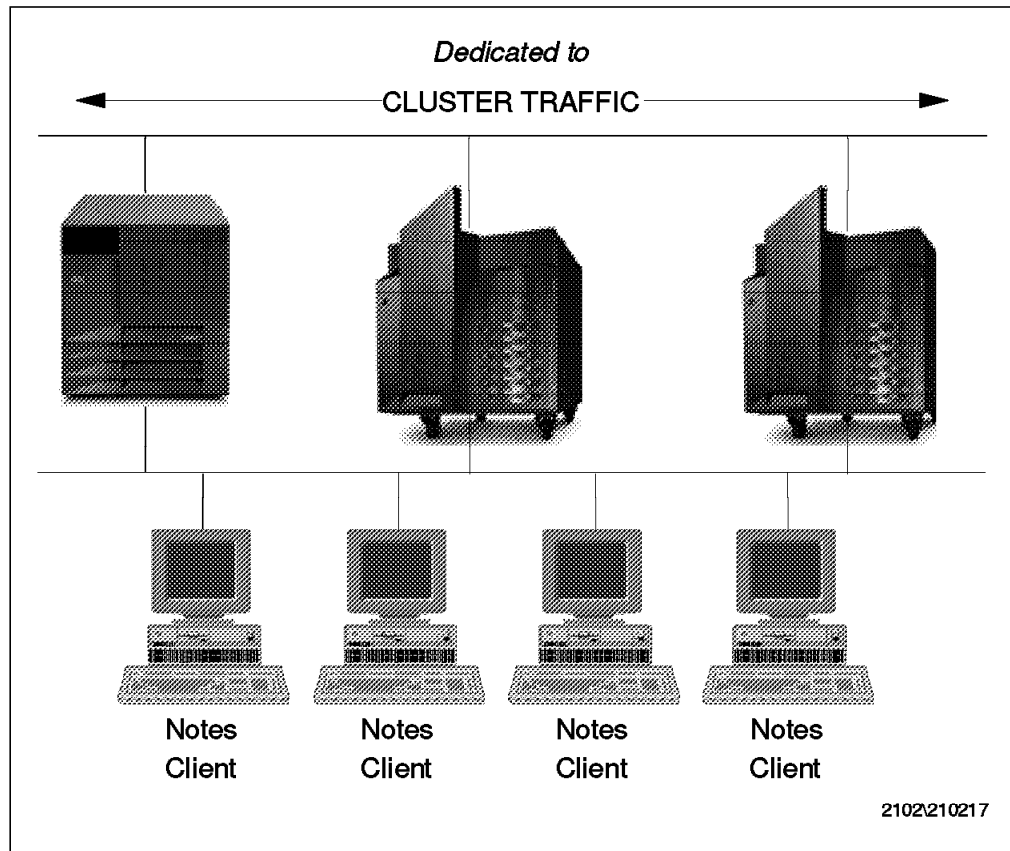


Figure 62. Domino Cluster Traffic

The cost? An additional network adapter in each server, a hub and wiring. There is a parameter in the NOTES.INI, `Server_Cluster_Default_Port`, which specifies the port used for intracenter network traffic.

5.7 Cluster Setup Strategies

After reviewing your Notes environment, you should plan a cluster strategy that tries to create a balance between your user's requirements for data availability and the physical ability of each computer in your cluster to manage additional workload. There are different strategies for optimizing database and server access.

The first strategy to be considered is the *high failover* strategy. This strategy attempts to keep important databases highly available to the clients. One obvious type of database considered in this strategy is the user mail database. You can make mail databases highly available by creating a replica of the databases on another server in the cluster. If the client accesses their mail database and the server is unavailable, the user fails over to the replica on another server.

The MailClusterFailover setting in the NOTES.INI enables Mail Router request failover. If users have replicas of mail files located on multiple servers, you can set this variable in the NOTES.INI file of all Notes Release 4.x servers in the domain to enable users to receive mail from servers within and outside the cluster when their home servers are down.

If the workload on your cluster is normally well-distributed, consider configuring the cluster for failover only, and not workload balancing:

The second strategy is *active workload balancing*. As we mentioned earlier, Domino clusters have an optional workload balancing feature that lets you distribute the workload of heavily used databases across multiple servers in the cluster. Again we should consider mail databases in balancing workload across servers in the cluster. It is recommended that you distribute the user mail files equally across the cluster so that no single server has a significantly larger number of mail files than any other server.

There are two NOTES.INI parameters that assist in workload balancing:

1. **Server_Availability_Threshold** - Specifies the acceptable level of system resources available to a server. By setting this value for each server in a cluster, you determine how the workload is distributed among cluster members. Domino compares this value against a server's availability index; when the availability index falls below the **Server_Availability_Threshold** value, the server becomes BUSY.
2. **Server_MaxUsers** - Specifies the maximum number of user sessions allowed on a server. When the server reaches this limit, the server goes into a MAXUSERS state. The Cluster Manager then attempts to redirect new user requests to other servers in the cluster. (You can use the **Server_MaxUsers** setting for any Domino server. This setting is not restricted to clusters.)

The third strategy is a combination of high failover and active workload balancing. When you have a database that users access continuously, such as a special discussion database, not only do you need high availability, but you also need to distribute within the cluster the heavy workload created by users accessing the database.

5.8 Domino Cluster Planning Guidelines

Listed below are some guidelines for planning a setup of a Domino cluster:

- All cluster members must run the Lotus Domino Release 4.6 Advanced Services license. Domino clusters are available for R4 servers only. Although the Advanced Services code ships with Lotus Domino Server 4.6; the Advanced Services license is separately purchasable for each server in the cluster.
- Only Notes R4 workstations can take advantage of the cluster failover feature. R4 workstations do not need a special license to use clusters. At this time failover features are not available to Web browser clients.
- All cluster members must use the same set of network protocols.
- In Release 4.6, Domino Advanced Services supports only TCP/IP environments for cluster and partitioned server configurations.

- Notes clients can access clustered Domino servers from any protocol the server supports. However, the advanced automatic failover and workload balancing features are only available for TCP/IP clients connected to Domino servers.
- Servers with clustered databases must replicate across TCP/IP networks.
- All cluster members must be in the same Domino domain and share a common Public Address Book.
- The Address Book for the Domino domain must have a specified administration server. If you do not specify one, the Domino Administration Process cannot change cluster membership. The server that is the administration server of the Public Address Book does not need to be a member of a cluster or be running the Advanced Services license. However, it must be running Lotus Domino Release 4.6.
- All cluster members must be connected using a high-speed local area network (LAN).
- Each server is a member of only one cluster at a time.
- Each server must have adequate disk storage space to function as a cluster member. Because clusters usually require more database replicas, Domino servers networked in clusters require more disk storage space than non-clustered servers.
- Each server must have adequate processor and memory capacity. In general, clustered servers require more computer power than non-clustered servers.
- Each server should be set up for optimum failover capability and workload sharing.

Note: Domino clustering was first introduced in Lotus Domino Release 4.5. Domino clusters are available for R4.5 and R4.6 servers only.

5.9 Domino Cluster Installation and Configuration Steps

In this section, we set up a Domino cluster on two servers running different operating systems, one running Microsoft's NT 4.0 server and the other running OS/2 Warp Server. We perform the following tasks:

1. Install and configure the NT Domino server, which we identify as RACK325_NT.
2. Register the additional OS/2 Domino server, RACK325_OS2.
3. Add both servers to a new Domino cluster, called DOMINOCUSTER.
4. Install and configure RACK325_OS2 (the OS/2 Domino server).
5. Replicate the CLDBDIR.NSF file on the two servers.
6. Replicate any Domino databases that you want to participate in Domino cluster replication.

5.9.1 Installation of the NT Domino Server

The first task is installing Lotus 4.6 server code on RACK325_NT. The Domino Server installation program, INSTALL.EXE, is located in the \W32INTEL\INSTALL subdirectory of the Lotus Domino CD. Refer to the Lotus Notes *Install Guide for Servers* for a complete description of the installation steps. In Figure 63, we point out the differences that should be noted for Advanced Services and clusters.

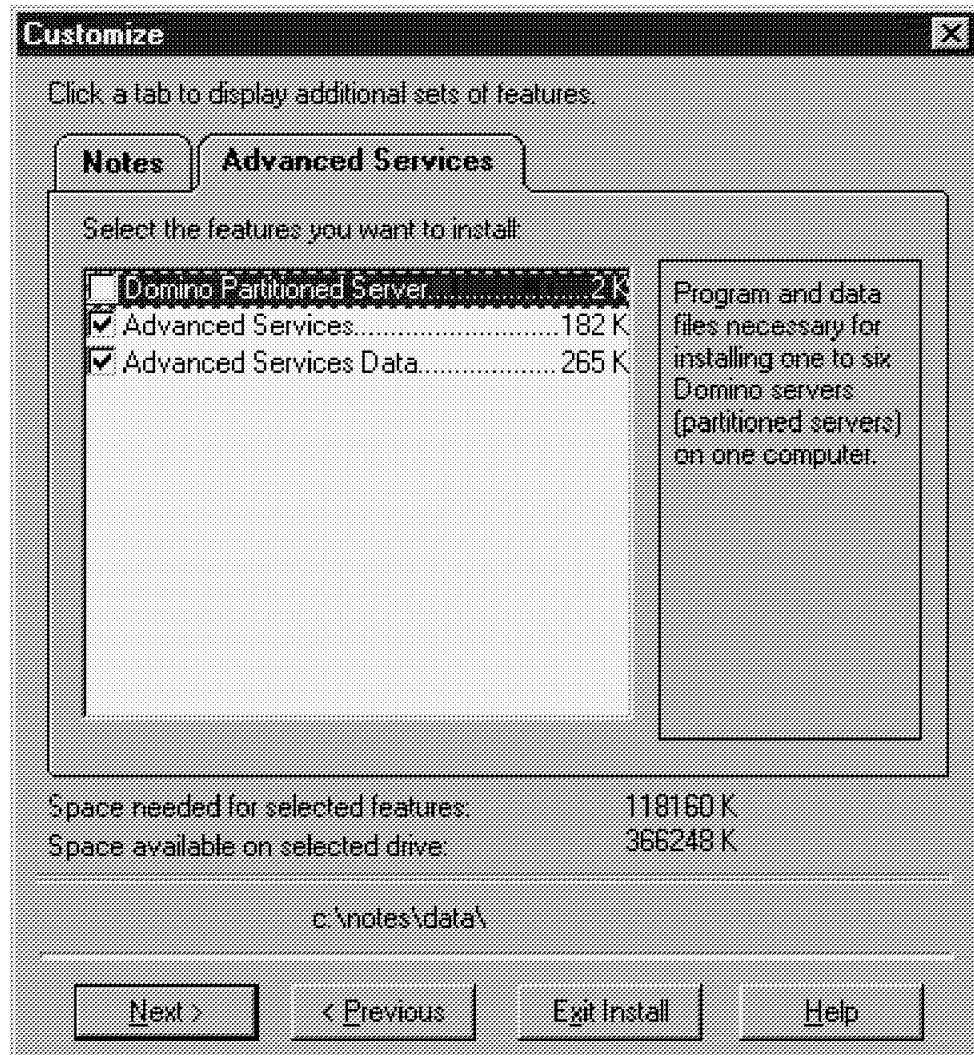


Figure 63. Advanced Services Customizing Page

Please complete the following steps to install and configure Domino clustering on two servers:

1. You will proceed through the start of the Domino Server installation as documented in the *Install Guide for Servers* until you get to the Advanced Options screen. You can select the location of your Notes program folder and Notes data folder plus three installation options. You need to select **Custom features - Manual install** to install Domino Advanced Services. Then select **Next>**.
2. The following screen is the Customize window. Notice that there are two tabs on the top of the page: Notes and Advanced Services. The default

choices on the Notes features page are sufficient for Domino clustering. If you require any additional features listed, select them. After you have completed the Notes feature page select the **Advanced Services** tab.

