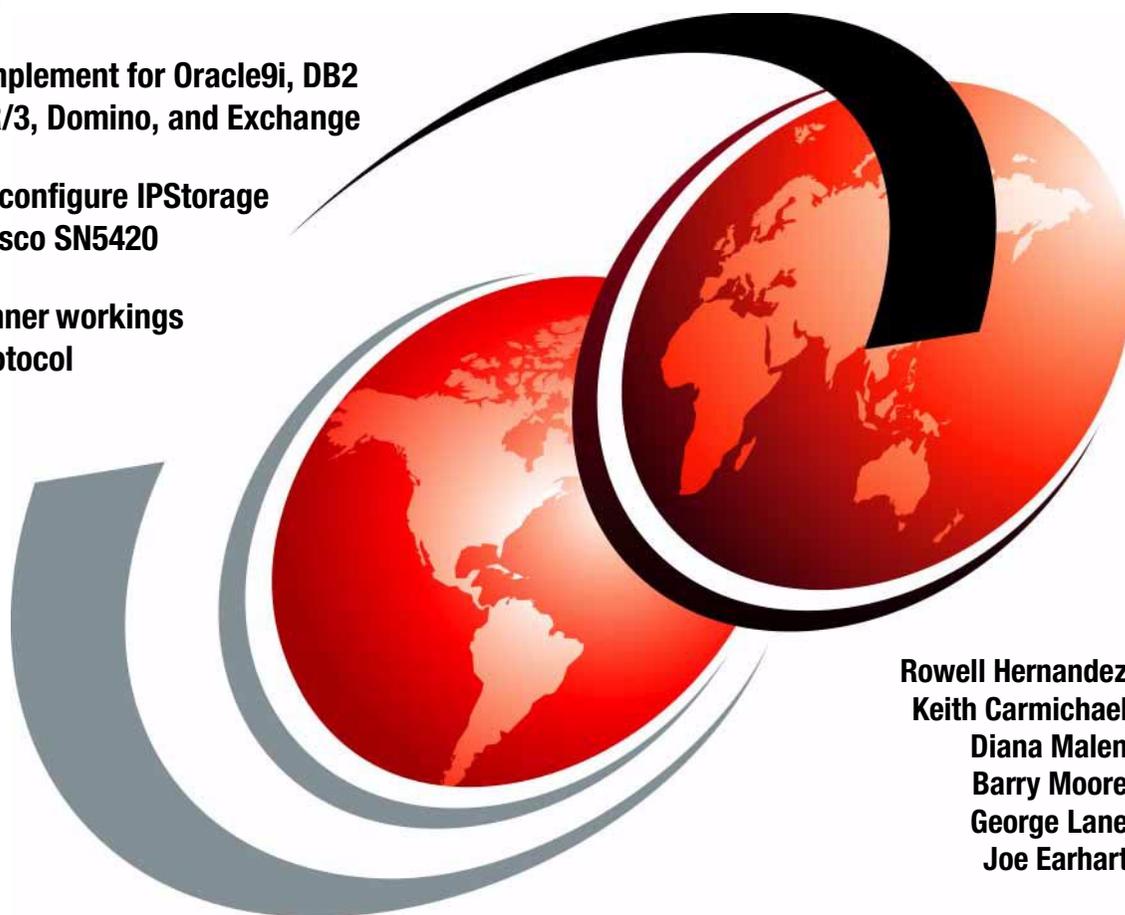


Using iSCSI Solutions' Planning and Implementation

Plan and implement for Oracle9i, DB2
UDB, SAP R/3, Domino, and Exchange

Install and configure IPStorage
200i and Cisco SN5420

Learn the inner workings
of iSCSI protocol



Rowell Hernandez
Keith Carmichael
Diana Malen
Barry Moore
George Lane
Joe Earhart



International Technical Support Organization

**Using iSCSI
Solutions' Planning and Implementation**

February 2002

Take Note! Before using this information and the product it supports, be sure to read the general information in “Special notices” on page 311.

First Edition (February 2002)

This edition applies to the IBM TotalStorage IPStorage 200i Models 210, 110 with microcode Release 1.2, Cisco SN5420 storage and initiator clients running on Redhat Linux 7.1, Windows 2000 and Windows NT.

Comments may be addressed to:
IBM Corporation, International Technical Support Organization
Dept. 471F Building 80-E2
650 Harry Road
San Jose, California 95120-6099

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 2002. 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

Figures	ix
Tables	xv
Preface	xvii
The team that wrote this redbook	xvii
Special notice	xix
IBM trademarks	xx
Comments welcome	xx
Chapter 1. Introduction to Internet SCSI	1
1.1 What Internet SCSI (iSCSI) is, in a nut shell	2
1.1.1 Internet Engineering Task Force draft	3
1.2 iSCSI comparison	4
1.2.1 SAN	5
1.2.2 NAS	6
1.2.3 iSCSI	8
1.2.4 iSCSI benefits	9
1.2.5 Considerations to make	11
1.2.6 When might this technology be used?	13
1.3 iSCSI components	13
1.3.1 iSCSI in detail	14
1.3.2 iSCSI technical introduction	15
1.3.3 How does iSCSI keep track?	15
1.3.4 iSCSI discovery	15
1.3.5 Starting iSCSI (iSCSI login)	15
1.3.6 iSCSI full feature phase	16
1.3.7 iSCSI data flows	16
1.3.8 iSCSI connection termination	17
1.3.9 iSCSI naming	17
1.3.10 iSCSI addressing	17
1.3.11 iSCSI message synchronization	18
1.3.12 iSCSI synch and steering	18
1.3.13 The IBM TotalStorage IP Storage 200i	18
Chapter 2. IBM TotalStorage IPStorage 200i	21
2.1 IBM TotalStorage IP Storage 200i series	22
2.1.1 IBM TotalStorage IP Storage 200i configurations	23
2.1.2 IBM TotalStorage IP Storage 200i technical specifications	24

2.1.3	IBM TotalStorage IP Storage 200i microcode	25
2.1.4	IBM TotalStorage IP Storage 200i features and profiles	26
2.1.5	IBM IP Storage high availability and serviceability	27
2.1.6	IBM IP Storage expandability and growth	27
2.1.7	IBM IP Storage 200i 4125-EXP expansion unit	28
2.1.8	IBM IP Storage 200i optional features	28
2.2	Implementation and setup	29
2.2.1	General hints and tips	31
2.2.2	Disk drive addressing	34
2.3	Step-by-step installation guide for 200i	38
2.3.1	Start up tips	38
2.3.2	Setup of the 200i	40
2.3.3	Logon to the appliance	41
2.3.4	Setting the service port IP addresses	42
2.3.5	Rebooting the 200i	43
2.3.6	Setting the data port IP address	46
2.3.7	Setting Up VLUNs	47
2.3.8	Client initiator setup procedure	53
2.3.9	Troubleshooting tips	61
	Chapter 3. Cisco SN5420 Storage Router	63
3.1	Cisco SN5420 Storage Router introduction	64
3.2	Features and benefits	66
3.2.1	Making disk subsystems IP aware	66
3.2.2	Interoperability	66
3.2.3	Scalability and reliability	66
3.2.4	Manageability	67
3.2.5	Security	67
3.2.6	Investment protection	67
3.3	Technical specifications	67
3.4	Installation	69
3.5	Configuration	71
3.5.1	SCSI router service	71
3.5.2	Clustering and high availability	71
3.5.3	Sample configuration	72
3.5.4	Configuration from the console port	73
3.5.5	Configuring management parameters	76
3.5.6	Configuring the high availability (HA) interface	81
3.5.7	Configuring SCSI routing	83
3.5.8	Configuring iSCSI targets	86
3.5.9	Configuring server access	89
3.5.10	Using the GUI on the Cisco SN5420 Storage Router	91
3.6	Host setup	93

3.6.1	iSCSI Windows 2000 driver installation and configuration	93
3.6.2	iSCSI Windows NT driver installation and configuration	96
3.6.3	iSCSI UNIX and Solaris driver installation and configuration	96
3.6.4	iSCSI Linux driver installation and configuration	98
3.6.5	Troubleshooting	101
3.7	Cisco SN5420 Storage Router implementation	106
3.7.1	IBM Enterprise Storage Server (ESS).	107
3.7.2	McData ED-5000 Enterprise Fibre Channel Director	107
3.7.3	Cisco SN5420 IP Storage Router	108
3.7.4	Cisco Catalyst 3500 Series XL switches.	108
3.7.5	Microsoft Windows AutoExNt-Service.	108
3.7.6	Creating a database on ESS disk LUNS.	108
Chapter 4. Implementing Enterprise Mail for iSCSI		111
4.1	Planning the Enterprise Mail installation	112
4.2	Prerequisites for installing Exchange 2000	114
4.2.1	Hardware requirements.	115
4.2.2	Software requirements	115
4.2.3	Before Exchange 2000 installation	118
4.2.4	Installing the Active Directory Connector	118
4.2.5	Advantage of using the forestprep switch	120
4.2.6	Preparing the forest with the /forestprep switch	120
4.2.7	Advantages of using the domainprep switch.	122
4.2.8	Preparing the domain with /domainprep switch	122
4.2.9	Installing Exchange 2000	125
4.2.10	Creating a new mail storage group	128
4.2.11	Creating an e-mail enabled user account	129
4.2.12	Configuring a users Outlook Express mail	131
4.2.13	Setting up startup scripts for Exchange	131
4.3	Installing Lotus Domino	135
4.3.1	Prerequisites for installing Lotus Domino	136
4.3.2	Hardware requirements.	136
4.3.3	Software requirements	137
4.3.4	Steps for installing Lotus Domino	137
4.3.5	Installation of Lotus Domino Administrator	141
4.3.6	Creating a new user in Domino.	146
4.3.7	Configuring startup scripts for Domino	148
4.4	General performance overview.	150
4.4.1	Performance monitors.	151
4.4.2	Available tools in Windows 2000.	151
4.4.3	Spotting a bottleneck.	152
Chapter 5. Implementing DB2 UDB EE V7 Database Server using iSCSI		155

5.1	Fundamentals of RDBMS	156
5.1.1	Database	156
5.1.2	Tables	156
5.1.3	Data dictionary	156
5.1.4	Indexes	157
5.1.5	Tablespaces	157
5.1.6	Partitioning options	159
5.1.7	Log files	159
5.1.8	Control files	160
5.1.9	Configuration parameters	161
5.2	DB2 UDB EE V7 concepts	161
5.2.1	DB2 UDB products	161
5.2.2	DB2 UDB EEE partitioning option	162
5.2.3	Instance	163
5.2.4	Database and database partitions	163
5.2.5	Nodegroups	164
5.2.6	Tablespaces	165
5.2.7	Recovery logs	166
5.3	Database planning considerations	167
5.3.1	Backup requirements	168
5.4	Creating a database on the iSCSI storage device	168
5.4.1	Creating a database	169
5.5	Implementation of the AutoExNt-Service	170
Chapter 6. Implementing Oracle9i on iSCSI		173
6.1	Database architecture	174
6.1.1	Oracle9i instance	176
6.1.2	Oracle9i database	178
6.1.3	The Oracle client	179
6.2	Oracle9i new features	180
6.2.1	New feature categories	183
6.3	Database planning	190
6.3.1	Requirements for the environments	193
6.3.2	Oracle9i installation and database types	196
6.4	Database implementation and configuration	197
6.4.1	Windows 2000 installation and configuration	197
6.4.2	Linux installation and configuration	220
6.4.3	Database creation on iSCSI disks	242
6.5	Hints and tips	283
6.6	Configuring startup scripts for Oracle	286
Chapter 7. SAP R/3 considerations for iSCSI		289
7.1	A general model of SAP R/3	290

7.2	SAP R/3 certification requirements (iXOS)	291
7.2.1	Certification categories	291
7.2.2	Hardware components	292
7.3	Use of Windows AutoExNt-Service and iSCSI	293
7.3.1	Windows AutoExNt-Service for SAP R/3 and DB2 UDB.	293
7.3.2	Implementation of the AutoExNt-Service	295
7.4	Information about disk layouts	296
7.4.1	Disk considerations for DB2 and SAP R/3	297
7.4.2	Recommended disk layout and configuration	299
7.5	SAPRouter configuration and recommendations.	300
7.6	Networking recommendations	301
7.7	Cisco SN5420 IP Storage Router and IBM ESS	302
	Appendix A. Data migration considerations for DB2.	303
	Moving database objects.	304
	DB2MOVE EXPORT and DB2MOVE IMPORT	304
	DB2MOVE EXPORT, DB2LOOK, DB2MOVE LOAD	305
	Related publications	307
	IBM Redbooks	307
	Other resources	307
	Referenced Web sites	307
	How to get IBM Redbooks	308
	IBM Redbooks collections.	309
	Special notices	311
	Index	313

Figures

1-1	SAN/NAS/iSCSI comparison	5
1-2	File I/O example	7
1-3	iSCSI uses block I/O	8
1-4	iSCSI appliance connection options	9
1-5	iSCSI components	14
1-6	The seven layers of the OSI reference model	16
1-7	Data encapsulation of iSCSI frame	17
2-1	IBM IP Storage 200i	23
2-2	Lab configuration	31
2-3	Netscape proxy setting	33
2-4	Internet Explorer proxy settings	33
2-5	200i operating panel	40
2-6	Log on window	41
2-7	200i storage home page	42
2-8	Network options	43
2-9	System options	44
2-10	Reset warning	44
2-11	Reset final warning	45
2-12	Reboot wait message	45
2-13	Data port selection	46
2-14	Edit data port option	47
2-15	Storage options	48
2-16	Access control page	49
2-17	VLUN ADD option	50
2-18	Configured VLUNs	51
2-19	Client setup	52
2-20	A single VLUN	53
2-21	Installation wizard	54
2-22	IBM license agreement	55
2-23	Summary page	55
2-24	Reboot option	56
2-25	IBM iSCSI configuration	56
2-26	IBM iSCSI configuration options	56
2-27	IBM iSCSI configuration utility	57
2-28	IBM iSCSI configuration logon	58
2-29	IBM iSCSI configuration	58
2-30	IBM iSCSI configuration advanced options	59
2-31	Windows disk management	60

2-32	Available iSCSI disk	60
3-1	Cisco SN5420 Storage Router	64
3-2	iSCSI routing	65
3-3	Fibre Channel port	69
3-4	Gigabit Ethernet port	69
3-5	Console port	70
3-6	Management port	70
3-7	HA port	70
3-8	Initial start up screen	74
3-9	Enter the management port IP address	75
3-10	Re-establish Admin mode by entering enable command	76
3-11	Example network	77
3-12	Initial configuration commands (screen 1)	78
3-13	Initial configuration commands (screen 2)	79
3-14	HA configuration commands	82
3-15	SCSI routing example	84
3-16	Commands to create SCSI routing and ACLs	85
3-17	Target and LUN mapping using WWPN	87
3-18	Command to apply ACL on the SCSI router	90
3-19	Initial login window	91
3-20	Management port configuration	92
3-21	System software display	92
3-22	Windows 2000 iSCSI configuration	94
3-23	Windows 2000 disk administrator	95
3-24	Reactivating a disk	95
3-25	System troubleshooting commands (screen 1)	102
3-26	System troubleshooting commands (screen 2)	103
3-27	System troubleshooting commands (screen 3)	104
3-28	System troubleshooting commands (screen 4)	105
3-29	System troubleshooting commands (screen 5)	106
3-30	iSCSI via Cisco SN5420, Fibre Channel to IBM ESS	109
4-1	Increasing page file size	116
4-2	IIS selection page	118
4-3	ADC setup selection	120
4-4	Exchange 2000 - ForestPrep	121
4-5	ForestPrep installation bar	122
4-6	Exchange 2000 - DomainPrep	124
4-7	Pre-Windows 2000 compatibility warning	124
4-8	DomainPrep progress bar	125
4-9	Exchange 2000 selection window	126
4-10	Exchange 2000 installation progress bar	127
4-11	MMC console	130
4-12	New user mail box options	130

4-13	Domino server installation option	139
4-14	Lotus Domino server setup	140
4-15	Server audience	140
4-16	Administrator settings	141
4-17	Notes installation selection options	143
4-18	Notes client configuration setup	144
4-19	Server connection error window	145
4-20	Domino administration	146
4-21	New user registration	147
5-1	Table relationship to tablespaces	157
5-2	Data storage hierarchy	158
5-3	Database logfiles	160
5-4	DB2 partitioning	162
5-5	DB2 instance hierarchy	163
5-6	Nodegroups	165
6-1	Oracle9i architecture	175
6-2	Listener process listens for incoming requests	179
6-3	Oracle Universal Installer welcome window	200
6-4	OUI File locations window	201
6-5	Available products window	202
6-6	Installation types window	203
6-7	Database configuration window	204
6-8	Upgrading and migrating an existing database window	205
6-9	Database identification window	206
6-10	Database file location window	207
6-11	Database character set window	208
6-12	Summary window	209
6-13	Install window	210
6-14	Configuration tools window	211
6-15	Configuration assistant window	212
6-16	ODCA database information window	213
6-17	Password management window	214
6-18	End of installation window	215
6-19	Linux OUI exit window	215
6-20	Services processes window	216
6-21	SQL Plus log on window	217
6-22	Oracle SQL*Plus window	218
6-23	Oracle SQL*Plus result set window	218
6-24	Linux OUI welcome window	224
6-25	Linux OUI file locations window	225
6-26	Linux OUI available products window	226
6-27	Linux OUI installation types window	227
6-28	Linux OUI database configuration window	228

6-29	Linux OUI privileged operating system groups window	229
6-30	Linux OUI database identification window	230
6-31	Linux database character set window	231
6-32	Linux OUI choose JDK home directory window	232
6-33	Linux OUI summary window	233
6-34	Linux OUI install window	234
6-35	Linux OUI error window	235
6-36	Linux Oracle fixed script window	236
6-37	Linux OUI setup privileges window	236
6-38	Linux OUI root.sh execution window	237
6-39	Linux OUI configuration tools window	237
6-40	Linux database configuration assistant window	238
6-41	Linux ODCA database information window	239
6-42	Linux OUI end installation window	240
6-43	Linux OUI exit window	241
6-44	DBCA welcome window	244
6-45	DBCA operations window	245
6-46	DBCA database templates window	246
6-47	Database identification window	247
6-48	Database connection options window	248
6-49	Initialization parameters window	249
6-50	Initialization parameters archive tab window	250
6-51	Initialization parameters db sizing tab window	251
6-52	Initialization parameters file location tab window	252
6-53	Database storage window	253
6-54	File name and file directory window	254
6-55	Change to iSCSI drive window	255
6-56	Summary window	256
6-57	Creation options windows	257
6-58	Create the database window	258
6-59	Linux database storage	259
6-60	Managed standby error window	260
6-61	Oracle not available error window	260
6-62	Shared memory error window	260
6-63	OUI client welcome window	262
6-64	OUI client file locations window	263
6-65	OUI client installation types window	264
6-66	OUI client summary window	265
6-67	OUI client install window	266
6-68	OUI client configuration tools window	267
6-69	ONCA client welcome window	268
6-70	NSNC database version window	269
6-71	ONCA complete window	270

6-72	OUI end of installation window	271
6-73	ONCA welcome window	272
6-74	ONCA net service name configuration window	273
6-75	NSNC database version window	274
6-76	NSNC service name window	275
6-77	NSNC select protocols window	276
6-78	NSNC TCP/IP protocol window	277
6-79	NSNC test window	278
6-80	NSNC connecting window	279
6-81	NSNC net service name window	280
6-82	NSNC another service name window	281
6-83	NSNC complete window	282
6-84	ONCA welcome window	283
6-85	Command prompt error window	285
7-1	SAP technical hierarchy	290

Tables

2-1	Product specification 200i	24
3-1	Cisco SN5420 Storage Router technical specifications	67
3-2	Subnet configuration example	70
3-3	Sample system parameters	72
3-4	Initial configuration commands	79
3-5	HA configuration commands	83
3-6	SCSI routing and ACL configuration commands	86
3-7	SCSI target configuration commands	87
3-8	Server access control commands	90
3-9	System troubleshooting commands	101
4-1	Exchange 2000 Server/Enterprise options	113
4-2	Setup information for Exchange 2000	114
4-3	Affected services for starting Exchange 2000	132
4-4	Setup information for Lotus Domino R5	135
4-5	Lotus Domino startup script	148
5-1	AutoExNt-Service	171
6-1	Oracle9i new features	180
6-2	Affected services that need to be changed	287
7-1	AutoExNt-Service	294
7-2	Filesystem layouts for SAP R/3 on Microsoft Windows NT/2000	300

Preface

This IBM Redbook gives a broad understanding of the upcoming iSCSI architecture standard. You can gain an understanding of the iSCSI technology and products that provide iSCSI solutions from IBM with the TotalStorage 200i, and from Cisco with the SN5420 Storage Router.

This redbook will guide you in planning, installing, implementing and configuring an iSCSI solution, and the installation of the iSCSI clients for both pieces of equipment on their supported operating systems.

The book will also cover, in detail, the installation of applications when using iSCSI storage. The applications include Oracle9i for Windows 2000 and Linux, IBM DB2, SAP R/3, Microsoft Exchange 2000 and Lotus Domino for Windows 2000.

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 and San Jose Centers.

Rowell Hernandez is a Project Leader for Network Attached Storage and Internet SCSI research and documentation at the International Technical Support Organization, San Jose Center. Before joining the ITSO in 2001, he worked as an IT Specialist for IBM Philippines, providing support for the IBM @server xSeries Server and Netfinity servers, Windows NT, Clustering and Linux. Rowell is also an IBM @server Certified Systems Expert - xSeries Windows, Microsoft Certified Systems Engineer + Internet, Citrix Certified Administrator, Comptia A+ Certified Technician and Red Hat Certified Engineer. He holds a bachelor's degree in Computer Science from AMA University, along with graduate work in Information Management at Ateneo De Manila University.

Keith Carmichael is an advisory IT Availability Professional from IBM South Africa. He's worked for IBM for the last five years in Integrated Technology Services Division. Keith's current responsibilities include Technical Support for PCs as well as managing the Parts Recovery Center. He is a Microsoft Certified Professional and also busy working toward his Windows 2000 MCSE certification. His areas of expertise are Windows NT, Windows 2000, xSeries and Netfinity servers, Desktop, ThinkPads, and thin clients. Keith holds a National Diploma in Electrical Engineering.

Diana Malen is a DBA Specialist for Strategic Outsourcing at IBM Limited New Zealand Global Services in Auckland. She has a bachelor's degree in Computer Science from the University of Tulsa, Oklahoma. She has more than 11 years of experience in the IT industry. Before moving to New Zealand, she worked as a Software Developer and as a Sybase and Oracle Database Administrator in the USA. Her expertise is in database administration, specifically for Oracle and Informix.

Barry Moore is a Consulting IT Specialist at IBM with 15 years of experience in the IT field. He holds a bachelor's degree in Business Computer Information Systems from the University of North Texas, along with graduate work in Computer Systems at Marist College, Poughkeepsie, New York. His areas of expertise include SAP R/3, Relational Database Management Systems (RDBMS), RS/6000 servers and storage, Storage Area Networks (SAN), and IBM Websphere.

George Lane is a Storage Networking Consultant at IBM. He started with IBM 23 years ago as a Customer Engineer. His focus has always been customer support. In previous positions he has worked as an IT Specialist in Networking and AIX. Recently George was part of IBM Global Services in Software Support Practice. Currently George is working with SSG in Pre-Sales Marketing Support for NAS and iSCSI products, and has worked extensively with the NAS and iSCSI development and test teams prior to these products coming to market. George's responsibilities now cover worldwide support for SSG.

Joe Earhart is a Systems Engineer based in Raleigh, North Carolina. He has 21 years of experience in the telecommunications field. He has worked at Cisco Systems for one year. His areas of expertise include TCP/IP, routing, switching and storage area networking.

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

Emma Jacobs
International Technical Support Organization, San Jose Center

Gabrielle Velez
International Technical Support Organization, Rochester Center

John Byrd, David Watts, Steve Russell, Tamikia Barrow, Gail Christensen,
Margaret Ticknor, and Jeanne Tucker
International Technical Support Organization, Raleigh Center

Sandy Albu, William Champion, Brian Tesar, Megan Kirkpatrick, Jim Blue, Paul Mattson, Matthew Lietzke, Karen Belding, Sandra Kipp, Margaret Lukowski, David Reich, Thomas Daniels, Jim Sedgwick, Bryan Winkleman, Garry Rawlins, Holly Tallon, John Dowdy, Hiromi Fukushima, and Jim Tuckwell
IBM Raleigh

Paul Fried
IBM Poughkeepsie

Pete Potosky, John Hartman, and Steve Miedema
IBM Rochester

Susan Grey, Mickey Jones, Bruce Allworth, and Lamar Reavis
IBM Advanced Technical Support - Storage Networking

Julian Satran
IBM Haifa

Wolfgang Singer
IBM Austria

Special notice

This publication is intended to help IBM Professionals, Customers and Business Partners to understand, plan and implement solutions using iSCSI. The information in this publication is not intended as the specification of any programming interfaces that are provided by IBM TotalStorage 200i and Cisco SN5420. See the PUBLICATIONS section of the IBM Programming Announcement for IBM TotalStorage 200i and Cisco SN5420 for more information about what publications are considered to be product documentation.

IBM trademarks

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

e (logo)® 	Redbooks Logo 
IBM ®	Redbooks™
AIX®	S/390®
AS/400®	ServeRAID™
DB2®	SP™
Enterprise Storage Server™	SP2®
ESCON®	ThinkPad®
Home Director™	Tivoli®
Hummingbird®	TotalStorage™
IBM®	Wave®
Informix™	xSeries™
Netfinity®	Lotus®
OS/2®	cc:Mail®
Perform™	Lotus Notes®
RS/6000®	Notes®
	Domino™

Comments welcome

Your comments are important to us!

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

- ▶ Use the online **Contact us** review redbook form found at:
ibm.com/redbooks
- ▶ Send your comments in an Internet note to:
redbook@us.ibm.com
- ▶ Mail your comments to the address on page ii.



Introduction to Internet SCSI

Storage is one of the fastest growing investment areas in the I/T arena. In the July 2000 issue of *Technology Investor* it states, “more storage capacity was sold in 1999 than in all previous years combined”. Industry experts estimate that data stored is doubling every year.

What is the driving force behind this phenomenal growth in storage? We find that dramatic growth in existing applications, such as online transactions, e-mail and so forth, plus the development of complex new e-commerce applications, including multimedia applications, play a big part. It is driven by systems like the Internet and intranets. The requirement for continuous availability of information in the e-business world encourages organizations to duplicate, even triplicate online copies of their data. And, it is ultimately enabled by the extraordinary ability of the disk drive industry to keep doubling the disk drive capacities almost yearly, while at the same time delivering 30% to 40% compound price reductions.

With the increased need for storage, the shortage of I/T professionals in the marketplace, the competition to get new products and technologies to market, and the increased budget constraints to accomplish these tasks, the need for more storage options becomes an important factor.

In this chapter we will discuss one of these new technologies, iSCSI. We will try to answer the questions:

- ▶ What is iSCSI?
- ▶ How does it compare to SAN and NAS?

- ▶ When would iSCSI be considered as a storage solution?
- ▶ How does it work?
- ▶ What products are available to implement an iSCSI solution?

IBM is a leader in the development and delivery of iSCSI technology and storage products. With the announcement in February 2001 of the TotalStorage IP Storage 200i, IBM is the first to market with an iSCSI solution. IBM, as well as other network and storage vendors, is working very closely with the Internet Engineering Task Force or IETF in developing the iSCSI standards.

In later chapters we will discuss some of the operating environments and applications that may be used with this technology. We will also take a closer look at the iSCSI hardware that was used in the development of this redbook.

1.1 What Internet SCSI (iSCSI) is, in a nut shell

iSCSI is a proposed industry-standard that allows SCSI block I/O protocols (commands, sequences and attributes) to be sent over a network using the popular TCP/IP protocol. This is analogous to the way SCSI commands are already mapped to Fibre Channel, parallel SCSI, and SSA media. (Do not confuse this with the SCSI cabling transport mechanism. We are talking about protocols).The iSCSI proposal was made to the Internet Engineering Task Force (IETF) standards body jointly by IBM and Cisco.

The question arises whether we can use TCP/IP, the networking technology of Ethernet LANs and the Internet, for storage. This could enable the possibility of having a single network for everything. This could include storage, data sharing, Web access, device management using SNMP, e-mail, voice and video transmission, and all other uses. IP SANs could leverage the prevailing technology of the Internet to scale from the limits of a LAN to wide area networks, therefore enabling new classes of storage applications. SCSI over IP could enable general purpose storage applications to run over TCP/IP. Moreover, an IP SAN would also automatically benefit from new networking developments on the Internet, such as quality of service and security. It is also widely anticipated that the total cost of ownership of IP SANs would be lower. This is due to larger volumes of existing IP networks and the wider skilled manpower base that is familiar with them.

The problem is that IP networking is based on design considerations different from those of storage concepts. Is it possible to merge the two concepts and yet provide the performance of a specialized storage protocol like SCSI, with block I/O direct to devices? The challenges are many. The TCP/IP protocol is

software-based and geared towards unsolicited packets, whereas storage protocol are hardware-based and use solicited packets. Also a storage networking protocol would have to leverage the TCP/IP stack without change and still achieve high performance.

IBM research, at its Almaden and Haifa research centers, has undertaken to solve these issues. The goal is to make the promise of IP SANs a reality. Efforts are concentrated along two different directions: the primary effort is to bridge the difference in performance between Fibre Channel and IP SANs. In parallel, there is an effort to define a standard mapping of SCSI over TCP/IP. The result is Internet SCSI (iSCSI).

In the simplest terms, iSCSI allows the transport of SCSI commands over TCP/IP. Yet the implication is tremendous. The proliferation of the Internet brought the explosion of the networking industry in the last decade. Still, the growth area for the new decade will be storage. At the higher end of the spectrum, the deployment of SANs will reflect the deployment of LANs in the earlier decade. However, this realization will not result from Fibre Channel (FC) alone. While FC has the advantage of throughput at the moment, it lacks the maturity and breadth compared to IP-related technology. This lack is reflected in cost, interoperability, and functionality of FC equipment. At the same time, this lack opens a window for IP-based SAN. FC's big advantage of throughput is being diminished by the rapid deployment of Gigabit Ethernet technology, with 10G on the horizon. iSCSI brings SAN onto the IP infrastructure that is familiar, affordable, mature, and rich in function.

Yet iSCSI as a technology is still in its infancy. The protocol still is a draft at this writing, by the Internet Engineering Task Force (IETF) working group.

1.1.1 Internet Engineering Task Force draft

The IETF is a large open international community of network designers, operators, vendors, and researchers concerned with the evolution of the Internet architecture and the smooth operation of the Internet. It is open to any interested individual. The actual technical work of the IETF is done in its working groups, which are organized by topic into several areas (for example, routing, transport, security, and so on). The IETF is under the Storage Networking Industry Association or SNIA's IP Storage Forum that promotes interoperability among storage products and vendors. SNIA is accepted as the primary organization for the development of SAN and NAS standards, with over 150 companies and individuals as its members, including all the major server, storage, and fabric component vendors. SNIA is committed to delivering architectures, education and services that will propel storage networking solutions into a broader market. IBM is one of the founding members of SNIA, and has senior representatives participating on the board and in technical groups. The SNIA mission is to

promote the use of storage network systems across the IT community. SNIA has become the central point of contact for the industry. It aims to accelerate the development and evolution of standards, to promote their acceptance among vendors and IT professionals, and to deliver education and information.

SNIA established the IETF in April 2001 joining approximately twenty companies together to establish the iSCSI standard. The announcement can be found at the following Web site:

<http://www.snia.org/English/Overview>

In that announcement SNIA made the following comments about the far reaching usage of iSCSI technology; they said, “iSCSI specifies a way to transport data residing on SCSI devices natively over TCP. iSCSI allows a company’s existing Ethernet network to transfer SCSI commands and data with total location-independence. Standard routers can be used to extend the access to a WAN or, if transmitted over the Internet, to any location in the world. The use of IP allows both small and large enterprises to leverage their investments on existing Ethernet infrastructure and IT staff knowledge. Hence, products built on the iSCSI standard will be welcomed by businesses of all sizes as a means to maximize their return on investment (ROI).

IBM has taken a leadership role in the development and implementation of open standards for iSCSI. As it is a new technology you can expect additional developments as iSCSI matures. Also we can anticipate extensions to other existing technologies, such as faster Ethernet (10GB/second), HBA’s with TCP/IP off load engines, and so on.

1.2 iSCSI comparison

How does iSCSI compare to SAN and NAS? Figure 1-1 is a simple diagram showing the three implementations side by side.

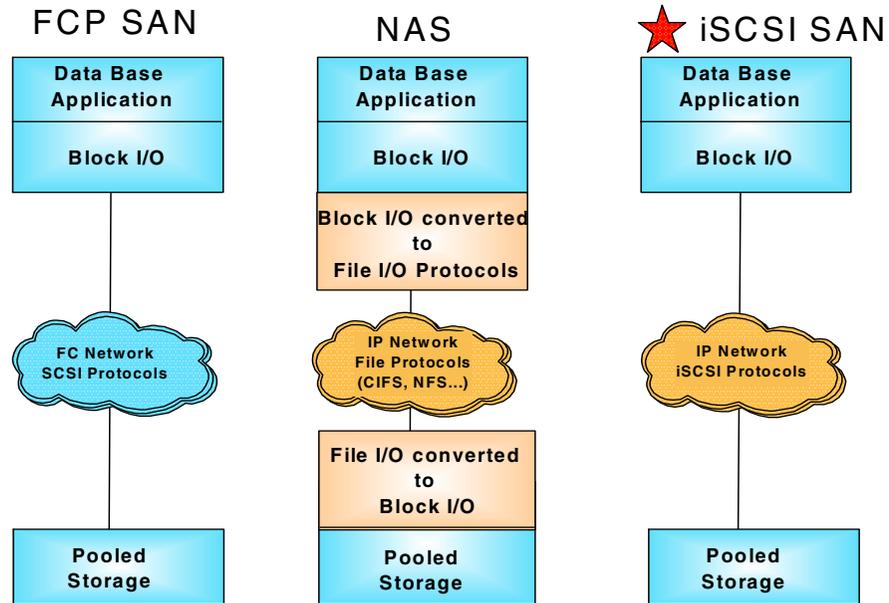


Figure 1-1 SAN/NAS/iSCSI comparison

1.2.1 SAN

A SAN is a specialized, dedicated high speed network. Servers and storage devices may attach to the SAN. It is sometimes called “the network behind the servers”. Like a LAN, a SAN allows “any to any” connection across the network, using interconnect elements such as routers, gateways, hubs and switches. Fibre Channel is the *de facto* SAN networking architecture, although other network standards could be used. In this book, when we refer to SANs we mean Fibre Channel SAN. Therefore, a decision to implement a SAN is a decision to develop a new storage network infrastructure. Large numbers of customers worldwide are implementing Fibre Channel SAN’s right now. Industry analysts view this as the storage network infrastructure with the most momentum during the next two or three years.

Fibre Channel is a multi-layered network, based on a series of American National Standards Institute (ANSI) standards. These define characteristics and functions for moving data across the network. Like other networks, information is sent in structured packets or frames, and data is serialized before transmission. But, unlike other networks, the Fibre Channel architecture includes a significant amount of hardware processing. This is oriented to storage block I/O protocols,

such as serial SCSI (known as FCP). It is capable of delivering very high performance, relative to a NAS device, which is optimized for network file I/O. The speed currently achieved is 100 MBps full duplex, with 200 MBps soon to be delivered.

Measured effective data rates of Fibre Channel has been demonstrated in the range of 60 to 80 MBps over the 1 Gbps implementation. This compares to less than 30 MBps measured over Gigabit Ethernet. The packet size of Fibre Channel is 2,112 bytes (rather large compared to some other network protocols). For instance an IP packet is 1,518 bytes, although normally IP transfers are much smaller. But for Fibre Channel, a maximum transfer unit sequence of up to 64 frames can be defined, allowing transfers of up to 128 MBps without incurring additional overheads due to processor interrupts. Therefore, today Fiber Channel is unsurpassed for efficiency and high performance in moving large amounts of data.

However, a SAN implementation does not come without a price. Because of the complexities involved, a SAN can be an expensive investment. Storage management becomes a consideration. A high level of skill is needed to maintain and manage a SAN. Costs can be as high as \$2500 per port for the hardware alone. Much planning must go into implementing a SAN. The decision to implement a SAN is not a trivial one.

1.2.2 NAS

Storage devices which optimize the concept of file sharing across the network have come to be known as NAS. NAS solutions utilize the mature Ethernet IP network technology of the LAN. Data is sent to and from NAS devices over the LAN using TCP/IP protocol. By making storage devices LAN addressable, the storage is freed from its direct attachment to a specific server, and any-to-any connectivity is facilitated using the LAN fabric. In principle, any user running any operating system can access files on the remote storage device. This is done by means of a common network access protocol, for example, NFS for UNIX servers, and CIFS for Windows servers. In addition, a task, such as back-up to tape, can be performed across the LAN, using software like Tivoli Storage Manager (TSM), enabling sharing of expensive hardware resources, such as automated tape libraries, between multiple servers. A storage device cannot just attach to a LAN. It needs intelligence to manage the transfer and the organization of data on the device. The intelligence is provided by a dedicated server to which the common storage is attached. It is important to understand this concept. NAS comprises a server, an operating system, plus storage which is shared across the network by many other servers and clients. So a NAS is a *device*, rather than a *network infrastructure*, and shared storage is attached to the NAS server.

One of the key differences of a NAS disk device, compared to Direct Attached Storage or other network storage solutions, such as SAN or iSCSI, is that all I/O operations use file level I/O protocols. File I/O is a high level type of request that, in essence, specifies only the file to be accessed, but does not directly address the storage device. This is done later by other operating system functions in the remote NAS appliance. A file I/O specifies the file. It also indicates an offset into the file. For instance, the I/O may specify “Go to byte ‘1000’ in the file, (as if the file was a set of contiguous bytes), and read the next 256 bytes beginning at that position”. Unlike block I/O, there is no awareness of a disk volume or disk sectors in a file I/O request. Inside the NAS appliance the operating system keeps tracks of where files are located on disk. The OS issues a block I/O request to the disks to fulfill the file I/O read and write requests it receives. In summary, the network access methods, NFS, CIFS, and Netware, can only handle File I/O requests to the remote file system. This is located in the operating system of the NAS device. I/O requests are packaged by the initiator into TCP/IP protocols to move across the IP network. The remote NAS file system converts the request to block I/O and reads or writes the data to the NAS disk storage. To return data to the requesting client application the NAS appliance software re-packages the data in TCP/IP protocols to move it back across the network. A database application which is accessing a remote file located on a NAS device, by default, is configured to run with File System I/O. It cannot utilize a ‘raw I/O’ to achieve improved performance.

N A S uses File I/O

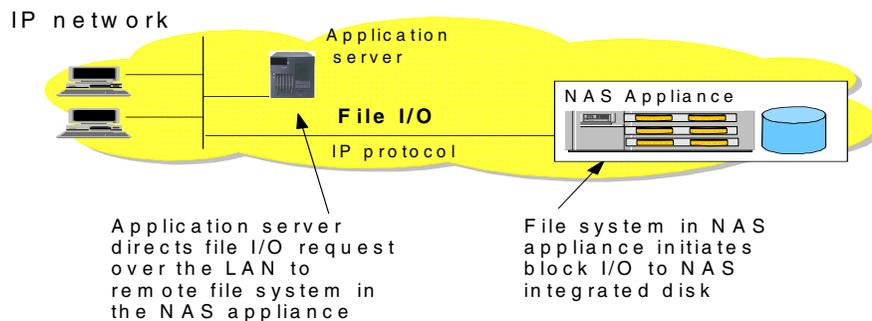


Figure 1-2 File I/O example

NAS offers a number of benefits, which address some of the limitations of directly attached storage devices and overcome some of the complexities associated with SANs.

- ▶ **Resource pooling:** A NAS appliance enables disk storage capacity to be consolidated and pooled on a shared network resource, at great distance

from the clients and servers which will share it. Therefore, a NAS device can be configured as one or more file systems, each residing on specified disk volumes. All users accessing the same file system are assigned space within it on demand. This contrasts with individual DAS storage, when some users may have too little storage, and others may have too much. Consolidation of files onto a centralized NAS device can minimize the need to have multiple copies of files spread across distributed clients. Therefore, overall hardware costs can be reduced. NAS pooling can reduce the need physically to reassign capacity among users. The results can be lower overall costs through better utilization of the storage, lower management costs, increased flexibility, and increased control.

- ▶ **Exploits existing infrastructure:** Because NAS utilizes the existing LAN infrastructure, there are minimal costs of implementation. Introducing a new network infrastructure, such as a Fibre Channel SAN, can incur significant hardware costs. In addition new skills must be acquired, and a project of any size will need careful planning and monitoring to bring to completion.
- ▶ **Simple to implement:** Because NAS devices attach to mature, standard LAN infrastructures, and have standard LAN addresses, they are, typically, extremely easy to install, operate and administer. This plug and play operation results in low risk, ease of use, and fewer operator errors, so it contributes to a lower cost of ownership.

1.2.3 iSCSI

As mentioned in 1.1, “What Internet SCSI (iSCSI) is, in a nut shell” on page 2, iSCSI is a network transport protocol for SCSI that operates on top of TCP. iSCSI encapsulates SCSI protocols into a TCP/IP frame, so that storage Controllers can be attached to IP networks.

i S C S I u s e s B l o c k I / O

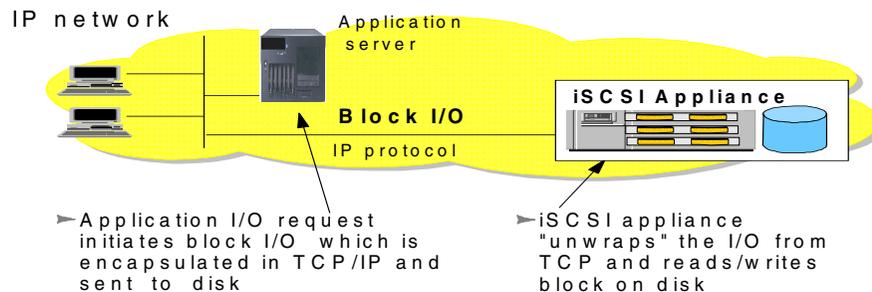


Figure 1-3 iSCSI uses block I/O

1.2.4 iSCSI benefits

The IBM TotalStorage IP Storage 200i appliance offers a number of benefits, and it can be viewed as a complementary solution to the other storage networking implementations already discussed.

Connectivity

iSCSI can be used for DAS or SAN connections. iSCSI capable devices could be placed on an existing LAN (shared with other applications) in a similar way to NAS devices. Also, iSCSI capable devices could be attached to a LAN which is dedicated to storage I/O (in other words an IP SAN), or even on a LAN connected to only one processor (like a DAS). These options are shown in Figure 1-4.

iSCSI Storage

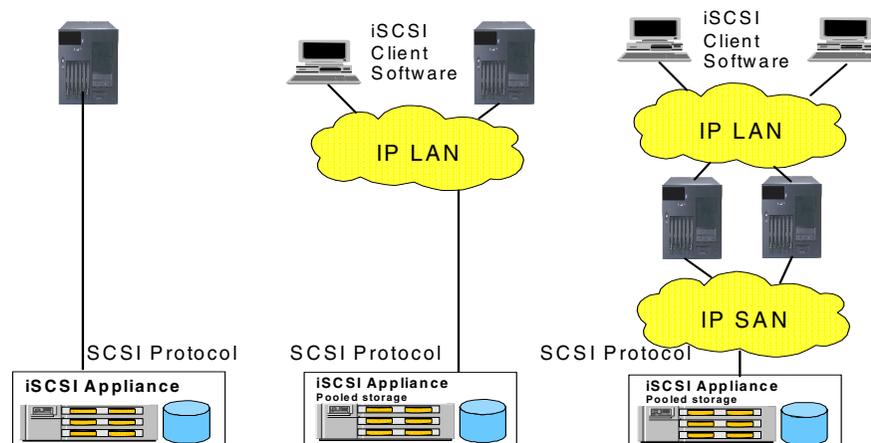


Figure 1-4 iSCSI appliance connection options

Extended distance

IP networks offer the capability easily to extend beyond the confines of a LAN, to include Metropolitan and Wide Area Networks (MANs and WANs). This gives greater flexibility, and at far less cost and complexity, compared to the inter connection of Fibre Channel SANs over wide areas.

Media and network attachments

iSCSI and NAS devices both attach to IP networks. This is attractive compared to Fibre Channel, because of the widespread use of IP networks. IP networks are already in place in most organizations and are supported by existing IT skills. TCP/IP-based networks can potentially support much longer distances than can pure Fibre Channel SANs, promising the possibility of scaling IP SANs to Storage Wide Area Networks (SWAN). Purchasing of attachment technologies is simplified, as they are the same as for LAN attachments.

Interoperability

The well-known, early-life interoperability problems of devices on Fibre Channel SANs disappear on networks using the familiar TCP/IP protocol.

Backup

Backup of data on the IP Storage is the same as for any direct attach storage; that is, via any method that supports SCSI-attached volumes. A backup application running on an external server, including the one hosting the iSCSI initiator code, will control the backup of data that is physically stored on the iSCSI appliance. In the future IBM plans to include more embedded storage management routines in the IP Storage 200i, including enterprise class solutions like Tivoli Storage Manager, and other popular industry solutions.

A NAS appliance, because it “hides” disk volumes from its clients, and often includes specialized backup facilities, may be easier to install and manage.

Management

iSCSI is managed like any direct-attach SCSI device. iSCSI-connected disk volumes are visible to attached processors. Compared to Fibre Channel SANs, iSCSI benefits from using IP networks for which there are established network management tools and people skills, such as Tivoli Netview or HP Openview. Such tools enable network administrators to coordinate provision of bandwidth among users and applications, traffic management, and overall network operations. Training in new networking skills is minimized.

Fibre Channel SANs currently have more storage-related management tools than iSCSI, such as support for tape sharing for backup. This advantage is likely to diminish as iSCSI matures and the market demands SAN-like management for iSCSI devices.

Low Cost

Cost comparisons are difficult to generalize and will probably depend on particular products. An iSCSI SAN is likely to have lower costs than a Fibre Channel SAN. For example, iSCSI network hardware, such as Ethernet host adapters, are generally lower cost than Fibre Channel host adapters. If iSCSI (or NAS) is attached to an existing LAN, no new host adapter cards may be needed at all. An iSCSI SAN can be built more quickly and with fewer new skills than a Fibre Channel SAN. An iSCSI disk device, all else being equal, may be lower cost than a NAS appliance since the iSCSI device does not need to support file systems, file sharing protocols, and other facilities often integrated into NAS products.

1.2.5 Considerations to make

Since the iSCSI appliance attaches to the existing Ethernet network, implementation is very similar to a NAS device. Administration and implementation are accomplished much more quickly than in a SAN environment. In this way, NAS and iSCSI are very similar. However, when you make the comparisons between iSCSI, SAN and NAS, it really comes down to Block I/O versus File I/O. Both iSCSI and SAN use Block I/O to transport data, whereas NAS uses File I/O. The determining factor might ultimately be what kind of application is to be used. Will the application work better in a Block I/O environment, or File I/O environment? Is the application a database or file sharing oriented? How much does performance play in this application?

All File I/Os result at a lower layer in Block I/Os. This means that iSCSI can also support File I/O applications. However, if “visibility” and “sharing” of files are needed, File I/O applications are better supported by NAS than by iSCSI devices.

SAN offers better performance, but is more expensive and requires a higher skill set to implement. iSCSI and NAS offer better pricing and skills may already be in place to implement them. However, both SAN and iSCSI offer the performance benefit of Block I/O. These are all considerations that must be taken into account when making the decision to add storage to the current environment.

Network congestion

In the case of NAS, and backup over IP networks, congestion may cause variable performance for other applications. This also remains a problem for iSCSI implementations. This should be alleviated with high speed Gigabit Ethernet, and largely overcome or masked, in time by the delivery of higher speed network transmission, such as 10 GB/sec.

Performance

A performance comparison is difficult to generalize because there are so many variables. That said, Fibre Channel at 100 MBps (1 Gigabit/second) is generally more efficient for I/O traffic than TCP/IP over Ethernet at equivalent bandwidth. iSCSI performs better than NAS (when both are on Ethernet) due to reduced protocol overhead. This is because it handles SCSI directly, rather than translating between file-I/O protocols and SCSI. This makes the IP Storage 200i appliance suitable for a variety of applications which might not deliver optimum performance on a NAS appliance. For instance, initial results indicate that database performance of IBM's IP Storage 200i is superior using raw I/O, compared to the equivalent database function executing through a file system.

TCP/IP is a software intensive network design which requires significant processing overheads. These can consume a substantial proportion of available processor cycles when handling Ethernet connections. This is a drawback for performance intensive storage applications.

A performance consideration, therefore, is the impact of the software protocol stack on processor utilization. Fibre Channel SANs support SCSI commands mapped directly to Fibre Channel media, and processor overhead for this mapping is low. In iSCSI, handling of the TCP/IP protocol requires processor cycles at both ends. Therefore, at this early time in the evolution of iSCSI, it is best suited for situations of relatively low I/O activity. This point generally applies to NAS as well. ("Low" in this case can be thousands of I/Os per second, but will be less than the very high performance levels which a Fibre Channel SAN could support.)

The IP Storage 200i, prior to general availability, is achieving its initial throughput targets. We recommend that early installations of the 200i are best suited to applications with low to moderate bandwidth requirements under 30 MBps throughput. Performance testing and tuning is in progress. Overall throughput is expected to achieve 60MB - 75MB throughput per second in later releases of microcode.

Security

An argument in favour of Fibre Channel SAN's is that data moves over a separate, secure and dedicated network. Many IT managers would have serious reservations about running mission critical corporate data on an IP network which is also handling other traffic. iSCSI introduces the possibility of an IP network SAN, which could be shared. If a separate IP network is implemented to minimize this risk, the cost advantage of IP SAN would be reduced.

Product maturity

iSCSI is very early in its development cycle. Industry standards are under discussion in the key industry organizations, and may take some time to emerge. Some IT managers will want to “wait and see” how things develop over the coming months. But, everything has a beginning. The advantages seem to be significant, and industry analysts have been extremely positive about the direction which IBM has taken with IP network solutions. The IP Storage 200i series solutions offer low entry points for IT departments to introduce open IP storage network solutions in order to test and establish the principles, and prepare for the future.

1.2.6 When might this technology be used?

The following applications for iSCSI are contemplated:

- ▶ Local storage access, consolidation, clustering and pooling (as in the data center)
- ▶ Client access to remote storage (for example, a “storage service provider”)
- ▶ Local and remote synchronous and asynchronous mirroring between storage controllers
- ▶ Local and remote backup and recovery

With these applications in mind, the IBM TotalStorage IP Storage 200i will be well suited for departments and workgroups within large enterprises, mid-size companies, and service providers, such as Internet Service Providers (ISP), and e-business organizations.

1.3 iSCSI components

iSCSI consists of several parts to provide a total solution. Figure 1-5 shows the different pieces that may appear in an iSCSI environment.

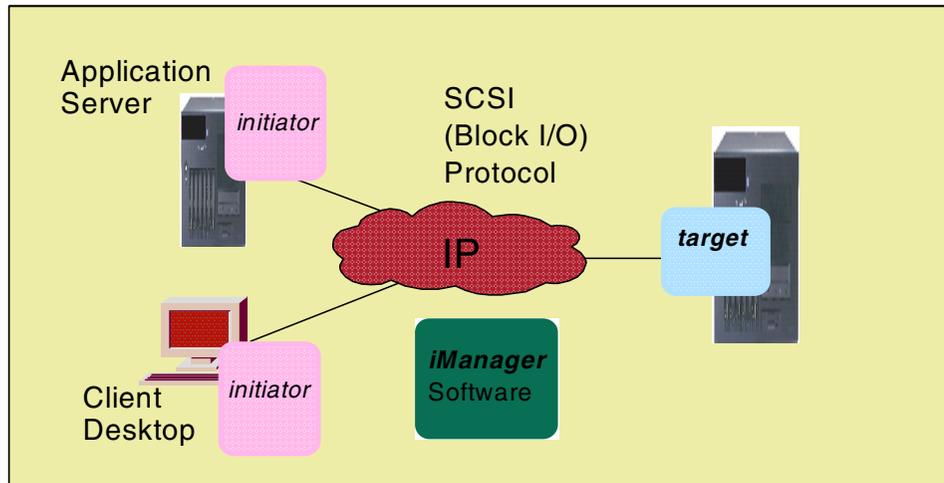


Figure 1-5 iSCSI components

- ▶ **iSCSI Initiator:** This consists of software drivers that initiate the SCSI request over IP to a target server. These drivers will coexist with existing SCSI devices without disruption. Current IBM drivers support Windows NT/2000 and Linux.
- ▶ **iSCSI Target:** This software manages the SCSI over IP requests on the target device.
- ▶ **iSCSI iManager:** The iManager software sits outside the datastream providing TCP/IP fabric services to aid device discovery, storage policy, LUN virtualization (partitioning, mapping, volume management).

The software to support iSCSI is provided with the IBM TotalStorage IP Storage 200i. Currently the initiator code is provided as a driver for the client workstations. As the product matures, iSCSI host bus adapters will be available to provide the iSCSI initiator code.

1.3.1 iSCSI in detail

We have discussed in general terms what iSCSI is suppose to be and how it compares to SAN and NAS. Now lets take a little closer look at some technical aspects of iSCSI.

1.3.2 iSCSI technical introduction

As we have mentioned more than once, iSCSI is the mapping of a SCSI remote procedure invocation model on top of the TCP protocol. A SCSI message is encapsulated inside the TCP frame and is called an “iSCSI protocol data unit”, or iSCSI PDU. iSCSI protocol is based on version three of the SCSI (SAM-2) protocol standard.

1.3.3 How does iSCSI keep track?

Command numbering is session wide and provides the means for ordered commands. Status numbering is per connection and aids recovery from connection failures. Data sequencing is per command and detects missing data packets.

Fields in the iSCSI PDU contain numbering information. With unidirectional traffic, special NOP-message PDUs may be utilized to synchronize communications between client and server. The status number is started after login. Data sequencing is required since multiple client-server paths may be used.

1.3.4 iSCSI discovery

Discovery allows an initiator to find the target(s) to which it has access. This requires a minimum of user configuration. Several methods of discovery may be used:

- ▶ A list of targets at the initiator
- ▶ Queries to known iSCSI server
- ▶ Queries to Storage Name Server (SNS)
- ▶ Use of the Service Location Protocol (SLP) as a type of “multicast” mechanism

1.3.5 Starting iSCSI (iSCSI login)

The purpose of login is to establish a TCP connection for iSCSI to use. This process authenticates the iSCSI end stations. IBM’s implementation of this process uses ID and password, because in a DHCP environment, the IP address is dynamic and could change from one boot to the next. The session parameters are set and a security protocol association is created. A TCP connection is mapped to an iSCSI session. iSCSI connections use layer 5 of the OSI stack. See Figure 1-6.

The OSI reference model Seven Layers

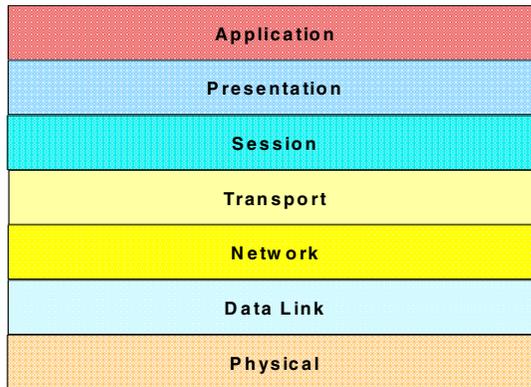


Figure 1-6 The seven layers of the OSI reference model

The Session layer establishes the initial logging on to a transmission between two nodes, plus the security and final termination of the session.

1.3.6 iSCSI full feature phase

After a successful login process, normal SCSI commands, data and messages can flow. Data and status for a given command must be sent over the same TCP connection. This is called connection allegiance and is based on port and IP number. Ordered command sets may use a different connection.

1.3.7 iSCSI data flows

Outgoing data (initiator to target) can be either solicited or unsolicited. Solicited data must be sent in response to Ready to Transfer PDUs (R2T) from the target. An initiator can send unsolicited data as part of a command or in a separate PDU. All subsequent data PDUs have to be solicited.

Figure 1-7 shows the iSCSI data packet frame.

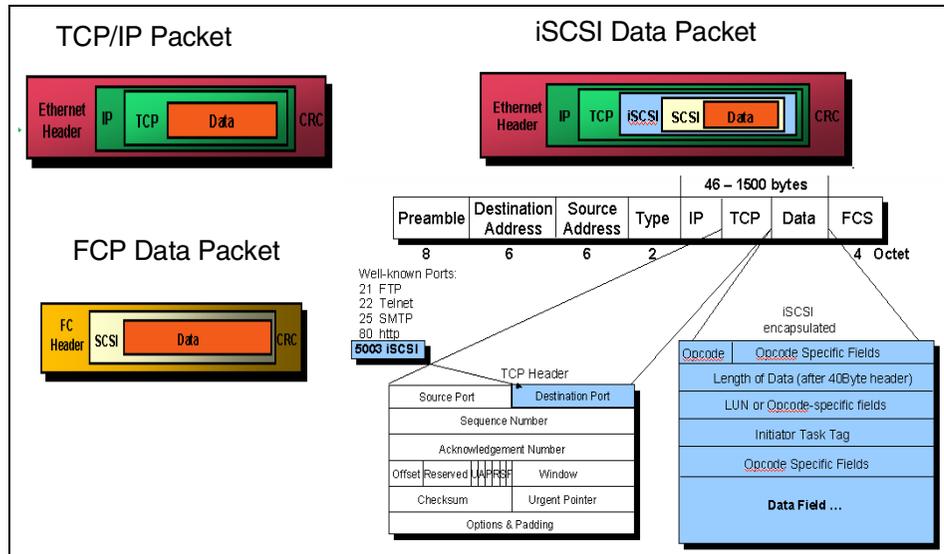


Figure 1-7 Data encapsulation of iSCSI frame

1.3.8 iSCSI connection termination

Connection terminations are exceptional events. Graceful connection shutdowns occur when there are no outstanding tasks with allegiances to the connection. If there are outstanding tasks, or tasks that have not yet sent status information, recovery actions may be needed.

Connection termination is also required as a prelude to recovery. Connection termination can avoid old PDUs from previous connections. (Logout is sent by the initiator.) Targets request the initiator to issue a Logout through an Asynchronous Event PDU.

1.3.9 iSCSI naming

Every initiator and target will be identified by its World-Wide Unique Identifier (WWUI) name. The WWUI name is not to be confused with an address. Targets are allowed to have a generic WWUI called “iSCSI”. The WWUI is used as a part of an iSCSI address.

1.3.10 iSCSI addressing

The iSCSI address has the format of:

<IP address>[:<port>]/<wwui>

The IP address can be either IPv4, IPv6, or the fully qualified domain name. The <port> is optional; it specifies the TCP port that the target is listening for connections on. If it is not used, the well-known iSCSI port is assumed. The <wwui> is the WWUI of the device.

The iSCSI address specifies a single path to an iSCSI target. The iSCSI address is primarily used during discovery. It ensures that the SCSI PDU gets delivered to the correct agent and does NOT effect any SCSI protocol specific addresses.

1.3.11 iSCSI message synchronization

iSCSI PDUs are of varying length, but TCP does not have a mechanism to identify these message boundaries within the TCP layer. iSCSI places message length information in the iSCSI header to delineate the end of the current message and the start of the next message. A potential problem still exists if iSCSI PDUs are delivered out of sequence or are missing.

There are several schemes that can be used to delineate the starting points of various messages. One of the easiest to implement is by using fixed interval markers. A special marker frame with a 32-bit pointer to the next iSCSI PDU start is used.

1.3.12 iSCSI synch and steering

An optional layer can be inserted between the iSCSI and TCP layers called “synch and steering”. This layer retains the ending address for every delivered iSCSI PDU. It includes the minimum header information as part of its encapsulation process. The information is used to deliver PDU contents to a final address and aid with recovery operations if TCP frames are lost during transmission.

1.3.13 The IBM TotalStorage IP Storage 200i

In February 2001 IBM announced the IBM TotalStorage IP Storage 200i, which was available at the end of June 2001.

This is a network appliance that uses the new iSCSI technology. The IP Storage 200i appliance solution includes client initiators. These comprise client software device drivers for Windows NT, Windows 2000 and Linux clients. These device drivers co-exist with existing SCSI devices without disruption. They initiate the iSCSI I/O request over the IP network to the target IP Storage 200i. The IP Storage 200i target appliance has both iSCSI target code and embedded storage

up to 3.52TB. IBM plans to add additional clients in response to customer feedback and market demands. IBM is committed to support and deliver open industry standard implementations of iSCSI as the IP storage standards in the industry are agreed.

Processors (*initiators*) supporting iSCSI can attach to IP Storage 200i over a TCP/IP network, such as (but not necessarily limited to) an Ethernet LAN. Attachment might be directly to the server or storage, or might be via a protocol converter at either end. The initiators are device drivers that intercept the low level SCSI commands, redirect them via TCP/IP encapsulated command, and map the SCSI command directly to its embedded storage.]

The IBM IP Storage 200i is a low cost, easy to use, native IP-based storage appliance. It integrates existing SCSI storage protocols directly with the IP protocol. This allows the storage and the networking to be merged in a seamless manner. iSCSI-connected disk volumes are visible to IP network attached processors, and as such are directly addressable by database and other performance oriented applications. The native IP-based 200i allows data to be stored and accessed wherever the network reaches, LAN, MAN or WAN distances.

Two options for attachment exist. You may choose to integrate the 200i directly into your existing IP LAN, combining storage traffic with other network traffic. This is a low cost solution for low activity storage applications. The alternative is SAN. Servers attach only to storage devices on the dedicated IP SAN. It acts as an extra network behind the servers, while the LAN in front of the servers remains dedicated to normal messaging traffic.

IBM TotalStorage IP storage 200i, comprises the 4125 Model 110 and 4125 Model 210 tower systems, and the 4125 Model EXP rack mounted system.

These are high-performance storage products, and they deliver the advantages of pooled storage, which FC SANs provide. At the same time they take advantage of the familiar and less complex IP network fabric.



IBM TotalStorage IPStorage 200i

In this chapter we discuss, in detail, the hardware that makes up IBM's entry into the iSCSI arena. The IBM TotalStorage IP Storage 200i is built around the xSeries 230 server. The operating system used is Linux Redhat 6.2.

We also look at the environment we used to develop this redbook. We will cover the network setup and discuss some recommendations and considerations for implementing iSCSI in your environment.

We cover the setup procedure for the 200i as well as step-by-step instructions for setting up the initiator on Windows NT and 2000.

In later chapters we discuss the Cisco SN5420 iSCSI Router and how it is installed in our network. You will also find details on the database implementations that we used.

2.1 IBM TotalStorage IP Storage 200i series

This is a network appliance that uses the new iSCSI technology. The IP Storage 200i appliance solution includes client initiators. These comprise client software device drivers for Windows NT, Windows 2000 and Linux clients. These device drivers co-exist with existing SCSI devices without disruption. They initiate the iSCSI I/O request over the IP network to the target IP Storage 200i. The IP Storage 200i target appliance has both iSCSI target code and embedded storage up to 3.52TB. IBM plans to add additional clients in response to customer feedback and market demands. IBM is committed to support and deliver open industry standard implementations of iSCSI as the IP storage standards in the industry are agreed.

The IBM IP Storage 200i is a low cost, easy to use, native IP-based storage appliance. The 200i is designed for workgroups, departments, general/medium businesses, and solution providers that have storage area network requirements across heterogeneous clients. It integrates existing SCSI storage protocols directly with the IP protocol. This allows the storage and the networking to be merged in a seamless manner. iSCSI-connected disk volumes are visible to IP network attached processors, and as such are directly addressable by database and other performance oriented applications. The native IP-based 200i allows data to be stored and accessed wherever the network reaches, LAN, MAN or WAN distances.

Two options for attachment exist. You may choose to integrate the 200i directly into your existing IP LAN, combining storage traffic with other network traffic. This is a low cost solution for low activity storage applications. The alternative is SAN. Servers attach only to storage devices on the dedicated IP SAN. It acts as an extra network behind the servers, while the LAN in front of the servers remains dedicated to normal messaging traffic.

IBM Total Storage IP storage 200i, is comprised of the 4125 Model 110 and 4125 Model 210 tower systems, and the 4125 Model EXP rack mounted system. All required microcode comes preloaded, minimizing time required to setup, configure and make operational the IP Storage 200i. There are only two types of connections to make, which are attaching the power cord(s) and the Ethernet connection(s) to the network. High speed, 133 MHz SDRAM is optimized for 133 MHz processor-to-memory subsystem performance. IBM IP Storage 4125 Model 110 and IP Storage 4125 Model 210 use the ServerWorks ServerSet III LE (CNB3.OLE) chipset to maximize throughput from processors to memory, and to the 64-bit and 32-bit Peripheral Component Storage (PCI) buses.

These are high-performance storage products, and they deliver the advantages of pooled storage, which PC SANs provide. At the same time they take advantage of the familiar and less complex IP network fabric.

After power on, the initial IP address configuration is a straightforward task which would be completed by the system administrator. The IBM TotalStorage IP Storage 200i provides a browser-based interface with which the system administrator can configure the network easily. RAID provides enhanced disk performance while minimizing storage failure. Adding disks and administering operations can occur while the system is online, providing excellent operational availability.

IBM will provide iSCSI initiator drivers for Linux, Windows NT, and Windows 2000. These drivers are available for download from the following Web site:

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

IBM will provide a user ID and password to authorized customers and users. The download package will extract all files, including a README, which explains how to build the initiator for particular hardware types and Linux versions. The Windows NT and 2000 install packages run under Install Shield, which will install drivers and update the registry. The information provided will explain how to configure the IP address of the iSCSI target. Once installed and configured (assuming the system administrator assigns access to storage for the initiator machine), the iSCSI initiator driver will open a connection to the iSCSI target on startup and will treat the assigned storage just like a locally attached disk. This is an important concept and has implications which will be discussed in 2.2, “Implementation and setup” on page 29.

2.1.1 IBM TotalStorage IP Storage 200i configurations

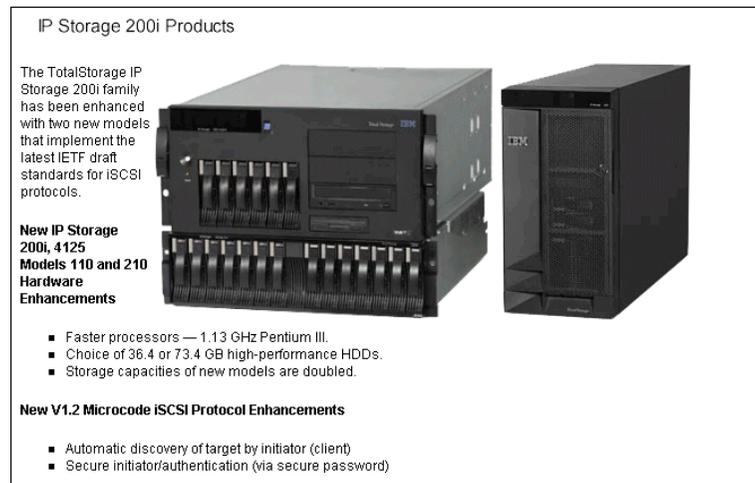


Figure 2-1 IBM IP Storage 200i

The workgroup model, IP Storage 200i, 4125 Model 110, is a compact tower design. It consists of the following components:

- ▶ One 1.13 GHz Pentium III Processor
- ▶ 512 KB Level 2 cache
- ▶ 512 MB of ECC 133 MHz System Memory
- ▶ ServeRAID-4MX - two channel RAID adapter
- ▶ 3/109GB of HDD Storage, expandable up to 6/440GB internal
- ▶ Three 250W, hot-swappable power supplies

For the departmental model, IP Storage 200i, 4125 Model 210, it is rack mounted and consists of the following components:

- ▶ Dual 1.13 GHz Pentium III Processors
- ▶ 1GB of ECC 133 MHz System Memory
- ▶ 512 KB Level 2 cache per processor
- ▶ ServeRAID-4H - high function, four-channel RAID adapter
- ▶ 3/109 GB of HDD Storage, expandable up to 6/440 GB internal
- ▶ 9/440 GB of HDD Storage, expandable up to 48/3.52 TB with three EXP units attached externally
- ▶ Three 250W, hot-swappable power supplies

The IBM TotalStorage IP Storage 200i 4125 Model EXP is a storage expansion unit that provides additional storage capability for the rack-based 4125. It provides up to 1.027 TB storage capacity per unit and up to three expansion units can be attached to a single 4125 Model 210, providing a maximum of 3.52TB of storage.

2.1.2 IBM TotalStorage IP Storage 200i technical specifications

Table 2-1 shows the technical specifications of the 200i model 110 and 210.

Table 2-1 Product specification 200i

Specifications	4125-110 Tower	4125-210 Rack
Form Factor	Tower	5U Rack drawer
Number of Processors (std./max)	1/2 1.13 GHz Pentium III	2/2 1.13 GHz Pentium III
Internal Disk	20 GB IDE	20 GB IDE

Specifications	4125-110 Tower	4125-210 Rack
L2 Cache	512 KB	512 KB Level 2 cache per processor
Memory (std./max)	512MB/1GB	1G/2GB
Expansion Slots	5	5
Capacity (std. max)	3/109GB, 6/440GB internal	3/109GB, 6/440GB internal, 9/440GB 48/3.52 TB with 3 EXP units attached externally
Network	10/100/1000Mbps or Intel Gigabit Fibre	10/100/1000 Mbps or Intel Gigabit Fibre
Integrated Advanced System Management Processor	Yes	Yes
Power Supply	3 x 250W Hot - Swap redundant power supplies	3 x 250W Hot - Swap redundant power supplies
Hot Plug Components	HDDs, Power supply's, fans	HDDs, Power supply's, fans
Light Path Diagnostics	Yes	Yes
Total PCI slots/ available	5 (4x64-bit and 1x32-bit)/ 1	5 (4x64-bit and 1x32-bit)/ 1
RAID Support	ServeRAID-4MX-Dual channel	ServeRAID-4H-4 channel

2.1.3 IBM TotalStorage IP Storage 200i microcode

The IBM Total Storage IP Storage 200i, 4125 Models 110 and 210, are preloaded with IETF Standard (V1.2) compliant machine code. This code is specifically designed to handle iSCSI initiators (clients) accessing varying amounts of data storage on the 4125 Models 110 and 210 appliances using block I/O requests.

The preloaded machine code includes:

- ▶ Linux OS optimized for iSCSI operations and functions
- ▶ iSCSI Target, iSCSI Initiator functions and drivers
- ▶ Gigabit Ethernet SX (Fibre) or Ethernet 10/100/1000 (Gigabit Copper) or NIC drivers
- ▶ ServeRAID-4 storage controller code

- ▶ ServeRAID Configuration and Monitoring — Web-based configuration tools and RAID Management using ServeRAID-4 controllers
- ▶ Web-based Configuration Utility — Manages the IBM Total Storage IP Storage 200i, 4125 Models 110 and 210, from a single client workstation

2.1.4 IBM TotalStorage IP Storage 200i features and profiles

The IBM TotalStorage IP Storage 200i serves out local storage to iSCSI initiator (client) requests. The IBM TotalStorage IP Storage 200i acts as a single server to the physical disk drives on behalf of the iSCSI clients. The physical disk drive storage is partitioned into multiple virtual logical units of storage (LUNs) for storage assignment to individual iSCSI clients.

The ability to access storage residing on the IBM TotalStorage IP Storage 200i is coordinated by access control logic in the Web-based User Interface (UI). iSCSI clients use an assigned client ID and password to access assigned LUNs.

Internal IP-Storage appliance system functions are integrated on top of a base Linux core (kernel level 2.4). Core architecture addresses initial IP addressing, product boot and recovery, and general box management.

Two methods are provided for “first boot” IP address assignment:

- ▶ The Ethernet NIC will default to address 192.9.200.1. By initially installing this product in a “private network”, where this address is reachable by a workstation having the same sub-net address, the administrator can point the workstation's browser to 192.9.200.1:1959 to access the user configuration.
- ▶ ARP (Address Resolution Protocol) find is a process listening on TCP port 3939. The function monitors the local network for “unanswered ARPs”. By pointing the administrator's browser to 3939, the administrator will cause an ARP into the network, where the 4125-200i has been installed. After seeing that no station is responding to the ARP, the 4125-200i will install the desired IP address and will respond to the ARP. If within 20 seconds a packet is received at port 3939, the 4125-200i will keep the address. At that point, the administrator can point the browser at the same address port 1959 to access user configuration panels.

The system disk is partitioned for multiple system images for upgrade and recovery. The system is booted from the primary partition. If the boot fails, the system is automatically booted from the recovery CD-ROM, which invokes failure recovery procedures. Through the service interface, the user can apply new system images from a local management station.

Network management is supported via SNMP and standard MIBs. SNMP agents and subagents support internal functions. A specific iSCSI MIB is not supported in this initial product release.

An Ethernet device driver supports Gigabit Ethernet SX (Fibre) or Ethernet 10/100/1000 (Gigabit Copper) connectivity.

The RAID levels supported are RAID 0, 1, 1E, 5, and 5E. Disk partitioning and management, as well as RAID arrays, are supported. Hot-spare disks can be defined for automatic failed disk replacement (with the exception of RAID 0).