As shown in Figure 63 on page 137, there are three options available on the Advanced Services page:

- Domino Partitioned Server - Program and data files necessary for installing one to six Domino logical servers (partitioned servers) on one computer. Each partitioned server is an independent Domino server with a unique Name and Address book and security naming tree.
 - Advanced Services - The program files necessary for Advanced Services: clusters, billing and partitioned servers. When installing partitioned servers, use this option only on the first partitioned server.
 - Advanced Services Data - The data files necessary for Advanced Services.
3. Select **Advanced Services** and **Advanced Services Data** for clustering.
 4. Select the **Next** button.
 5. After you have selected in which folder to place Domino, accept the default subdirectory for Domino. The Domino software is installed on the server.

This completes the installation stage for the NT Domino Server. Next, we configure the server.

5.9.2 Configuring the NT Domino Server

After you have successfully installed Lotus Domino 4.6 server code on RACK325_NT, the Domino server needs to be configured. The installation program placed the Lotus Domino Server and Notes Workstation icons in the program group specified during the installation phase (Lotus Notes). RACK325_NT is the first Domino Server in our Notes domain. (For those familiar with the server configuration procedure, there is nothing special required for clusters.)

1. Start the Lotus Notes workstation program.
2. Select **The first Lotus Notes server in your organization** and click **OK**. The First Server Setup window is displayed in Figure 64:

The screenshot shows the 'First Server Setup' dialog box with the following fields and values:

- Server name (e.g. Acme Server1): Rack325_NT
- Organization name (e.g. Acme Corp): IBM
- Administrator's last name: Doe
- First name: John
- MI: Q
- Administration password (case sensitive): password
- Network type: TCP/IP
- Serial port: (None)
- Modem type: Auto Configure (for unlisted modems only)
- Server is also administrator's personal workstation:

Figure 64. First Server Setup

3. Type Rack325 NT in the Server name field.
4. Type IBM in the Organization name field. This is your organization or company's name.
5. Type Doe in the Administrator's last name field. This is your administrator's last name.
6. Type John in the First name field and Q in the MI field. This is your administrator's first name and middle initial.
7. Type password in the Administration password field. This is the administrator's password.
8. Select **TCP/IP** from the pull-down list in the Network type field. This is the network protocol used for Lotus Notes and Domino network communications.

Leave Serial port and Modem type at the default (optional - configure if you plan to use an installed modem on your server).

Leave Server is also administrator's personal workstation unchecked (optional - check if you plan to use the server as a workstation)

9. Select **Advanced Options...** Select all the default values plus:
 - a. Select **Log all replication events** (optional).
 - b. Select **Log all client sessions** (optional).
10. Click **OK** on the Advanced Server Setup Options window.
11. Click **OK** on the First Server Setup window.
12. Select the appropriate time zone and check Observe Daylight Savings Time (if appropriate). Select **OK**.
13. Click **OK** when Notes setup is complete.
14. Exit the Notes Workstation program.

This completes the initial part of the NT Domino server configuration. There are several other customization steps that must be completed before we can start setting up the OS/2 Domino server. The first thing we want to do is start the NT Domino server.

5.9.3 The First NT Domino Server Startup

Starting the NT Domino server and viewing the server's console listing provides several clues to the required customization tasks.

1. Click on the **Domino Server** icon in the Lotus Notes folder to start the Lotus Domino Server program. You will see the following screen:

```
Lotus Domino Server, Release 4.6, September 19, 1997
Copyright c 1985-1997, Lotus Development Corporation, All Rights Reserved

09/24/97 11:09:42 AM Informational: The registry value HKEY_LOCAL_
MACHINE\System\CurrentControlSet\Control\Session Manager\Memory
Management\LargeSystemCache is currently set to 1. For better server
performance, change the setting to 0
09/24/97 11:09:42 AM Informational: The registry value HKEY_LOCAL_
MACHINE\System\CurrentControlSet\Control\PriorityControl\
Win32PrioritySeparation is currently set to 2. For better server
performance, change the setting to 0
09/24/97 11:06:26 AM Database Replicator started
09/24/97 11:06:30 AM Mail Router started for domain IBM
09/24/97 11:06:30 AM Router: Internet SMTP host rack325_nt in domain
raleigh.ibm.com
09/24/97 11:06:35 AM Index update process started
09/24/97 11:06:41 AM Stats agent shutdown
09/24/97 11:06:45 AM Agent Manager started
09/24/97 11:06:48 AM AMgr: Executive '1' started
02/24/97 11:06:51 AM No Administration Server has been designated
for the Public Address Book.
09/24/97 11:06:51 AM Creating Administration Requests database.
09/24/97 11:06:52 AM Administration Process started
09/24/97 11:06:55 AM Schedule Manager started
09/24/97 11:06:55 AM SchedMgr: Validating Schedule Database
09/24/97 11:06:56 AM SchedMgr: Done validating Schedule Database
09/24/97 11:06:56 AM SchedMgr: Calculating hourly statistics
09/24/97 11:06:56 AM SchedMgr: Done calculating hourly statistics
09/24/97 11:07:00 AM Calendar Connector started
09/24/97 11:07:05 AM Database Server started
09/24/97 11:15:48 AM The Administration Request database -
admin4.nsf - on RACK325_NT/IBM is currently a wildcard replica.
09/24/97 11:15:48 AM The Administration Process on RACK325_NT/IBM
will not work until an Administration Server has been designated
for this domain's Public Address Book.
```

A Review of the RACK325_NT/IBM console listing (note highlighted lines) shows some additional configuration work that must be completed on the server.

2. Click on the **Lotus Notes** icon in the Lotus Notes folder to start the Lotus Notes Client program.

We need to create the Certification Log. Before we can register additional servers or create additional certifiers, you must create the Certification Log (CERTLOG.NSF) file. You need the Certification Log if you plan to use the Administration Process to rename, recertify, or upgrade users and servers.

- a. Choose **File - Database - New**.
- b. Select the name of the server you're using to register users in the Server field.
- c. Type Certification Log in the Title field.
- d. Type CERTLOG.NSF in the File field. (You must name this file CERTLOG.NSF.)
- e. Click **Show Advanced Templates**, select **Certification Log** as the template, and then click **OK**.

- f. Assign Editor access to all administrators who will register users and servers and re-certify IDs.
3. We need to designate an administration server for the Public Address Book. One of the cluster planning guidelines stated that the Notes domain must have a specified administration server. If you do not specify one, the Notes Administration Process cannot change cluster membership. To use the Administration Process, you must assign the Public Address Book an administration server. The following is the procedure to set up RACK325_NT as the Administration Server for the Public Address Book (PAB):
 - a. Click **File - Tools - Server Choose Administration**.
 - b. Click **Database Tools**.
 - c. In the Server box, select the server storing databases (RACK325_NT) for which you want to select an administration server.
 - d. In the Databases box, select the Public Address Book database (Acme's Address Book) for which you want to select an administration server.
 - e. In the Tool box, select **Administration Server**.
 - f. In the Administration server box, select an administration server (RACK325_NT/IBM) for the databases.
 - g. Click **Update**.
 - h. You should see the informational window stating 1 database processed - 0 errors. Click **OK**, then **Done**.
 - i. Exit Notes.
 - j. Quit the server.
4. Delete the ADMIN4.NSF file.
 - a. This file is located in the \NOTES\DATA\ subdirectory. (Do not delete ADMIN4.NTF.)
 - b. Restart the server. (A new ADMIN4.NSF is created during this server startup.) The following lines will display on the console at the next server startup.

```

02/24/97 02:23:45 PM RACK325_NT/IBM is the Administration Server of
the Public Address Book.
02/24/97 02:23:45 PM Creating Administration Requests database.
```

Next, we need to register the OS/2 Domino server, which we name RACK325_OS2.

5.9.4 Registering an Additional Server - the OS/2 Domino Server

Before you can install additional servers you must register them. The server registration process performs two tasks:

- Creates a server document in the Public Address Book
- Creates a server ID

The registration server is RACK325_NT. It must be up and running on the network. Registration creates an entry for the new server in the Certification Log (CERTLOG.NSF), which we created earlier. The following is the procedure required to register RACK325_OS2. (You will need a blank formatted diskette.)

1. Start the Notes workstation program.
2. Choose **File - Tools - Server Administration**.
3. Click the **Servers** icon and choose **Register Server**.
4. Click **Yes** if you purchased a server license.
5. Type the certifier password and click **OK**.
6. Choose applicable security type, North American or International, in the Security field. Click **Continue**.
7. Type RACK325_OS2 as the name for the new server.
8. Enter a password for the server and enter a number for the minimum password length. The password is case-sensitive.

Each time you start the server, Notes prompts you at the additional server's console for this password. The server does not start up until the password is entered. If this is a problem you can set up the server without a password (assuming the server is located in a secure location). To set up the server without a password prompt:

- Type 0 for minimum password length.
 - After clicking the **Other** icon, store the server ID in a file and not in the Public Address Book.
9. The Domain name IBM should already be displayed in the Domain field.
 10. The name of the administrator should be entered in the Administrator field. You can leave the default, Mark Twain/IBM.
 11. Select the **Other** icon in the left-hand column.
 12. In the Store Server ID field, select **In file** to store the server ID on the blank diskette (in addition to storing it in the Public Address Book). Click **Set ID** file. RACK325_OS2.id should be displayed in the file name. Select **Save**.
 13. Click **Register** and insert a blank formatted diskette when requested.

We now have two servers, RACK325_NT and RACK325_OS2, registered and set up in our Public Address Book.

5.9.5 Configuring a Domino Cluster

In our next step, we configure the Domino Cluster and add our servers to the cluster. Afterwards, we install and configure the OS/2 Domino server.

1. Open the Public Address Book of the server NAMES.NSF.
2. Select the **Server - Servers** view.
3. Select servers **RACK325_NT/IBM** and **RACK325_OS2/IBM** and choose the **Add to Cluster** button (see Figure 65 on page 144).

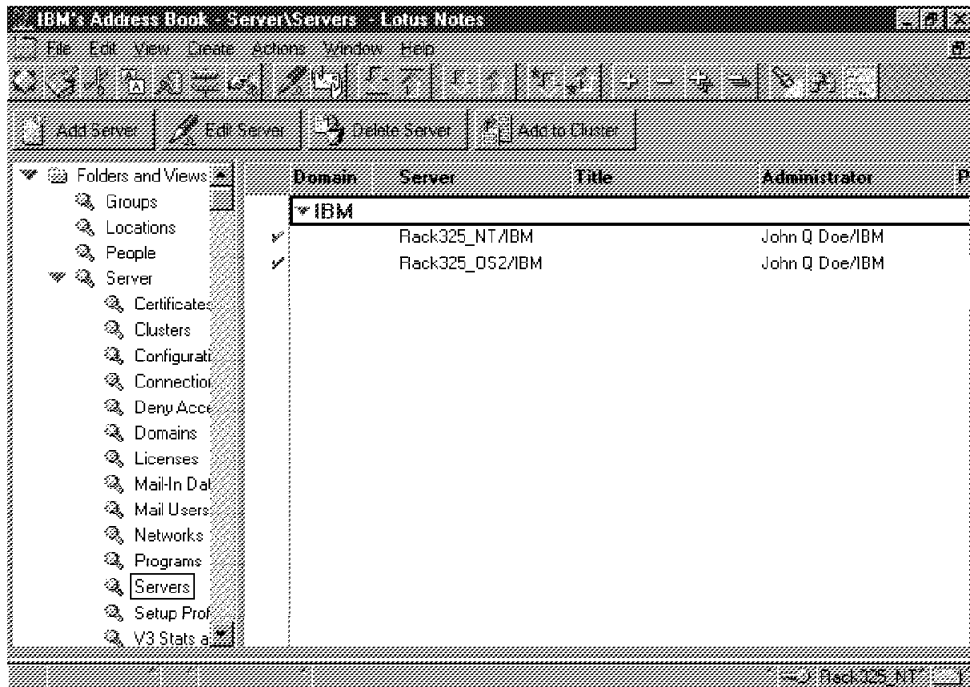


Figure 65. Public Address Book Server-Server Document

4. This is a new installation; therefore, we are creating a new cluster. The system prompts you to choose the cluster you want to add the server(s) to or create a new cluster. Select **Create** and click **OK**. The new cluster name is called DOMINOCLUSTER.