2.1.5 IBM IP Storage high availability and serviceability

IBM TotalStorage IP Storage 200i delivers economical reliability and serviceability:

- ▶ Six hot-swap HDD bays with SCA-2 connectors — support SAF-TE functions
- ▶ Standard ServeRAID-4H or 4MX controllers — support RAID levels 0, 1, 5, 1E, 5, and 5E
- ▶ ECC DIMMs combined with an integrated ECC memory controller — corrects soft and hard single-bit memory errors, while minimizing disruption of service to LAN clients
- ▶ Memory hardware scrubbing — corrects soft memory errors automatically without software intervention
- ▶ ECC L2 cache processors — ensures data integrity while reducing downtime
- ▶ Three worldwide, voltage-sensing 250-watt power supplies — auto restart and redundancy
- ▶ Information LED panel — visual indications of system well-being
- ▶ Easy access to system board, adapter cards, processor, and memory

2.1.6 IBM IP Storage expandability and growth

The IP Storage 4125 Model 110 and IP Storage 4125 Model 210 mechanical packages are available in tower or rack models. The rack model is engineered to meet the compactness of a 5 U rack drawer. A conversion kit (feature #3601 5600 Tower-to-Rack Kit) is available to convert a tower mechanically for rack mounting. It features the following:

- ▶ Standard 512 MB (Model 110) or 1 GB (Model 210) of system memory — expandable to 2 GB
- ▶ Five full-length adapter card slots — 5 (4x64-bit and 1x32-bit)

- ▶ ServeRAID-4MX Ultra160 SCSI Controller (IP Storage 4125 Model 110) — dual channel supports internal RAID storage
- ▶ ServeRAID-4H Ultra160 SCSI Controller (IP Storage 4125 Model 210) — four channels support internal and three external channels
- ▶ Ten drive bays:
 - Six 3.5-inch slim-high, hot-swap drive bays, three 5.25/3.5-inch half-high device bays, and one 3.5-inch slim-high drive bay
 - Up to 440 GB of internal data storage
 - A 40x-17x IDE CD-ROM and 1.44 MB diskette drive

These servers have the flexibility to handle applications for today and expansion capacity for future growth.

2.1.7 IBM IP Storage 200i 4125-EXP expansion unit

This highly available external storage expansion unit is supported with the IP Storage 4125. It ships with three slim-high 10K-4 Ultra160 SCSI.

The 4125-EXP contains two hot-swap, redundant power supply/fan assemblies. Key features of the storage expansion unit include:

- ▶ Supports 14 slim-high HDDs, maximum capacity 1.027 TB (14 X 73.4 GB)
- ▶ Supports data transfer speeds of up to 160 MB
- ▶ 3U form factor for minimum rack space usage
- ▶ Accommodates single or dual SCSI bus configurations
- ▶ Dual hot-swap, 250 W redundant power supplies with integrated fan assemblies
- ▶ Includes two line cords and publications

2.1.8 IBM IP Storage 200i optional features

Here are the optional features for the IBM TotalStorage IP Storage 200i.

- ▶ **Gigabit Ethernet SX Adapter (#3302)**

The xSeries Gigabit Ethernet SX Adapter provides 1000BASE-SX connectivity to a Gigabit Ethernet network for servers over a 50 or 62.5 micron multimode fiber optic link attached to its duplex SC connector. Its 1000 Mbps data rate and 32- or 64-bit PCI bus mastering architecture enables the highest Ethernet bandwidth available in an adapter. It is compliant with IEEE 802.3z Ethernet and PCI 2.1 standards, ensuring compatibility with existing

Ethernet installations. It also supports 802.1p packet prioritization and 802.1q VLAN tagging. Either this feature (#3302-Ethernet SX Adapter) or feature #3303-10/100/1000 (Gigabit) Ethernet Copper Adapter must be selected for the IP Storage 4125 Model 110 and IP Storage 4125 Model 210.

▶ **10/100/1000 (Gigabit) Ethernet Copper Adapter (#3303)**

The Adapter delivers up to 1000 Mbps over existing Category 5 twisted pair cables. Existing Fast Ethernet cabling infrastructure can be used for Gigabit throughput. No re-cabling is necessary. Link speed auto sensing is supported, so the adapter can operate at 10, 100, or 1000 Mbps depending on the configuration of the switch or hub. This supports the migration to Gigabit without having to replace or reconfigure the adapter. Either this feature (#3303-10/100/1000 (Gigabit) Ethernet Copper Adapter) or feature #3302-Gigabit Ethernet SX Adapter must be selected for the IP Storage 4125 Model 110 and IP Storage 4125 Model 210.

▶ **512 MB 133 MHz ECC SDRAM RDIMM (#3403)**

This is memory for the IP Storage 4125 Model 110 and IP Storage 4125 Model 210. It can be used to increase memory to further tune performance for the intended environment. ECC SDRAM RDIMM is a special type of memory module which is recommended for use on servers.

▶ **1 GB 133 MHz ECC SDRAM RDIMM (#3404)**

This is memory for the IP Storage 4125 Model 110 and IP Storage 4125 Model 210. The IBM TotalStorage IP Storage 200i supports up to 2 GB when populated.

▶ **5600 Tower-To-Rack Kit (#3601)**

This may be used to rack-mount the IP Storage 4125 Model 110.

2.2 Implementation and setup

In preparing for this redbook, it was our intent to come up with several real world applications to use featuring the iSCSI hardware and software. Performance was not our primary consideration. To this end, we used both IBM and Cisco hardware to set up various environments. At the time of writing this redbook, the standards for iSCSI have not been finalized by the IETF as previously mentioned, so we tried some things that may not be necessary once the standard is established.

We used the following hardware to create our iSCSI LAN:

- ▶ IBM TotalStorage IP Storage 200i
- ▶ Cisco SN-5420 iSCSI Router

- ▶ WS-C3508G Catalyst 3500 Series XL GigE Switch
- ▶ WS-C2924M-XL-EN 24 port 10/100 W/2 GigE Slots

We did not use a router, although it should be considered in a normal setup, because it is recommended that the control port and data port of the IBM 200i and the Cisco SN-5420 data and maintenance ports are on different subnets. There would be no problem connecting the 200i and the Cisco 5420 maintenance ports on a “maintenance subnet” and the data ports on a different subnet. It is up to the user as to how this would be accomplished, but to avoid conflicts, this needs to be considered.

Our client stations on which we loaded our initiator code consisted of two IBM 760 Thinkpads and one Netfinity 7600 server. We tested it with Windows NT, Windows 2000 and Linux. One of the interesting discoveries we made during our testing is that we were able to use both the Cisco and IBM initiators on a single workstation simultaneously. As a result, we were able to access both VLUNS from the 200i as well as LUNs from the Enterprise Storage Server from the same desktop. Database applications were run from the Netfinity 7600 while the VLUNS and ESS LUNs were used for database storage. The Thinkpads were used as client machines accessing the database.

Figure 2-2 is a high level overview of our lab setup.

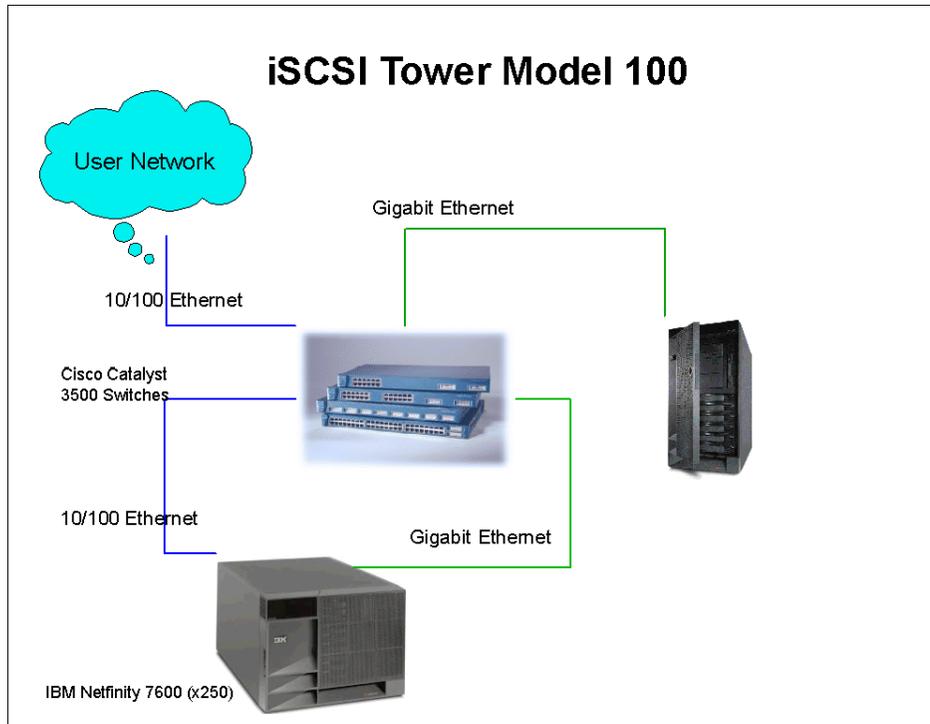


Figure 2-2 Lab configuration

The preceding diagram is a simplified layout of the lab we used to develop this redbook. It is not a recommendation of how an iSCSI setup may be accomplished. Through this testing we have been able to come up with ideas and recommendations for iSCSI do's and don'ts.

Each chapter contains hints and tips for successful iSCSI installation and implementation. In the next section we give some observations based on the general setup of iSCSI.

2.2.1 General hints and tips

Much of the material you will need to setup the IBM TotalStorage IP Storage 200i can be found in the *Installation and Administrator's Guide* that comes with the product. It is recommended that you become familiar with the material found there, if any questions arise while going through this book.

IP suggestions

The *Installation Guide* describes two different methods for setting up IP addresses for the 200i. We found that using the GUI accessible through either Netscape or Internet Explorer is the simplest.

The *Installation Guide* also indicates that management can be performed through either the control port or data port addresses. As mentioned earlier, it is recommended that the maintenance port and data port be on separate IP segments. Though management can be accomplished from either port, storage access should only be performed through the data port (GigE port). In 2.3, “Step-by-step installation guide for 200i” on page 38, we show how to set up and configure the 200i. This was accomplished through the control port address.

Management console

We recommend that if a management workstation is chosen with multiple NIC cards (multiple Ethernet or Token Ring cards or a mixture of each) that the routing functions are turned off. It is also a good idea to use different IP addresses that have different subnets so that any requests sent to the management port will not be misdirected to another subnet.

When accessing the browser for the purpose of management, turn off the proxy server information. Make sure your browser is using a direct connection. Proxy server information will interfere with access to the management GUI.

Figure 2-3 shows the proxy settings for using Netscape as your management GUI.

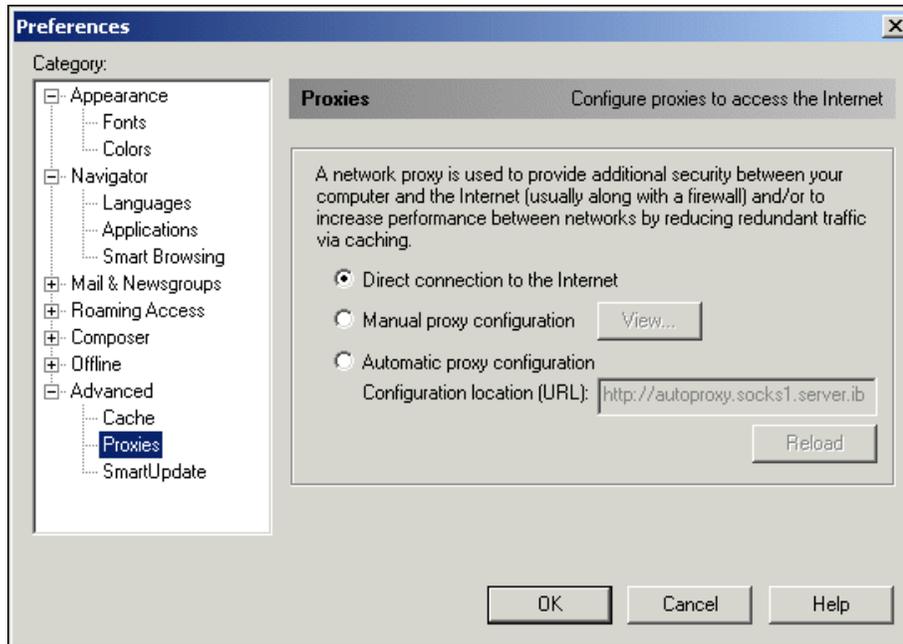


Figure 2-3 Netscape proxy setting

It is also possible to use Internet Explorer as your management GUI. Figure 2-4 shows how the proxy settings should appear when using Internet Explorer. Notice that there is nothing checked under Local Area Network Settings.

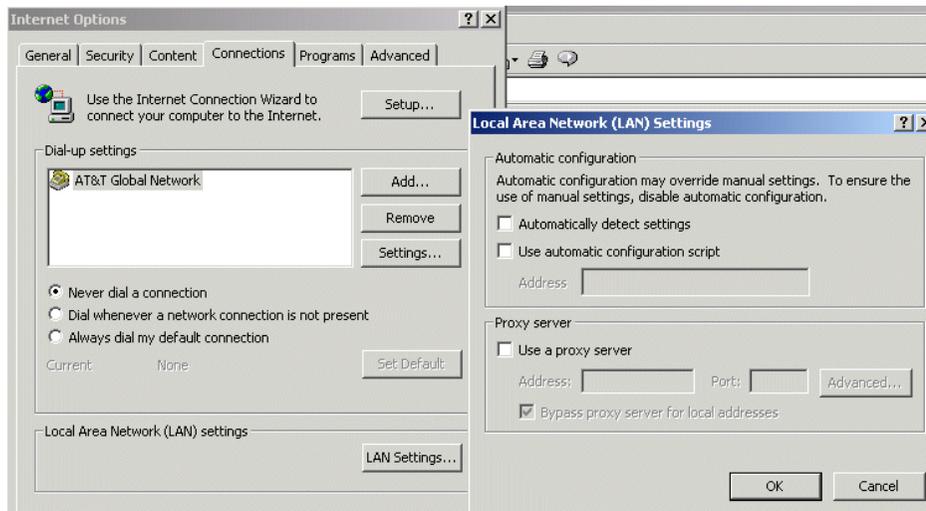


Figure 2-4 Internet Explorer proxy settings

The default address for the maintenance port of the 200i is 192.9.200.1. In order to set up the 200i using the GUI, you must first be sure you can ping that address from your management station; **ipconfig** can be used to check the IP address of your management station on Windows NT and 2000, and **ifconfig** can be used to check the IP address under Linux.

2.2.2 Disk drive addressing

We found it advisable to setup the iSCSI and 5420 drives toward the end of the alphabet. That way there would be no conflicts with drive letters in the beginning of the alphabet. Since iSCSI drives are actually network drives (though they appear as local drives to the user), they are added to the system later than the local drives. It is possible to overwrite a drive letter with a device and not see the iSCSI drives.

Another interesting anomaly that may or may not be seen is the timing relationship between defining the local drives and network drives on power up. Sometimes the iSCSI drives will be seen, but no drive letters will be assigned, because the system finishes its power up sequence before Disk Management assigns the drive letters. A common issue with the IBM iSCSI Windows client is that certain automatic services, such as those associated with popular database programs or the LanmanServer service, can start before the iSCSI drives are available. If these services depend on being able to access files on the iSCSI disk, and the disk is not yet available when these services start, then the services generate errors and manual user intervention is required.

The solution for this is to use a Microsoft tool that is provided in the Windows 2000 resource kit. The following files should be loaded into the **C:\WINNT\SYSTEM32** directory:

- ▶ autoexnt.bat
- ▶ instexnt.exe
- ▶ autoexnt.exe
- ▶ Sleep.exe

AutoExNt is a service which starts automatically each time the machine starts. As a service, it can do anything that an administrator of the machine can do. It can start and stop services or applications. It will work on both Windows NT and Windows 2000.

The **LanmanServer** service is responsible for restoring network shares on your computer. The **IBMiSCSIcfg** service is responsible for restoring your iSCSI target configuration after rebooting, but this service does not start until after the **LanmanServer** starts. The reason for this is because Microsoft system services always start before third-party services start (like **IBMiSCSIcfg**). Therefore, **LanmanServer** service tries to restore the shares on your iSCSI disks before the iSCSI disks are available.

The solution is to use the **AutoExNt** service to start the **IBMiSCSIcfg**, stop the **LanmanServer** service and restart the **LanmanServer** service, guaranteeing that any shares on the iSCSI disks will be restored. Here are the steps to perform this:

1. Install the **AutoExNt** service:
 - a. Copy the files `instexnt.exe`, `autoexnt.exe`, `sleep.exe` to the system 32 directory (usually `c:\winnt\system32`)
 - b. From a command prompt, change to the system32 directory and type:

```
instexnt.exe install
```

Note: `instexnt.exe install /interactive` will display bat file running and the program can be stopped from this DOS window.

2. Create a blank text file 'AutoExNt.bat' in the system32 directory — this file will tell the **AutoExNt** service what you want it to do.
3. From the 'Services' control panel applet:

- **Control Panel -> Services on Windows NT**
- **Control Panel -> Administrative Tools -> Services on Windows 2000**

Change the 'Startup type' setting from automatic to manual of these:

- **IBMiSCSIcfg**
- **server**
- **Computer browser**

Also change the Startup type setting from automatic to manual of any application-specific services that are affected. For information on affected services refer to the event viewer.

4. Edit the `AutoExNt.bat` file that was created in step 2. Here is where you start and stop services manually, as needed, and in the order that you specify. For this example (restoring shares to iSCSI disks) the batch file may look like this:

Example 2-1 Single iSCSI drive checking script

```
net stop "Computer browser" /yes  
net stop "server" /yes
```

```

:drive_offline
net start "IBMiSCSIcfg"

sleep 10

net start "server"
net start "computer browser"
net start "Distributed File System"

if exist g:\ goto iscsi_yes

net stop "IBMiSCSIcfg" /yes
net stop "computer browser" /yes
net stop "server" /yes

sleep 5

goto drive_offline

:iscsi_yes

net start "The required services 1"
net start "The required services 2"

```

Note: This Batch File will keep trying to reconnect the iSCSI box before starting any other services, as it checks for the availability of G: drive before starting any other services.

Important: There have been situations where Dynamic iSCSI drive does not go online after a system reboot. To overcome this problem, select **Reactivate Disk** under Disk Management, as shown in Figure 3-24 on page 95.

For checking multiple iSCSI drives the following script can be used:

Example 2-2 Multiple iSCSI drive checking script

```

net stop "Computer browser" /yes
net stop "server" /yes

:drive_offline
net start "IBMiSCSIcfg"

sleep 10

net start "server"

```

```
net start "computer browser"
net start "Distributed File System"

if exist g:\ goto iscsi_First_drive

:net_stop
net stop "IBMiSCSIcfg" /yes
net stop "computer browser" /yes
net stop "server" /yes

sleep 5

goto drive_offline

: iscsi_First_drive
if exist h:\ goto iSCSI_yes

goto net_stop

:iscsi_yes

net start "The required services 1"
net start "The required services 2"
```

Note: the 'sleep xx' above specifies a number of seconds to do nothing, which ensures that the IBMiSCSIcfg service has enough time to find and log into the target. It is also used between the net stop and net start, because sometimes the server service will crash if you stop it and then immediately restart it. The Browser service is also involved in this case, because it is dependent on the server service.

Tip: The drive does not have to be specified. A file on a specific drive can also be selected to prevent the services from starting in the case where a new drive was added and a wrong drive letter was allocated; add the file name to the command line:

```
if exist h:\file.txt goto.....
```

(This can be any file, even an executable or database file, because this part of the script just checks for the availability of the file, it does not open or launch the file.)

Important: Make sure that the AutoExNt service is set to automatic. Failure to do so will result in the script not running!

5. Reboot the computer.

If everything was set up correctly, the AutoExNt.bat file will be executed automatically by the Service Control Manager when the computer boots. There is no need to log into the computer for the AutoExNt.bat file to execute, because it runs as a system service.

2.3 Step-by-step installation guide for 200i

In this section we show how the 200i can be set up. Remember, for any questions not answered here, refer to the *Installation Guide* that comes with the product. It will cover some items that we will not discuss in this redbook.

2.3.1 Start up tips

Before attempting to set up the 200i there are some preliminary checks that need to be performed:

- ▶ Ensure that all power supplies are plugged in. The 200i comes standard with three 250 Watt power supplies. It will run on two power supplies for redundancy.
- ▶ Wait at least 20 seconds after plugging in the power before turning on the 200i. This gives the system-management processor time to initialize before applying power.
- ▶ Since the IP Storage 200i is considered an appliance, it does not require a display, keyboard, and mouse to run. Therefore, you will hear several beeps right after powering on the system. It is NORMAL to hear one long and three short beeps after pressing the power button.

Note: If you hear anything other than one long and three short beeps, there may be a system problem and you will need to attach a display, keyboard and mouse to see what the problem is. Power off the system, attach the components, and power the system back on.

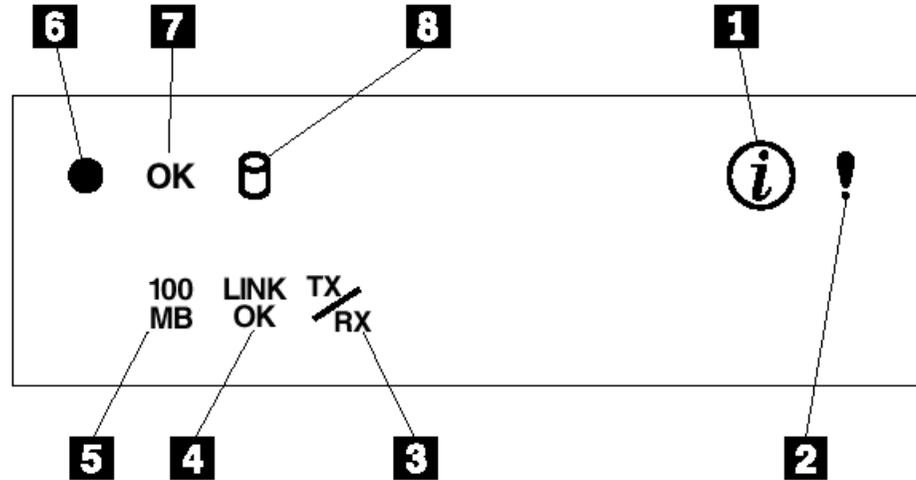
After attaching a display, keyboard, and mouse the beep pattern will change but diagnostic information will now be displayed in the window.

For specific troubleshooting information, please refer to “Appendix A” in the *Administrator’s Guide*.

- ▶ It will serve no purpose to continue to have a display attached to the 200i, There is no usable information that will be displayed other than to say that the Linux kernel is being uncompressed. During normal operation of the 200i all usable information will be displayed from the GUI.

Figure 2-5 shows the display panel lights and their functions known as light path diagnostics. For more information on the operating and function of the light path diagnostic, refer to “Troubleshooting” in the *Administration Guide*.

The Ethernet connection portion of the display refers to the Ethernet port on the system board only. If the indicators are not lit, the on-board Ethernet adapter is either not working or an Ethernet cable is not plugged in.



1. Information LED
2. System Error LED
3. Service Ethernet Transmit/Receive Activity (TX/RX) LED
4. Service Ethernet Link Status (LINK OK) LED
5. Service Ethernet Speed (100MB) LED
6. Power-on LED
7. POST Complete (OK) LED
8. SCSI Hard Drive In-use LED

Figure 2-5 200i operating panel

2.3.2 Setup of the 200i

Setup of the 200i is a relatively simple operation. However, as we have previously discussed, there are some actions you can take to ensure a smooth install. The default address of the control or maintenance port of the 200i is 192.9.200.1. All setup functions are performed through this port.

The easiest way to set up this appliance is to access the control port. Setup the workstation that will be used for maintenance to be able to ping this address. Once you can ping this address, access to the 200i is relatively easy.

Important: The administrations console requires:

- ▶ Java Swing 2
- ▶ Sun Java Plug-in 1.3 or later

The administration console will not operate properly if the required software is not installed. The software is available from:

<http://java.sun.com/getjava>

2.3.3 Logon to the appliance

Point the Web browser to the control port address (Figure 2-6):

192.9.200.1:1959/

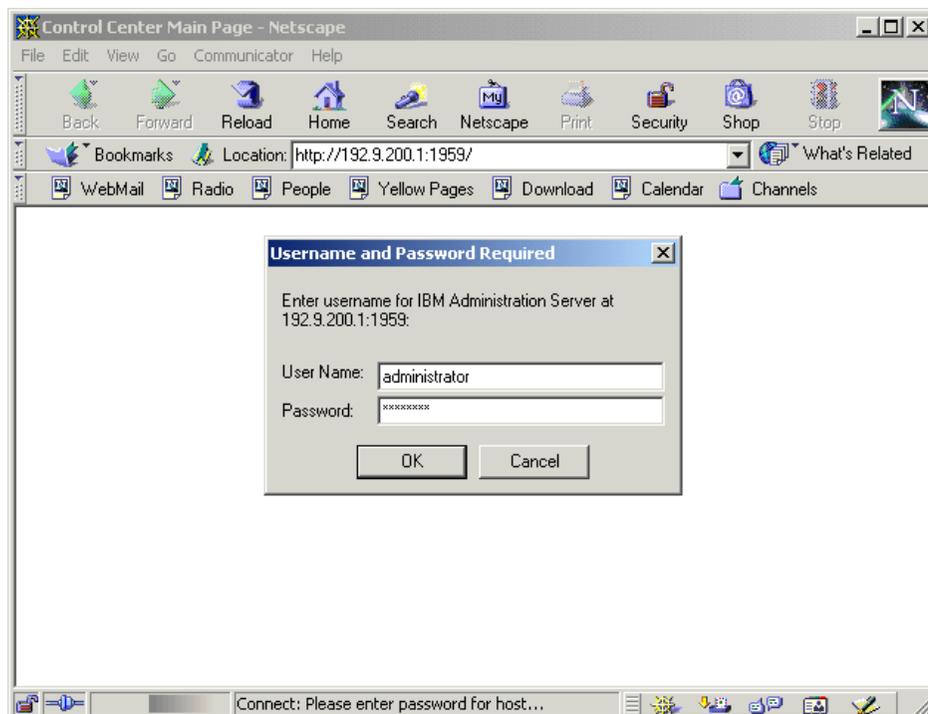


Figure 2-6 Log on window

- ▶ Enter the User Name as: administrator
- ▶ Enter the Password as: password

Note: The User Name and Password fields are case sensitive.

Figure 2-7 is the GUI Management window. From this window, you will be able to set the IP address, assign VLUNs and clients, review logs, reboot the system, and many other administrative tasks.



Figure 2-7 200i storage home page

2.3.4 Setting the service port IP addresses

The next steps will show how to set the service port and data port addresses.

- ▶ Select **Network** from the left hand side of the window (Figure 2-7).
- ▶ Select **Service port**.

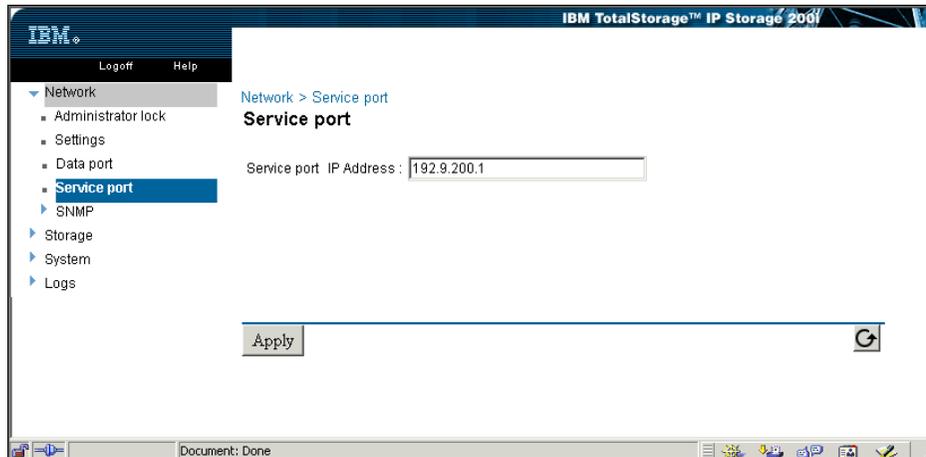


Figure 2-8 Network options

In this window, you are able to set a new IP address for the service port (or maintenance port). Simply key in your new address and press **Apply** in the lower left of your window.

Note: The system will tell you that a reboot will be required to apply the new IP address.

2.3.5 Rebooting the 200i

- ▶ Select **System**.
- ▶ In the window, select **Reset** (Figure 2-9).

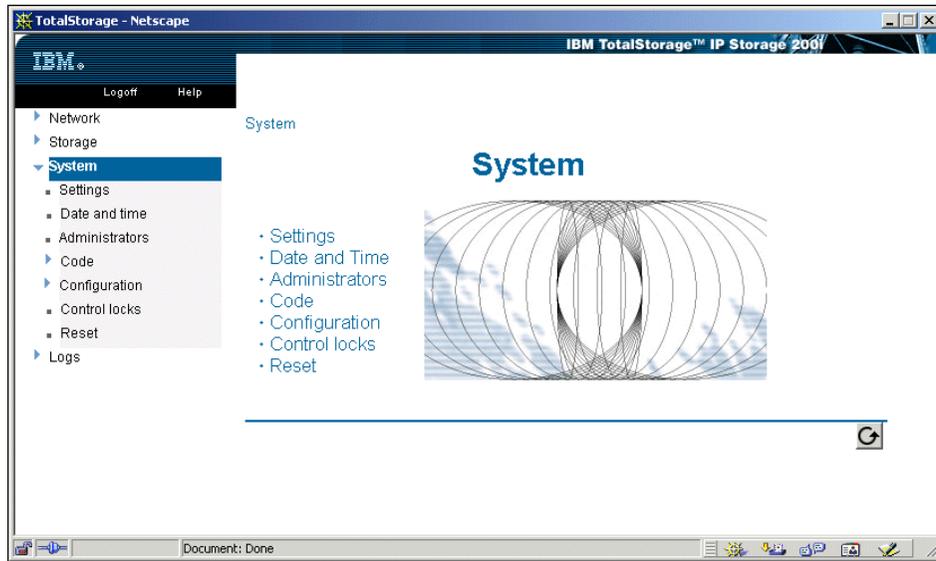


Figure 2-9 System options

- ▶ When **Reset** is selected, you will be presented with the next window (Figure 2-10). Select **Restart** to restart the 200i.

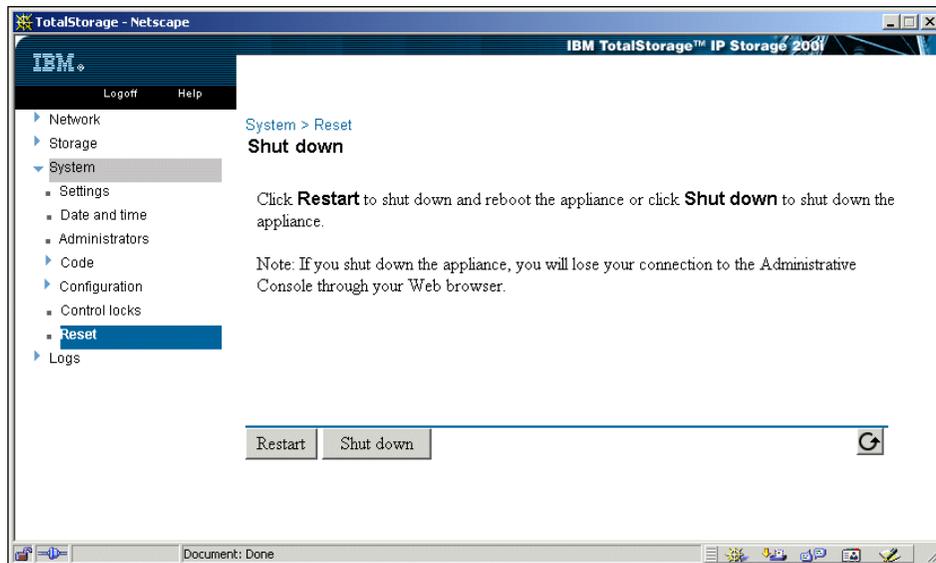


Figure 2-10 Reset warning

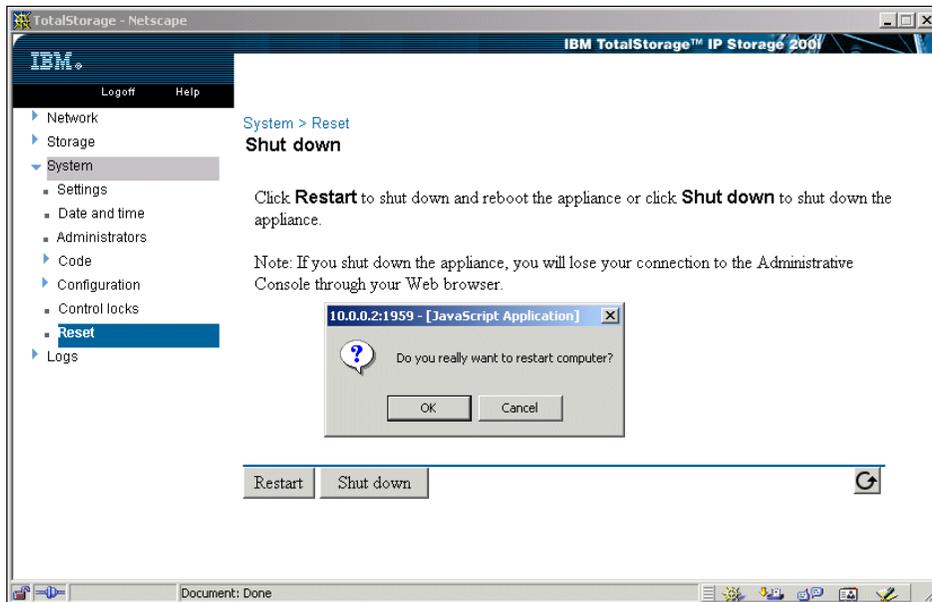


Figure 2-11 Reset final warning

You will be prompted, “Do you really want to restart the computer?”

- ▶ Select **OK**.



Figure 2-12 Reboot wait message

The system will show please wait while rebooting (Figure 2-12). This may take several minutes. Ping the control port address, once you receive a response, and you will be able to log back onto the management GUI.

Note: If you changed the control port address, you will now need to log on with your browser pointing it to the new address.

2.3.6 Setting the data port IP address

The next steps will show you how to change the data port address for the GigE Ethernet adapter card. This is the port you will use to access the storage over the network.

Note: You can also access the management GUI from this port, but you cannot access storage through the control port.

- ▶ From the home page as shown in Figure 2-7 on page 42, select **Network**.
- ▶ Select **Data Ports** (Figure 2-13).
- ▶ Click the select tick box for the adapter that will be edited.
- ▶ Click **Edit**.

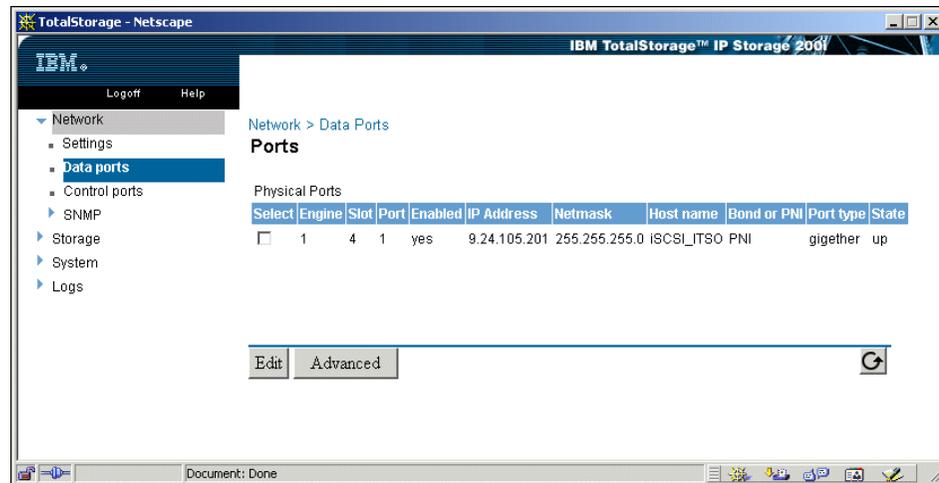


Figure 2-13 Data port selection

- ▶ Make the required changes and select **Apply** (Figure 2-14).

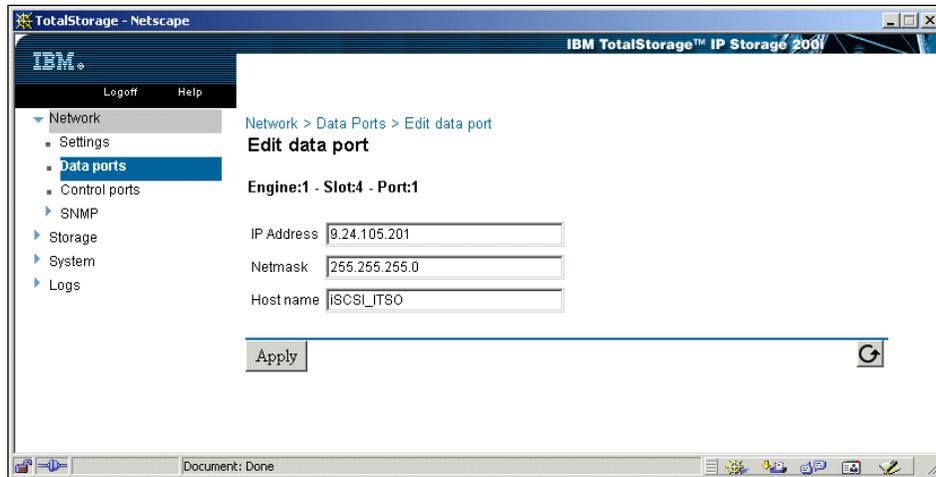


Figure 2-14 Edit data port option

Note: Remember a reboot will be required to activate the new IP address.

After Apply is selected, the system will display a message that activation is complete.

2.3.7 Setting Up VLUNs

From the home page window as shown in Figure 2-7 on page 42:

- ▶ Select **Storage**.
- ▶ Select **iSCSI Clients**.

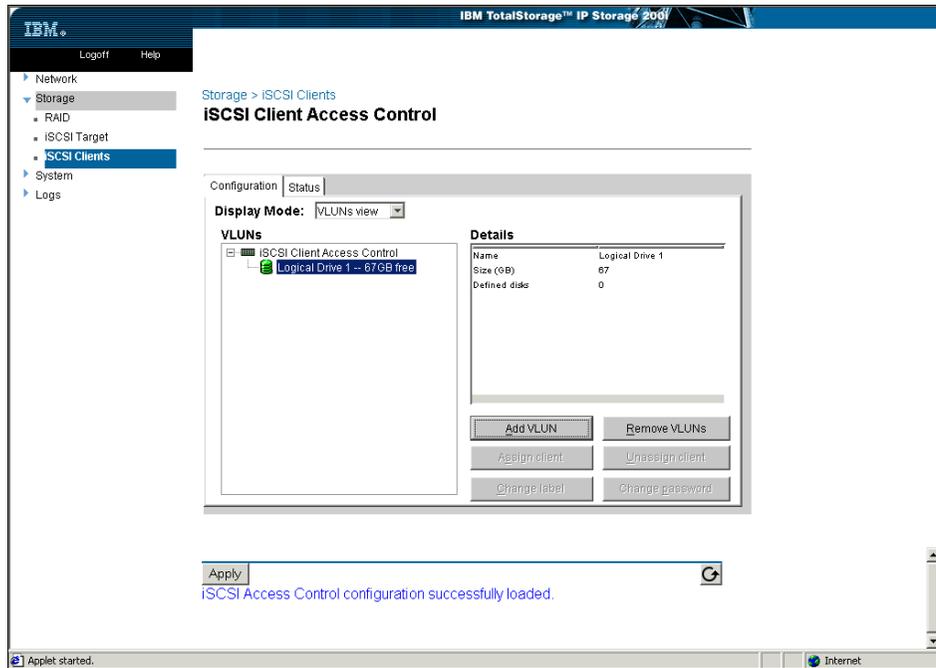


Figure 2-15 Storage options

- ▶ After iSCSI clients are selected, the management GUI will load the iSCSI access control page (Figure 2-15), this may take several minutes for you to complete.

The next window will display the disks and amount of storage that is available to use on this 200i. Changes to storage sizes can be made from this page.

- ▶ Highlight the disk that will house the new VLUN as shown in (Figure 2-16).
- ▶ On the lower right of the display, two selection blocks will appear — Add VLUN and Remove VLUNs.

Note: A single client can be assigned to multiple VLUNs, but a VLUN can only be assigned to one client at a time.

- ▶ Select **Add VLUN**.

The Add VLUN window will appear (Figure 2-17).

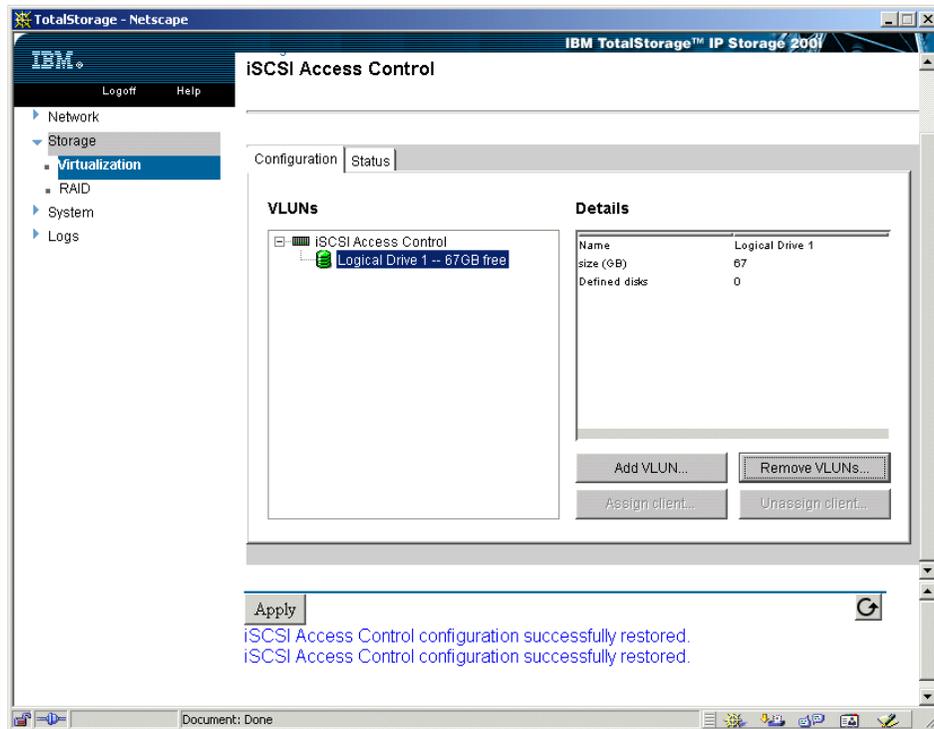


Figure 2-16 Access control page

- ▶ Key in the VLUN size, number and label.
- ▶ Select **Add**.

Restriction: A VLUN must be the same size or smaller than the size of the logical drive.

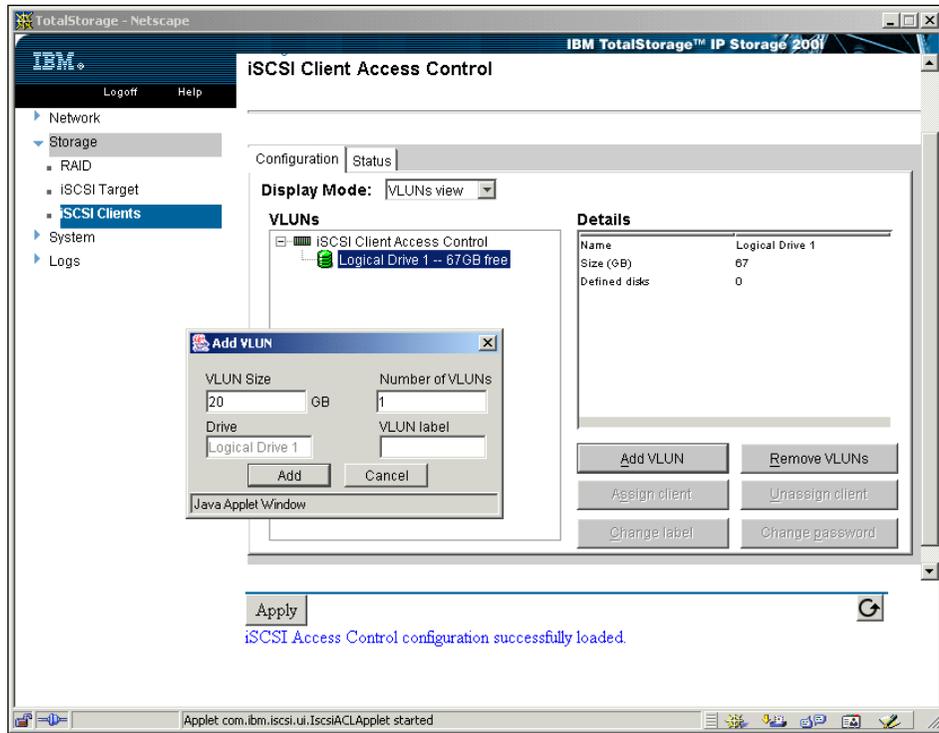


Figure 2-17 VLUN ADD option

- ▶ Once Add is selected, the VLUNs will be created (Figure 2-18).

In this example we created two 30 GB VLUNs. We utilized the rest of the disk and added a 7 GB VLUN.

Once finished creating the VLUNs:

- ▶ Select **Apply** in the lower left of the window for the assignments to be completed.

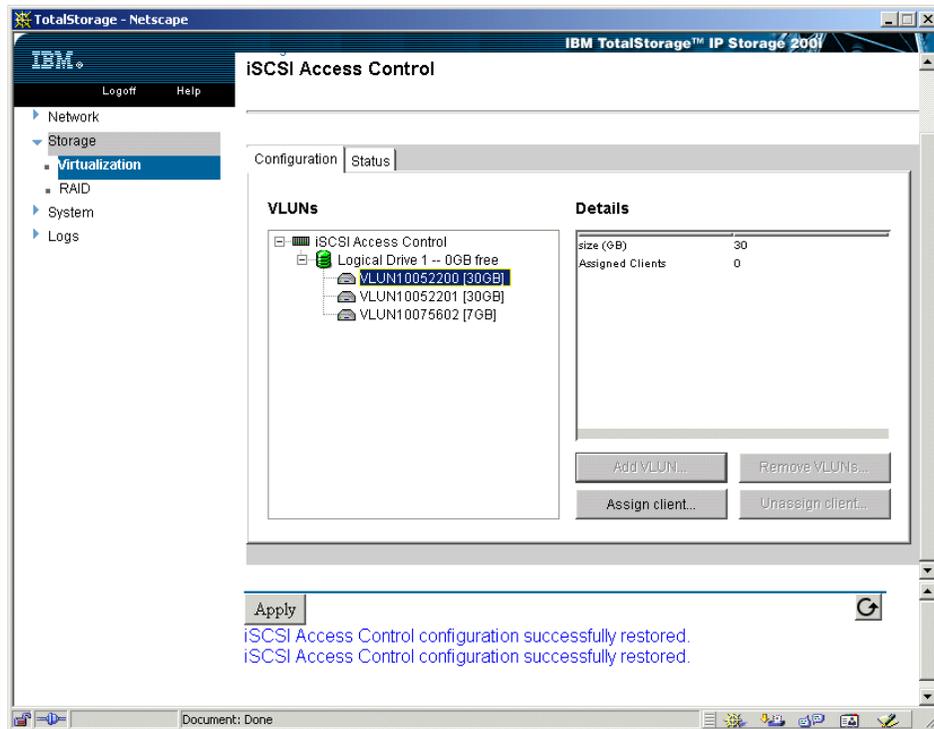


Figure 2-18 Configured VLUNs

The next step will be to assign the VLUNs to clients. Once the VLUNs are created, highlight one of the VLUNs. The Assign Client box becomes active:

- ▶ Select the **Assign Client** option.
- ▶ A window will appear allowing you to create the client providing the Login ID and Password (Figure 2-19).
- ▶ Key in a user ID and password, and confirm the password that the client will be using. This is not associated with the computer name and network name of the client machine.
- ▶ Select **OK**.

The client name appears under the VLUN that was previously created (Figure 2-20). You may continue assigning a single client to all of the VLUNs or a different client to each VLUN. In this example we used all of the VLUNs for a single client.

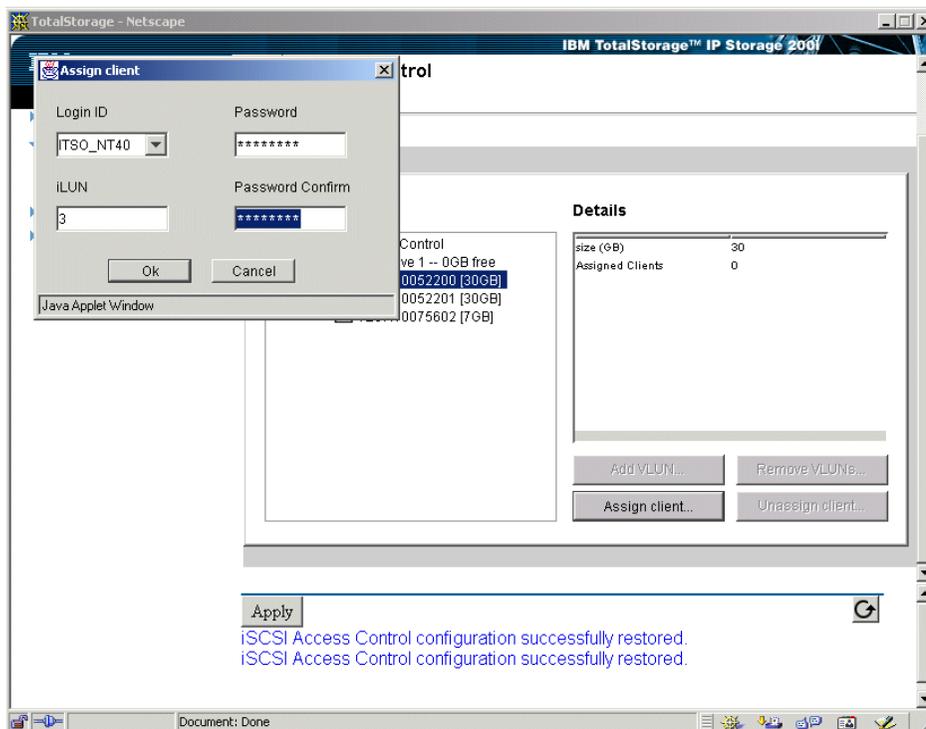


Figure 2-19 Client setup

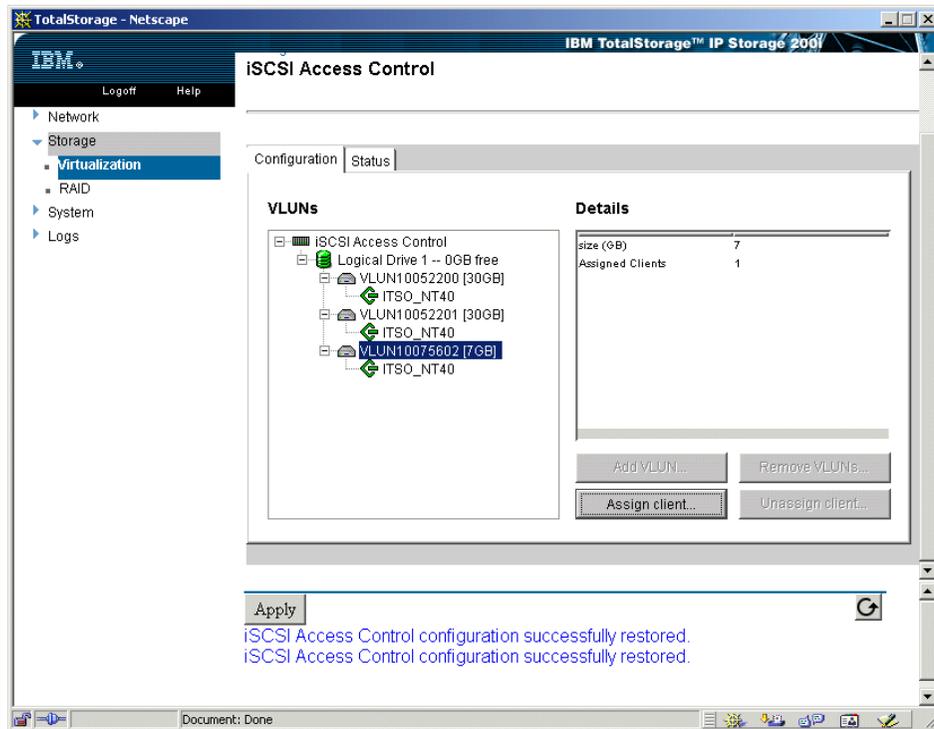


Figure 2-20 A single VLUN

- ▶ Once the clients are created, select the **Apply** button to complete the task. The management window will return confirming that the tasks are completed.

Once the iSCSI client disk configuration has been completed the iSCSI is ready for client connection.

2.3.8 Client initiator setup procedure

The next steps will show you the procedure for installing and setting up the Initiator code as well as the process for attaching to the iSCSI target machine.

The initiator code is a zip file that contains several smaller files. It should be unzipped into a directory from where it will be installed. In this setup the following directory was used:

```
C:\install\
```

From the Install directory run the setup.exe file:

- ▶ Double-click the setup file to start the initiator setup procedure.

The InstallShield Wizard for the IBM iSCSI Client for Windows will be launched (Figure 2-21).



Figure 2-21 Installation wizard

- ▶ Select **Next**.
- ▶ Read the license agreement.
- ▶ Select **I accept the License Agreement**.
- ▶ Select **Next** (Figure 2-22).

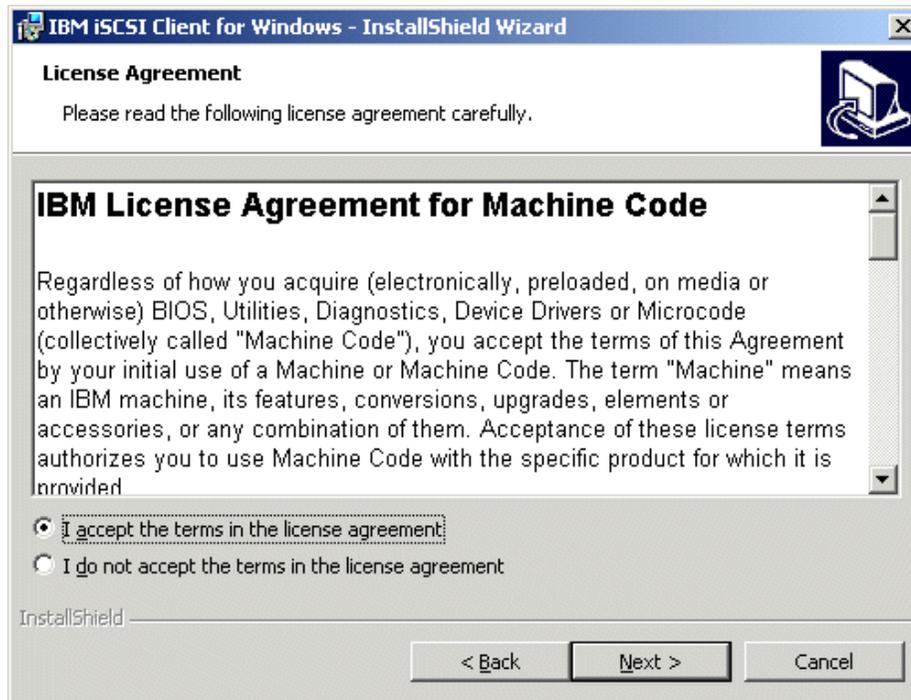


Figure 2-22 IBM license agreement

- ▶ Select **Install** to start the installation process (Figure 2-23).

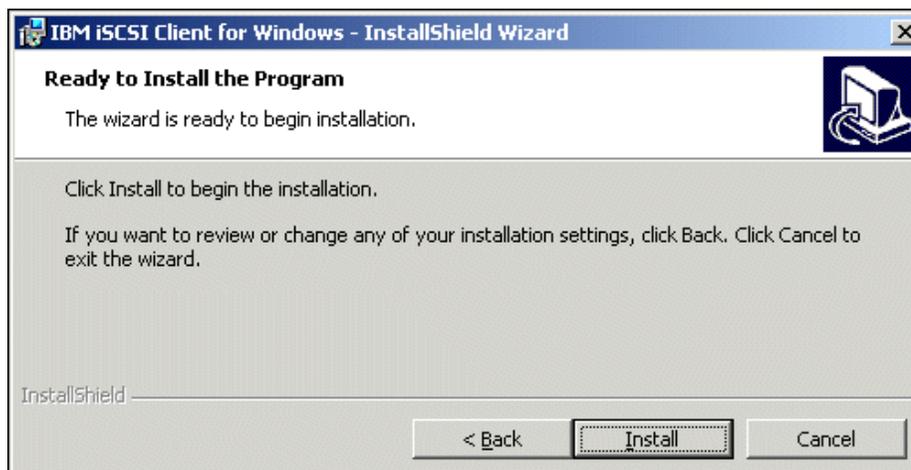


Figure 2-23 Summary page

- ▶ After the installation is complete, select **Finish**.

The iSCSI Initiator machine will need to be rebooted for the driver to become active (Figure 2-24).

- ▶ Select **Yes**.



Figure 2-24 Reboot option

After system reboot the IBM iSCSI Configuration icon will appear in the lower right hand corner (Figure 2-25).

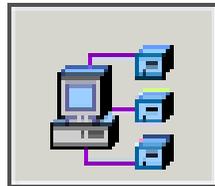


Figure 2-25 IBM iSCSI configuration

- ▶ Right-click on the **IBM iSCSI configuration icon**.
- ▶ The three menu items will appear (Figure 2-26).
- ▶ Select **IBM iSCSI Client Configuration** to configure the initiator.

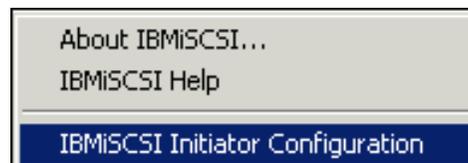


Figure 2-26 IBM iSCSI configuration options

The IBM iSCSI configuration utility window will appear and allow you to connect to your iSCSI target.

- ▶ Select **Add** (Figure 2-27).

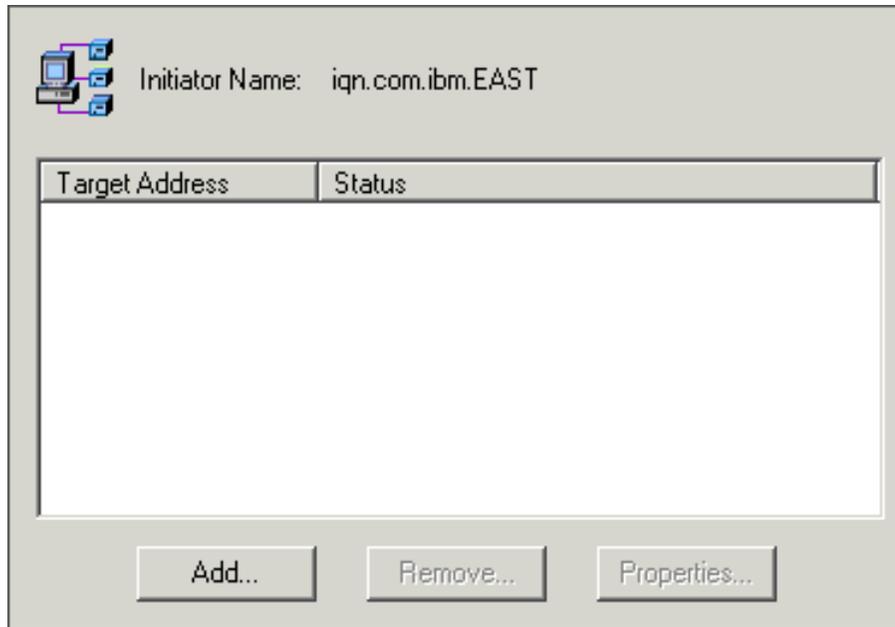


Figure 2-27 IBM iSCSI configuration utility

Once Add is selected the fields for IP address or hostname, port number, login ID and login password become active.

- ▶ The login ID and password is entered as configured in setting up the VLUNs (Figure 2-28).

The name can be acquired by logging in with a null modem cable, connecting with the HyperTerminal, enter:

- ▶ **iscsi**
- ▶ **summary**

Important: The target name must be the fully qualified name, the name is case sensitive, the format is iqn.com.ibm.Model.Serial no:

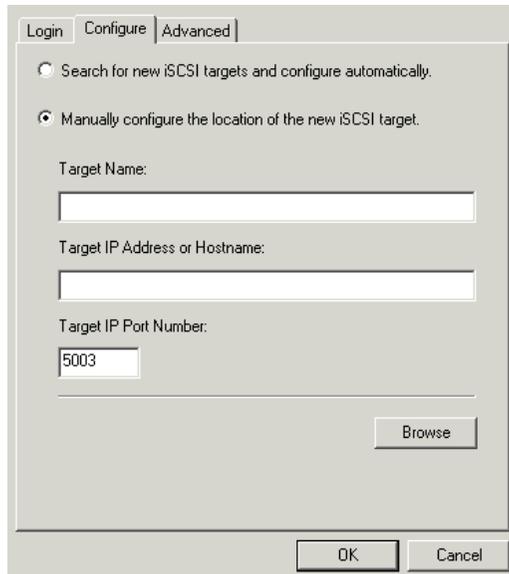
```
iqn.com.ibm.4125.23XXXXX
```

- ▶ Enter the IP address and target name of the data port/GigE Ethernet card under the configure option (Figure 2-29).
- ▶ Select **OK** to complete the process (Figure 2-30).



The image shows a dialog box with three tabs: "Login", "Configure", and "Advanced". The "Login" tab is selected. Inside the dialog, there is a section titled "Secure Login" with a checked checkbox labeled "Perform secure login.". Below this checkbox are two text input fields: "Username:" and "Password:". At the bottom of the dialog are "OK" and "Cancel" buttons.

Figure 2-28 IBM iSCSI configuration logon



The image shows a dialog box with three tabs: "Login", "Configure", and "Advanced". The "Configure" tab is selected. Inside the dialog, there are two radio button options: "Search for new iSCSI targets and configure automatically." (which is unselected) and "Manually configure the location of the new iSCSI target." (which is selected). Below these options are three text input fields: "Target Name:", "Target IP Address or Hostname:", and "Target IP Port Number:". The "Target IP Port Number" field contains the value "5003". To the right of the "Target IP Port Number" field is a "Browse" button. At the bottom of the dialog are "OK" and "Cancel" buttons.

Figure 2-29 IBM iSCSI configuration

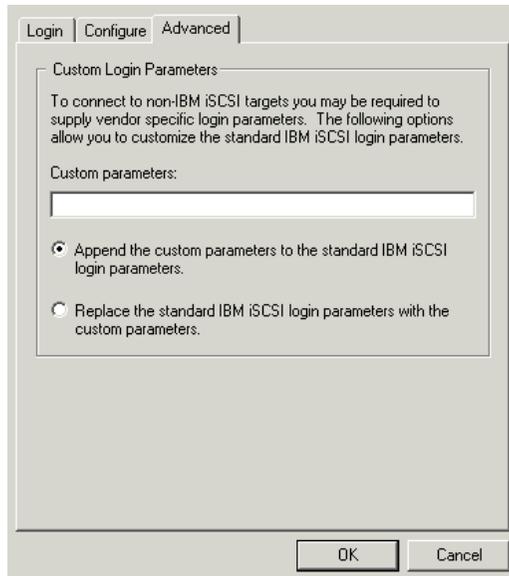


Figure 2-30 IBM iSCSI configuration advanced options

After OK is selected, the process to identify the iSCSI target, and also connecting to the assigned VLUNs begins.

Once the target has been identified and the VLUNs have been attached:

- ▶ Select **Yes**.

Windows Disk Management will start:

- ▶ Select **Yes** and proceed to set up the disk that has been made available.

Attention: It is recommended that iSCSI disks are setup as **Basic** disks and not Dynamic Disks.

Set up the drive using the Disk Management system. Write a signature, format the drive and assign a drive letter and name to the drive.

After completion of this process, you will be presented with the available drive (Figure 2-31).

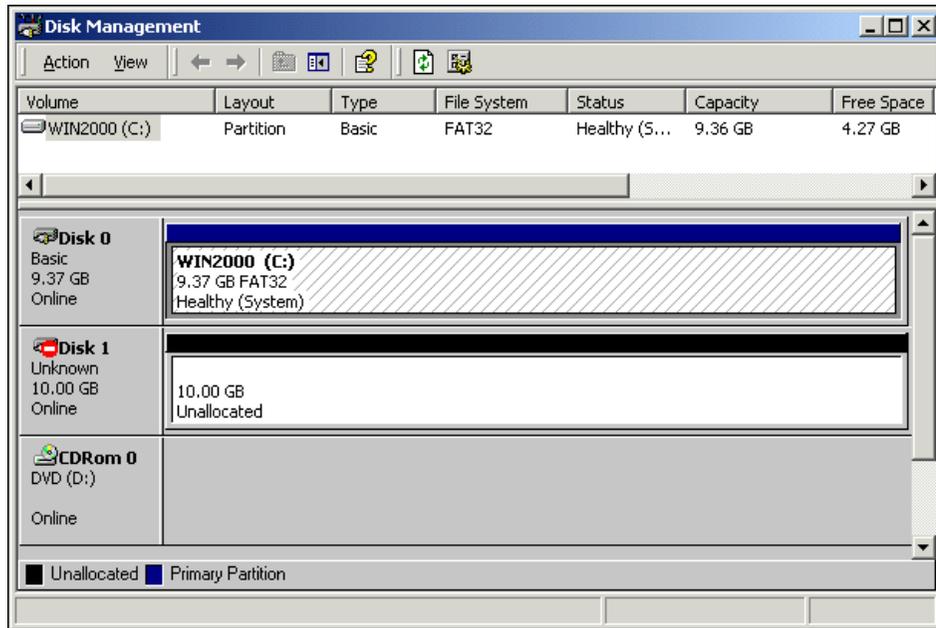


Figure 2-31 Windows disk management

Open My Computer from the desktop. The new drive is available in this example as drive E: (Figure 2-32).



Figure 2-32 Available iSCSI disk

The iSCSI disk is now ready for use as a locally attached drive on the system.

2.3.9 Troubleshooting tips

A guide to troubleshooting can be found in “Appendix A” of the *Administrator's Guide*, which is included as part of the ship group for the IP Storage 200i. For detailed troubleshooting procedures, please refer to that manual.

A few troubleshooting tips are suggested here for quick checks of the system:

- ▶ Be sure all three power cords are installed on the system.
- ▶ Check your hub or switch to ensure you have connectivity from the 200i on each of the LAN connections.
- ▶ Listen for the one long and three short beeps when the power button is pressed.
- ▶ Check for the “OK” LED on the front panel after power up to ensure the POST tests were run without error.
- ▶ If necessary, attach a display, keyboard, and mouse to the 200i to check for system errors on power up. After power up, these devices will not be needed.
- ▶ Be sure the control port and data port can be pinged from another workstation during setup. **Note:** This is very important.
- ▶ If after booting the initiator workstation an iSCSI drive is missing, run Disk Management to see if the drive exists. If it is in Disk Management, but does not show up in your list of drives under “My Computer”, try assigning a drive letter to it that is at the upper end of the alphabet. It is possible a new device has been installed locally on the system that has taken that drive letter. Remember, the iSCSI drives are network drives and as such are added later in the startup process.
- ▶ During initiator setup an error may be displayed indicating that the target cannot be found. Ensure the data port address is correctly typed into the “IP Address or Hostname” field in the iSCSI configuration utility. Also, make sure you can still ping the data port address from the initiator workstation.
- ▶ The diagnostic procedures in “Appendix A” of the *Administrator's Guide* indicates that the system will beep once on power up if no system problems are detected. This is true if a display, keyboard and mouse are attached when powering up. Otherwise you will hear one long and three short beeps.



Cisco SN5420 Storage Router

The goal of this chapter is to provide an introduction to the benefits and features of the Cisco 5420 Storage Router. A review of the hardware features along with an example software configuration is provided for your consideration. This configuration is typical of what you would encounter in many of your SAN and IP environments.

3.1 Cisco SN5420 Storage Router introduction

In April 2001, Cisco, IBM's partner in presenting the iSCSI protocol to the IETF, announced the Cisco SN5420 Storage Router, which offers a gateway between an IP network and a Fibre Channel network. This allows an IP network attached client to access Fibre Channel SAN storage via the gateway. IBM International Global Services has a re-marketing agreement with Cisco for the sale and support of SN5420.

The Cisco SN5420 Storage Router (Figure 3-1) provides access to Fibre Channel attached SCSI storage over IP networks. With the 5420 you can directly access storage anywhere on an IP network just as easily as you can access storage locally.

The Cisco SN5420 Storage Router provides servers with IP access to storage by means of SCSI routing using the iSCSI protocol. With SCSI routing, servers use an IP network to access storage as if the servers were directly attached to the storage devices.



Figure 3-1 Cisco SN5420 Storage Router

The Cisco SN5420 Storage Router is ideal when using both Fibre Channel and TCP/IP protocols. It combines the high performance data transfer capabilities of Fibre Channel with the interoperable, and widely understood, TCP/IP.

The Cisco SN5420 Storage Router uses the TCP/IP protocol suite for networking storage supporting the level of interoperability inherent to IP networks. It leverages existing management and configuration tools that are already well known and understood. And it is based on industry standards which maximize your investment by allowing you to leverage the existing TCP/IP experience and equipment.

The Cisco SN5420 Storage Router is based on both IP and Storage Area Network (SAN) standards, providing interoperability with existing local area network (LAN), wide-area network (WAN), optical and SAN equipment. The Cisco SN5420 Storage Router is a high performance router designed to allow block level access to storage regardless of your operating system or location. The Cisco SN5420 Storage Router accomplishes this by enabling Small Computer Systems Interface over IP (iSCSI). The Cisco SN5420 Storage Router

connects to both the FC SAN network and the IP network via Gigabit Ethernet. This allows the Cisco SN5420 Storage Router to perform gateway functions between environments and allows IP routing intelligence to be leveraged with storage networking technologies.

Each server that requires IP access to storage via the Cisco SN5420 Storage Router needs to have the Cisco iSCSI driver installed. Using the iSCSI protocol, the iSCSI driver allows a server to transport SCSI requests and responses over an IP network. From the perspective of a server operating system, the iSCSI driver appears to be a SCSI or Fibre Channel driver for a peripheral channel in the server. Figure 3-2 shows a sample storage router network. Servers with iSCSI drivers access the storage routers through an IP network connected to the Gigabit Ethernet interface of each 5420 storage router. The storage routers access storage devices through a storage network connected to the Fibre Channel interface of the management interface of each storage router.

For high availability operations, the storage routers communicate with each other over two networks: the HA network connected to the HA interface of each storage router, and the management network connected to the management interface of each storage router. Note that there are three IP interfaces on the Cisco 5420 storage router; Gigabit Ethernet, Management Ethernet and HA Ethernet. All three interfaces must have unique IP networks (subnets) defined. An example of this is explained in Figure 3-11 on page 77.

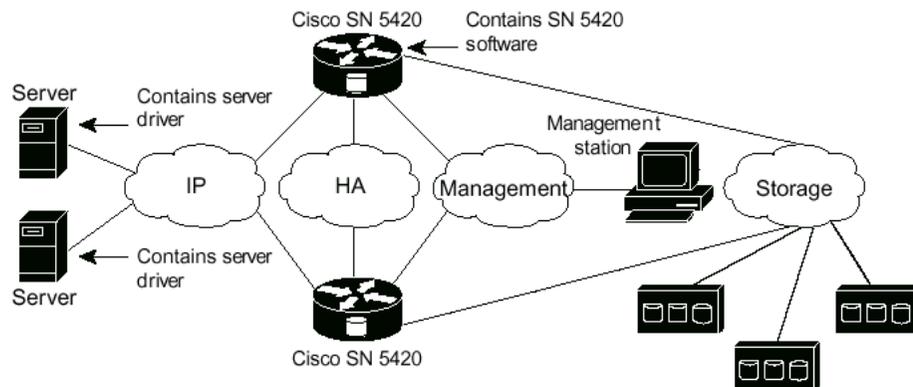


Figure 3-2 iSCSI routing

3.2 Features and benefits

Many features listed in this section make the Cisco SN5420 Storage Router a cost effective and reliable method of implementing IP based access to SAN storage devices.

3.2.1 Making disk subsystems IP aware

Implementing iSCSI within the server allows seamless block level access for all applications. When using the Cisco SN5420 Storage Router, servers now have the ability to access Fibre Channel attached storage with inexpensive Ethernet adapters instead, over an IP network. This capability allows existing application software to operate without modifications. The iSCSI driver, which resides on the host server, is a key component of the Cisco SN5420 Storage Router. The iSCSI driver intercepts SCSI commands, encapsulates them as IP commands, and redirects them to the Cisco SN5420 Storage Router. The iSCSI driver is supported on a variety of operating systems. Section 3.6, “Host setup” on page 93 lists the specifics of driver installation and configuration.

3.2.2 Interoperability

The Cisco SN5420 Storage Router fits seamlessly into existing storage and data networks. The Cisco SN5420 Storage Router uses the well-known TCP/IP protocol suite for network storage, supporting the level of interoperability inherent to mature IP networking protocols. The Cisco SN5420 Storage Router is based on current SAN standards as well and is compatible with existing SAN deployments, point-to-point, switched or arbitrated loop.

3.2.3 Scalability and reliability

The Cisco SN5420 Storage Router provides optimal performance and reliability. Additional Cisco SN5420 Storage Routers can be easily added to the network to match performance requirements. Reliability is accomplished by using the Cisco SN5420 Storage Router High Availability (HA) and Clustering features. If one Cisco SN5420 Storage Router fails, another Cisco SN5420 Storage Router automatically takes over for the failed component. Section 3.5, “Configuration” on page 71 lists details on installing and configuring HA features.