If this were not a new installation we possibly could have had existing clusters defined. If that were the case, we could have added one or more servers by selecting the cluster name from the list.

5. You get a message stating that Your request to add the server(s) to a cluster has been successfully submitted to the Administration Requests database. Click **OK**.

5.9.6 Verifying the Domino Cluster

In the previous step we successfully created a cluster called DOMINOCLUSTER which consists of our two servers. Select the **Server - Clusters** view to verify the servers are members of the cluster. Close the Public Address Book.

5.9.7 Restarting the NT Server

Below is the console log of the NT server after we have created the cluster and added the two servers. At this point we still have not installed the OS/2 Domino server known as RACK325_OS2.

Again, start the NT Domino server and view the server's console listing.

1. Click on the **Domino Server** icon in the Lotus Notes folder to start the Lotus Domino Server program. You will see the following screen:

Lotus Domino r Server, Release 4.6, September 19, 1997
Copyright c 1985-1997, Lotus Development Corporation, All Rights Reserved

```
09/26/97 08:20:15 AM Begin scan of databases to be consistency checked
09/26/97 08:20:15 AM End scan of databases: 1 found 09/26/97 08:20:16
AM Informational: The registry value
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\Session
Manager\Memory Management\LargeSystemCache is currently set to 1. For
better server performance, change the setting to 0 09/26/97 08:20:17 AM
Informational: The registry value
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\PriorityControl\
Win32PrioritySeparation is currently set to 2. For better server
performance, change the setting to 0
09/26/97 08:20:18 AM Cluster Administration Process started <---- (1)
09/26/97 08:20:18 AM Database Replicator started
09/26/97 08:20:23 AM Mail Router started for domain IBM 09/26/97
08:20:27 AM Index update process started 09/26/97 08:20:33 AM Stats
agent shutdown 09/26/97 08:20:36 AM Router: Internet SMTP host
rack325_nt in domain raleigh.ibm.com 09/26/97 08:20:39 AM Agent Manager
started_09/26/97 08:20:44 AM AMgr: Executive '1' started 09/26/97
08:20:45 AM RACK325_NT/IBM is the Administration Server of the Public
Address Book. 09/26/97 08:20:45 AM Administration Process started
09/26/97 08:20:49 AM Schedule Manager started 09/26/97 08:20:50 AM
SchedMgr: Validating Schedule Database 09/26/97 08:20:51 AM SchedMgr:
Done validating Schedule Database 09/26/97 08:20:51 AM SchedMgr:
Calculating hourly statistics 09/26/97 08:20:51 AM SchedMgr: Done
calculating hourly statistics 09/26/97 08:20:53 AM Calendar Connector
started
09/26/97 08:20:58 AM Cluster Database Directory started <----- (2)
09/26/97 08:21:03 AM Cluster Replicator was unable to configure using
Cluster Database Directory clbdir.nsf: File does not exist <----- (3)
09/26/97 08:21:59 AM Cluster Replicator started
09/26/97 08:21:07 AM Database Server started
09/26/97 08:21:50 AM Adding server to cluster DOMINOCUSTER
09/26/97 08:21:59 AM Error processing request: No cluster mates
found <----- (4)
09/26/97 08:21:59 AM Cluster Administration Process shutdown
09/26/97 08:22:05 AM Created database Cluster Database Directory:
clbdir.nsf <----- (5)
09/26/97 08:22:09 AM Finished initialization of Cluster Database Directory
```

2. The Cluster Administration (CLADMIN) process starts. It is responsible for the correct operation of all cluster components. On clustered servers, the process runs automatically at server startup and whenever the cluster membership changes. When the added server detects a change to its Address Book, the Cluster Manager and the Cluster Administration Process (CLADMIN) start.
3. The Cluster Database Directory Manager (CLDBDIR) task keeps the Cluster Database Directory (CLDBDIR.NSF) up to date with the most current database information. It also manages databases with cluster-specific attributes such as databases marked out of service or pending delete.
4. Cluster Replicator was unable to configure using the Cluster Database Directory, CLDBDIR.NSF. This is the first startup of the server as a cluster member. This database does not exist yet, but it will subsequently be created (see note #5).

The Cluster Database Directory (CLDBDIR.NSF) database resides on each server in a cluster and contains information about all the databases and replicas within a cluster.

5. The RACK325_NT server is added to the cluster but at this time there are no other cluster members found. At this point we have registered but not installed the RACK325_OS2 server.

6. The Cluster Directory Database (CLDBDIR.NSF) file is created. It uses the CLDBDIR.NTF template.

5.9.8 Installing the OS/2 Domino Server

Now we are going to install Domino 4.6 on our second server in the cluster, the OS/2 server. The installation program is started in the usual fashion, except you use the INSTPM.EXE command for OS/2.

Recommendation: OS/2 Warp Server 4.0 with TCP/IP (required) installed.

1. Select **Advanced Services** and then **OK**.
2. Read the License Warning window. The Advanced Services license must be purchased for every Domino server using Advanced Services features. If you have proper licensing, select **Yes**.
3. On the Install window update CONFIG.SYS, and click **OK**.
4. In the Install - Directories window, place the \NOTES directory and \NOTES\DATA directory on the desired drive. In the Notes components section choose **Select All** and select the **Install** button. The installation program will then load the Lotus Domino for OS/2 Advanced Services code on your server.
5. At the conclusion of the loading process, you need to shut down and restart the OS/2 server.

Before you shut down, though, we recommend that you test your TCP/IP configuration to make sure that TCP/IP is configured correctly on the two servers. From the OS/2 command prompt, try pinging the RACK325_NT server:

```
C:> PING RACK325_NT
```

If correctly configured, RACK325_NT should provide a response. If no response, open the **TCP/IP Configuration (LAN)** folder. Select the **Hostnames** tab. On page 1, RACK325_OS2 should be in the This machine's hostname field. Select page 2, and make sure that there is an entry for RACK325_OS2 and RACK325_NT with their associated TCP/IP addresses. Select **Save** and exit. Again try pinging RACK325_NT. If RACK325_NT responds, ping RACK325_OS2 from the RACK325_NT server. (If it still fails, ping using only the TCP/IP addresses. For example, ping 9.67.240.122 to ensure TCP/IP and wiring are correct.) When TCP/IP is correctly set up, shut down and restart OS/2.

6. After the system has restarted, open the Lotus Applications window and start the Lotus Notes Workstation program.

The Domino server code is now on the OS/2 server. This process also placed the Notes Server and the Notes Workstation icons in a Lotus Applications program group on your desktop. The first time you double-click the Notes workstation icon, the Notes setup program automatically starts. You must complete the server setup program before starting the server and working with Notes.

If you try to start the Notes server without first completing the Notes server setup program, the server exits with the message: You must first type "Notes" to set up your system. To restart the set up process, double-click the Notes workstation icon.

7. Refer to the online Server Installation document for complete instructions on the Notes setup program.

After completion of the Notes setup program, exit the workstation program and start the Domino server.

8. After a successful OS/2 Domino Server startup, you will need to replicate the Administration database (ADMIN4.NSF) from RACK325_NT on to RACK325_OS2. After replicating this file, exit the workstation program and stop the OS/2 Domino server. (Type QUIT at the Domino server console.) After the Domino server shutdown completes, restart the OS/2 Domino server.
9. Next start the Notes workstation program. Replicate CLDBDIR.NSF with the RACK325_NT server.
 - The CLDBDIR.NSF is an important Domino cluster file. The cluster database records and stores information about the databases on a Domino cluster. The cluster database directory (CLDBDIR.NSF) is updated regularly by the server's cluster database directory (CLDBDIR) task. It is used by the servers within a cluster when failover occurs. Each time you add a database to a server, the Cluster Database Directory Manager creates a document that profiles the database and adds it to the cluster database directory. This database document contains information such as the database name, server, path name, replica ID, and other replication and access attributes (see Figure 66).
 - The information in the CLDBDIR.NSF file is used by the cluster to determine failover paths and to control access to a database.

We suggest performing the replication from the NT workstation program.

- a. Open CLDBDIR.
- b. Select the **Replicate - With** options.
- c. Select **To and From Servers**.
- d. After replication the CLDBDIR.NSF should show entries in databases from both servers. (You might want to compare this CLDBDIR.NSF with the diagram of this database displayed in the earlier discussion of the cluster database directory.)

File Name	Server	Replica ID	Cluster Rep
* ACluster.nsf	RACK325_NT	85256482:002D4B0E	Enabled
* ACluster.nsf	Rack325_OS2	85256482:002D4B0E	Enabled
* admin4.nsf	RACK325_NT	85256481:035EE921	Enabled
* admin4.nsf	Rack325_OS2	85256481:035EE921	Enabled
* admin4.ntf	RACK325_NT	852561D4:006BAEBA	Enabled
* ADMIN4.NTF	Rack325_OS2	852561D4:006BAEBA	Enabled
* alog4.ntf	RACK325_NT	852561E2:003C5FC8	Enabled
* ALOG4.NTF	Rack325_OS2	852561E2:003C5FC8	Enabled
* approve4.ntf	RACK325_NT	8525623B:004C4E91	Enabled
* APPROVE4.NTF	Rack325_OS2	8525623B:004C4E91	Enabled
* billing.ntf	RACK325_NT	852561B7:0075AB18	Enabled
* BILLING.NTF	Rack325_OS2	852561B7:0075AB18	Enabled
* busines.nsf	RACK325_NT	85256481:00FE90CB	Enabled

Figure 66. The Cluster Database Records

- e. Files must be replicated between the clustered systems before any cluster replication for that file occurs. The next step, therefore, is to replicate any of the files that you want automatically replicated by the cluster replication service.
 - 1) From the RACK325_OS2 server's menu (in the N&A Book), select **File - Replication - New Replica**.
 - 2) Next, select a server, (in our case, RACK325_NT) and click **Open**.
 - 3) Highlight the database you want to replicate, then click **Select**.
 - 4) On the New Replica screen, choose the **Create: Immediately** button.
 - 5) Choose **OK**.
- f. To verify whether cluster replication is working, make some changes to the replicated database. You should see the changes replicated to the other cluster member almost immediately. Also review the console on both systems and at the database on the other system.

5.9.9 Changes to a Cluster Server's NOTES.INI File

After the server is added to a cluster there are new entries added to the server's NOTES.INI file. The next time the server starts, two new server tasks are automatically started - the Cluster Database Directory Manager and the Cluster Replicator tasks. The ServerTask statement in the NOTES.INI file of the server(s) is appended with these two new task parameters:

- CLDBDIR - All Domino servers within a cluster run a Cluster Database Directory Manager (CLDBDIR) task that creates and maintains the cluster database directory. When you add or remove a database on a server, the CLDBDIR task immediately changes the information in the cluster database directory to show the addition or deletion of that database.

All members of a cluster share a common cluster database directory. When one server updates its list of databases, the Cluster Replicator replicates the changes to the other servers in the cluster. In this way each cluster member has an up-to-date directory of databases in the cluster.

- CLREPL - The Cluster Replicator (CLREPL) task is responsible for the tight synchronization of data among replicas in a cluster. Each server in a cluster runs one Cluster Replicator by default although administrators may want to increase the number of CLREPL tasks on a server.

Whenever a change occurs to a local database, the Cluster Replicator pushes the change to the other replicas in the cluster. This behavior ensures that each time you access a database, you are looking at the most up-to-date version. The Cluster Replicator also replicates changes in private folders that are stored in a database.

The Cluster Replicator task only pushes replications to other servers in the cluster. The standard Notes replicator task (REPLICA) is still responsible for replicating changes to and from servers outside of the cluster.

In addition to these parameters that are automatically added there is the MailClusterFailover parameter that is important if you want MailRouter failover. This parameter enables or disables Mail Router request failover. If users have replicas of mail files located on multiple servers, you can set this variable in the NOTES.INI file of all Notes Release 4.x servers in the domain to enable users to receive mail from servers within and outside the cluster when their home servers are down. The format of the parameter is (MailClusterFailover=value).

There is no default value for this parameter. Use the following values to set this variable:

- 0 - Disables Mail Router request failover
- 1 - Enables Mail Router request failover

For example, MailclusterFailover=1, will enable a MailRouter request failover.

5.10 Recovering from Failures

What happens in the cluster when one of the members is unavailable? What happens to database changes in the cluster when one of the members is unavailable? When the Cluster Replicator task in a member of a cluster cannot replicate database changes, it retains the information in its memory and tries to replicate the changes until it is successful. To conserve resources that might be wasted on unsuccessful replication attempts and to increase the probability of achieving successful replication on the next retry, the Cluster Replicator increases the interval between retry attempts from one minute up to one hour for each failed replication.

When the destination (failed) server eventually restarts and replication succeeds, all modifications performed on the source database that have not yet been replicated are pushed to the destination database. That is, the Cluster Replicator brings the destination database into synchronization with the source database regardless of how many changes occurred on the source while the destination database was unavailable.

5.11 How to Set Up Failover and Workload Balancing to Work Together

As the cluster administrator, you should closely monitor your system for usage patterns during peak periods and adjust cluster resources accordingly. Provide failover for databases by creating replicas on two or more servers. You can balance server workloads by adjusting the availability thresholds and maximum user limits of each server. The Cluster Manager then monitors each server in the cluster, triggers failover, and redirects requests accordingly. It also monitors the number of active users and indicates to other cluster servers when a server is in the MAXUSERS state and is not accepting new user requests.

To see how often requests are being redirected, check the Notes Log for failover events. If redirection of the user request is unsuccessful, the user receives a message, and is not allowed access to the server.

5.12 Summary

In summary, Lotus Domino clusters, introduced in Lotus Domino Server 4.5, provide an application clustering solution for the Lotus Notes environment. This implementation combines the important basic clustering functions of high availability, scalability, and workload balancing.

In general, Domino clusters currently allow you to group up to six servers to form a cluster. Because it is an application clustering solution it provides you with a degree of operating system independence. In other words, if it is supported by Advanced Services it can participate in the cluster. Besides the

flexibility of which operating system is supported, you can also mix and match operating systems within a cluster.

Domino clusters also provide cluster scalability. It is easy to add or remove servers to the cluster. As the workload increases simply add another server to the cluster.

Another key clustering function is workload balancing. This necessary feature lets the system distribute the workload as the demand fluctuates or in the event a server fails. As an administrator you don't want to have to continually monitor which server is getting overloaded. Workload balancing lets the clustering technology try to distribute the workload dynamically, based on performance thresholds set in each server.

Domino clustering is also easy to implement. There are no special hardware requirements or rewriting of your Domino applications. If you currently have Lotus Domino 4.6, basically all you have to do is install the Advanced Services code, configure the cluster, and replicate your important databases throughout the cluster.

As we mentioned earlier, there are many new Intel-based clustering solutions being touted in the marketplace today. If you review the clustering functionality that really exists in each of these solutions, you will probably notice that many of these solutions provide only a subset of true clustering. For example, some solutions merely provide for high availability. Although high availability is certainly a key clustering function, quite simply high availability by itself is not a clustering solution. It is a high availability solution and does not address scalability or workload balancing.

After comparing the current marketplace, you will see that Domino clusters provide a comprehensive and advanced clustering solution for the Lotus Notes environment. In many ways, you should conclude at this time that the Lotus Domino clustering solution is leading the Intel-based server clustering pack.

Appendix A. Top 10 Ways to Improve Your Server Performance

This information has been provided by Iris.

By analyzing a variety of NotesBench reports, published over the last two years by NotesBench Consortium members, we came up with a list of the top 10 ways you can improve the performance of your server. The list shows you how to improve your server capacity and response time.

1. Make sure your server memory matches the number of users you want to support. Most NotesBench vendors use 300K-400 KB per active user. They also set their NSF_BUFFER_POOL_SIZE to the maximum for their memory configuration. This setting isn't necessary, because the Domino server initially obtains a quarter of available memory and grows only if necessary (depending on the load). You should use published physical memory configurations as a ceiling for memory configuration decisions.

2. Distribute I/O among separate devices.

For example, you can put the OS kernel on one drive, the page file on another, the Domino executable on a third, and finally the Domino data files on a fourth drive. In some cases, NotesBench vendors point their log.nsf file to a location different from the default data directory (using the log= setting in the server's NOTES.INI file).

3. Make I/O subsystem improvements.

For example you can:

- Move from EISA-based systems (such as, controllers) to PCI-based systems.
- Exchange EISA/PCI boards in favor of PCI-only boards (this way, lower speed EISA devices won't decrease the I/O throughput).
- Use striping to improve performance.
- Use multiple I/O controllers to distribute logical volumes (and use file pointers to databases across separate controllers). Make sure you have the latest BIOS for your I/O subsystem. This is an inexpensive way to remove a likely throughput bottleneck.

4. Use faster disk drives.

You can improve disk drive speeds from 5,400 rpm to 7,200 rpm. For most Windows NT systems, NotesBench vendors use 2 GB disk drives. For Solaris and IBM Netfinity systems, the drives were larger: 4 GB. For AS/400, the drives were even larger: 8 GB.

5. Increase the strip size.

NotesBench vendors use a strip size of 8 KB (Digital's systems) or 16 KB (IBM Netfinity reports). (The IBM Netfinity report provides additional information on I/O settings such as IOQ Depth, Outbound Posting, PCI Line Prefetch, and Address Bit Permitting.)

6. Use faster CPUs.

NotesBench vendors have moved beyond the Pentium, Sparc, and PowerPC processors, which were in the 100-200Mhz range, to higher-speed processors. However, they consistently use P6-based systems over the Pentium II systems for high-end Domino server loads. The size of your Level

2 cache should match your expected user loads and the response time you want. Vendors have moved from 256 KB to 512 KB, 1 MB to 2 MB Level 2 cache systems, especially on their greater than two-CPU configurations.

7. Improve your network.

NotesBench vendors have:

- Moved from 10 Mbps cards and networks to 100 Mbps configurations
- Used multiple LAN segments (one for each partition) to isolate network traffic, at the high-end user loads

8. Change your network protocol to IP.

Vendors were initially (two years ago) using NetBIOS and SPX internally, but have unanimously moved to IP for their performance publishing efforts.

9. Upgrade to a newer release of Domino.

NotesBench vendors have moved from Domino release 4.5a SMP version to Domino Release 4.52B SMP version for higher capacity results. The first Domino Release 4.6a result (AS/400) on a RAID5 configuration indicates a reliable configuration can still provide competitive response time with a properly designed I/O architecture.

10. Use Domino partitioned servers.

NotesBench vendors have increased scaling of active user loads and leveraged their more powerful configurations (faster clock cycles, fiber-connected I/O subsystems, OS kernel to CPU binding, and multiple I/O controllers) by using partitioned servers.

To understand how Iris came up with this top 10 list, they take you through the performance analysis of item number 2 in the list -- to distribute I/O among separate devices. Initially, many vendors placed the kernel, page, and Domino executables on one volume and the Domino data files on another. However, both volumes were on the same controller. Lately, the NotesBench reports show improvements in performance when the volumes are separated across multiple controllers, and individual volumes are separated across disks. What this means is that we found that vendors put the OS kernel on one drive, page file on another, Domino executable on a third, and finally the Domino data files on a fourth drive. In some cases, they pointed their log.nsf file to a location different from the default data directory (using the log= setting in the server's NOTES.INI file). Vendors who distributed the I/O over several disk drives had better server performance overall, and could support a higher capacity of users.

For example, in a NotesBench report published in May of 1996, Digital Equipment Corporation set up a server with the following specifications:

CPU: 133Mhz CPUs
Memory: 512 MB
Domino: Release 4.1

They placed the operating system and the Domino executable on drive C:\, the page file on drive D:\, and the Notes\data directory on drive E:\. They could support a maximum capacity of 1,500 users with this configuration.

In a NotesBench report published in September of 1997, IBM Corporation set up a server with the following specifications:

CPU: Three 200MHz Intel Pentium Pro processors
Memory: 1 GB
Domino: Release 4.51

They placed the operating system on drive C:\, the page file on drive C:\, the Notes\data directory on drive E:\, and the Domino executable on drive E:\. They supported a Mail-only workload of 3,500 active mail users. In a four-processor configuration, they supported a MailDB workload of 2,900 active users.

These examples led us to the conclusion that distributing I/O over several disk drives had better server performance overall, and could support a higher capacity of users. We went through many other NotesBench reports to collect the data shown in our top 10 list. You can visit the NotesBench Web site yourself to view published data and test results. Visiting the site may help you to come up with other ways to improve your server's performance.

Appendix B. What's New in Domino and Notes Release 5

In this section, we discuss the key features and enhancements in each of the R5 products, Domino, Lotus Notes and Domino Designer.

B.1 Lotus Domino R5

The Domino R5 family of servers delivers messaging, applications and on-line collaboration fast and reliably for organizations from small businesses to the largest enterprises. Domino R5 helps organizations reduce costs by making the server easier to administer and the desktop much easier to manage. Domino R5 provides businesses with the flexibility and openness they need to harness the power of the Web, along with the security they need to keep systems running smoothly and to prevent unauthorized access.

For an in-depth technical look at Lotus Domino R5, be sure to check out Domino R5 Preview at:

<http://www.notes.net/today.nsf/>

The following is the latest features in Domino R5:

- Enterprise-grade messaging server for companies of all sizes

- Internet mail

Domino messaging is Internet messaging; there is no distinction between them anymore.

For example: ensure easy to use, full fidelity Internet mail with native Internet addressing, native MIME content support and native SMTP routing.

- Messaging management

Domino R5 includes new messaging management features that make for easier, more efficient deployment, monitoring and management.

For example: detailed message tracking and monitoring, improved Message Store Management and routing control management, native SMTP routing, easy-to-set message quotas and built-in upgrade tools for MSMail, cc:Mail and Exchange mean you get the best messaging system, one that you can manage proactively at the lowest cost of ownership.

- Directories

R5 Domino Directory addresses the requirements of even the largest enterprise.

For example: complete support for LDAP V3, a highly compressed Directory Catalog that scales to support multiple millions of entries and replicates to your desktop so information is available even when you're on the road, makes it easy to quickly locate any resource within your enterprise or beyond.

- Secure infrastructure for messaging and applications

R5 provides the most flexible and pervasive security model for e-mail and applications while promoting information sharing.

For example: fully integrated X.509 certificates, stronger cryptography, support for Common Data Security Architecture (CDSA), S/MIME, SSL V3 for IIOp and LDAP.