3.2.4 Manageability

The Cisco SN5420 Storage Router leverages existing management and configuration tools that are already well known and understood. The Cisco SN5420 Storage Router provides full network management support through Simple Network Management Protocol (SNMP), WEB-based GUI and command line interface (CLI) access.

3.2.5 Security

The Cisco SN5420 Storage Router uses Access Control Lists (ACLs) to limit only specific IP addresses access to SAN-based storage. This controls server access to specific logical unit numbers (LUNs). The Cisco SN5420 Storage Router is also password protected to further control security.

3.2.6 Investment protection

Total cost of ownership (TCO) is a growing concern for most system administrators and management. The Cisco SN5420 Storage Router helps reduce the costs by leveraging your existing TCP/IP networking infrastructure while maintaining your current and near-term investments in storage systems and Fibre Channel infrastructure. The Cisco SN5420 Storage Router simplifies the cost of management, deployment and support issues, given the fact that technical skills in TCP/IP support are more widely available than SAN experience.

3.3 Technical specifications

Table 3-1 provides the technical specifications of the Cisco SN5420 Storage Router.

Table 3-1 Cisco SN5420 Storage Router technical specifications

Specifications	Description
Environmental	
Ambient operating temperature	32 to 104 F (0 to 40 C)
Humidity (RH)	10 to 95 percent non-condensing
Altitude	-500 to 10,000 ft. (-152.4 to 3,048 m)
Physical characteristics	
Dimensions (H x W X D)	1.75 x 17 x 15.5 in. (4.4 x 43.2 x 39.4 cm) 1 Rack Unit in height

Specifications	Description
Weight	7.5 lb (3.4 kg)
AC Power	
Output	70W
Power dissipation	35W
Current	1.0A maximum @ 100 to 120 VAC 0.5A maximum @ 200 to 240 VAC
Frequency	50 to 60 Hz
Airflow	Right and left side in, rear out
Gigabit Ethernet Port	
Connector	Duplex SC
Type	Short wavelength
Wavelength	850 nanometers
Fiber type	Multimode
Core size, modal bandwidth, maximum length	62.5, 160, 722 ft. (220 m) 62.5, 200, 902 ft. (275 m) 50.0, 400, 1640 ft. (500 m) 50.0, 500, 1804 ft. (550 m)
Fibre Channel Port	
Connector	Duplex SC
Type	SN (Shortwave laser without Open Fiber Control)
Wavelength	850 nanometers
Fiber type	Multimode
Core size, modal bandwidth, maximum length	62.5, 160, 984 ft. (300 m) 50.0, 400, 1640 ft. (500 m)

3.4 Installation

This section describes the physical installation and port overview of the Cisco SN5420 Storage Router. The router is a 1 RU rack-mountable chassis that has one Gigabit Ethernet port, one Fibre Channel port, two management ports and one high availability port.

Fibre Channel port

The FC port is labeled FIBRE CHANNEL 1 and is a 100-M5-SN-I/100-M6-SN-I interface using a duplex SC connector (Figure 3-3). It is used to connect to storage controllers on the FC network. The FC port supports point-to-point, loop or fabric topologies and functions either as a Fibre Channel N_Port or NL_Port.



Figure 3-3 Fibre Channel port

Gigabit Ethernet port

The Gigabit Ethernet port is labeled GIGABIT ETH 2 and is a 1000Base-SX (short-wavelength) interface used to connect to a switch (Figure 3-4). It is from this port that servers, requiring storage, connect using the iSCSI driver. The port uses a duplex SC connector and has status LEDs. While testing, we connected the Gigabit Ethernet port to a Cisco Catalyst 3508-G. This switch was also connected to a Cisco Catalyst 3524-XL fast Ethernet switch with two gigabit uplink ports. These switches provided the network connectivity for all of the device data, management and HA Ethernet connections.



Figure 3-4 Gigabit Ethernet port

RS-232 management port

The RS-232 management port is labeled CONSOLE and is an RS-232 serial interface used for local console access (Figure 3-5). The command line interface (CLI) is accessed using a terminal emulator setup for 9600 bps, 8 data bits, 1 stop bit and no parity. A rolled cable is required when connecting to a terminal emulator. The supplied Cisco cable is a rolled cable.

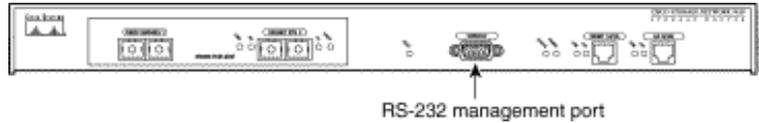


Figure 3-5 Console port

Management Ethernet

The port labeled MGMT 10/100 is for Ethernet network management access (Figure 3-6). Through a management network you can manage the storage router using telnet to enter CLI commands, a Web based graphical user interface (GUI), or SNMP commands. This port uses a modular RJ-45 straight through UTP cable for a connection.

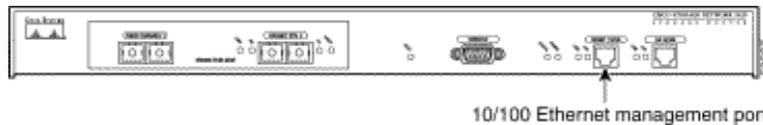


Figure 3-6 Management port

High Availability (HA) port

The HA port is labeled HA 10/100 and uses a modular RJ-45 connector (Figure 3-7). This port is used to join other Cisco SN5420 Storage Routers providing fault operation. If connecting directly to another HA port, use rolled unshielded twisted pair (UTP) cable.



Figure 3-7 HA port

Important: For the Cisco SN5420 Storage Router to operate properly, the Management Ethernet, Gigabit Ethernet and High Availability Ethernet ports must be defined using separate and unique networks or subnets. See Table 3-2.

Table 3-2 Subnet configuration example

Description	Network	Subnet
Management network	192.168.1.0	255.255.255.0

Description	Network	Subnet
HA network	192.168.2.0	255.255.255.0
Data network	192.168.3.0	255.255.255.0

3.5 Configuration

This section explains the SCSI router service that runs in the Cisco SN5420 Storage Router and describes the concept of High Availability/Clustering and includes a sample configuration.

3.5.1 SCSI router service

Access for SCSI routing is controlled in the servers and a storage router. In a server, the IP address of each storage router in which the server is to transport SCSI requests and responses is configured in the iSCSI driver. In the Cisco SN5420 Storage Router a SCSI router service is defined with an access list that identifies which servers can access storage devices attached to it.

Important: Up to four SCSI routing services can be defined per storage router.

Note: At the time of writing this redbook, the latest SN5420 software release was V1.1.5.

3.5.2 Clustering and high availability

Clustering is used in conjunction with high availability and allows storage routers to back each other up in case of failure. A storage router cluster consists of two storage routers configured with the same cluster name and connected as follows:

- ▶ Connected to the same servers
- ▶ Connected to the same storage systems
- ▶ Connected to each other through their management and HA interfaces

In a cluster, storage routers continually communicate HA and configuration information between each other by balancing the exchange of information through both the HA and Management Ethernet ports. In the event of a hardware or software failure in the primary 5420, the secondary will take over SCSI router operations. The HA interface for each storage in a cluster should be on the same IP network or subnet. All SN5420s that participate in a cluster must have the same cluster name.

3.5.3 Sample configuration

Table 3-3 lists the parameters that have been used in our sample configuration. We highly recommend that you fill in this table with your site-specific information, before attempting to configure the router.

Table 3-3 Sample system parameters

System parameter	Value
High availability (HA) cluster name	Mt_Olympus
System name (if using DNS this is the name you will associate with the management interface IP address).	SN_5420-MG1
Monitor-level password	sanfran
Administrator-level password	cisco@123
Management interface IP address and network mask	10.1.10.244/24 (255.255.255.0)
High availability (HA) interface IP address and network mask	10.1.20.56/24 (255.255.255.0)
Primary DNS IP address	10.1.40.243
Secondary DNS IP address	10.1.50.249
Domain name	mystoragenet.com
Local gateway IP address	10.1.10.201
NTP server IP address	10.1.60.86
Primary SNMP trap manager IP address	10.1.30.17
Secondary SNMP trap manager IP address	10.1.30.18
SNMP get community name	public
SNMP set community name	mynetmanagers

System parameter	Value
System administrator contact name	Pat J. Smith
System administrator contact e-mail	pjsmith@mystoragenet.com
System administrator contact phone	763-555-7766
System administrator contact pager	763-555-7766
SCSI routing services instance name (maximum of four).	zeus
SCSI routing services server interface IP address and network mask	10.1.0.45/24 (255.255.255.0)
SCSI routing services device interface	fc1

3.5.4 Configuration from the console port

The command line interface, CLI, method described here is used to initially configure the 5420. Other methods such as Web based GUI or SNMP are available, but are not covered. The CLI is very much like Cisco IOS. There are however still many differences and the exact syntax is noted here.

Follow this procedure to configure the 5420 from the console port:

1. Connect the com port of a laptop or PC to the console port of the 5420.
2. Using HyperTerminal in Windows, create a new connection with the setting of 9600 bits per second, parity of none, 8 data bits and 1 stop bit.
3. Power up the 5420. Once the boot process completes you will see the message:

```
 “[SN5420]$date time:CLI:AS_NOTICE :Started CLI on console”
```

 appear.
4. Press the **Enter** key and a monitor prompt SN5420> will appear.
5. Enter **enable** to enter the configuration mode. The prompt will change to SN5420#.

Attention: Since the release of SN5420 software V1.1.3, you do not have to enter a password from the local console port.

```
Cisco Console.r2w - Reflection for UNIX and Digital
File Edit Connection Setup Macro Window Help

Jul 02 16:56:30:HA:AS_ALERT :HA failed to initialize: Neither MGMT/HA interface is valid
Jul 02 16:56:30:HA:AS_ALERT :HA has been Idled: no network interfaces
Jul 02 16:56:30:ConfMgmt:AS_NOTICE :Cluster Application registered with Configuration Management
Jul 02 16:56:30:ConfMgmt:AS_NOTICE :SCSI Router registered with Configuration Management
Jul 02 16:56:30:ClusterConf::createConfNode:AS_DEBUG :failed to create configuration node -- null node name
Jul 02 16:56:30:ge2:AS_NOTICE :Valid signal detected
Jul 02 16:56:30:ge2:AS_CRIT :Interface Up

Done executing startup script /ata0/WuSpeed.start
-> reset stdin to /pipe/dummy fd = 10
                                shell restarted.

                                ->

*[SN5420]$ Jul 02 16:56:37:CLI:AS_NOTICE :Started CLI on console

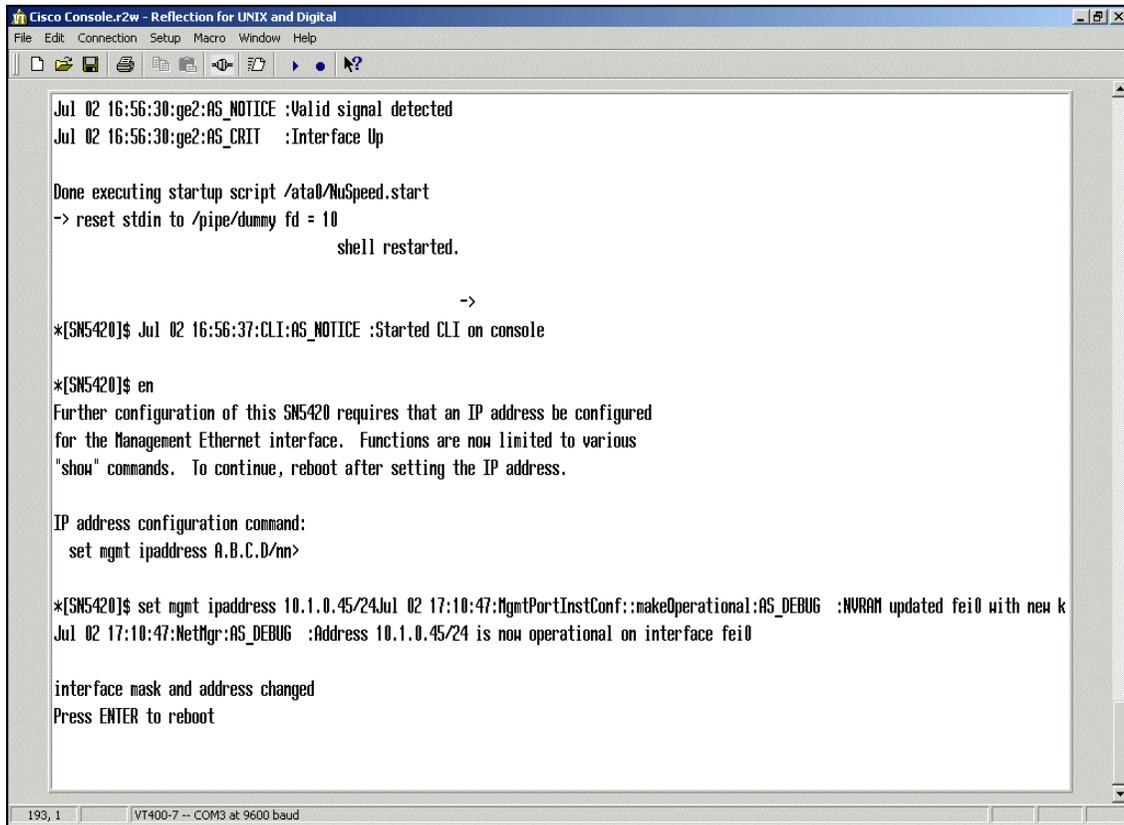
*[SN5420]$ en
Further configuration of this SN5420 requires that an IP address be configured
for the Management Ethernet interface. Functions are now limited to various
"show" commands. To continue, reboot after setting the IP address.

IP address configuration command:
set mgmt ipaddress A.B.C.D/mn>

*[SN5420]$
```

Figure 3-8 Initial start up screen

Since this is the first configuration of the router, the IP address and mask must be entered next. Once this is done you will be required to reboot the system for this change to take effect.



```
Cisco Console.r2w - Reflection for UNIX and Digital
File Edit Connection Setup Macro Window Help

Jul 02 16:56:30:ge2:AS_NOTICE :Valid signal detected
Jul 02 16:56:30:ge2:AS_CRIT  :Interface Up

Done executing startup script /ata0/MuSpeed.start
-> reset stdin to /pipe/dummy fd = 10
                                shell restarted.

                                ->

*[SN5420]$ Jul 02 16:56:37:CLI:AS_NOTICE :Started CLI on console

*[SN5420]$ en
Further configuration of this SN5420 requires that an IP address be configured
for the Management Ethernet interface. Functions are now limited to various
"show" commands. To continue, reboot after setting the IP address.

IP address configuration command:
set mgmt ipaddress A.B.C.D/nn>

*[SN5420]$ set mgmt ipaddress 10.1.0.45/24Jul 02 17:10:47:MgmtPortInstConf::makeOperational:AS_DEBUG :NVRAM updated fei0 with new k
Jul 02 17:10:47:NetMgr:AS_DEBUG :Address 10.1.0.45/24 is now operational on interface fei0

interface mask and address changed
Press ENTER to reboot

193, 1 | VT400-7 -- COM3 at 9600 baud
```

Figure 3-9 Enter the management port IP address

Once the router reboots, re-establish the configuration mode by issuing the **enable** command again.

```

Cisco Console.r2w - Reflection for UNIX and Digital
File Edit Connection Setup Macro Window Help

ClusterApp initing...

[ClusterApp::init]register type cluster succeeded
nuEWtask starting...
SystemApp starting...
ClusterApp starting...
Jul 02 17:13:51:ConfMgmt:AS_NOTICE :Cluster Application registered with Configuration Management
tmEmWeb active - Port 80 (buffers:640 size:128)
Jul 02 17:13:51:ConfMgmt:AS_NOTICE :SCSI Router registered with Configuration Management
Jul 02 17:13:51:ClusterConf::createConfNode:AS_DEBUG :instance id: cluster/6440ec40, registered
Jul 02 17:13:51:ClusterApp::makeOperational:AS_DEBUG :cluster registration succeeded
value = 0 = 0x0

Done executing startup script /ata0/NuSpeed.start
-> reset stdin to /pipe/dummy fd = 10
                                shell restarted.

                                ->

[SN5420]$ Jul 02 17:13:58:CLI:AS_NOTICE :Started CLI on console

[SN5420]$ en
[Entering Administrator mode]

[SN5420]#

```

345, 11 | VT400-7 -- COM3 at 9600 baud

Figure 3-10 Re-establish Admin mode by entering enable command

3.5.5 Configuring management parameters

A configuration wizard is available that will prompt you to enter management, HA and SCSI routing parameters. By typing **setup** you will invoke the script. To better understand parameter options we have listed each command individually with their description in Table 3-4 on page 79. The commands are listed in specific order to properly configure the system. Many commands are listed as optional and are only necessary depending on your specific network installation. Figure 3-11 is a diagram of our example network.

Important: When an asterisk is added to the beginning of the prompt, this means that the system configuration has been modified but not saved. If the system were restarted these entries would not be retained when the system is reloaded.

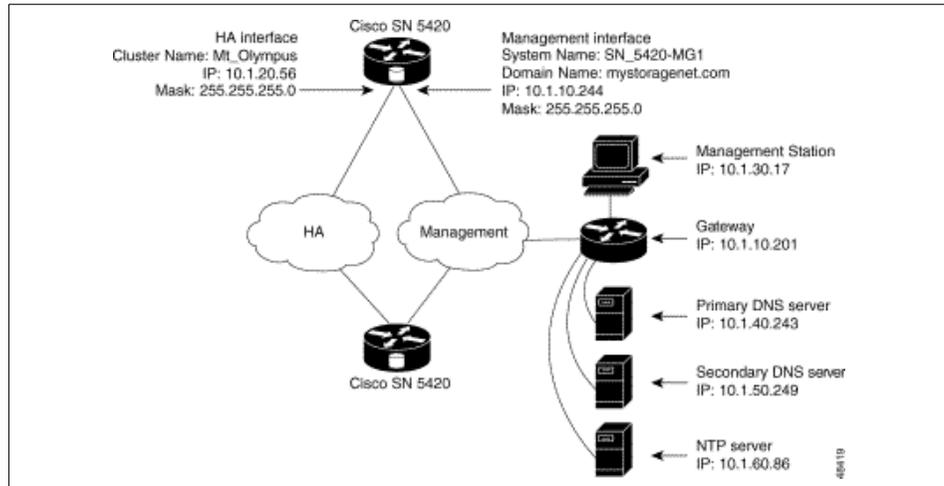


Figure 3-11 Example network

Figure 3-12 and Figure 3-13 show how each screen looks if you type and test the commands in Table 3-4 on page 79.

```
Cisco Console.r2w - Reflection for UNIX and Digital
File Edit Connection Setup Macro Window Help

*[SN_5420-MG1]$ en
[Entering Administrator mode]

*[SN_5420-MG1]# set systemname SN_5420-MG1
*[SN_5420-MG1]# add route 10.1.30.0/24 gw 10.1.10.201
*[SN_5420-MG1]# Jul 02 17:55:26:NetMgr:AS_DEBUG :Route to 10.1.30.0/24 using gateway 10.1.10.201 via interface fei0 is now operati

*[SN_5420-MG1]# set nameserver 10.1.40.243 domain mystoragenet.com
*[SN_5420-MG1]# add route 10.1.40.243/32 gw 10.1.10.201
*[SN_5420-MG1]# Jul 02 17:56:42:NetMgr:AS_DEBUG :Route to 10.1.40.243/32 using gateway 10.1.10.201 via interface fei0 is now operati

*[SN_5420-MG1]# set secnameserver 10.1.50.249
*[SN_5420-MG1]# add route 10.1.50.249/32 gw 10.1.10.201
*[SN_5420-MG1]# Jul 02 17:57:37:NetMgr:AS_DEBUG :Route to 10.1.50.249/32 using gateway 10.1.10.201 via interface fei0 is now operati

*[SN_5420-MG1]# set clock 13:00:00 07/02/2001
*[SN_5420-MG1]# set ntpserver 10.1.60.86
*[SN_5420-MG1]# Jul 02 13:00:24:CLISetSysNTP:AS_ALERT :NTP server set to 10.1.60.86

*[SN_5420-MG1]# Jul 02 13:00:34:NTP task:AS_DEBUG :Unable to get current time from NTP server

*[SN_5420-MG1]# add route 10.1.60.86/32 gw 10.1.10.201
*[SN_5420-MG1]# Jul 02 13:01:26:NetMgr:AS_DEBUG :Route to 10.1.60.86/32 using gateway 10.1.10.201 via interface fei0 is now operati

25, 1 | VT400-7 -- COM3 at 9600 baud
```

Figure 3-12 Initial configuration commands (screen 1)

```

Cisco Console.r2w - Reflection for UNIX and Digital
File Edit Connection Setup Macro Window Help

*[SN_5420-MG1]# add route 10.1.50.249/32 gw 10.1.10.201
*[SN_5420-MG1]# Jul 02 17:57:37:NetMgr:AS_DEBUG :Route to 10.1.50.249/32 using gateway 10.1.10.201 via interface fei0 is now operat

*[SN_5420-MG1]# set clock 13:00:00 07/02/2001
*[SN_5420-MG1]# set ntpserver 10.1.60.86
*[SN_5420-MG1]# Jul 02 13:00:24:CLISetSysNTP:AS_ALERT :NTP server set to 10.1.60.86

*[SN_5420-MG1]# Jul 02 13:00:34:NTP task:AS_DEBUG :Unable to get current time from NTP server

*[SN_5420-MG1]# add route 10.1.60.86/32 gw 10.1.10.201
*[SN_5420-MG1]# Jul 02 13:01:26:NetMgr:AS_DEBUG :Route to 10.1.60.86/32 using gateway 10.1.10.201 via interface fei0 is now operat

*[SN_5420-MG1]# set snmp getcommunity public
*[SN_5420-MG1]# Jul 02 13:03:05:CLISetSNMPcommunity:AS_WARNING:getcommunity changed to public

*[SN_5420-MG1]# set snmp setcommunity mynetmanagers
*[SN_5420-MG1]# Jul 02 13:03:37:CLISetSNMPcommunity:AS_WARNING:setcommunity changed to mynetmanagers

*[SN_5420-MG1]# set snmp traphosts primary 10.1.30.17 secondary 10.1.30.18Jul 02 13:04:09:CLISetSNMPtrapinfo:AS_WARNING:Primary tra7
Jul 02 13:04:09:CLISetSNMPtrapinfo:AS_WARNING:Secondary trap host changed to 10.1.30.18

*[SN_5420-MG1]# set admin contactinfo name "Pat J. Smith" email pjsmith@mystoragenet.com phone "763 555-7766" pager "763 555-7766"
*[SN_5420-MG1]# save system
[SN_5420-MG1]#

37, 16 | VT400-7 -- COM3 at 9600 baud

```

Figure 3-13 Initial configuration commands (screen 2)

Table 3-4 Initial configuration commands

Command	Description
set systemname <i>SN_5420-MG1</i>	Configures a name for the management interface. If using DNS this is the name you will associate with the management interface IP address.
set mgmt ipaddress <i>10.1.10.244/24</i>	Configures the IP address of the management interface. This will be prompted automatically once the router is brought up for the first time.

Command	Description
add route <i>10.1.30.0/24 gw 10.1.10.201</i>	(Optional) This adds a gateway IP address so that the SN5420 can be managed from a station outside of the local subnet.
set nameserver <i>10.1.40.243 domain mystoragenet.com</i>	(Optional) Specifies the address and domain of the primary DNS server.
add route <i>10.1.40.243/32 gw 10.1.10.201</i>	(Optional) This adds a gateway IP address so that the SN5420 can access the remote DNS server. note that this is a host route (/32).
set secnameserver <i>10.1.50.249</i>	(Optional) Specifies the secondary DNS server.
add route <i>10.1.50.249/32 gw 10.1.10.201</i>	(Optional) This adds a gateway IP address so that the SN5420 can access the remote secondary DNS server. note that this is a host route (/32).
set clock <i>05:20:00 07/15/2001</i>	Set current time and date.
tzselect	This command will display a menu to select the proper time zone. Enter the correct time zone designation from the list.
set ntpserver <i>10.1.60.86</i>	(Optional) A Network Time Protocol server can be identified with this command.
add route <i>10.1.60.86/32 gw 10.1.10.201</i>	(Optional) This adds a gateway IP address so that the SN5420 can access the NTP server. note that this is a host route (/32).
set snmp getcommunity <i>public</i>	(Optional) Specifies the name of the SNMP community that will have read access to the storage router.
set snmp setcommunity <i>mynetmanagers</i>	(Optional) Specifies the name of the SNMP community that will have write access to the storage router.
set snmp traphosts primary <i>10.1.30.17 secondary 10.1.30.18</i>	(Optional) Specifies the primary and secondary SNMPv1 TRAPS. If this host is outside of the Management Interfaces IP subnet then an add route command is required.

Command	Description
set monitor password <i>sanfran</i>	Sets the monitor password.
set admin password <i>cisco@123</i>	Sets the administrator password.
set admin contactinfo name <i>"Pat J. Smith"</i> email <i>pjsmith@mystoragenet.com</i> phone <i>"763 555-7766"</i> pager <i>"763 555-7766"</i> quotes are not working...showing up as "	(Optional) Adds Contact information.
save system	Save system parameters.

3.5.6 Configuring the high availability (HA) interface

Configuring HA allows a secondary 5420 to take over routing in the event the primary Cisco SN5420 Storage Router fails. Please refer to 3.2.3, "Scalability and reliability" on page 66 for information about the high availability feature.

Important: Both Cisco SN5420 Storage Routers that will participate in HA must be defined with the same cluster name. The management interface as well as the HA interface must be defined on the same network or subnet.

Figure 3-14 and Table 3-5 on page 83 show the HA configuration commands.

```
Cisco Console.r2w - Reflection for UNIX and Digital
File Edit Connection Setup Macro Window Help

[SN_5420-M61]# set ha ipaddress 10.1.20.56/24 Jul 02 14:02:50:NetMgr:AS_DEBUG :Attaching interface 'fe1'
0x60a128 (tNetTask): fe1: link down
                Jul 02 14:02:50:NetMgr:AS_DEBUG :Address 10.1.20.56/24 is now operational on interface fe1

interface mask and address changed
*[SN_5420-M61]# Jul 02 14:02:50:HA:AS_NOTICE :HEARTBEAT SEND FAILURE (62)

*[SN_5420-M61]# set cluster Mt_Olympus

The current configuration of the SN5420 may conflict with the new cluster.
If you choose to merge the configurations, the applications configured on
the SN5420 will be replicated to other SN5420s in the new cluster. This may
cause conflicts in the new cluster.

** Enter CTRL-C to cancel. **

Merge or delete applications ? [merge/delete (delete)] Last message repeated 30 times
delete
                Jul 02 14:03:21:HA:AS_NOTICE :HEARTBEAT SEND F

#####
Please confirm that you want to delete the configuration.

321, 1 | VT400-7 -- COM3 at 9600 baud
```

Figure 3-14 HA configuration commands

Table 3-5 HA configuration commands

Command	Description
set ha ipaddress <i>10.1.20.56/24</i>	Sets the HA port IP address. Must be on the same network/subnet as the other 5420s that will participate in HA.
set cluster <i>Mt_Olympus</i>	This sets the cluster name. This will force a reboot of the storage router. You will have the option to either merge or delete. Merge means that the configuration of this storage router (including SCSI routing services) will be propagated to the other storage routers in the named cluster. Deleting means that the existing configuration will be deleted from the storage router. The cluster name must be the same as the other 5420s that will participate in HA.

3.5.7 Configuring SCSI routing

Storage routing is configured by creating a SCSI router instance. In our example we named the instance “zeus”. See Figure 3-15. Additional steps to configure SCSI routing require us to apply an IP address to the server interface (Gigabit Ethernet port) and identify the topology on the Fibre Channel port. We can then define access control lists (ACL) which limit storage access to specific servers on the IP network. The Figure 3-16 and the command examples in Table 3-6 show how to properly configure SCSI routing in the Cisco SN5420 Storage Router.

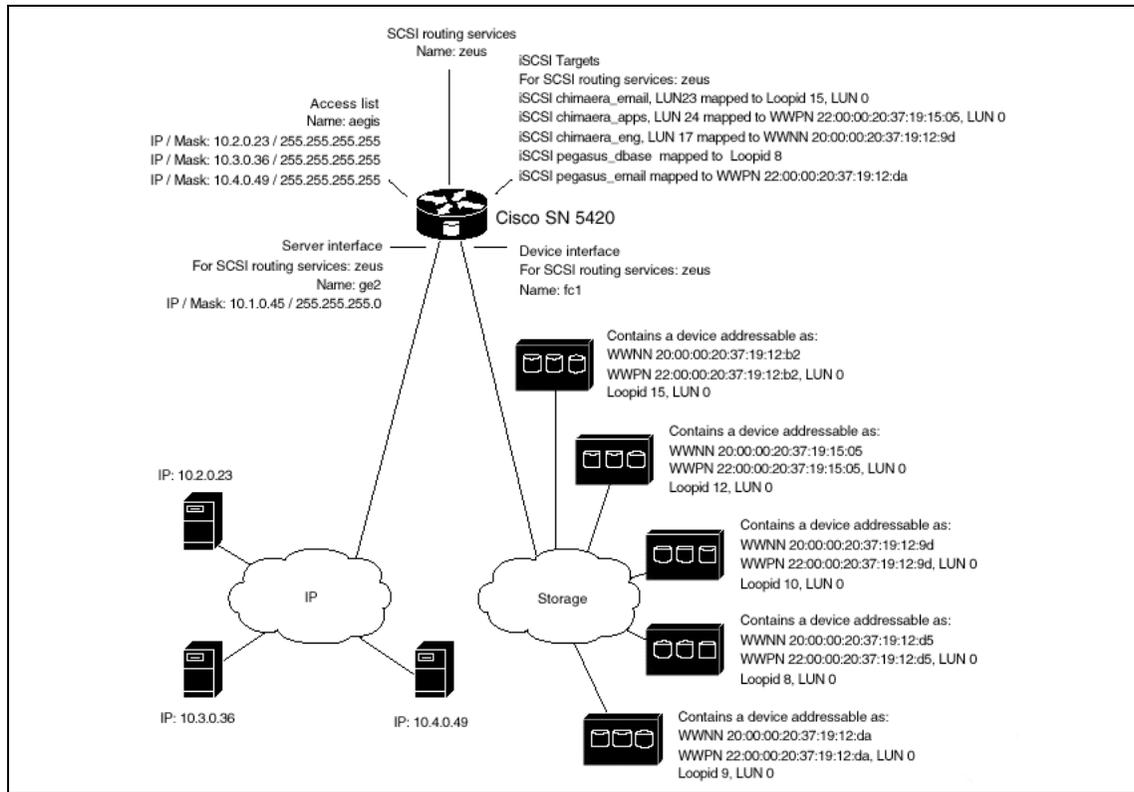


Figure 3-15 SCSI routing example

The following settings for configuration have been taken from the *Cisco SN5420 Software Installation* manual and are reprinted with permission.

When configuring the storage router for SCSI routing, specify these parameters:

- ▶ **Access list:** Specifies the IP addresses of servers allowed to access a common set of storage resources via the storage router. You can create any number of access lists.
- ▶ **Name of SCSI routing services:** Specifies the name of the instance of the SCSI routing services in the storage router. This becomes a binding point for the association of other configuration parameters such as the network IP address, targets, and access lists. You can define up to four SCSI routing service instances per stand-alone SN5420 Storage Router, or per storage router cluster.
- ▶ **Network IP address:** Specifies the storage router IP address that servers will address for access to storage. Specifies the iSCSI target names and LUNs

and the physical storage addresses (controllers and LUNs) mapped to the iSCSI targets and LUNs.

- ▶ **Storage (device) interface:** Specifies the SN5420 interface to be used to interface with storage devices.
- ▶ **Server interface:** Specifies the SN5420 interface to be used to interface with servers.

A mapping of physical devices to iSCSI targets must be completed for configuration. Once the IP address of the storage router has been set in the driver residing on the server, the storage router will forward SCSI requests and responses between servers and the mapped storage devices. That is when the server will recognize its iSCSI storage device.

```
[SN_5420-MG1]# create scsirouter zeusJul 02 13:24:11:ScsiRouter::haCreate:AS_DEBUG :Received scsirouter/zeus from HA
Jul 02 13:24:11:ScsiRouter::haCreate:AS_DEBUG :ScsiRouter config already has scsirouter/zeus
Jul 02 13:24:11:ScsiRouter::createConfNode:AS_DEBUG :registered app instance with HA -- instance: scsirouter/zeus

*[SN_5420-MG1]# Jul 02 13:24:16:ScsiRouter::haStart:AS_NOTICE :starting zeus
Jul 02 13:24:16:ScsiRouter::haStart:AS_NOTICE :zeus has been started

*[SN_5420-MG1]# add scsirouter zeus serverif ge2 10.1.0.45/24
*[SN_5420-MG1]# add scsirouter zeus deviceif fc1Jul 02 13:25:48:srsCreate:AS_NOTICE :zeus started!
Jul 02 13:25:48:NetMgr:AS_DEBUG :Address 10.1.0.45/24 is now operational on interface ge2

*[SN_5420-MG1]# set interface fc1 topology ptp
Topology changed from ptpref to ptp on interface fc1
You must save your configuration and reboot the system for this change to take affect
*[SN_5420-MG1]# create accesslist aegisJul 02 13:26:38:newACL:AS_INFO :Creating accesslist aegis

*[SN_5420-MG1]# add accesslist aegis 10.2.0.23/32Jul 02 13:27:21:newACL:AS_INFO :Adding 10.2.0.23/32 to accesslist aegis

*[SN_5420-MG1]# add accesslist aegis 10.2.0.36/32Jul 02 13:27:45:newACL:AS_INFO :Adding 10.2.0.36/32 to accesslist aegis

*[SN_5420-MG1]# add accesslist aegis 10.2.0.49/32Jul 02 13:28:02:newACL:AS_INFO :Adding 10.2.0.49/32 to accesslist aegis

*[SN_5420-MG1]# save all
[SN_5420-MG1]#
```

Figure 3-16 Commands to create SCSI routing and ACLs

Table 3-6 SCSI routing and ACL configuration commands

Command	Description
create scsirouter <i>zeus</i>	Create a scsirouter instance named zeus.
add scsirouter <i>zeus serverif ge2</i> <i>10.1.0.45/24</i>	Binds the server interface and IP address to the scsirouter service zeus. This is the IP address that iSCSI clients will use to access SCSI routing services.
add scsirouter <i>zeus deviceif fc1</i>	Binds the device interface fc1 to the scsirouter service zeus.
set interface fc1 topology ptp Or set interface fc1 topology loop	Configure the device interface topology to arbitrated loop or point to point.
create accesslist <i>aegis</i>	Create an access list.
add accesslist <i>aegis 10.2.0.23/32,</i> <i>10.3.0.36/32, 10.4.0.49/32</i>	Add IP addresses to the access list. /32 denotes that only the specific IP address is provided access.
save all	Save configuration before rebooting.
reboot	Reload the router to enable the interface topology.

3.5.8 Configuring iSCSI targets

Once SCSI routing and ACLs are defined in the Cisco SN5420 Storage Router you need to add storage target definitions (Table 3-7). There are several ways to map iSCSI targets within a 5420. Target-only and Target-and-LUN mapping are supported using either LoopID, World Wide Port Name (WWPN) or World Wide Node Number (WWNN). An example of Target-and-LUN mapping using WWPN addressing is shown in Figure 3-17.

```

Cisco Console.r2w - Reflection for UNIX and Digital
File Edit Connection Setup Macro Window Help
[SN_5420-M61]# add scsirouter zeus target chimaera_apps lun 31 wwpn
wwpn      wwpn      loopid lun vendor  product
serial number
1: 0000000000000000 10000000c922112f 129  0 IBM   2105F20
   20114830
2: 0000000000000000 10000000c922112f 129  1 IBM   2105F20
   20414830
3: 0000000000000000 10000000c922112f 129  2 IBM   2105F20
   30214830

[SN_5420-M61]# add scsirouter zeus target chimaera_apps lun 31 wwpn #2
Succeeded adding scsirouter zeus target chimaera_apps lun 31 wwpn 10:00:00:c9:22:11:2f lun 1
*[SN_5420-M61]#

```

14, 17 | VT400-7 -- COM3 at 9600 baud

Figure 3-17 Target and LUN mapping using WWPN

Table 3-7 SCSI target configuration commands

Command	Description
add scsirouter zeus target <i>chimaera_email lun 23 loopid 15 lun 0</i> Or add scsirouter zeus target <i>chimaera_email lun 23 loopid ?</i> And add scsirouter zeus target <i>chimaera_email lun 23 loopid #5</i>	Example of Target-and-LUN mapping using LoopID addressing. Specifies iSCSI target and LUN as <i>chimaera_email</i> , LUN 23 and maps it to a physical address LoopID 15, LUN 0. The ? will create a list of physical addresses seen by the 5420, the index number of the corresponding physical address can then entered.