- Open, secure Web application server

R5 has enhancements for serving high value, secure Web applications that integrate your business.

For example: enhanced HTTP engine performance, option to use Microsoft IIS as the HTTP engine for Domino, access control for HTML files, Internet cluster management; deliver applications that can change based on the browser capabilities, provide live access to relational database fields; virtual server logging, support for simple easy-to-search URLs, CORBA/IIOp protocol support.

- Reliable, available and scaleable server infrastructure

- Database and Architecture Improvements

R5 improvements to the Domino core architecture ensure 24x7 service levels as well as unsurpassed scaleability and performance.

For example: increased database storage (the 4GB limit on database size has been increased to 32 GB), memory and I/O optimization, transactional logging, quicker restart after failure, and the Domino R5 directory scales to support at least one million registered users.

- 24x7 availability & reliability

R5 provides 24x7 service levels.

For example: Internet Cluster Manager extends failover and load balancing to Web browser clients, integration with OS clusters, clustering enhancements, unlimited partitions per server for greater scaleability and availability, on-line indexing, auto-restart and recovery enhancements, on-line in-place compaction, and full and incremental on-line backup support.

- Simple, flexible administration

- Install/setup improvements

Domino R5 server can be up and running within minutes.

For example: flexible, wizard-driven install programs, user upgrade at registration, server setup without a client, automated client setup, enhanced server setup profiles provides 'cookie cutter' approach to server setup.

- Domino Administrator

The Domino Administrator makes administration easy with a task-oriented interface.

For example: drag-and-drop functionality and tabs for easy access to frequently performed tasks will free up administrators to spend more time on proactive management activities.

- Desktop management

Administrators have centralized control over user desktops.

For example: centralized client configuration and centralized upgrades reduce support, deployment and training costs.

- Innovations in collaboration, knowledge management

Domino R5 search allows users to efficiently search any information type to find what they need.

For example: domain search across all databases and any file systems, universal filters to index documents and attachments, results returned in customizable pages and summaries.

- Localization Features

Domino provides is the premier platform for multi-national organizations.

For example: inbound/outbound message support for more languages and character sets, easier administrator configuring for multilingual messaging needs, Multilingual Directory support.

B.2 Lotus Notes R5

The latest release of Notes is an easy-to-use, Internet messaging client that works with any server. Customizable and lightweight, with R5 you work the way you want, from anywhere, with all the power of Notes. Marries the Internet with the unmatched power of Notes. Delivers to you new and unprecedented search and navigation capabilities that help you quickly and easily find what you need to deliver on your commitments.

For an in-depth technical analysis of Lotus Notes R5, be sure to check out Notes R5 Preview at:

<http://www.notes.net/today.nsf/>

The new features of Lotus Notes R5 are mentioned here:

- Works with any server

With the addition of standard protocol support, Notes R5 becomes a powerful tool for users no matter what their infrastructure.

For example: You can use Notes to get to your ISP-hosted mail! With support for POP3, IMAP4, LDAP v3, MIME and S/MIME, X.509 certificates and native Internet addressing, Notes R5 is a true Internet e-mail client.

- Easy to use Internet client

R5 builds on the best features of Notes -- e-mail, calendaring, group scheduling, contact management, task management -- and makes them better.

For example: You can create your own best start to every day with a user-customizable Headlines page. Bookmark important information for easy access, and find new information quickly with browser-like navigation.

- Mobile support

The industry-leading features that make Notes users just as productive, whether they working connected or disconnected, have been made even better.

For example: Now you can access Notes through your ISP connection, take information with you in an optimized local directory catalog so you can address memos and find intranet resources, change your computer settings with point-and-click ease so you can work from wherever you happen to be.

- Innovations in collaboration, knowledge management

Notes R5 is an intuitive, easy-to-use environment for collecting, synthesizing and transforming data into knowledge.

For example: start each day with the information you need to know brought to your attention. Headlines that you can customize alert you to the meetings, deadlines, urgent messages, news bulletins that you need to know about. Find information faster and more easily with powerful search capabilities, paging and summarization of search results, full text index of file system, and file filters.

B.3 Domino Designer R5

Domino Designer R5 enables application developers to be more productive and to build more powerful solutions for the Domino Web Application Server faster. Domino Designer R5 offers a more open environment than its predecessor, Lotus Notes Designer for Domino R4.6. With Designer R5, developers use the languages and methods they're most comfortable, create applications that work with multiple clients and Web browsers and can be delivered to a worldwide audience.

For an in-depth analysis of Domino Designer R5, be sure to check out Domino Designer R5 Preview at:

<http://www.notes.net/today.nsf/>

The new features of Domino Designer R5 are mentioned here:

- Intuitive, open environment for developers

- Task-Oriented IDE

Designer R5 has a task-oriented visual environment for designing and developing Notes and Web applications.

For example: intuitive toolbars and tool palettes, a tabbed environment for multiple work sessions, improved programming environment supporting Web standard languages.

- Site Design and Navigation

New tools make it easy and faster to design Web sites and navigation, reducing Web site maintenance costs.

For example: Outline Designer provides a visual tool for designing the structure and navigation of a Web site; with Frameset Designer developers easily create multi-pane interfaces to their applications.

- Page/Form Layout and Design

Visual authoring tools making it faster and easier to lay out and design pages and forms.

For example: Page Designer provides a WYSIWYG HTML authoring, applets add features previously only available to Notes clients, new table options, enhanced image support, greater control over positioning of elements, improved color support, templates.

- Easy access to enterprise data and applications

New and enhanced tools make it easy to access enterprise data and applications.

For example: Domino Enterprise Connection Services (DECS) provides live access to relational databases including DB2, Oracle, Sybase files, EDA/SQL ODBC; use Java agents to access relational data.

- Industry Standards Support

Designer R5 supports the industry standards that businesses need in order to develop eBusiness applications. New choices of languages are available for writing and scripting code.

For example: Designer supports HTML 4, Java, JavaScript, and CORBA/IIOP; users can mix and match WYSIWIG and HTML code, bring in existing HTML code by hand, or import HTML pages; new application objects have been added to the Domino Object Model.

- Choice of development tools

With Designer R5, it's faster than ever before to develop applications that converge enterprise data and applications with automated business processes. Developers can use leading development tools and standard technologies for building Domino applications.

For example: HTML Web pages and links created in HTML editors like NetObjects Fusion can be saved to the Domino object store; integrates with 3rd party Java tools such as VisualAge for Java, Symantec Visual Cafe or Borland JBuilder).

- Multi-client application support

Developers can preview their work in the clients and browsers that end-users demand, develop applications once that run on any client type. Give the average browser the rich, full experience of Notes.

For example: developers can preview their applications in the most popular Web browsers; applets make it easy to provide Notes functionality to browsers.

- Multilingual application support

With Domino Global WorkBench now completely integrated with Designer, developers can create multilingual Web sites right out of the box.

For example: generate and maintain different language versions of Web page, support multiple currency, support long dates, and provide more flexible database sorting options.

Appendix C. Domino and the Internet

In this section, we discuss the various ways that Domino works with the Internet to provide you with a complete Internet solution. Information regarding Domino and the Internet is provided in much greater detail in *Working with Lotus Notes and the Internet*. This documentation shipped with releases of Domino 4.5.

C.1 The Internet and the World Wide Web

The Internet is comprised of a wide area network (WAN) collection of networks across the world. Those networks include regional networks, local networks at universities and research institutions, local networks at business locations, and many military networks. The term Internet refers to this entire set of networks.

The most popular application on the Internet is the World Wide Web. The Web is an extensive hypermedia system that connects millions of pages of information stored on Internet servers located all over the world. To access and display Web pages, you use Web browser software. A Web browser will allow you to access information that resides on various Internet servers that run different protocols.

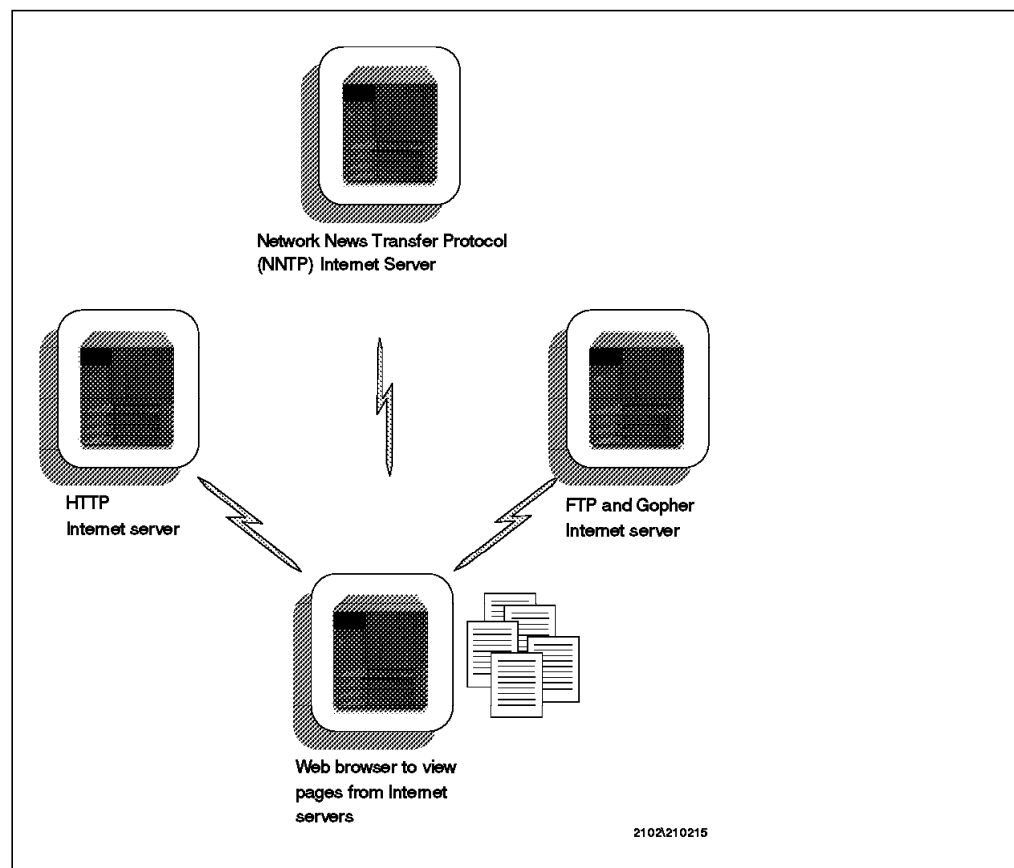


Figure 67. A Web Browser Accessing Various Internet Servers

Web pages are identified by a unique address, known as a Uniform Resource Locator (URL), such as:

<http://www.pc.ibm.com/us/netfinity/>

The first part of the URL (http) specifies the network protocol used to access the page. The second part of the URL (www.pc.ibm.com) is the Internet server name. The third part of the URL (us/netfinity/) is the actual file name of the page on that server.

C.2 Types of Internet Connections

Today, you must make a decision to determine the type of connection that best suits your organizational needs. Questions you need to ask yourself are how many people will be connecting, how often they will be connecting, the speed with which they will be communicating, and so on. Next, you need to contract with an Internet Service Provider (ISP) to obtain your actual Internet connection. An ISP is a company that offers different types of network connections to the Internet. These connections range from simple dial-up connections to fully dedicated leased-line connections. With Domino, you can use any type of connection to hook into the Internet through an ISP. To locate an ISP in your area that suits your needs, visit the following Web pages:

- <http://www.isps.com/>
- <http://thelist.internet.com/>
- <http://www.thedirectory.org/>
- http://www.herbison.com/herbison/iap_meta_list.html

The basic types of Internet connections you can obtain through an ISP are:

C.2.1 Dial-Up Connection

A dial-up connection is considered a direct connection to the Internet. If you have a dial-up connection, you connect your computer (for example, your Domino server) to your corporate LAN and your computer communicates through a dial-up modem that connects to an ISP.




Figure 68. Dial-Up Connection

C.2.2 Leased-Line Connection

A leased-line connection is considered a direct connection to the Internet. If you have a leased-line connection, you connect your computer to your corporate LAN and communicate through a router to a leased phone line that connects to an ISP.




Figure 69. Leased-Line Connection

C.2.3 Proxy Connection

A proxy is a server that provides indirect access to the Internet. A proxy server usually runs in conjunction with firewall software to pass incoming and outgoing requests between servers on both sides of a firewall. Typically, firewalls are placed at the entry point to a public network such as the Internet. If your organization uses a proxy server for its Internet connection, you connect your computer to your corporate LAN that communicates through the proxy and firewall servers and on to an ISP.




Figure 70. Proxy Server Connection

For more information on ways to implement Internet security for your Domino server, you should obtain the IBM redbook *The Domino Defense: Security in Lotus Notes and Internet*, SG24-4848.

C.3 Extending Your Notes Network Over the Internet

This section describes ways to put your Domino servers on the Internet and how to access other Domino servers over the Internet.

You can extend your corporate network to allow your Domino servers to communicate over the Internet. Your first step is to connect your Notes environment to the Internet. Then, once your Internet connection is established, you can:

- Put a Domino server on the Internet.
- Connect from one server to another server over the Internet.
- Use NotesNIC to locate and communicate with servers on the Internet.

C.3.1 Putting a Domino Server on the Internet

To put your Domino server on the Internet:

1. Install the Domino Release 4.6 server software. Refer to the *Notes Release 4.6 Install Guides for Servers*.
2. Obtain your Internet connection from an ISP.
3. Install one of these types of network connections:
 - TCP/IP protocol software for leased-line connections
 - SLIP/PPP (dial-up versions of TCP/IP) for modem connections
4. Set up a proxy server to connect your internal LAN to the Internet (optional, but highly recommended).
5. Test your Internet connection using a ping utility.

C.3.2 Connecting Servers to Servers over the Internet

You can connect a server to a server over the Internet so that you can perform normal server tasks, such as replication and mail routing.




Figure 71. Domino Server-to-Server Connections Over the Internet

To establish the connection:

1. Find out the IP address of the destination Domino server with which you want to connect.
2. Create a Server Connection document in your server's Public Address Book.
3. Edit the Server document in your server's Public Address Book if you are connecting through a proxy server (optional, but recommended with proxy server).
4. Establish the connection to the Internet outside of Domino.
5. Connect to the Domino server.

C.3.3 Using NotesNIC to Communicate with Other Domino Servers

NotesNIC is a Notes directory service on the Internet that provides a way for Notes sites to locate and communicate with each other. To accomplish this, all NotesNIC servers belong to the "NET" domain. As a result, all NotesNIC servers share a replica copy of the "NET" domain's Public Address book, which provides the information needed to locate and communicate with Notes Internet sites. Any Notes site can register to become a part of the NotesNIC domain (NET). Once registered, your Notes site can:

- Use scheduled replication to share and update information among sites around the world.
- Use Domino Mail to send and receive mail to other Notes sites over the Internet with no need for SMTP gateways.
- Use anonymous access/unauthenticated server access features to access Domino servers on the Internet without proper prearranged cross-certified authentication.
- Administer your global connections from the familiar environment of a standard Notes Public Address Book.

The following figure illustrates how the NotesNIC "NET" domain on the Internet works along with internal corporate domains. In the figure, both ABC and XYZ corporations physically host and maintain all their own Domino servers and Notes workstations, including the Domino server that is in the "NET" domain. Using a passthru connection through their "NET" server, they can communicate not only with other servers in the "NET" domain, but with workstations and servers in other corporate domains.

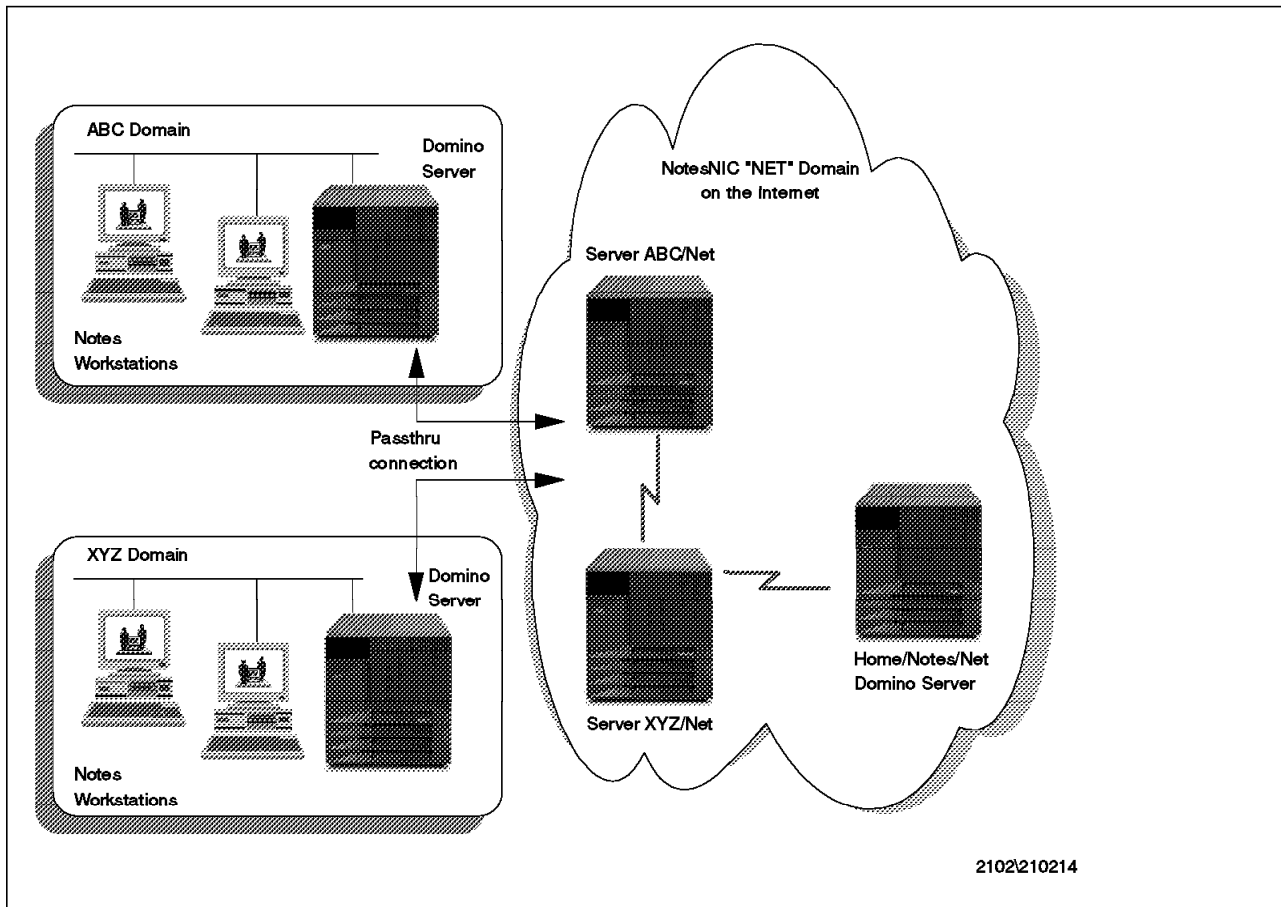


Figure 72. The NotesNIC Servers Network

For more information on NotesNIC, visit the home site for the NotesNIC:

Notes site: welcome.nsf at home/notes/net (205.159.212.10)

Web site: <http://www.notes.net/notesnic.nsf>

Appendix D. The Tivoli Enterprise Module for Domino/Notes

As part of Tivoli's applications management initiative, Tivoli is working to create application management modules, called Tivoli Enterprise modules, for large client/server applications. When installed on top of a Tivoli Management Environment, these Tivoli Enterprise modules will allow for automatic full life-cycle management of the client/server application, by making use of the functionality available from the Tivoli Enterprise core applications. The Tivoli Enterprise Module for Domino/Notes is one such module.

In this section, we introduce the functions and features of the Tivoli Enterprise Module for Domino/Notes. Details of these functions and features are provided in the redbook *Managing Domino/Notes with Tivoli Manager for Domino*, SG24-2104. This redbook assists with the installation and administration of the Tivoli Enterprise management module for Domino and Notes and the optional Admin Extension for Tivoli Enterprise with clients using different operating system platforms. This publication discusses the experiences of IBM's International Technical Support Organization working with Tivoli Enterprise, the Tivoli Enterprise for Domino/Notes, Lotus Domino Server Release 4.5 and Lotus Notes Workstation Release 4.5 software.

Note: Because the Domino/Notes environment and the Tivoli Enterprise environment are prerequisites for the Tivoli Enterprise Module for Domino/Notes, a working knowledge of these environments is important in thoroughly understanding some of the management and customization shown throughout the book.

D.1 Overview of the Tivoli Enterprise Module for Domino/Notes

The Tivoli Enterprise Module for Domino/Notes allows for the management of Domino/Notes environments by providing the following functionality:

- Security enhancements
- Task administration
- Availability management
- Automated Notes client installation
- User management provided with the Admin Extension for Tivoli Enterprise Module for Domino/Notes

Since the Tivoli Enterprise Module for Domino/Notes utilizes the Tivoli Enterprise Framework and other core applications, the functionality the module provides is fully customizable and extendable.

D.1.1 Security Enhancements

Each Notes domain is configured within a Tivoli Management Region (TMR) as its own policy region. In addition, two resource roles are created for each configured Notes domain. Giving a Tivoli Administrator one or both of these resource roles in a particular policy region authorizes the administrator to perform certain Notes tasks.

D.1.2 Task Administration

The Tivoli Enterprise Module for Domino/Notes provides tasks and jobs for configuring, operating and maintaining Notes domains and users. In the Tivoli Enterprise Module for Domino/Notes environment, the term task refers to a defined action which is executed against a given group of managed nodes. A job refers to a resource consisting of a task and its predefined execution parameters, which includes the group of managed nodes against which the task will be executed. Jobs can be scheduled to run using the Scheduler. A task library is simply a container that a Tivoli Enterprise Module for Domino/Notes administrator can create and store tasks and jobs. Here we use the term *task* in a general sense to refer to Tivoli Enterprise Module for Domino/Notes tasks or jobs.

The tasks in the Tivoli Enterprise Module for Domino/Notes are pre-organized into three task library categories: utilities, tasks, and user tasks:

- Utilities** This task library is made up of tasks used mostly for configuration and are seldom used once the initial configuration enabling the Tivoli Enterprise Module for Domino/Notes to manage a Notes installation is complete.
- Tasks** This task library contains common operational tasks that are used in the day-to-day administration and management of Notes servers. This task library contains both tasks and jobs.
- User Tasks** This task library contains common operational tasks that are used in the day-to-day administration and management of Notes users. This task library contains both tasks and jobs.

Table 12 presents each supported task.