Command	Description
<p>add scsirouter zeus target <i>chimaera_apps lun 24 wwpn</i> <i>22:00:00:20:37:19:15:05 lun 0</i></p> <p>Or</p> <p>add scsirouter zeus target <i>chimaera_apps lun 31 wwpn ?</i></p> <p>And</p> <p>add scsirouter zeus target chimaera_apps lun 31 wwpn #4</p>	<p>Example of Target-and-LUN mapping using WWPN addressing.</p> <p>Specifies iSCSI target and LUN as <i>chimaera_apps</i>, LUN 24 and maps it to a physical address <i>wwpn 22:00:00:20:37:19:15:05</i>, LUN 0.</p> <p>The ? will create a list of physical addresses seen by the 5420, the index number of the corresponding physical address can then entered.</p>
<p>add scsirouter zeus target <i>chimaera_eng lun 17 wwnn</i> <i>22:00:00:20:37:19:12:9d</i></p> <p>Or</p> <p>add scsirouter zeus target <i>chimaera_eng lun 17 wwnn ?</i></p> <p>And</p> <p>add scsirouter zeus target chimaera_eng lun 17 wwnn #3</p>	<p>Example of Target-and-LUN mapping using WWNN addressing.</p> <p>Specifies iSCSI target and LUN as <i>chimaera_eng</i>, LUN 17 and maps it to a physical address <i>wwnn 22:00:00:20:37:19:12:9d</i>.</p> <p>The ? will create a list of physical addresses seen by the 5420, the index number of the corresponding physical address can then entered.</p>
<p>add scsirouter zeus target <i>pegasus_dbase loopid 8</i></p> <p>Or</p> <p>add scsirouter zeus target <i>pegasus_dbase loopid ?</i></p> <p>And</p> <p>add scsirouter zeus target pegasus_dbase loopid #1</p>	<p>Example of Target only mapping using LoopID addressing.</p> <p>Specifies iSCSI target <i>pegasus_dbase</i>, and maps it to LoopID 8.</p> <p>The ? will create a list of physical addresses seen by the 5420, the index number of the corresponding physical address can then entered.</p>

Command	Description
add scsirouter zeus target <i>pegasus_email wwpn</i> <i>22:00:00:20:37:19:12:da</i> Or add scsirouter zeus target <i>pegasus_email wwpn ?</i> And add scsirouter zeus target pegasus_email wwpn #2	<p>Example of Target only mapping using WWPN addressing.</p> <p>Specifies iSCSI target <i>pegasus_email</i>, and maps it to WWPN <i>22:00:00:20:37:19:12:da</i>.</p> <p>The <i>?</i> will create a list of physical addresses seen by the 5420, the index number of the corresponding physical address can then entered.</p>

3.5.9 Configuring server access

Access to iSCSI targets on the Cisco SN5420 Storage Router is controlled by access lists. The access lists contain the IP address of the servers and are enabled by using the set of commands in Table 3-8. We previously created the ACL Aegis to limit access to (3) IP hosts: 10.2.0.23, 10.2.0.36 and 10.2.0.49. See Figure 3-16 on page 85. This next step applies that ACL to the specific targets (Figure 3-18).

```

Cisco Console.r2w - Reflection for UNIX and Digital
File Edit Connection Setup Macro Window Help

[SN_5420-M61]# add scsirouter zeus target chimaera_apps lun 31 wwpn
wwpn      wwpn      loopid lun vendor  product
serial number
1: 0000000000000000 10000000c922112f 129  0 IBM    2105F20
   20114830
2: 0000000000000000 10000000c922112f 129  1 IBM    2105F20
   20414830
3: 0000000000000000 10000000c922112f 129  2 IBM    2105F20
   30214830

[SN_5420-M61]# add scsirouter zeus target chimaera_apps lun 31 wwpn #2
Succeeded adding scsirouter zeus target chimaera_apps lun 31 wwpn 10:00:00:c9:22:11:2f lun 1
*[SN_5420-M61]# set scsirouter zeus target chimaera_apps accesslist aegis
*[SN_5420-M61]#

39, 17 | VT400-7 -- COM3 at 9600 baud

```

Figure 3-18 Command to apply ACL on the SCSI router

Table 3-8 Server access control commands

Command	Description
set scsirouter zeus target chimaera_email accesslist aegis	The iSCSI target chimaera_email can be accessed by servers listed in the access list named aegis using the SCSI routing service zeus.
set scsirouter zeus target chimaera_apps accesslist all	The iSCSI target chimaera_apps can be accessed by all servers using the SCSI routing service zeus.
set scsirouter zeus target all accesslist aegis	All iSCSI targets defined to the SCSI routing service zeus can be accessed by servers listed in the access list named aegis.

Command	Description
set scsirouter zeus target all accesslist all	All iSCSI targets defined to the SCSI routing service zeus can be accessed by all servers using the SCSI routing service zeus.
set scsirouter zeus target chimaera_apps accesslist none	No server can access the iSCSI target named chimaera_apps using the SCSI routing service zeus.
set scsirouter zeus target all accesslist none	No server can access any iSCSI target using the SCSI routing service zeus.

3.5.10 Using the GUI on the Cisco SN5420 Storage Router

Another management option that can be used to configure and control a Cisco SN5420 Storage Router is the graphical user interface. By using a Web browser, insert the IP address that was assigned to the management port of the Cisco SN5420 Storage Router. All functions, except for the initial assignment of the management IP address, can be performed using the GUI. The initial login window, main admin and system maintenance, are shown here for reference (Figure 3-19 through Figure 3-21).

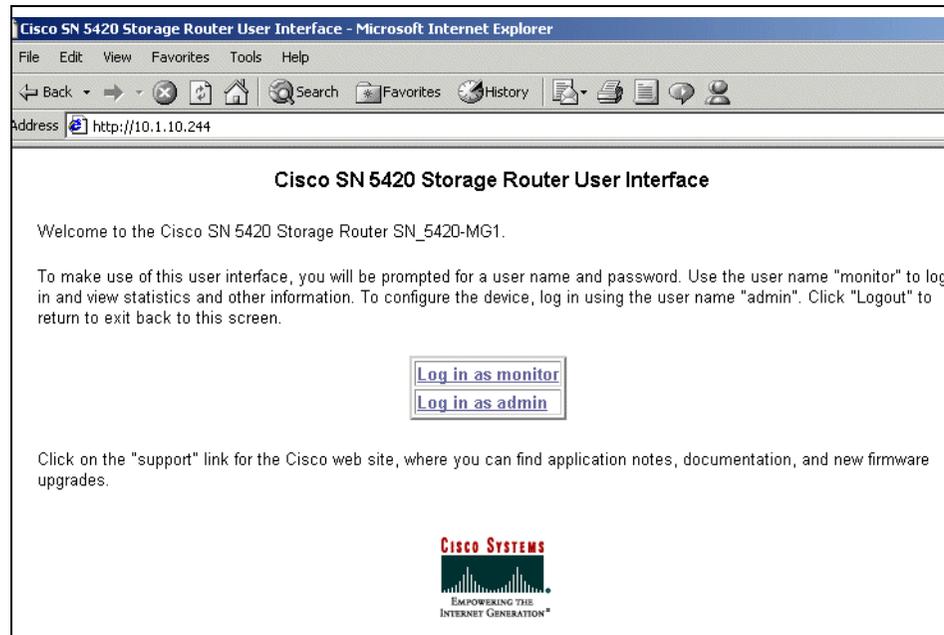


Figure 3-19 Initial login window

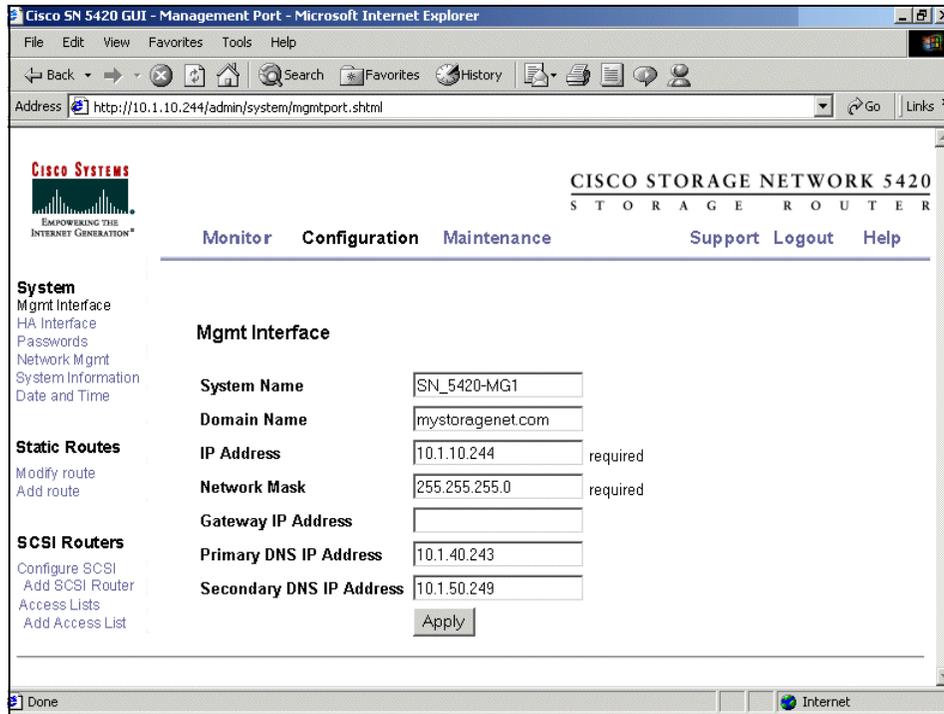


Figure 3-20 Management port configuration

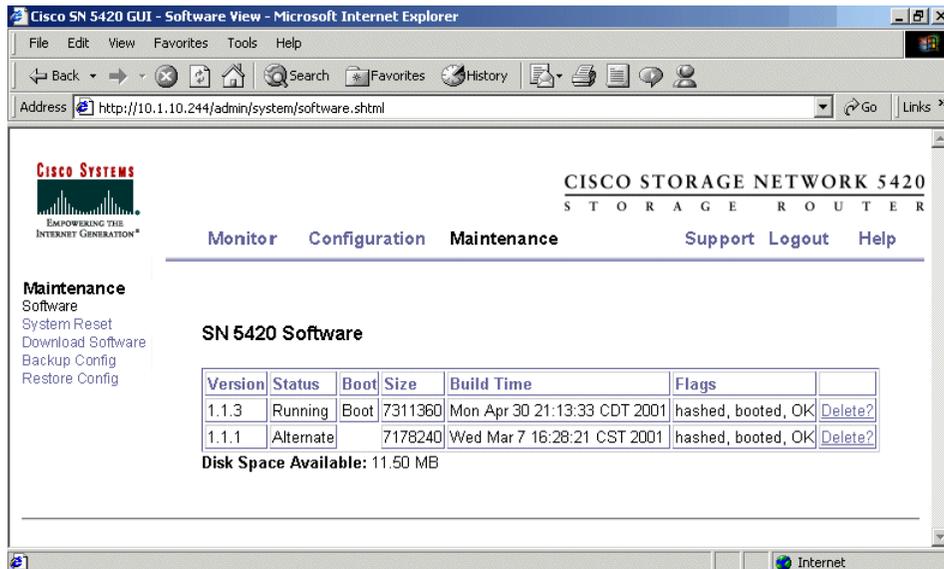


Figure 3-21 System software display

3.6 Host setup

This section describes the steps required to configure IP host access to the Cisco SN5420 Storage Router and storage targets. For the IP hosts to access storage via the Cisco SN5420 Storage Router you must first install and configure client iSCSI drivers specific to their host operating system. Cisco and Cisco Partners have developed or are currently working on iSCSI drivers that support these operating systems (the first four are currently working):

- ▶ Linux
- ▶ Solaris
- ▶ Windows NT
- ▶ Windows 2000
- ▶ AIX (under development by IBM)
- ▶ HP UX (under development by HP)
- ▶ Netware (under development by Novell)

Note: The latest Cisco iSCSI drivers for all hosts is available at:

<http://www.cisco.com/cgi-bin/tablebuild.pl/sn5420>

3.6.1 iSCSI Windows 2000 driver installation and configuration

This section covers the installation of the Cisco iSCSI driver for Windows 2000.

Requirements

The Cisco Storage Networking iSCSI Driver for Microsoft Windows 2000 requires Microsoft Windows 2000 Professional or Server with Service Pack 2 or higher. If Service Pack 2 is installed, Microsoft hotfix Q302895 is required. This fix does not appear on the Microsoft hotfix Web site; you must specifically request it. To obtain these hot fixes you will need to contact Microsoft. This fix will eventually be bundled into Microsoft Windows 2000 Service Pack 3. If Windows 2000 Service Pack 3 has been installed, then these hot fixes are not needed.

When installing the Windows 2000 iSCSI driver, we needed the following hot fixes from Microsoft:

- ▶ Q302895 - Program that runs multiple I/O operations to multiple Fibre Channel disks may hang.
- ▶ Q297961 - SCSI bus reset update.
- ▶ Q280114 -ScsiPort hangs after it fails a request, because the target was busy.

To install the iSCSI driver, follow these instructions:

1. Run the Setup.exe file to install the iSCSI driver.

2. Type in the IP address or host name of the Gigabit Ethernet interface on the Cisco 5420 and click **Add**. The address should appear in the list (Figure 3-22).
3. Click **OK**, and you will be prompted to restart the computer. Restart the computer.

The configuration is now complete.

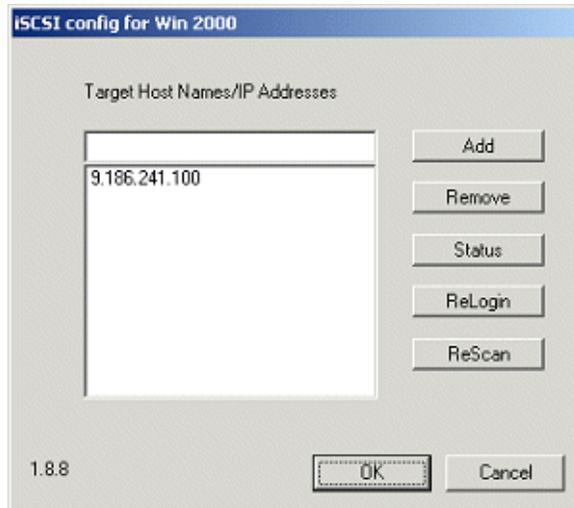


Figure 3-22 Windows 2000 iSCSI configuration

Assigning a drive letter to the iSCSI device

1. Follow this sequence to assign a drive letter from Disk Management (Figure 3-23):

Click **Start -> Programs -> Administration Tools -> Computer Management -> Storage**

Figure 3-23 Windows 2000 disk administrator

You should now see Disk Management.

2. There will be a disk that is offline which is the iSCSI partition. In order to bring it online, right click on the offline disk and choose Reactivate Disk (Figure 3-24).

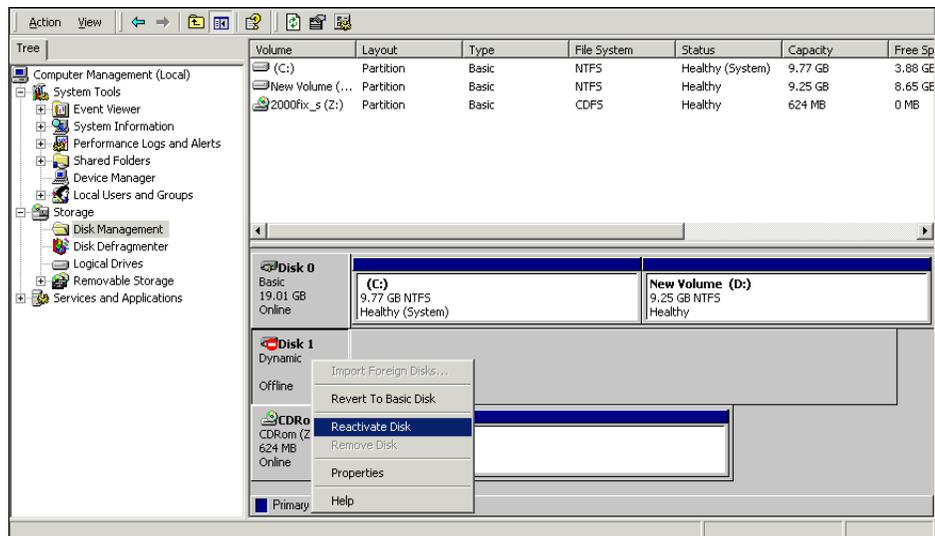


Figure 3-24 Reactivating a disk

3. The disk is now mapped to the next available drive letter. In order to use the disk, the drive must be formatted. We suggest you change the drive letter to

be further down in the alphabet, so that if another local drive is connected it will not interfere with the iSCSI drive.

Note: As of updating this document the latest Windows 2000 iSCSI driver from Cisco available is release 1.8.10.

3.6.2 iSCSI Windows NT driver installation and configuration

In this example we will walk through a Windows NT installation and configuration. All available drivers and complete installation documentation are available for download at Cisco Connection Online (CCO):

<http://www.cisco.com/pgi-bin/tablebuild.pl/sn5420>

These instructions are based on Cisco Windows NT iSCSI driver version 1.8.6.

The Cisco SN iSCSI server driver for Microsoft Windows NT requires Windows NT Server (or Workstation) Version 4.0 with Service Pack 6a or later.

iSCSI driver installation steps for Windows NT

1. Download the latest iSCSI driver files. Store and unzip these files into a directory. Follow the installation instructions in the readme.txt file.
2. Once the iSCSI drivers are installed on the Windows NT system, open the Control Panel and double-click the iSCSI config icon. At the iSCSI config dialog box, enter the IP address of the SN5420 Storage Router and click Add.
3. Restart the server.
4. Once the server has restarted, run Disk Administrator. This will cause Windows NT to discover the iSCSI devices. Disk Administrator is found under:
Start->Programs->Administrative Tools->Disk Administrator
5. Using Disk Administrator, configure hard disk partitions and drive letters for the discovered iSCSI devices.
6. At the iSCSI Config dialog box, click Save Boot Info. This will save the device partition and drive letter assignment in the registry.

Note: At the time of writing this redbook, the latest Windows NT iSCSI driver available from Cisco was Release 1.8.9.

3.6.3 iSCSI UNIX and Solaris driver installation and configuration

In this example we will walk through a UNIX/Solaris iSCSI driver installation and configuration.

These instructions are based on Cisco Solaris iSCSI driver Version 1.8.7.

You must be a super-user (root) to install the iSCSI package. The Cisco iSCSI driver for SUN servers runs only on SPARC machines. The driver requires either Solaris Version 7 or 8 with all Sun recommended patches installed for the specific version used. Because the network connections usually go down first in a system shutdown, any iSCSI file systems should be unmounted prior to shut down.

iSCSI driver installation steps for UNIX/Solaris

1. Download the iSCSI driver file. Copy the `solaris_iscsi_<ver>.tar.Z` file to an empty working directory, such as `/usr/local/NuScsiTcp`. The `<ver>` is the three digit version, for example, 0.8.1.
2. Change to the working directory created in Step 1 and untar the file using the tar command.

```
cd /usr/local/NuScsiTcp
uncompress solaris_iscsi_<ver>.tar
tar xvf solaris_iscsi_<ver>.tar
```

The working directory now contains the iSCSI package.

3. Add the package.

```
cd /usr/local/NuScsiTcp
pkgadd -d . CSC0iscsi
Adding the package will install all the files to their correct locations.
/etc/rc0 . d/K08iscsi (iSCSI shut down script)
/etc/rc3.d/S78iscsi (iSCSI start up script)
/etc/NuScsiTcp.conf (iSCSI configuration file)
/kernel/drv/sd.conf (iSCSI SCSI disk configuration)
/kernel/drv/nupseudo (iSCSI psuedo driver)
/kernel/drv/nuscsitcp (iSCSI HBA driver)
/kernel/strmod/nustrmo (iSCSI stream module)
/usr/local/bin/nuscsitcpd (iSCSI daemon)
```

The package install will note that the system needs to be rebooted, perform the following steps before rebooting.

4. Edit the `/etc/NuScsiTcp.conf` file to insert the correct SN5420 IP address. There are two types of entries in the configuration file, "TargetIpAddr" and "TargetHiWater".

TargetIPAddr specifies the SN5420's IP address. The configuration file can have up to eight different remote connections. The value of `TargetIpAddr` can either be an IP address or a hostname. The format of this entry is:

```
TargetIpAddr=remote_address
```

For example: `TargetIpAddr=10.1.0.45`

TargetHiWater specifies an optional configuration for the TCP transmit and receive window size. The format of this entry is:

```
TargetHiWater:remote_address=high_water_value
```

For example: TargetHiWater:10.1.0.45=1048576

5. The file, /kernel/drv/sc.conf, was modified with the pkgadd command. A generic set of targets and LUNs were added for the iSCSI driver. Edit this file to include only the applicable targets and LUNs.
6. The order of the boot script may need to be modified because the iSCSI daemon tries to initiate a TCP connection to each configured SN5420 and if the network interface has not started by the time the boot script runs, the connection will not occur.
7. Reboot the system to start and reload the iSCSI driver.

If you do not wish to reboot the machine, the following command will manually start the iSCSI devices immediately:

```
/etc/rc3.d/S78iscsi start
```

To manually stop the iSCSI driver, enter:

```
/etc/rc0.d/K08iscsi stop
```

Start and stop scripts are Bourne shell executables. If you are running csh, use "sh /etc/rc3.d/S78iscsi start".

8. The first reboot should reconfigure devices. Otherwise, the configuration commands (drvconfig, disks) may be run manually. Normal disk commands like format, newfc, mkfs, and fsck will work on the iSCSI devices just like they work on a local hard drive.
9. To remove the iSCSI package enter:

```
pkgrm CSC0iscsi
```

Note: At the time of writing this redbook, the latest Solaris iSCSI driver available from Cisco was Release 1.8.9.

3.6.4 iSCSI Linux driver installation and configuration

In this example we will walk through a Linux iSCSI driver installation and configuration. These instructions are based on Cisco Linux iSCSI driver Version 1.8.5.

This driver requires Linux kernel Version 2.2.16-3 or higher. Compiling the driver requires some include files from the kernel. The appropriate "kernel-headers" rpm package, if using Redhat Linux, must be installed on the compiling host. If you do not have the kernel-headers package or the kernel source package

installed, you can download the source code from the kernel that is currently running from <http://kernel.org>. You must ensure that the kernel source is extracted to “/usr/src/linux” so that the driver will compile correctly. To determine the current kernel version that is running on your system run the following command:

```
#cat /proc/version
Linux version 2.2.16-22 (root@porky.devel.redhat.com) (gcc version egcs-2.91.66
19990314/Linux (egcs-1.1.2 release)) #1 Tue Aug 22 16:49:06 EDT 2000
#
```

Once compiled, the objects and executables can be moved to another host running the same level of operating system. Since the network connections usually go down first in a system shutdown, any iSCSI file systems should be unmounted prior to shut down.

iSCSI driver installation steps for Linux

1. Download the iSCSI driver file. Copy the linux_iscsi_<ver>.tgz file to an empty working directory, such as /usr/src/NuScsiTcp. The <ver> is the three digit version, for example, 0.8.1.

```
mkdir /usr/src/NuScsiTcp
cp /tmp/linux_iscsi_<ver>.tgz /usr/src/NuScsiTcp
```

2. Change to the working directory created in Step 1 and untar the file using the tar command.

```
cd /usr/src/NuScsiTcp
tar xzvf linux_iscsi_<ver>.tgz
```

The working directory now contains the iSCSI driver and associated files.

3. Compile the iSCSI driver.

```
make clean
make
```

4. Install the iSCSI driver. You have to be a super-user (root) to install the driver.

```
make install
```

The install process does not over write an existing /etc/NuScsiTcp.conf. Make sure that any additional features, from the latest release, are added to the existing configuration file, by reviewing the new packages NuScsiTcp.conf file.

5. Edit the /etc/NuScsiTcp.conf file to insert the correct SN5420 IP address. There are two types of entries in the configuration file, “TargetIPAddr” and “Target,Lun”.

TargetIPAddr specifies the SN5420’s IP address. The configuration file can have up to eight different remote connections.

The value of TargetIpAddr can either be an IP address or hostname. The format of this entry is:

```
TargetIpAddr=remote_address
```

For example: TargetIpAddr=10.1.0.45

Target,Lun are optional entries and are needed only if you want to specifically activate the specified LUNs. The first number is the target and the second is the LUN. Targets start at zero and increment by one, LUNs start at zero and increment by one. If the configuration file does not contain any “Target,Lun” entries then all the Targets and LUNs are probed.

The format of the entry is:

```
Target,Lun=0,0
```

6. If you are using a non-Redhat Linux distribution you may need to edit your boot scripts to properly run the iSCSI setup script. You may also need to change the order of the boot script, because the iSCSI daemon tries to get a TCP connection to each configured SN5420, and if the network interface has not started by the time the boot script runs, the connection will not occur.
7. Reboot the system to start and reload the iSCSI driver. If you do not want to reboot your machine, you may use the following command to start the iSCSI drivers immediately:

```
/usr/src/NuScsiTcp/rc.NuScsiTcp start
```

To manually stop the iSCSI driver enter:

```
/usr/src/NuScsiTcp/rc.NuScsiTcp stop
```

Normal disk commands like fdisk, mkfs, and fsck will work on the iSCSI devices just like they work on a local hard drive.

8. To optionally change the TCP window for the iSCSI connection modify the following values prior to starting the iSCSI daemon. You may want to put them in the SCSI rc script. Normally, the larger the window the better the performance.

```
echo "1048576" > /proc/sys/net/core/wmem_default
echo "1048576" > /proc/sys/net/core/wmem_max
echo "1048576" > /proc/sys/net/core/rmem_default
echo "1048576" > /proc/sys/net/core/rmem_max
```

9. To remove the iSCSI driver, enter:

```
make remove
```

This deletes the files from /lib/modules, /etc and /bin.

Back up one directory and delete the working directory.

```
cd ..
rm -fr /usr/src/NuScsiTcp
```

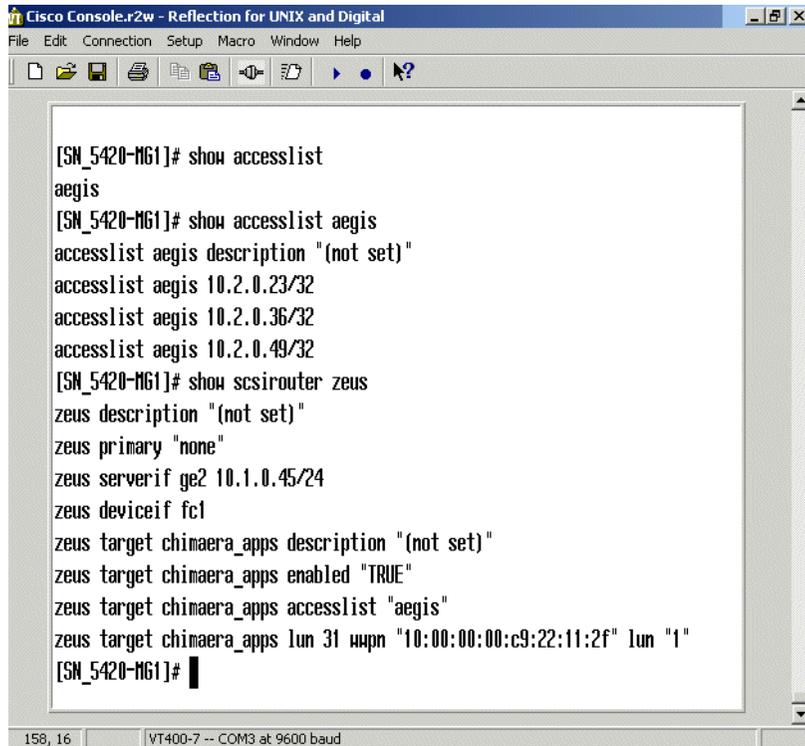
Note: At the time of updating this redbook, the latest Linux iSCSI driver from Cisco was Release 1.8.7.

3.6.5 Troubleshooting

Table 3-9 lists several show commands that are available to display system variables and verify proper operation of the Cisco SN5420 Storage Router. Also see the accompanying figures (Figure 3-25 through Figure 3-29).

Table 3-9 System troubleshooting commands

Command	Description
show accesslist	Verify an access list exists.
show accesslist <i>aegis</i>	Display entries in the access list named aegis.
show scsirouter <i>zeus</i>	Display configuration parameters for the SCSI routing instance named zeus.
show system	Displays system management and configuration variables.
show net route	Displays the system route table.
show snmp	Displays SNMP variables.
show admin contact info	Displays contact information.
show cluster	Displays cluster information.



The screenshot shows a Cisco console window titled "Cisco Console.r2w - Reflection for UNIX and Digital". The window contains the following text:

```
[SN_5420-M61]# show accesslist
aegis
[SN_5420-M61]# show accesslist aegis
accesslist aegis description "(not set)"
accesslist aegis 10.2.0.23/32
accesslist aegis 10.2.0.36/32
accesslist aegis 10.2.0.49/32
[SN_5420-M61]# show scsirouter zeus
zeus description "(not set)"
zeus primary "none"
zeus serverif ge2 10.1.0.45/24
zeus deviceif fc1
zeus target chimaera_apps description "(not set)"
zeus target chimaera_apps enabled "TRUE"
zeus target chimaera_apps accesslist "aegis"
zeus target chimaera_apps lun 31 wwpn "10:00:00:00:c9:22:11:2f" lun "1"
[SN_5420-M61]#
```

At the bottom of the window, the status bar shows "158, 16" and "VT400-7 -- COM3 at 9600 baud".

Figure 3-25 System troubleshooting commands (screen 1)

```

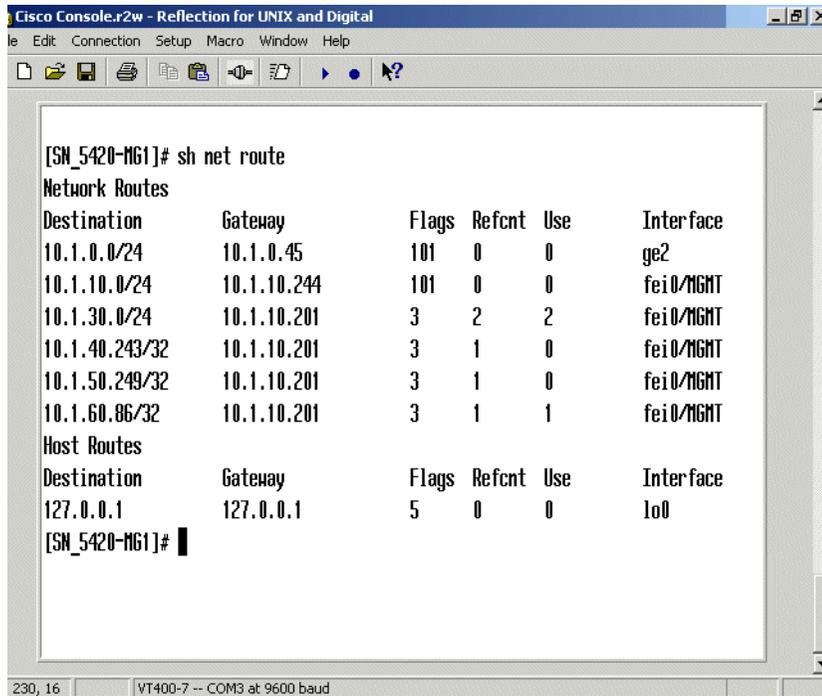
Cisco Console.r2w - Reflection for UNIX and Digital
le Edit Connection Setup Macro Window Help
[SN_5420-M61]# sh system
  System Name: SN_5420-M61
  System Model: Cyclone
  Disk Capacity: 30.50 Mbytes
  Free Disk Space: 11.50 Mbytes
  Software Version: 1.1.3
  Last Reset: Mon Jul  2 14:07:19 GMT 2001
  Current Time: Mon Jul  2 14:09:39 GMT 2001
  Time Zone: none
  NTP Server: 10.1.60.86
  Name Servers: 10.1.40.243(Pri) 10.1.50.249(Sec)
  Domain: mystoragenet.com

          Model Number  Rev  Serial Number
System    SN5420-FCM-GEM    01  SAD05030PAU
Processor CYCLONE-E4          01  SAP05065C9D
Fibre Channel      FCM    01  VMT01030136
Gigabit Ethernet   GEM    15  SAD045105BY

Device  IP/Netmask          MAC
fei0/MGMT 10.1.10.244/24      00:01:64:40:ec:40
lo0      127.0.0.1/8         00:00:00:00:00:00
fei1/HA  10.1.20.56/24       00:01:64:40:ec:41
fc1
ge2
644, 1 | VT400-7 -- COM3 at 9600 baud

```

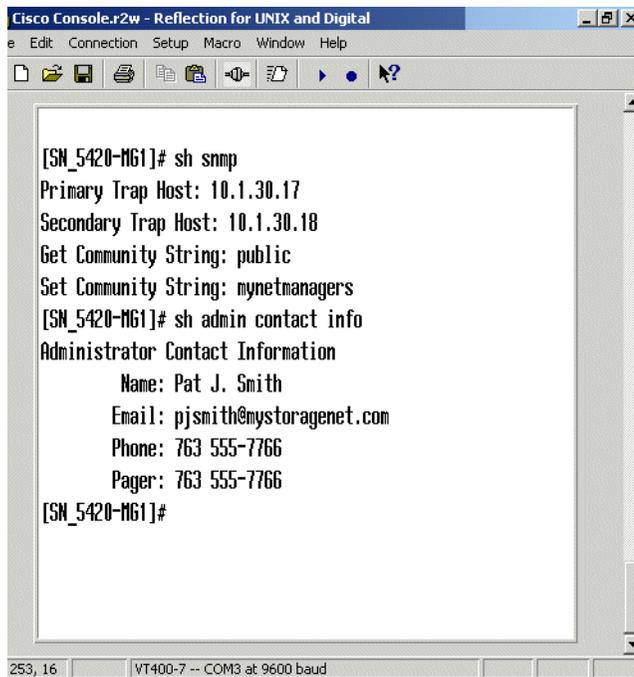
Figure 3-26 System troubleshooting commands (screen 2)

The image shows a screenshot of a Cisco console window titled "Cisco Console.r2w - Reflection for UNIX and Digital". The window has a menu bar with "le", "Edit", "Connection", "Setup", "Macro", "Window", and "Help". Below the menu bar is a toolbar with various icons. The main area of the window displays the output of the command "[SN_5420-M61]# sh net route". The output is divided into two sections: "Network Routes" and "Host Routes".

```
[SN_5420-M61]# sh net route
Network Routes
Destination      Gateway          Flags  Refcnt  Use    Interface
10.1.0.0/24      10.1.0.45       101    0        0      ge2
10.1.10.0/24     10.1.10.244    101    0        0      fei0/MGMT
10.1.30.0/24     10.1.10.201    3      2        2      fei0/MGMT
10.1.40.243/32  10.1.10.201    3      1        0      fei0/MGMT
10.1.50.249/32  10.1.10.201    3      1        0      fei0/MGMT
10.1.60.86/32   10.1.10.201    3      1        1      fei0/MGMT
Host Routes
Destination      Gateway          Flags  Refcnt  Use    Interface
127.0.0.1        127.0.0.1       5      0        0      lo0
[SN_5420-M61]#
```

At the bottom of the window, there is a status bar showing "230, 16" and "VT400-7 -- COM3 at 9600 baud".

Figure 3-27 System troubleshooting commands (screen 3)



The screenshot shows a Cisco console window titled "Cisco Console.r2w - Reflection for UNIX and Digital". The window contains the following text:

```
[SN_5420-M61]# sh snmp
Primary Trap Host: 10.1.30.17
Secondary Trap Host: 10.1.30.18
Get Community String: public
Set Community String: mynetmanagers
[SN_5420-M61]# sh admin contact info
Administrator Contact Information
      Name: Pat J. Smith
      Email: pjsmith@mystoragenet.com
      Phone: 763 555-7766
      Pager: 763 555-7766
[SN_5420-M61]#
```

At the bottom of the window, the status bar displays "253, 16" and "VT400-7 -- COM3 at 9600 baud".

Figure 3-28 System troubleshooting commands (screen 4)

```

Cisco Console.r2w - Reflection for UNIX and Digital
le Edit Connection Setup Macro Window Help

[SN_5420-M61]# sh cluster
Cluster Name: Mt_Olympus
HA State: RUNNING
Cluster Changes: 2
Last Change: Mon Jul 2 14:08:00 GMT 2001

Sent 410 heartbeats
Rcvd 0 heartbeats

Node ID      fei0/MGMT IP      fei1/HA IP  Last Heard From
6440ec40    10.1.10.244      10.1.20.56  Self

Application Name      IP  State      Last Config Update

[SN_5420-M61]#
703, 16 | VT400-7 -- COM3 at 9600 baud

```

Figure 3-29 System troubleshooting commands (screen 5)

3.7 Cisco SN5420 Storage Router implementation

In this implementation model, we describe the various components used to implement DB2. The components and topics are:

- ▶ IBM Enterprise Storage Server (ESS)
- ▶ McData ED-5000 Enterprise Fibre Channel Director
- ▶ Cisco SN5420 IP Storage Router
- ▶ Cisco Catalyst 3500 Series XL switches
- ▶ Microsoft Windows AutoExNt-Service
- ▶ Creating a database on ESS disk LUNS

Figure 3-30 on page 109 displays a high level logical view of the Cisco SN5420 IP Storage Router model implementation used for this redbook.