<i>Table 12. Tivoli Enterprise Module for Domino/Notes Module for Domino/Notes Supported Tasks</i>	
Tivoli Enterprise Module for Domino/Notes Module for Domino/Notes Supported Tasks	
Type of Task	Supported Tasks
Utilities Task Library	
Tivoli Enterprise Module for Domino/Notes Configuration Tasks	Configure Event Server
	Configure Tivoli Enterprise Module for Domino/Notes for a Notes Domain
	Configure and Deconfigure the Notes Admin Extension
	Enable and Disable a User Profile for Notes
Notes Domain Policy Region Configuration Tasks	Configure Client Install
	Configure Event Console
	Configure Remote Console
	Configure Tivoli Enterprise Module for Domino/Notes for a DomainAlias (Notes Domain policy region) Server
	Create DomainAlias (Notes domain policy region) Client List
	Create a DomainAlias (Notes domain policy region) Server List
Tasks Task Library	
Notes Server Tasks	Start and Stop the Notes Server
	Start and Stop the Notes Server Monitors
	Start and Stop the Notes Event Adapter
	Execute Remote Console Command for a Notes Server
	Show Notes Server Statistics
	Start the Notes Server's Client
User Tasks Task Library	
Notes User Tasks	Add Users to a Notes Domain
	Delete Users from a Notes Domain
	Get User Information from a Notes Domain
	Modify User Information in a Notes Domain
	Resend a Notes User ID File
	Install Notes on a Notes Domain Client

D.1.3 Availability Management

Managing the availability of a Notes environment with the Tivoli Enterprise Module for Domino/Notes includes the following features:

Tivoli Enterprise Module for Domino/Notes Event Adapter for Alert Monitoring

The Tivoli Enterprise Module for Domino/Notes provides a Notes add-in task in the form of an event adapter that monitors Notes alerts and translates them into Tivoli/Enterprise Console (T/EC) events.

Tivoli Enterprise Module for Domino/Notes Notes Monitor Collections

The Tivoli Enterprise Module for Domino/Notes supplies two customized monitor collections: central and remote monitor collections. These collections check such things as:

- Server availability
- Notes statistics
- Daemons, replicators, routers and add-in tasks that are running

Proactive monitoring and automatic creation of different severity levels for T/EC events is enabled through configurable thresholds and responses.

Event Filtering

The Tivoli Enterprise Module for Domino/Notes uses a configuration file to allow the administrator to filter selected events to prevent them from being forwarded to the event server rulebase for processing.

Tivoli Event Server Rulebase

The Tivoli Event Server Rulebase provides intelligent event processing and correlation. The Tivoli Enterprise Module for Domino/Notes provides two rule sets to simplify and enhance Notes management:

1. The T/EC adapter rule set captures and processes Notes events from the Tivoli Enterprise Module for Domino/Notes Event Adapter and forwards them to the Tivoli Enterprise Module for Domino/Notes Console.
2. The Tivoli Enterprise Module for Domino/Notes Distributed Monitoring rule set receives and processes information from the Notes Central and Remote Monitoring collections and forwards them to the Tivoli Enterprise Module for Domino/Notes Console.

These rule sets enable such features as automatic closing of harmless events, event correlation, and multi-region support.

D.1.4 Automated Notes Client Installation

The Tivoli Enterprise Module for Domino/Notes also provides automated Notes client installations. Features included in the support include a task for creating Notes client installation file packages, administrator defined file pack organization, and transparent heterogeneous platform support. Heterogeneous Notes client file packages allow the administrator to organize the file package into profile managers any way that they wish, for example, organizationally, as the determination of the platform is performed as part of the pre-installation processing.

D.1.5 User Management with the Admin Extension

If you have the Tivoli Enterprise Module for Domino/Notes Admin Extension installed, you can manage Notes users by distributing user profiles via profile managers to subscribers (Notes users). Included in the management of users through the Admin Extension is adding/modifying/deleting user information from the Notes Name and Address Books.

D.2 The Tivoli Enterprise Module for Domino/Notes Module for Domino/Notes Managed Environment

The Tivoli Enterprise Module for Domino/Notes management product suite allows you to manage an environment from network devices, to the systems connected across these networks, to the database that the applications running on these systems depend on, as well as the applications themselves. The managed and supported environment include the following:

- Managed Server Platforms
AIX, HP-UX, Solaris, and Windows NT Server
- Managed Client Platforms
AIX, HP-UX, Solaris, Windows 3.X, Windows NT, Windows 95 and OS/2 Warp

By correlating all the logistical elements in a management paradigm, Tivoli Enterprise Module for Domino/Notes Module for Domino/Notes provides the industry's most extensive depth and breadth of management for Domino 4.5 and 4.6, Powered by Notes.

Appendix E. Lotus Notes NetBIOS Port Drivers

There are two Notes-supplied NetBIOS port drivers available for the OS/2 Domino server. These are called NETBIOS.dll and IBMEENB.dll. These port drivers use resources from the underlying NetBIOS protocol stacks provided by the OS/2 LAN Requester and the OS/2 Multi-Protocol Transport Services (MPTS) respectively.

We recommend that you use the NETBIOS.dll, if possible, because of its ease of configuration and its efficient use of resources.

The NETBIOS.dll driver is limited to a system total of approximately 100-125 clients. Those familiar with OS/2 LAN Server might ask themselves "Why is Notes limited to only 125 NetBIOS clients? OS/2 Warp Server can support up to 1000 NetBIOS clients."⁴ The reason for this limitation is in the design of the port driver, not a limitation of the NetBIOS protocol itself. The port driver uses 16-bit OS/2 system semaphores, two per client, and there is a system limit of 256 16-bit system semaphores in OS/2.

Now that we have explained the capacity limitation, let's discuss the two NetBIOS implementations available for OS/2-based Domino servers. In the OS/2 environment, there are two NetBIOS APIs, the submit interface and the NB30 interface. The submit interface is supplied by the OS/2 LAN Requester and the Notes driver NETBIOS.dll. gets its network resources through this interface. That is the reason for the requirement that the OS/2 LAN Requester must be active in order to use this driver.

NetBIOS 3.0 (NB30) is supplied by Multiprotocol Transport Services (MPTS) and it is this interface that is used by the IBMEENB.dll driver.

Since it uses the resources available to the OS/2 LAN Requester, the NETBIOS.dll driver is dependent on the resources allocated to the requester that are allocated in the IBMLAN.INI file. In other words, the network resources (sessions, commands, and names) that both the Notes port driver and the OS/2 LAN Requester use come out of the same pool. In effect, the Notes port driver and OS/2 LAN Requester must share these resources. Consequently, when you configure the IBM LAN Requester, you should make sure that you configure sufficient resources for the OS/2 LAN Requester and the Notes port driver. You will probably need to increase the number of sessions, commands and names from their default values in order to provide sufficient quantities for both components.

The startup sequence of OS/2 LAN Requester and Notes is also important when using the NETBIOS.dll driver. Since the Notes driver is dependent on LAN Requester, you must start LAN Requester before starting Notes.

As stated earlier, the IBMEENB.dll driver makes calls to the standard NetBIOS 3.0 (NB30) interface instead of using the NetBIOS Submit interface. This means you do *not* need to install or run the LAN Requester. However, this driver is more difficult to configure and does not allocate resources as efficiently as the NETBIOS.dll driver. This is explained below:

⁴ Using four LAN adapter cards.

1. Notes has several processes that perform communications. They are the database server, database client, replicator, background replicator on the client, and the periodic macro agent. Each of these uses the port driver to handle the lower level tasks associated with the configured protocol.
2. When the process is being initialized, it makes a call through the port driver to NetBIOS asking for NetBIOS resources to be allocated to the process. The number of resources can be explicitly specified or it can be left to default.
3. In the case of NetBIOS using the IBMEENB.dll driver, Notes uses the default value of 15 sessions per process for all cases except for the database server. This means that 15 sessions are allocated whether they are needed or not. This results in an inefficient use of NetBIOS sessions, which are a limited commodity anyway.
4. The database server process on the Notes server is the one exception; it uses the value specified on the port line in NOTES.INI. For example, if the port line is LANO=IBMEENB,0,50,2000,,12288, then the database server process gets 50 sessions. This value can be changed if required.

As an example, let's say that you want to perform the following sequence on your Notes server:

1. Start the replication process
2. Start the Notes client interface
3. Access a database on another server
4. Start a background replication
5. Start a timed macro to access another server

In order to perform the preceding sequence of events, you will need $50+15+15+15+15+15 = 125$ sessions! These sessions must be available or the process will fail at initialization time. (For the IBMEENB.dll driver using the NB30 interface, the maximum number of sessions is defined in the NETBEUI section in the PROTOCOL.INI file.)

Contrast this resource allocation with the NETBIOS.dll driver. In this case, the OS/2 LAN Requester allocates all the resources specified in IBMLAN.INI when the requester service starts. Notes and any other application that uses the NetBIOS interface will share those resources with the LAN requester. So with our earlier example, Notes would only need $X+1+1+1+1+1 = X+5$ sessions, where 'X' equals the number of currently attached Notes users. As you can see this is much more resource-efficient.

Note: To modify the default resources, you need to modify the NOTES.INI file on servers using IBMEENB.DLL or the PROTOCOL.INI and IBMLAN.INI files on servers using NETBIOS.DLL.

For more information on setting up Domino on OS/2 Warp Server, please review Chapter 4 "Running Domino on NetBIOS" in Lotus Domino's *Configuring the Domino Network*.

Appendix F. Special Notices

This publication is intended to help IS professionals perform server selection and implementation for a Notes network environment. The information in this publication is not intended as the specification of any programming interfaces that are provided by Lotus Notes, Microsoft NT Server, OS/2 Warp Server, and Novell NetWare. See the PUBLICATIONS section of the IBM Programming Announcements for more information about what publications are considered to be product documentation.

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact IBM Corporation, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The information about non-IBM ("vendor") products in this manual has been supplied by the vendor and IBM assumes no responsibility for its accuracy or completeness. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Any pointers in this publication to external Web sites are provided for convenience only and do not in any manner serve as an endorsement of these Web sites.

Any performance data contained in this document was determined in a controlled environment, and therefore, the results that may be obtained in other

operating environments may vary significantly. Users of this document should verify the applicable data for their specific environment.

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

AIX	AIX/6000
C + +	EtherStreamer
LANStreamer	IBM
LAN Distance	NetFinity
OS/2	OS/2 Warp
PC Server	Netfinity
Tivoli Enterprise	

The following product names are trademarks of Lotus Development Corporation:

Lotus	Notes Mail
Domino	Domino Intranet Starter Pack
Domino.Merchant	Notes Net
LotusSphere	Lotus
Notes	LotusScript
Lotus Notes	

The following product names are trademarks of Tivoli Systems, Inc.:

Tivoli Enterprise	Tivoli
Tivoli Management Framework	Tivoli Global Enterprise Manager
Tivoli NetView	NetView
Tivoli Enterprise Module for Domino/Notes	

The following terms are trademarks of other companies:

C-bus is a trademark of Corollary, Inc.

Java and HotJava are trademarks of Sun Microsystems, Incorporated.

Microsoft, Windows, Windows NT, and the Windows 95 logo are trademarks or registered trademarks of Microsoft Corporation.

NoteSizr, GroupSizr and WebSizr are trademarks or registered trademarks of Technovations, Incorporated.

PC Direct is a trademark of Ziff Communications Company and is used by IBM Corporation under license.

Pentium, Pentium II, Pentium II Xeon, MMX, ProShare, LANDesk, and ActionMedia are trademarks or registered trademarks of Intel Corporation in the U.S. and other countries.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

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

Appendix G. Related Publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

G.1 International Technical Support Organization Publications

For information on ordering these ITSO publications see "How to Get ITSO Redbooks" on page 181.