Important: DB2 Binaries were not installed on external disk in this model. The Cisco solution is not implemented as a service, so Microsoft Windows AutoExNt-Services could not initiate the disk mounts prior to DB2 services startup.

3.7.1 IBM Enterprise Storage Server (ESS)

The IBM Enterprise Storage Server (ESS) is the ultimate SAN utility, providing the information “fuel” that runs the e-business “engine”. Extensive heterogeneous server connectivity makes the ESS a natural fit for server consolidation requirements. The ESS supports rapid universal access to vast quantities of data through many advanced functions and features, making it a workhorse for support of business intelligence and other business-critical applications. Because of its enterprise-wide support and management scope, the ESS is tailor made to help provide consistent, efficient and effective enterprise resource planning. Using the ESS to address any or all strategic and tactical business initiatives will give an organization the business advantage needed to survive and thrive in the e-world.

For more information about this product, please refer to:

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

3.7.2 McData ED-5000 Enterprise Fibre Channel Director

To help provide high data-availability across a SAN, IBM offers the McDATA Enterprise Fibre Channel Director, which provides the scalability demanded by rapidly growing e-business and other mission-critical applications. The Director design is based upon the IBM ESCON Director (designed and manufactured for IBM by McDATA), which has provided the level of industry-leading data availability, performance, and data integrity required by today's most demanding data centers.

The Director features redundant power supplies and cooling fans, along with redundant active components that support automatic failover. Other high-availability capabilities include hot swapping for all field-replaceable units, nondestructive firmware updates, and automatic fault detection and isolation. In addition, “call-home” and e-mail capabilities automatically alert support and maintenance personnel to accelerate problem resolution.

Directors with the High Availability Option provide redundancy for all active electronic components, helping to enable continuous data access and high performance in the event of a single component failure. Together, these features help enable nondestructive maintenance and upgrades without loss of data access. You can use multiple Directors to create a scalable enterprise SAN backbone that supports consolidated storage management applications, such as disk sharing, tape pooling, and enterprise-wide data sharing.

For more information about this product, please refer to:

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

3.7.3 Cisco SN5420 IP Storage Router

The Cisco SN5420 Storage Router provides access to storage over IP networks. With the SN5420 Storage Router, you can directly access storage via an IP network. The robustness of the (TCP/IP) protocol suite, including security, availability, manageability, and quality of service (QoS), is available when you access storage. The SN5420 Storage Router from Cisco Systems enables the Small Computer Systems Interface over IP (iSCSI), which is the first storage implementation based on IP standards and interoperability.

For more information about this product, please refer to:

<http://www.cisco.com>

3.7.4 Cisco Catalyst 3500 Series XL switches

The Cisco Systems Catalyst 3500 series XL is a scalable line of stackable 10/100 and Gigabit Ethernet switches. This line of switch solutions is appropriate for midsize networks deploying business applications, such as IP based networking.

For more information about this product line, please refer to:

<http://www.cisco.com>

3.7.5 Microsoft Windows AutoExNt-Service

The Microsoft Windows AutoExNt-Service may be implemented in the model. We noticed that at times there was a delay in the acquisition of the logical drives. A `sleep.exe` was the only inserted command in the `AutoExNt.bat` file. Specific implementations may require additional testing.

For more details, please refer to 2.2.2, “Disk drive addressing” on page 34.

3.7.6 Creating a database on ESS disk LUNS

Since the ESS LUNS are presented to the Windows disk administrator as a logical drive, the database creation procedure in Example 5-1 on page 169 may be used, as well as on any logical drive.

iSCSI via Cisco SN5420

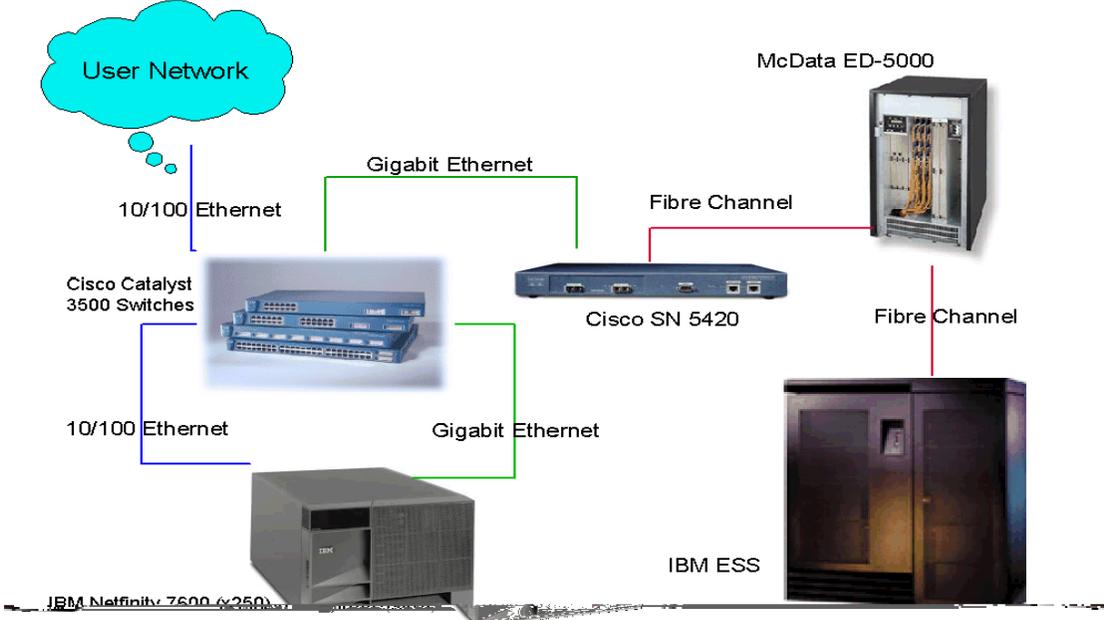


Figure 3-30 iSCSI via Cisco SN5420, Fibre Channel to IBM ESS



Implementing Enterprise Mail for iSCSI

One of the most common solutions being used in companies today is Enterprise mail. The Enterprise mail functions have expanded to more than just messaging or e-mail. Enterprise mail is now used for collaboration and day-to-day processes in replacement of paper signatures and approval. There are many different Enterprise mail applications in the market. In this chapter we will focus on Lotus Notes Domino and Microsoft Exchange.

We will limit this chapter to providing information about installation, configuration, and administration of Microsoft Exchange 2000 and Lotus Domino on Windows 2000 Server, in conjunction with the IBM TotalStorage IP Storage 200i storage unit.

4.1 Planning the Enterprise Mail installation

The following factors need to be considered in the planning/implementing stage of Enterprise mail.

- ▶ Validate that the required hardware and software is installed on the target server, because this will impact on the success/failure of the installation of Exchange 2000/Lotus Domino.
- ▶ Validate that all the required Windows 2000 components and services are installed and are functioning correctly.

- ▶ Test period

In this period, it is very important that all the applications and servers get tested to their full capability to minimize problems in a live environment.

- ▶ Pilot Site

This is an important phase, because a lot of potential problems and installation errors can be picked up and corrected before cutting high profile users over to the new installation. From this phase, product documentation can be generated, such as:

- Deployment parameters
- Training material
- Required user training
- Required support

- ▶ Cut over period

- ▶ System administration and tools

In this phase, we recommend that the system administrator's skills are identified and the required training given for use of system and user support. These need to be addressed:

- Mail box administration/creation
- System policies for Exchange 2000 (max. mail box size)
- Mail box size limitation for Lotus Domino
- Connection to the Internet for mail routing
- Windows 2000 or Pre-Windows 2000 security

- ▶ Check the components that will be selected for Exchange 2000 or Exchange 2000 Enterprise Server, as shown in Table 4-1.
- ▶ Decide which Lotus Domino software package to install: Domino, Domino Mail Server, Domino Application Server, Domino Enterprise Server.
- ▶ Provide support to end users

This requires that the skill levels of the employees are identified, because this will determine the required support and training that will be given to the support personnel and end users. This may include:

- Telephone support
- On-site support
- Time frame when support will be available

▶ The amount of servers

It is important that service, availability, and scalability are included when planning the amount of servers that will be deployed. This is dependent on customer need and downtime impact on the business.

▶ Whether or not to use RAID and what level of RAID

A cost versus performance study needs to be done to determine what level of RAID, if implemented, will be used. This must be completed in the test phase.

▶ Required hard disk space

The following factors need to be carefully planned, because this has a big impact on the server performance and the environment in which it operates:

- Amount of RAM
- Number of CPUs and speed
- Amount of network cards and speed
- Performance monitoring
- Performance tuning

Table 4-1 Exchange 2000 Server/Enterprise options

Function	Exchange 2000	Exchange 2000 Enterprise
Interrogation with Active Director	X	X
Messaging and Collaboration	X	X
Outlook 2000 SR release	X	X
Outlook for Macintosh 8.2.2	X	X
Microsoft Web Storage System	X	X
Connectors to Lan based e-mail systems, MSMail, cc:Mail, Notes/Domino /GroupeWise	X	X
Distributed Configuration (Front-end Back-end Server)		X
POP3 and IMAP4	X	X

Function	Exchange 2000	Exchange 2000 Enterprise
HTTP Outlook Web Agent	X	X
SMTP Connector		X
X.400 Connector	X	X
Unlimited Message Storage		X
Multiple Stores per Server		X
Windows Clustering		X
Instant Messaging	X	X
Chat		X

4.2 Prerequisites for installing Exchange 2000

In this section we will discuss the minimal hardware and software requirements as well as some factors that need to be considered when installing Exchange 2000.

Test environment

The environment was set up accordingly, as shown in Table 4-2.

Table 4-2 Setup information for Exchange 2000

DNS Name	TONGA.ISCSI.RED	LOCHNESS.ISCSI.RED
Operating System	Windows 2000 Advanced Server SP2	Windows 2000 Advanced Server SP2
Domain Role	Domain Controller, DNS Server	Member Server
Additional Software	Window 2000 IIS Hot Fix	Exchange 2000 enterprise server, Window 2000 IIS Hot Fix, ADC
Additional Windows Components	DNS	NNTP
CPU Speed	P III 800MHZ	P III 800MHZ
Memory	768MB	768MB

DNS Name	TONGA.ISCSI.RED	LOCHNESS.ISCSI.RED
Network Card	10/100MB Ethernet	10/100MB Ethernet
Additional Hardware		200i Network attached storage device

Important: It is important that the IIS Hot Fix is installed because there is security exposure to viruses; in this case, we used:

`Q301625_W2K_SP3_x86_en.EXE`

4.2.1 Hardware requirements

Hardware requirements are diverse for each organization. These are based on minimal installation requirements, the amount of users, scalability, and availability. Here are the factors:

- ▶ Minimum of a Pentium-compatible 300 MHz or higher (500 MHz or above recommended)
- ▶ Minimum of 528 MB of RAM (512 MB or higher recommended up to 4GB)
- ▶ 500 MB of free hard disk space on the target drive
- ▶ 200 MB of free hard disk space on the system drive (recommended 4 GB of hard disk space, in total, is made available)
- ▶ CD-ROM
- ▶ VGA monitor

Tip: The total amount of hard disk space required is equal to the total number of users times the max. mail box size set by the Exchange administrator.

4.2.2 Software requirements

Exchange 2000 will only run on Windows 2000 server, Windows 2000 advanced server and Windows 2000 Datacenter Server.

- ▶ Increase the page file to double the amount of RAM on the server that will host Exchange 2000, as shown in Figure 4-1.
- ▶ TCP/IP Protocol
- ▶ Internet Information Services (IIS) with Network News Protocol (NNTP) selected as a additional selection

Tips:

- ▶ We recommended that the page file is moved from the system drive to a separate local drive, for performance enhancement.
- ▶ Download the latest Service Pack, IIS and NNTP fix's for Windows 2000, along with service pack one for Exchange 2000 at:

<http://www.microsoft.com/servers>

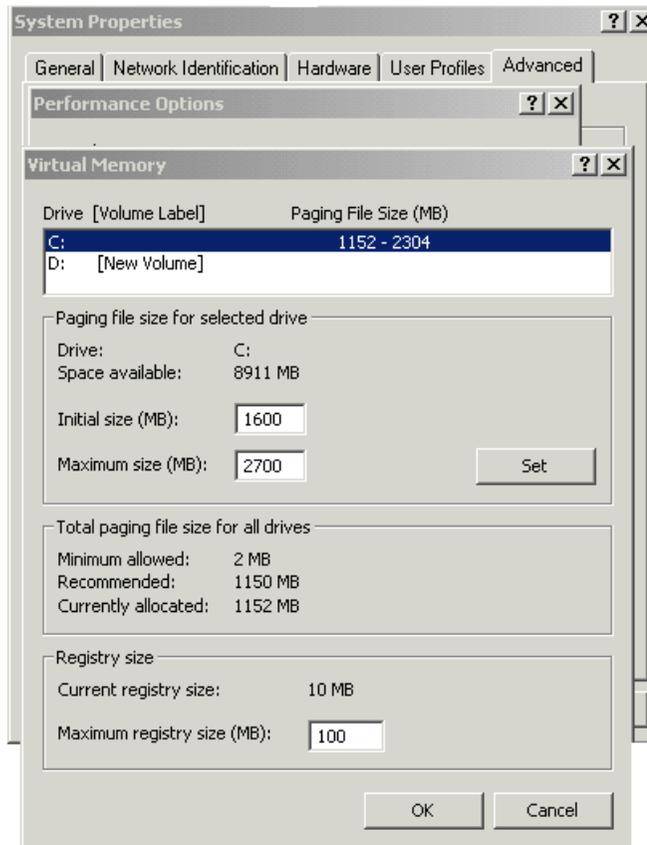


Figure 4-1 Increasing page file size

We recommend that Exchange 2000 installation is done on a NTFS partition for performance and security enhancement reasons.

Make sure that service pack two, (minimum of service pack one) is installed on the server that is going to host Exchange 2000. The installation will not run if the minimum service packs are not installed.

The server that will host Exchange 2000 must be a member server of an Active Directory Domain; if not, it will exit with an error.

Internet Information Services (IIS) is a very important for the installation of Exchange 2000, it provides all of the Internet message transport protocols for Exchange 2000. Simple Mail Transfer Protocol (SMTP), Hypertext Transfer Protocol (HTTP) and Network News Protocol (NNTP) protocols are extended to support the advanced features of Exchange 2000.

Make sure that SMTP and HTTP services are installed; the services are part of the default installation. NNTP services are an additional selectable components and have to be installed and verified that they are working, if NNTP services are not installed, Exchange 2000 will exit with a warning.

Installation of NNTP can be done by running the Add/Remove Windows components:

- ▶ Selecting IIS
- ▶ Select NNTP Service, as shown in Figure 4-2, for Windows 2000 servers.
- ▶ Verify that the NNTP service is running; this can be accomplished by opening a command prompt and telneting in to the local host on port 119:

```
Telnet localhost 119
```

A telnet session will open; if it fails to open make sure that NNTP services are installed and running by checking that NNTP services have started. To end the telnet session type:

```
Quit
```

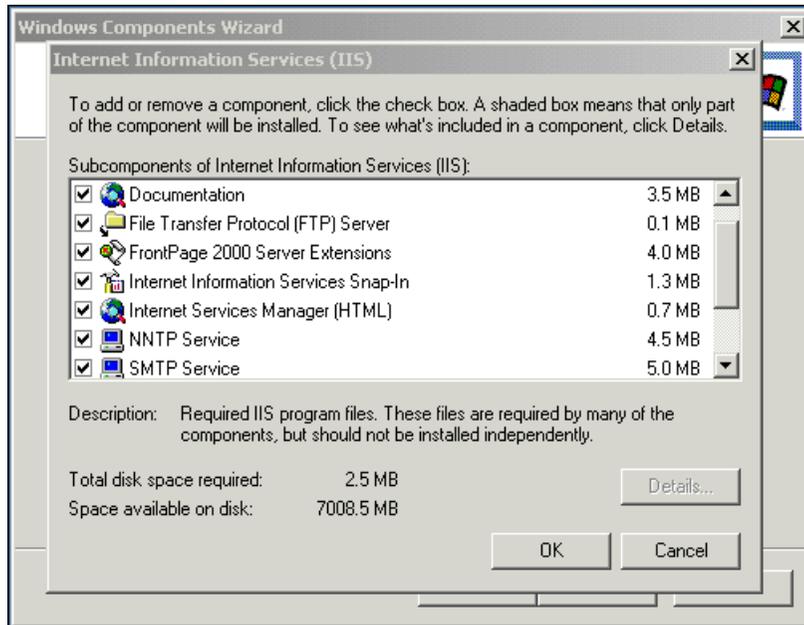


Figure 4-2 IIS selection page

4.2.3 Before Exchange 2000 installation

Exchange 2000 can be installed from the installation guide that runs when the Exchange CD is installed into the drive. This is not advisable for the first Exchange server in the domain. The forestprep and domainprep command needs to be run on this server.

Keep the installation and system files local and create the mail stores on the iSCSI unit.

4.2.4 Installing the Active Directory Connector

When installing a server running Exchange 2000 into an Exchange 5.5 site, it is important to install the Active Directory Connector (ADC) and verify that it is working.

The server/workstation needs the ADC from where the Exchange 2000 management and mail box creation tasks will be done, and this includes the Domain Controller. If the ADC is not installed or not working, the Exchange tasks will not be available in the Active Directory Users and Computers Snap for the Microsoft Management Console (MMC).

Use the ADC from the Exchange service pack one, instead of the one from the Exchange 2000 CD-ROM. This will enable support for Microsoft Exchange 2000 and Microsoft Exchange 5.5 to coexist on the same Active Directory.

During the installation of Exchange 2000, the installation of the ADC is confirmed by checking the configuration partition of the Active Directory.

Tip: Install the ADC from Exchange service pack one.

Log on to the server using the Domain Administrator account. This can be accomplished by extracting all the downloaded files of Exchange 2000 service pack one, to a local folder. In this case c:\temp was used. Run setup from the:

```
C:\Temp\adc\i386\setup.exe
```

Alternatively, if service pack one is not available, ADC can be launched from the Exchange 2000 CD under:

- ▶ **Z:\adc\i386\setup.exe** (for this installation the CD-ROM letter was Z:)
- ▶ Select **Next**.
- ▶ Select the required components as shown in Figure 4-3, and select **Next**.

The selection of the Service and Management component is very important. Failure to do so will result in non-manageability of the Exchange 2000 server.

- ▶ Select the installation drive, and select **Next**.
- ▶ Type in the Account name and password, and select **Next**.
- ▶ The installation of the ADC will launch.
- ▶ Select **Finish** to complete the installation.

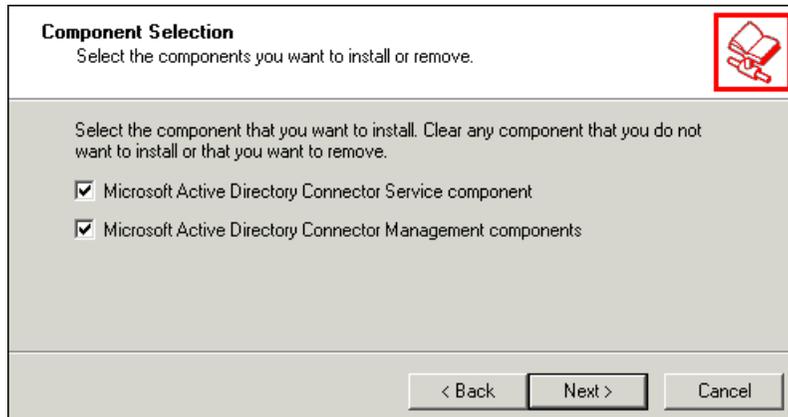


Figure 4-3 ADC setup selection

4.2.5 Advantage of using the forestprep switch

Forestprep prepares the Active Directory forest for Exchange 2000. Several modifications are made to the Active Directory schema and configuration without installing Exchange 2000. Forestprep also grants the first Exchange administrator user group permission to the Exchange organization.

Attention: This is a one way change to the Active Directory, and changes to the schema are not reversible.

4.2.6 Preparing the forest with the /forestprep switch

Forest preparation is accomplished by changing to a command prompt, and switching to the CD-ROM drive. Change the directory to G:\setup\i386 (for this installation the CD-ROM letter was G:); from the command prompt type:

```
G:\setup\i386\setup /forestprep
```

Note: If forestprep is misspelled a normal installation will run. Close the installation and check the spelling of forestprep on the command line.

The installation welcome screen will be launched:

- ▶ Select **Next**.
- ▶ Read the Licence agreement, and select **I Agree**.
- ▶ Select **Next**.
- ▶ Enter the CD key for Exchange 2000.

- ▶ Select **Next**.

Make sure that ForestPrep appears as shown in Figure 4-4. Select the installation drive, and select **Next**.

- ▶ Select **Create a new Exchange Organization**.
- ▶ Select **Next**.
- ▶ Type in the organization name.
- ▶ Select **Next**.
- ▶ Type in the account name.
- ▶ Select **Next**.

The component installation progress bar will appear (Figure 4-5). This part of the installation may take some time to complete.

- ▶ Select **Finish** to complete the installation.

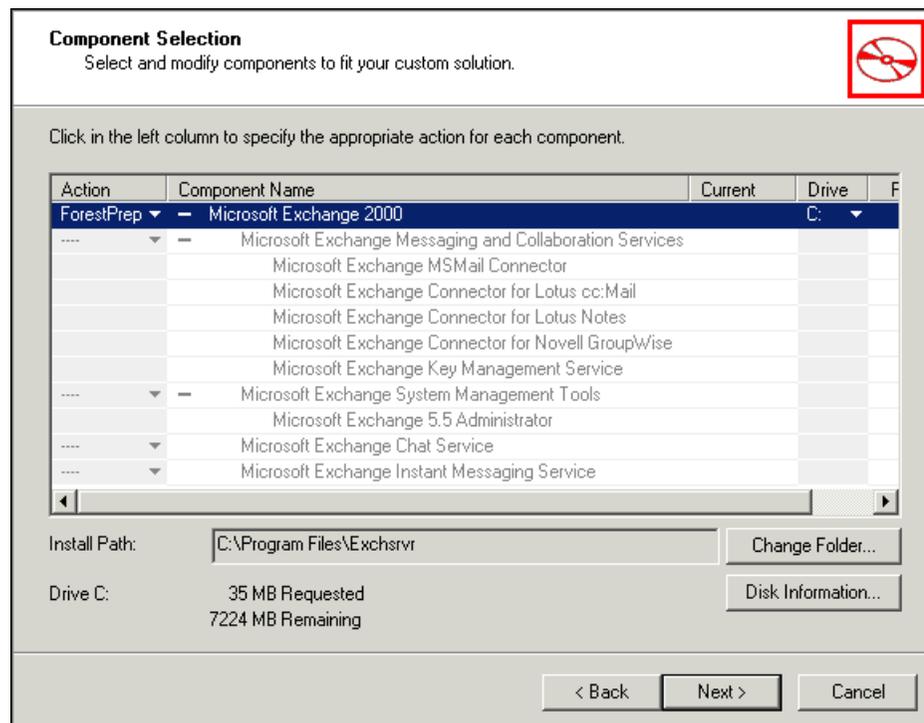


Figure 4-4 Exchange 2000 - ForestPrep

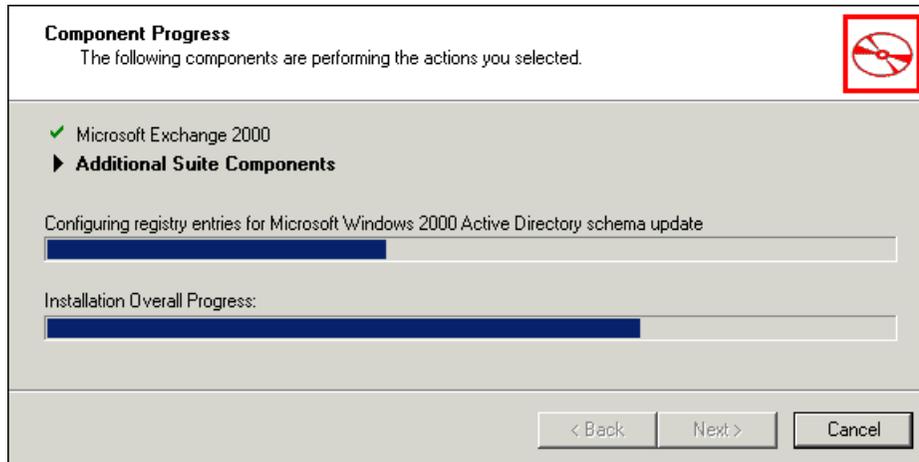


Figure 4-5 ForestPrep installation bar

4.2.7 Advantages of using the domainprep switch

Domainprep prepares the domain by creating an Exchange Domain Servers domain global group that will contain all the servers running Exchange 2000 in the domain.

By default the first Exchange 2000 server installed in the domain is the Recipient Update Server (RUS). The advantages are:

- ▶ Creates a global security group, Exchange Domain Servers.
- ▶ Creates a domain local security group, Exchange Enterprise Servers.
- ▶ Adds the Exchange Domain Servers group to the Exchange Enterprise Servers group.
- ▶ Creates the container for the public folder proxy object.
- ▶ Creates a user account EUSER_EXSTOREEVENTS for use with the script events host. This account has less access than a guest account.
- ▶ Grants the appropriate permissions for Exchange 2000 administrators and Exchange servers.

4.2.8 Preparing the domain with /domainprep switch

Domain preparation is accomplished by changing to a command prompt, and switching to the CD-ROM drive. Change the directory to G:\setup\i386 (for this installation the CD-ROM letter was G:); from the command prompt, enter:

```
G:\setup\i386\setup /domainprep
```

Note: If domainprep is misspelled a normal installation will run. Close the installation and check the spelling of domainprep on the command line.

The installation welcome screen will be launched.

- ▶ Select **Next**.
- ▶ Read the Licence agreement, and select **I Agree**.
- ▶ Select **Next**.
- ▶ Enter the CD key for Exchange 2000.
- ▶ Select **Next**.
- ▶ Make sure that the DomainPrep option appears, as shown in Figure 4-6.
- ▶ Select the target drive, and select **Next**.

If there is a Pre-Windows 2000 domain running, the following error might be displayed as shown in Figure 4-7. Take note of the warning, correct it later if applicable for the installation.

- ▶ Select **OK** to continue.

The installation/preparation will now start (Figure 4-8).

- ▶ Select **Finish** to complete the installation.

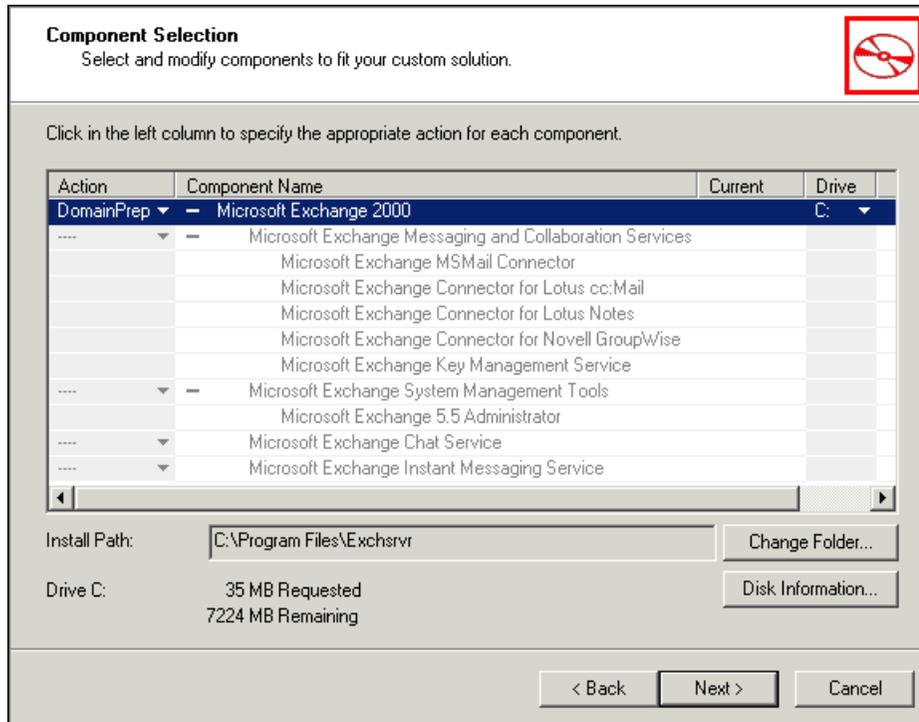


Figure 4-6 Exchange 2000 - DomainPrep



Figure 4-7 Pre-Windows 2000 compatibility warning

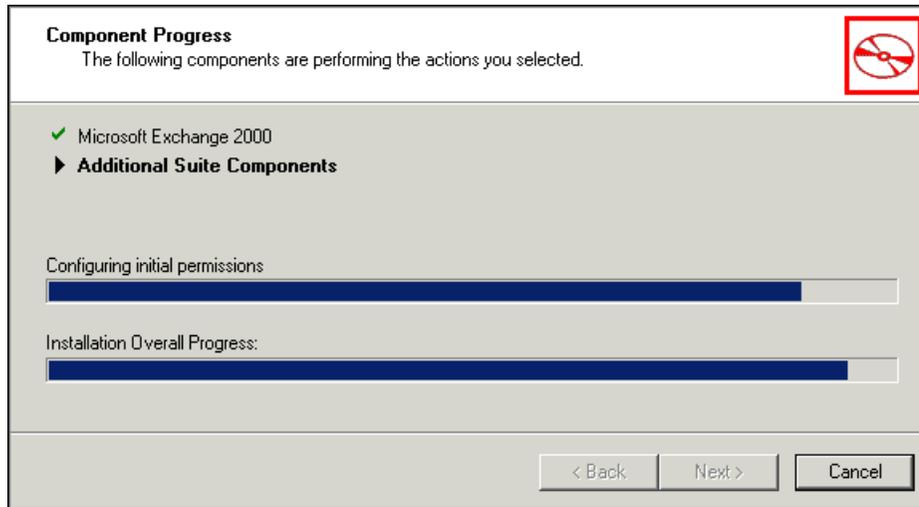


Figure 4-8 DomainPrep progress bar

4.2.9 Installing Exchange 2000

This section covers the installation and possible pitfalls for Exchange 2000.

Tip: If this is a reinstallation make sure that the C:\Program Files\Exchsrvr\MDBDATA is *empty*, because Exchange 2000 will not install, if there is file in this folder. For this example, C: was used as the target drive.

You can install Exchange 2000 from the application that launched when the CD was installed into the CD-ROM drive; alternatively it can be installed from:

```
G:\setup\i386\setup.exe
```

- ▶ Select **Next**.
- ▶ Read the license agreement, and select **I Agree**.
- ▶ Select **Next**.
- ▶ Enter the CD key.
- ▶ Select **Next**.
- ▶ Changes are made by selecting the drop-down box and then making a required selection (Typical Installation is suggested), as shown in Figure 4-9. Select the target drive.
- ▶ Select **Next**.

Take note that this is a per seat license agreement.

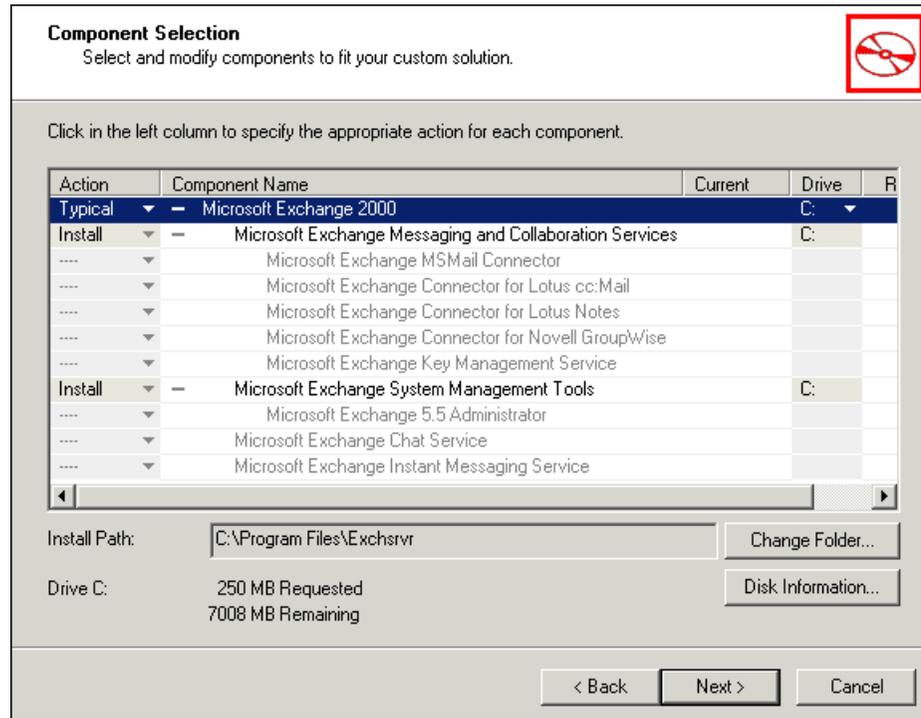
- ▶ Select **Next**.

On the summary page, make sure that the desired components and target drive are selected.

- ▶ Select **Next**.

The installation process will now begin (Figure 4-10). This may take some time to complete.

- ▶ Select **Finish** to complete the installation.



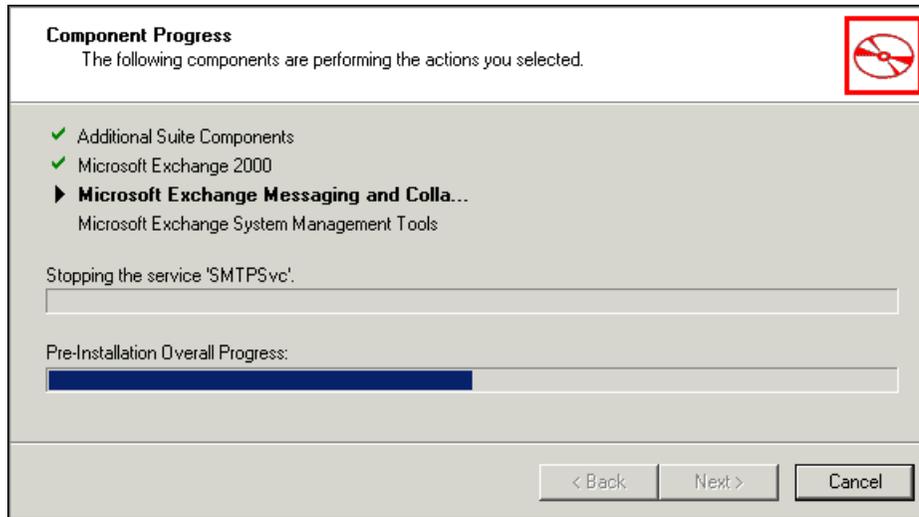


Figure 4-10 Exchange 2000 installation progress bar

Note: Before loading the Active Directory Schema you must register the dynamic-link library by typing this at the command prompt:

```
> Regsvr32 schmmgmt.dll
```

Tip: The Exchange System Manager security tab is not displayed by default.

To display the Exchange System Manager security tab, use Regedit:

- ▶ Add a DWORD VALUE to HKEY_CURRENT_USER\Software\Microsoft\Exchange\EXAdmin.
- ▶ Right-click EXAdmin point to new, click DWORD Value. The name value is ShowSecurityPage.
- ▶ Double-click ShowSecurityPage, type 1 in the value data box, and click OK.
- ▶ The installation of Exchange 2000 is now complete.

Rebooting of the Exchange 2000 Server and Domain Controller is not required, but is recommend at this point!

Tip: Validate that the POP3 and SMTP services has installed and started. Telnet can be used to validate these services.

To validate the POP3 service at the command line, type:

```
Telnet localhost 110
```

To validate the SMTP service on a command line, type:

```
Telnet localhost 25
```

- ▶ Type **quit** to escape from the telnet session.

Note: For complete installation and configuration guidelines, refer to the Exchange 2000 Server Resource Kit; alternatively, data is available from:
<http://www.microsoft.com/exchange>

4.2.10 Creating a new mail storage group

In this section we discuss how to create a new storage group and a new Mailbox store on the iSCSI unit.

From the Exchange System Manager tool, expand the current work group:

- ▶ Expand the servers option.
- ▶ Expand the applicable server.
- ▶ Select the applicable server.
- ▶ Right-click on the applicable server.
- ▶ Select **New Storage Group**.
- ▶ Enter the New Storage Group name.
- ▶ Enter the location destination of the New Transaction log (select the iSCSI drive that was set up for the Log Files).
- ▶ Enter the System Path destination Location (select the iSCSI drive that was set up for the System Path Location).
- ▶ Select **Apply**.
- ▶ Select **OK** to close the window.

The new storage group will now be available under the selected server.

Creating a new mailbox store

- ▶ Select the newly created storage group.
- ▶ Right-click the new storage group.

- ▶ Select **New Mailbox Store**.
- ▶ Name the new mailbox store.
- ▶ Select the **Database** tab.
- ▶ Enter the location destination of the Exchange database and Exchange streaming database (select the iSCSI drive that was set up for the database files).
- ▶ Changes to storage limits and so forth can be made here as well, under the different selection tabs.
- ▶ Select **Apply**.
- ▶ Select **OK** to close the window.