- *Netfinity Tape Solutions*, SG24-5218
- *Netfinity Server Management*, SG24-5208
- *Lotus Domino Intranet Starter Pack 2.0 Cookbook for IBM Netfinity Servers and PC Servers*, SG24-5202
- *Case Study: Domino.Merchant Server for IBM Netfinity Servers*, SG24-5217
- *The Domino Defense: Security in Lotus Notes and the Internet*, SG24-4848
- *IBM PC Servers and Windows NT 4.0 Integration Guide*, SG24-4763
- *IBM Netfinity and PC Server Technology and Selection Guide*, SG24-4760
- *Secrets to Running Lotus Notes: The Decisions No One Tells You How to Make*, SG24-4875
- *Managing Domino/Notes with Tivoli Manager for Domino, Enterprise Edition, Version 1.5*, SG24-2104
- *Implementing ServeRAID SCSI and SSA RAID Disk Subsystems*, SG24-2098
- *OS/2 Warp Server Integration Guide for IBM Netfinity and IBM PC Servers*, SG24-2125

G.2 Redbooks on CD-ROMs

Redbooks are also available on CD-ROMs. **Order a subscription** and receive updates 2-4 times a year at significant savings.

CD-ROM Title	Subscription Number	Collection Kit Number
System/390 Redbooks Collection	SBOF-7201	SK2T-2177
Networking and Systems Management Redbooks Collection	SBOF-7370	SK2T-6022
Transaction Processing and Data Management Redbook	SBOF-7240	SK2T-8038
Lotus Redbooks Collection	SBOF-6899	SK2T-8039
Tivoli Redbooks Collection	SBOF-6898	SK2T-8044
AS/400 Redbooks Collection	SBOF-7270	SK2T-2849
RS/6000 Redbooks Collection (HTML, BkMgr)	SBOF-7230	SK2T-8040
RS/6000 Redbooks Collection (PostScript)	SBOF-7205	SK2T-8041
RS/6000 Redbooks Collection (PDF Format)	SBOF-8700	SK2T-8043
Application Development Redbooks Collection	SBOF-7290	SK2T-8037

G.3 Other Publications

These publications are also relevant as further information sources:

- *Lotus Domino Web Site Development*, SR23-8686
- *Lotus Notes and Domino Network Design*, SR23-8120

- *Lotus Notes and Domino Server 4.6 Unleashed*, ISBN 0-67-231142-9
- *The Lotus Domino Server: Integrating Lotus Notes 4.6 with the Internet*, ISBN 1-55-851574-7

How to Get ITSO Redbooks

This section explains how both customers and IBM employees can find out about ITSO redbooks, redpieces, and CD-ROMs. A form for ordering books and CD-ROMs is also provided.

This information was current at the time of publication, but is continually subject to change. The latest information for customers may be found at <http://www.redbooks.ibm.com/>. and for IBM employees at <http://w3.itso.ibm.com/>.

How IBM Employees Can Get ITSO Redbooks

Employees may request ITSO deliverables (redbooks, redpieces, materials, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **Redbooks Web Site on the IBM Intranet**

Search for, view or download redbooks, redpieces, or materials or download a printable copy of a redbook, from the intranet Redbooks Web Site at:

<http://w3.itso.ibm.com/>
or <http://inews.ibm.com>

Register for information on workshops, residencies, and redbooks, by selecting the "ITSO Mailing List" button.

- **Hardcopy** —

Order hardcopies in the United States with PUBORDER.

IBMers outside the United States should use the normal hardcopy ordering process for their country.

Redpieces and Updates to Current Books

For information so current it is still in the process of being written, look at "Redpieces" on the Redbooks Web Site (<http://www.redbooks.ibm.com/redpieces.html>). Redpieces are redbooks in progress; not all redbooks become redpieces, and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

Redpieces will also be a repository for updates to current redbooks. Simply search on the book title and if you find a later "dash level" (-01, -02, etc.) check the redpiece for new information.

Materials

Materials are a repository of workshops, presentations, papers, and Web pages developed and written by the ITSO technical professionals. Look at all Materials on the intranet Redbooks Web Site by selecting the "Additional Materials" button at: (<http://w3.itso.ibm.com/>).

How Customers Can Get ITSO Redbooks

Customers may request ITSO deliverables (redbooks, redpieces, and CD-ROMs) in the following ways:

- **Redbooks Web Site on the World Wide Web**

Search for, view or download redbooks and redpieces at:

<http://www.redbooks.ibm.com/>

- **Online Orders** — send orders to:

In United States:
In Canada:
Outside North America:

IBMMAIL
usib6fpl at ibmmail
caibmbkz at ibmmail
dkibmbsh at ibmmail

Internet
usib6fpl@ibmmail.com
lmannix@vnet.ibm.com
bookshop@dk.ibm.com

- **Telephone Orders**

United States (toll free)
Canada (toll free)

1-800-879-2755
1-800-IBM-4YOU

Outside North America
(+45) 4810-1320 - Danish
(+45) 4810-1420 - Dutch
(+45) 4810-1540 - English
(+45) 4810-1670 - Finnish
(+45) 4810-1220 - French

(long distance charges apply)
(+45) 4810-1020 - German
(+45) 4810-1620 - Italian
(+45) 4810-1270 - Norwegian
(+45) 4810-1120 - Spanish
(+45) 4810-1170 - Swedish

- **Mail Orders** — send orders to:

IBM Publications
Publications Customer Support
P.O. Box 29570
Raleigh, NC 27626-0570
USA

IBM Publications
144-4th Avenue, S.W.
Calgary, Alberta T2P 3N5
Canada

IBM Direct Services
Sortemosevej 21
DK-3450 Allerød
Denmark

- **Fax** — send orders to:

United States (toll free)
Canada
Outside North America

1-800-445-9269
1-403-267-4455
(+45) 48 14 2207 (long distance charge)

- **1-800-IBM-4FAX (United States) or (+1)001-408-256-5422 (Outside USA)** — ask for:

Index # 4421 Abstracts of new redbooks
Index # 4422 IBM redbooks
Index # 4420 Redbooks for last six months

- **On the World Wide Web**

Redpieces and Updates to Current Books

For information so current it is still in the process of being written, look at "Redpieces" on the Redbooks Web Site (<http://www.redbooks.ibm.com/redpieces.html>). Redpieces are redbooks in progress; not all redbooks become redpieces, and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

Redpieces will also be a repository for updates to current redbooks. Simply search on the book title and if you find a later "dash level" (-01, -02, etc.) check the redpiece for new information.

IBM Redbook Order Form

Please send me the following:

Title	Order Number	Quantity

First name Last name

Company

Address

City Postal code Country

Telephone number Telefax number VAT number

• Invoice to customer number _____

• Credit card number _____

Credit card expiration date Card issued to Signature

We accept American Express, Diners, Eurocard, Master Card, and Visa. Payment by credit card not available in all countries. Signature mandatory for credit card payment.

Index

A

- administration 131
- Advanced Services 6
- alert 99
- application management modules 169
- application responsiveness 87
- application servers 65
- ATM 15
- authentication 10
- availability 130
- availability management 171

B

- bibliography 179

C

- calendar 5, 55
- cc:Mail 4
- CERT.ID 10
- chart 99
- cluster mail 55
- cluster replication 126, 128
- clustering 6, 15
- configuration steps 136

D

- database access control lists 10
- database indexing 92
- database servers 50
- databases cached 96
- dial-up 24
- directory services 89
- disaster recovery 68
- disk bottleneck 115
- disk partition 80
- disk space 64
- DISKPERF command 98
- Domino 3, 7
 - billing 125
 - optimization 90
 - partitioned servers 125
 - server clusters 125
- Domino.Action 5

E

- Ethernet 15
- event filtering 172
- examples of Domino server selections 69

F

- failover 130
- failover protection 126
- file system cache size 84, 87

G

- gateway services 89
- Groupware_A 54

H

- hierarchical name 9
- HTML 4
- HTTP 3, 4
- hub server 21, 52

I

- IBMEENB 90
- IBMEENB.DLL 176
- idle usage 54
- installation steps 136
- Internet 161
- Internet Connection Server 4
- IPX/SPX 8, 89
- ITSO environment 173

L

- link considerations 19
- log 99
- LOG.NSF 92
- Lotus Notes replication 127

M

- mail 4
 - database 54
 - hubs 65
 - poll time 96
 - routing hub 54
 - servers 50
 - threads 95
 - users 65
- Mail Only 54
- memory bottlenecks 112
- memory buffers 97
- Microsoft mail 4
- MTAs 5
- multiple replicators 96

N

- naming conventions 9
- NetBIOS 8, 89, 175
- NETBIOS.DLL 176
- Netfinity 97
- Netfinity servers
- NetWare 89
- NetWare Directory Services (NDS) 89
- network
 - analyzer 16
 - bottleneck 116
 - card 67
 - card driver 82
 - protocol 7
- Notes
 - client installation 172
 - log file 92
 - mail 4
 - monitor collections 172
- NOTES.INI 90, 91, 92
- NotesBench 54

O

- OLE 2.0 4
- operating system 7, 81, 82
- overloaded 24

P

- PAGEFILE.SYS 86
- paging 64
- paging file 86
- Passthru servers 52
- Performance Monitor data 99
- planning 131
- planning guidelines 135
- POP3 5
- processor 67
- processor bottlenecks 114
- protocol optimization 88

R

- RAID 66
- RAM 63
- random access memory 63
- recovering from failures 149
- Redundant Array of Inexpensive Disks 66
- Remote Access (Dial-In) 11
- Remote Access Servers 51
- remote Domino server 17
- Replica task 24
- replication 5, 20
- replication hub 54
- replication hubs 65
- replication schedules 20

- replicator 91
- router 91

S

- scalability 127
- scheduling 5, 55
- security considerations 10
- server
 - access control lists 10
 - configuration 79
 - optimization 79
 - recommendations 63
 - tasks 91
- SERVER.ID 10
- Server.Planner 57, 58
- ServerTasks 24
- sessions 93
- shared database 54, 55
- shared discussion 55
- SMP 6
- SMTP MTA 5
- SSL 6
- swap file 87
- symmetric multiprocessing 6
- system installation 80
- System Performance Monitor 97
- systems management 27

T

- task administration 170
- tasks task library 170
- TCP/IP 8, 89
- Tivoli Enterprise Module for Domino/Notes 169
 - overview 169
 - security enhancements 170
- Tivoli Enterprise Module for Domino/Notes event adapter 171
- Tivoli Enterprise Module for Domino/Notes0 Notes monitor collections 172
- Tivoli Enterprise modules 169
- Tivoli event server rulebase 172
- token-ring 15

U

- UPS 67
- URL 4
- user management with Admin Extension 172
- user tasks task library 170
- users 94
- utilities task library 170

V

- virus protection 68

W

- Warp server 87
- Web browsers 3
- Web Buyer 55
- Web server 4, 51
- Web Walker 54
- wide area networking 16
- Windows NT Performance Monitor 97
- Windows NT Server 83
- workload balancing 126, 131
 - maximum number of users 131
 - performance index threshold 131
- workloads 54

X

- X.500 naming standard 9
- X.PC 8

ITSO Redbook Evaluation

Lotus Domino Integration Guide for IBM Netfinity Servers
SG24-2102-02

Your feedback is very important to help us maintain the quality of ITSO redbooks. **Please complete this questionnaire and fax it to: USA International Access Code + 1 914 432 8264 or:**

- Use the online evaluation form found at <http://www.redbooks.ibm.com>
- Send your comments in an Internet note to redbook@us.ibm.com

Which of the following best describes you?

Customer **Business Partner** **Solution Developer** **IBM employee**
 None of the above

Please rate your overall satisfaction with this book using the scale:
(1 = very good, 2 = good, 3 = average, 4 = poor, 5 = very poor)

Overall Satisfaction _____

Please answer the following questions:

Was this redbook published in time for your needs? Yes____ No____

If no, please explain:

What other redbooks would you like to see published?

Comments/Suggestions: **(THANK YOU FOR YOUR FEEDBACK!)**