A notification will be shown and asking to mount the store.

- ▶ Select **Yes** to mount the new store.

A notification will be shown after the successful mounting of the new store.

- ▶ Select **OK**.

The new mail store is now ready for use.

Important: Do not create the log files and stores on the same physical hard disk drive, because this will impact the server performance.

4.2.11 Creating an e-mail enabled user account

Note: Creating a e-mail enabled user account option is only available on the Exchange 2000 server or workstation/server where the Exchange 2000 management tools are installed.

- ▶ Open a Microsoft Management Console (MMC)
Add Active Directory users and computers (Figure 4-11). Select the domain, select **New Users**, right-click the user folder and select action -> **New User**.
- ▶ Select **Next**.
- ▶ Follow the instructions, as shown in Figure 4-12. Make sure that the create an Exchange mailbox tic box is selected, and verify that the appropriate mail box store is selected.
- ▶ Select **Next**.
- ▶ Click **Finish** to create the new user.

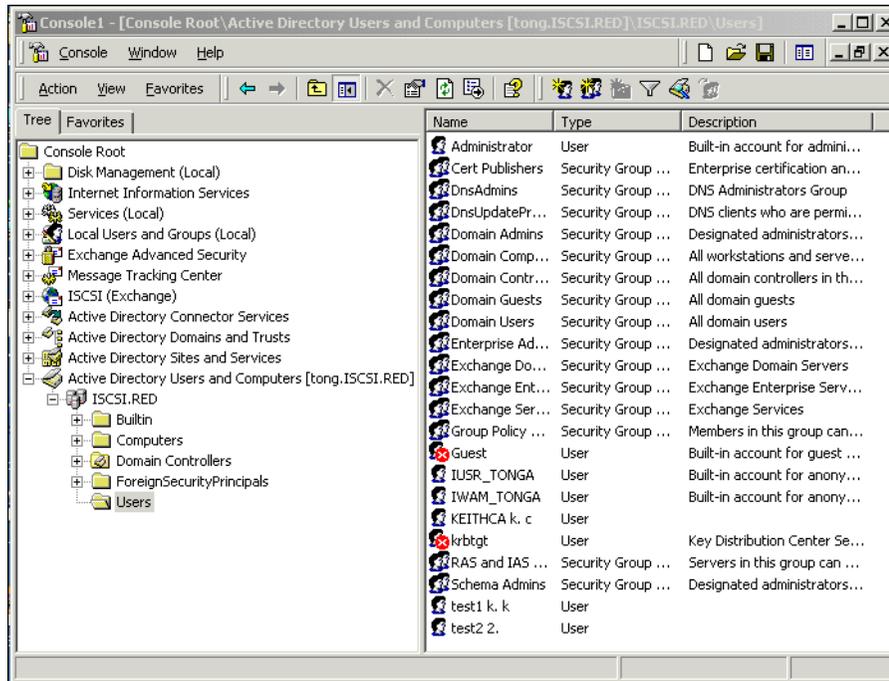


Figure 4-11 MMC console

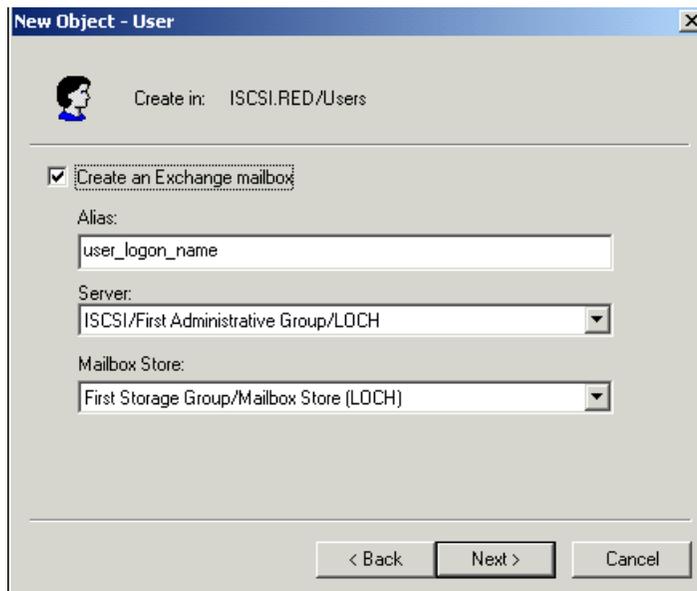


Figure 4-12 New user mail box options

4.2.12 Configuring a users Outlook Express mail

In this section we discuss setting up a new Outlook Express user e-mail account.

When starting Outlook Express for the first time, the Outlook Express configuration window will be launched. Enter the new user's:

- ▶ Name, Last Name.
- ▶ Select **Next**.
- ▶ Enter the users e-mail address.
- ▶ Select **Next**.
- ▶ Enter the POP3 and SMTP server address.
- ▶ Select **Next**.

Note: For POP3 and SMTP server entries, enter the IP address of the Exchange 2000 server, if a DNS record has not been created for the Exchange 2000 server.

- ▶ Enter the user name and password.
- ▶ Select **Next**.
- ▶ Select **Finish** to complete the installation.

The user is now ready to send and receive e-mail.

4.2.13 Setting up startup scripts for Exchange

This section will provide a complete script for starting Microsoft Exchange 2000, safely, after the iSCSI HDD has been attached as a local drive.

We strongly recommend that the installation of AutoExNt services are interactive, to note any errors in the start up. This is accomplished by running the installation of the instextn.exe with the interactive command:

```
Instextn.exe install /interactive
```

Table 4-3 has all the affected services that need to be changed to start Exchange 2000, safely.

Table 4-3 Affected services for starting Exchange 2000

Service name	Startup type
AutoExNt	Automatic
Computer browser	Manual
Server	Manual
Distributed File System	Manual
Microsoft Exchange System Attendant	Manual
Microsoft Exchange MTA Stacks	Manual
Microsoft Exchange Information Store	Manual
Microsoft Exchange POP3	Manual
Microsoft Exchange IMAP4	Manual
Microsoft Exchange Routing Engine	Manual

The script in Example 4-1 will attempt to connect the iSCSI drives. If it fails to connect, it will not start the Exchange 2000 services, and intervention is then required.

Example 4-1 Single iSCSI drive startup script

```
net stop "Computer browser" /yes
net stop "server" /yes

:drive_offline
net start "IBMiSCSIcfg"

sleep 10

net start "server"
net start "computer browser"
net start "Distributed File System"

if exist g:\ goto iscsi_yes
Goto ISCSI_Error
```

```

        sleep 5

goto drive_offline

:iscsi_yes
net start "Microsoft Exchange System Attendant"
SLEEP 2
net start "Microsoft Exchange MTA Stacks"
SLEEP 2
net start "Microsoft Exchange Information Store"
SLEEP 2
net start "Microsoft Exchange POP3"
SLEEP 2
net start "Microsoft Exchange IMAP4"
SLEEP 2
net start "Microsoft Exchange Routing Engine"

:iscsi_error

```

The script in Example 4-2 will keep attempting to reconnect the G: drive until it is successful. After connection of the G: drive, the required Microsoft Exchange 2000 services will be started. We recommend that this script is implemented at the end of the instigation, after the drive letter has been selected.

Example 4-2 Multiple iSCSI drive startup script

```

net stop "Computer browser" /yes
net stop "server" /yes

:drive_offline
net start "IBMiSCSIcfg"
sleep 10
net start "server"
net start "computer browser"
net start "Distributed File System"

if exist g:\ goto iscsi_yes

net stop "IBMiSCSIcfg" /yes
net stop "computer browser" /yes
net stop "server" /yes

sleep 5

goto drive_offline

:iscsi_yes

net start "Microsoft Exchange System Attendant"

```

```
SLEEP 2
net start "Microsoft Exchange MTA Stacks"
SLEEP 2
net start "Microsoft Exchange Information Store"
SLEEP 2
net start "Microsoft Exchange POP3"
SLEEP 2
net start "Microsoft Exchange IMAP4"
SLEEP 2
net start "Microsoft Exchange Routing Engine"
```

Note: For more information on scripts and multiple iSCSI drives, refer to 2.2.2, "Disk drive addressing" on page 34.

Shut down script for Exchange 2000

To shut down and restart Exchange 2000 safely, a shutdown script needs to be installed. This is not required, but this can be accomplished by creating a bat (shutdown.bat) file in this folder:

```
C:\WINNT\system32\GroupPolicy\Machine\Scripts\Shutdown\
```

Edit the file and insert the following statements:

Example 4-3 Shut down script for Exchange 2000

```
net stop "Microsoft Exchange POP3" /yes
net stop "Microsoft Exchange IMAP4" /yes
net stop "Microsoft Exchange Routing Engine" /yes
net stop "Microsoft Exchange Information Store" /yes
net stop "Microsoft Exchange MTA Stacks" /yes
net stop "Microsoft Exchange System Attendant" /yes
net stop "Computer browser" /yes
net stop "server" /yes
net stop "IBMiSCSIcfg" /yes
```

Save the file. Then open the Microsoft Management Console (MMC) and add the Local Computer Policy from the MMC:

- ▶ Select **Computer Configuration**.
- ▶ Select **Windows Settings**.
- ▶ Select **Scripts**.
- ▶ Select **Shut down**.
- ▶ Right-click on shut down.
- ▶ Select **Properties**.
- ▶ Select **Add**.

- ▶ Select **Browse**.
- ▶ Select:
C:\WINNT\system32\GroupPolicy\Machine\Scripts\Shutdown\shutdown.bat
- ▶ Select **Open**.
- ▶ Select **OK**.
- ▶ Select **Apply**.
- ▶ Select **OK** to close the shutdown script window.

4.3 Installing Lotus Domino

In this section we discuss installing Lotus Domino R5 on Windows 2000 Advanced server in an Active Directory.

Lotus Domino is not Domain dependent and can be installed on a member server as well as a stand alone server in the network. For more information on configuring and implementing Domino, refer to:

- ▶ *Using LDAP for Directory Integration*, SG24-6163
- ▶ *Getting the Most From Your Domino Directory*, SG24-5986
- ▶ *LDAP Implementation Cookbook*, SG24-5110

Test environment

The test environment was set up as shown in Table 4-4.

Table 4-4 Setup information for Lotus Domino R5

DNS name	TONGA.DOMINO.RED	CHARCOT.DOMINO.RED
Operating System	Windows 2000 Advanced Server SP2	Windows 2000 Advanced Server SP2
Domain Role	Domain Controller, DNS Server	Member Server
Additional Software	Window 2000 IIS Hot Fix, Lotus Notes, Lotus Domino Administrator	Lotus Domino R5, Window 2000 IIS Hot Fix, ADC
Additional Windows Components	DNS	
CPU	P III 800MHZ	P III 800MHZ
Memory	768MB	768MB

DNS name	TONGA.DOMINO.RED	CHARCOT.DOMINO.RED
Network Card	10/100MB Ethernet	10/100MB Ethernet
Additional Hardware		200i Network attached storage device

Important: The IIS Hot Fix is installed, because there is security exposure to viruses; in this case we used:

Q301625_W2K_SP3_x86_en.EXE

4.3.1 Prerequisites for installing Lotus Domino

In this section we discuss the software and hardware requirements for installing Lotus Domino on a Windows 2000 Server — the installation process and setting up a mail box for a user.

Note: For a detailed description of Domino deployment, refer to *A Roadmap for Deploying Domino in the Organization*, SG24-5617.

4.3.2 Hardware requirements

Hardware requirements are diverse for each organization. These are based on minimal installation requirements, the amount of users, scalability, and availability. Here are the factors:

- ▶ For Windows 2000 server a minimum of 128 MB of memory is required, and 256 MB of memory and more is recommended.
- ▶ Windows 2000 requires 2 GB with a minimum of 1 GB of free space.
- ▶ A minimum of a Pentium-compatible 133 MHz or higher (300 MHz or above is recommended).
- ▶ The minimum amount of memory you need in your Domino server is the sum of two parts, the base RAM required to run the server application and the operating system, plus memory for each client. The formula is:

Minimum memory = 128 + (number of concurrent users/5)MB

This is a minimum value and should be used as a starting point for memory requirements. If the server runs additional tasks, such as file/print, backup and recovery, or anti-virus software running on your server, you will need to configure additional memory for them.

Notes: Lotus Domino only addresses up to 2 GB of RAM. If you install extra memory, it will only be used by the operating system and other additional tasks. This means that installing more than 2.5 GB of RAM will not significantly improve server performance.

For additional information about system requirements or hardware platforms, visit the Lotus Domino Web site at:

<http://www.Lotus.com/domino>

4.3.3 Software requirements

Lotus Domino can be installed on other systems, such as IBM @server iSeries Servers and AS/400, IBM @server zSeries Servers and S/390, IBM AIX, IBM OS/2, Linux, Sun Solaris, Windows NT/2000, and HP-UX.

- ▶ We recommend that the latest Service Pack and Hot fixes for Windows 2000 is installed for the applicable services that will be installed.
- ▶ TCP/IP protocol needs to be installed.
- ▶ We recommend that the Lotus Domino installation is done on a NTFS partition for performance and security enhancement reasons.

For additional information on system requirements and supported operating systems, visit the Lotus Domino Web site at:

<http://www.lotus.com/domino>

Tip: Download the latest Service Pack, IIS and NNTP fixes for Windows 2000, along with service pack one for Exchange 2000 from:

<http://www.microsoft.com/servers/downloads>

4.3.4 Steps for installing Lotus Domino

You can install Lotus Domino from the installation CD-ROM, or from the file that is available for download at:

<http://www.lotus.com/domino>

Keep the installation and system files local and create the mail stores on the iSCSI box. Both installation methods will bring you to the installation setup window.

For this installation we downloaded the file from the Web site above. The file was saved to local disk and the installation started by browsing to the Lotus Domino installation file and then launched from the source location:

`C:\temp\w32n_domino_entepriase.exe`

- ▶ The File Extraction process will launch (this may take some time to finish).
- ▶ In the Lotus Domino Installation window, select **Next**.
- ▶ Read the License agreement and select **Yes**.
- ▶ Enter your Name And Company.

Attention: Partitioned Server Installation should be selected if applicable, for the installation.

- ▶ Select **Next**.
- ▶ Select the target drive and folder for the program folder and data folder.

Tip: If Partitioned Server Installation was selected the data folder destination option is not available here.

- ▶ Select **Next**.
- ▶ Select the type of Domino Server to install as shown in Figure 4-13.

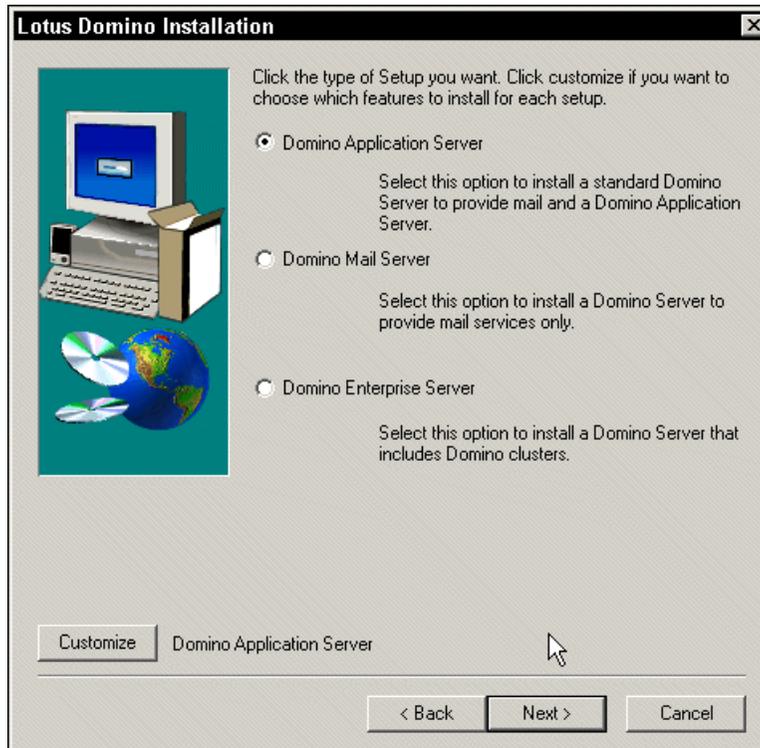


Figure 4-13 Domino server installation option

Select **Customize** for a customized component selection with Lotus Domino.

If the **customize** option was selected, select the additional components:

- ▶ Select **Next** to complete the installation.
- ▶ Select **Next** if a standard installation will be done.
- ▶ Select **Next** to start the installation of Lotus Domino (this may take some time to complete).
- ▶ Select **Finish** to complete the installation.

After the installation of Lotus Domino, select:

- ▶ Start -> Programs -> Lotus Applications -> Lotus Domino Server

The first time Lotus Domino is launched, it will bring up the Domino Server Setup window (Figure 4-14).

For this installation, we used First Domino Server.

- ▶ Click the > (right arrow, on top of the window).

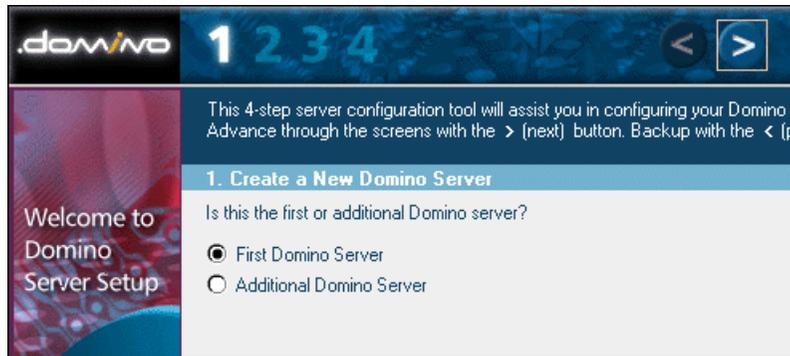


Figure 4-14 Lotus Domino server setup

We recommend that you select Quick and Easy installation.

- ▶ Click the > selection box.
- ▶ Select the components for installation in the Server Audience selection page as shown in Figure 4-15.
- ▶ Click the > box.

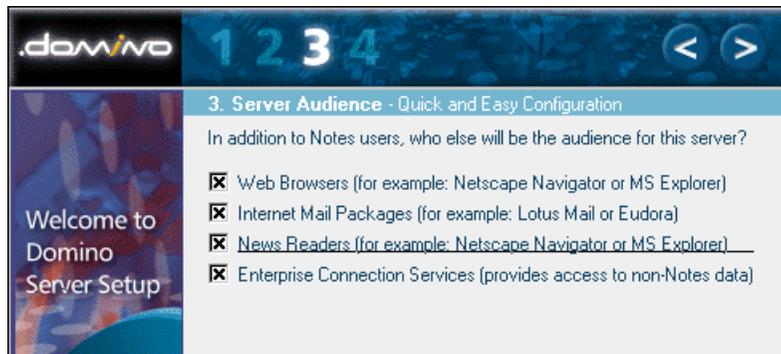


Figure 4-15 Server audience

Make sure to select edit to make changes to the Domino configuration (Figure 4-16), before selecting the **Exit** option.

Attention: Change the certifier password along with the administrator password, and change the server name and certifier options if required.

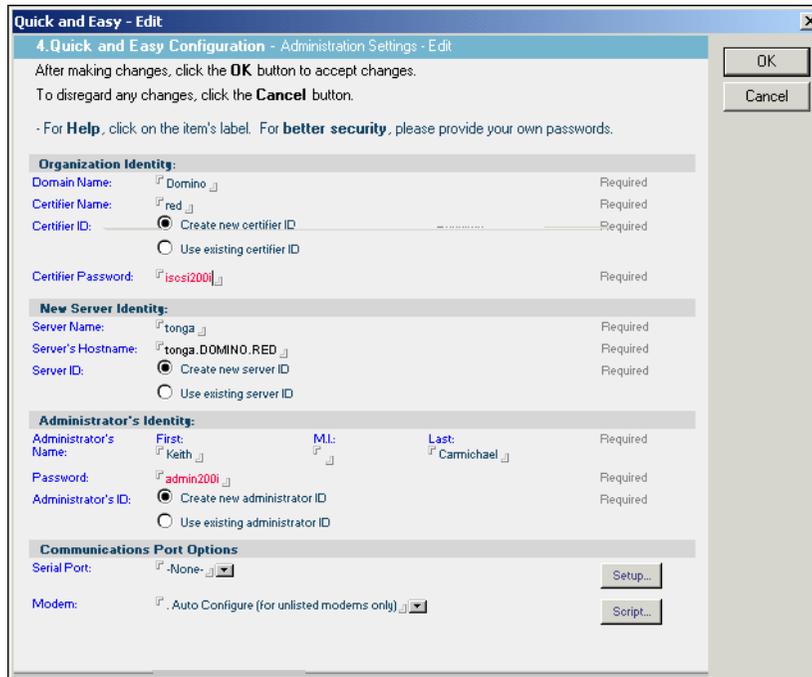


Figure 4-16 Administrator settings

- ▶ Select **OK** on the settings page.
- ▶ Select **Exit**.
- ▶ Select **Finish** to complete the installation.

After the setup of Lotus Domino, click:

- ▶ Start -> Programs -> Lotus Applications -> Lotus Domino Server

The Domino server application will be launched and displayed in an OS window.

4.3.5 Installation of Lotus Domino Administrator

You can install Lotus Domino from the installation CD-ROM, or from the file that is available for download at:

<http://www.lotus.com/domino>

Both installation methods will bring you to the installation setup window.

For this installation we downloaded Lotus Domino Administrator from the Web site above. The file was saved to local disk and the installation started by browsing to the Lotus Domino Administrator installation file and then launched from the source location:

```
C:\temp\w32n_allclient.exe
```

Important:

- ▶ Wait for the Domino server to start up all the processes before opening Lotus Domino Administrator.
 - ▶ Do not install the Domino Administrator on the same server as the Domino server; this configuration is not advisable and not supported.
-
- ▶ The File Extraction process will launch (this may take some time to finish).
 - ▶ In the Lotus Notes Installation window, select **Next**.
 - ▶ Read the License agreement and select **Yes**.

Enter your Name and Company, and select the shared installation if applicable.

- ▶ Select **Next**.

Note: If the shared installation option is selected, the data folder destination option is not available.

Select the program and data folder destination, as applicable.

- ▶ Select **Next**.
- ▶ Select the type of installation: Notes Client, Domino Designer, Domino Administrator, and all Clients to select a complete installation as shown in Figure 4-17.

Select Customize for a customized component installation.

- ▶ Select **Next** to complete the installation.
- ▶ Select **Next** if a standard installation will be done.
- ▶ Select **Next** to start the installation of Lotus Notes (this may take some time to complete).

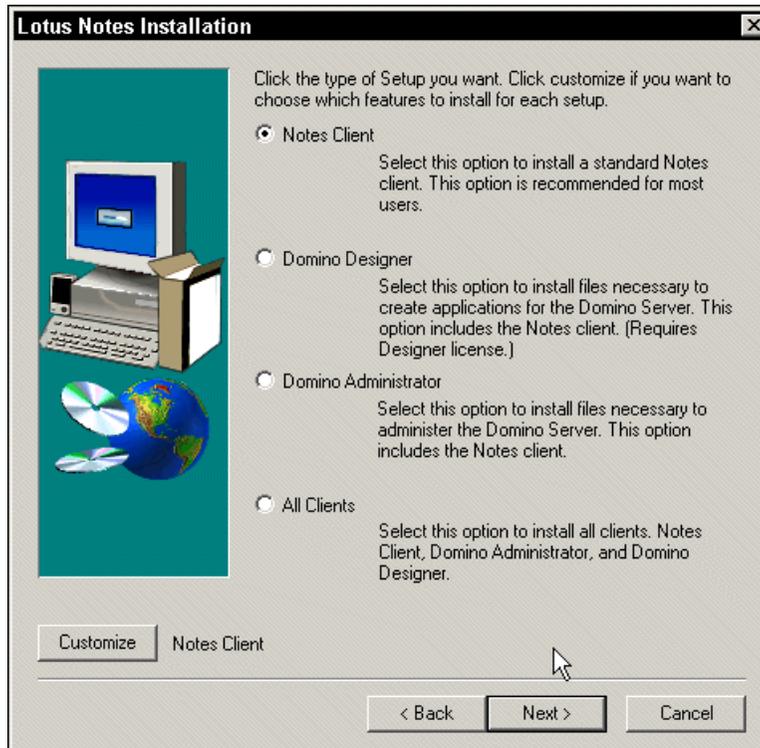


Figure 4-17 Notes installation selection options

After the installation of Lotus Notes completes, select:

- ▶ Start -> Programs -> Lotus Applications -> Lotus Domino Administrator

The first time that Lotus Domino Administrator is launched it will bring up the Notes Client Configuration window (Figure 4-18).

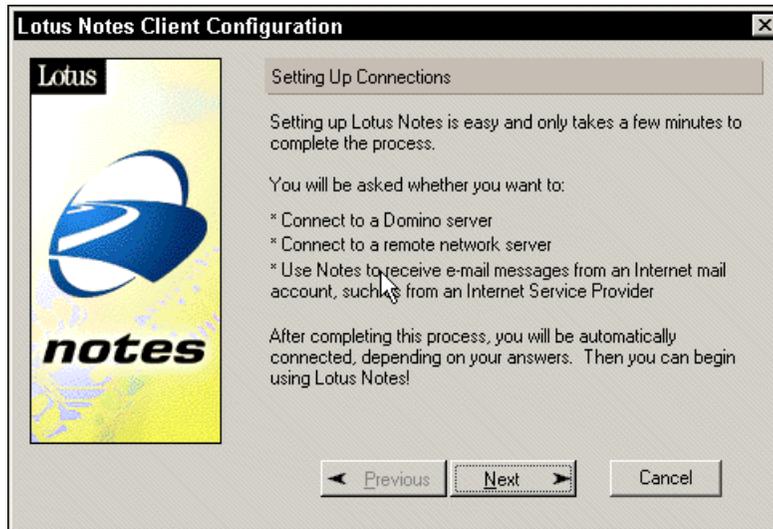


Figure 4-18 Notes client configuration setup

- ▶ Select **Next**.
- ▶ Select **I want to connect to a Domino Server**.
- ▶ Select **Next**.

For this installation we used a standard LAN connection.

- ▶ Select **Setup a connection to a Local area network**.
- ▶ Select **Next**.
- ▶ Type in the server name as noted when the Domino Server was installed.
- ▶ Select **Next**.
- ▶ Select **Use My Name as Identification**, and type in the administrator name as noted in the Domino Server setup.
- ▶ Select **Next**.

Note: If there is a error connecting to the server, the connection error window will be displayed as in Figure 4-19.

Connect to the Domino server by selecting a network type and a server address.

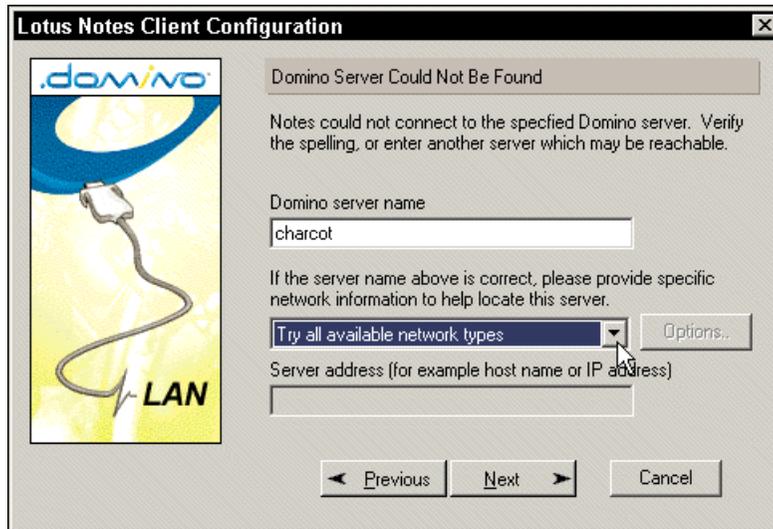


Figure 4-19 Server connection error window

- ▶ Select **Next**.
- ▶ Select **I want to create a Internet mail account**, if you want to set up incoming mail from outside of the organization.
- ▶ Select **Next**.

If the Internet mail option is selected, follow the instructions for setting up the mail account.

- ▶ Select **Next**.

If the News server option is selected follow the instructions for setting up the news server.

- ▶ Select **Next**.

For our scenario, we did not connect to another directory device.

- ▶ Select **Next**.

Make the appropriate selection for your connection to the Internet. In our scenario, we have a direct connection to the Internet and have not used a proxy server.

- ▶ Select **Next**.
- ▶ Select your connection type to the LAN.
- ▶ Select **Next**.

- ▶ Select **Finish** to complete the installation.

4.3.6 Creating a new user in Domino

This section guides you through creating a new e-mail account in Lotus Domino from the Domino Administrator.

Open Domino Administrator; this can be accomplished by selecting:

- ▶ Start -> Programs -> Lotus Application -> Lotus Domino Administrator

Enter the Domino Administrator password that was selected during the setup phase of Domino.

- ▶ Select the **Administration** tab.
- ▶ Select **People & Groups**.
- ▶ Select **Tools** to expand the tools options.
- ▶ Select **People** to expand the available functions, as shown in Figure 4-20.

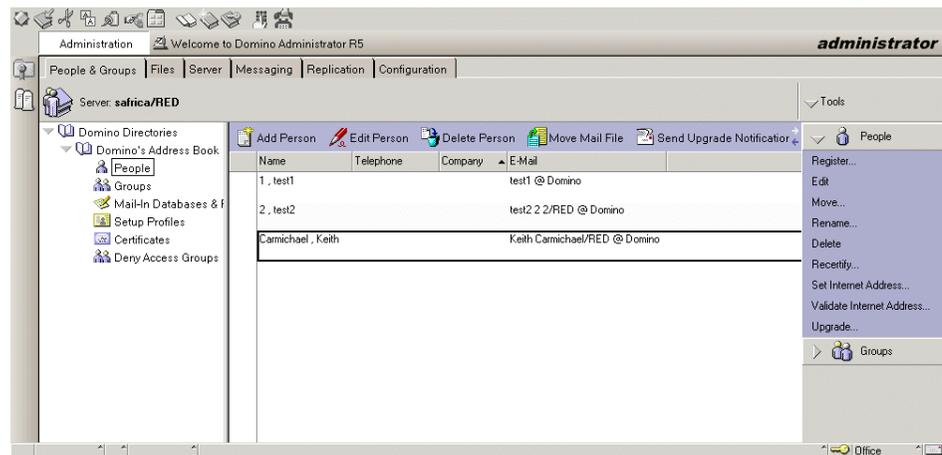


Figure 4-20 Domino administration

- ▶ Select **Register**.

A window will open that will prompt for the Certifier ID. This ID file was created in the installation portion of Lotus Domino, and is available as the cert.id file on the Lotus Domino Server:

C:\lotus\domino\data\ (in this scenario c:\ is the installation directory)

Tip: Copy the cert.id file from the Lotus Domino Server to the administration workstation.

- ▶ Select the **cert.id** file.
- ▶ Select **Open**.
- ▶ Enter the password as noted in the installation portion of the Lotus Domino server.

The Registration window will now open (Figure 4-21).

Advanced

Registration Server... safrica/RED

First name MI Last name Short name

Password

Mail server safrica/RED

Mail file name

Setup profile (None)

Registration status

Provide name and password information for the new person. To view or edit additional registration settings, check the 'Advanced' checkbox above.

Add person Import Text file... Migrate people...

Registration queue:

User Name	Registration Status	Date

Register All Register Delete Options... Done

Figure 4-21 New user registration

- ▶ Enter the required user data.
- ▶ Select **Add Person** (the user data will appear in the Registration queue).
- ▶ Select **Register** to complete the registration of a new user.
- ▶ Select **Done** to exit the Registration window.

- ▶ Press **F9** to refresh the People options page, and the new users will be present.

To configure the new user mail account options, select the user account as shown in Figure 4-20 on page 146.

- ▶ Select **Edit Person**.
- ▶ On the Basic window, the users Internet password can be set for POP3 and SMTP use, and the userid file can be detached and renamed to match the username.id format.

The detached ID file can be distributed to the user for setting up Lotus Notes. After the required changes have been made to the user account and to exit the user account option window, select:

- ▶ **Save and close Tab**

The user is now ready to send and receive e-mail messages.

4.3.7 Configuring startup scripts for Domino

This section will provide a complete script for starting Lotus Domino, safely, after the iSCSI HDD has been attached as a local drive.

Table 4-5 has all the affected services that need to be changed to start Lotus Domino, safely.

Table 4-5 Lotus Domino startup script

Service name	Startup type
AutoExNt	Automatic
Computer browser	Manual
Server	Manual
Distributed File System	Manual
Lotus Domino Server (LotusDominoData)	Manual

The script in Example 4-4 will attempt to connect the iSCSI drives, if it fails to connect, it will not start the Domino services, and intervention is then required.

Example 4-4 Connect the iSCSI drive

```
net stop "Computer browser" /yes
net stop "server" /yes

:drive_offline
```

```

net start "IBMiSCSIcfg"

sleep 10

net start "server"
net start "computer browser"
net start "Distributed File System"

if exist g:\ goto iscsi_yes
Goto ISCSI_Error

    sleep 5

goto drive_offline

:iscsi_yes

net start "Lotus Domino Server (LotusDominoData)"

:iscsi_error

```

The script in Example 4-5 will keep attempting to reconnect the G: drive until it is successful. After connection of the G: drive, the required Lotus Domino services will be started. We recommended that this script is implemented at the end of the instigation, after the drive letter has been selected.

Example 4-5 Reconnect the drive

```

net stop "Computer browser" /yes
net stop "server" /yes

:drive_offline
net start "IBMiSCSIcfg"

sleep 10

net start "server"
net start "computer browser"
net start "Distributed File System"

if exist g:\ goto iscsi_yes

net stop "IBMiSCSIcfg" /yes
net stop "computer browser" /yes
net stop "server" /yes

sleep 5

goto drive_offline

```

```
:iscsi_yes
```

```
net start "Lotus Domino Server (LotusDominoData)"
```

Note: For more information on scripts, refer to Example 2-1 on page 35.

4.4 General performance overview

E-mail servers act as repositories and routers of electronic mail. They handle the transfer of e-mail to its destination. E-mail servers generate communication workload without user request, such as directory replication, communication links between sites, and its protocol conversion process. You may need to connect your server to third-party e-mail servers. This generates additional workload on your e-mail servers.

The important subsystems for e-mail servers are:

- ▶ Network
- ▶ Memory
- ▶ CPU
- ▶ Disk

If there is not enough bandwidth between your users and servers, all the other performance efforts will be unsuccessful. Adequate network bandwidth is an essential performance factor for e-mail servers.

E-mail servers use system memory to support database buffers and e-mail server services. Sufficient size of the memory and effective disk subsystem are very important performance factors for e-mail server performance. If the size of the memory is sufficient, valuable user data will stay in the system memory.

E-mail servers use log files to transfer modified data to an information store. These log files are written sequentially. New transactions are added at the end of the transaction files. Log files and database files have different usage patterns. Log files provide better performance with separate physical disks, and database files provide high performance with striped disk arrays, because of random workload. Using several drives instead of a single drive doubles the e-mail throughput. Read-ahead disk caching disk subsystems can also offer performance benefits.

Users mailboxes can be stored either on the server or on each user's local hard drive or both. In each case, you need high network performance, because the users retrieve their mail over the network. If the average size of the messages are big, you need more network bandwidth. Also, server-to-server replication traffic can be a significant load on the network. Using multiple LAN adapters can help its network performance.

When an e-mail message arrives at the server, the destination is determined. If the address is local, it is stored in the database of the e-mail server. If the address is not local, the message is passed on to an appropriate server for processing. If the address is a distribution list, the server checks the addresses in the list and routes the message accordingly. These processes require CPU cycles. Sufficient amounts of memory must also be provided for these processes.

If your server supports directory replication and connectors between sites, your server will experience high distribution list usage, and your CPU cycles will be a more important factor in e-mail server performance.

4.4.1 Performance monitors

Performance Monitor is one of the most valuable monitoring tools available to Windows 2000 administrators. It is commonly used to monitor server performance and to isolate bottlenecks. The tool provides real-time information about server subsystem performance. The data collection interval can be adjusted based on your requirements.

The logging feature of Performance Monitor makes it possible to store, append, chart, export, and analyze data captured over time. Exchange provides additional monitors that allow Performance Monitor/Console to extend its usefulness beyond the operating system level.

Performance Monitor can send alerts when predefined threshold levels are reached. This is useful especially when you want to perform actions as soon your preset threshold conditions are met. With Capacity Manager, you can only view exceeded thresholds once you gather data from the test systems and generate a report.

4.4.2 Available tools in Windows 2000

These are the available tools:

- ▶ The Process Resource Monitor (**PMON.EXE**, part of Windows 2000 Support Tools in the **\SUPPORT\TOOLS** directory of the Windows 2000 CD-ROM) is a command-line tool that monitors process resource usage by tracking CPU and memory usage.

- ▶ The Process Viewer (**PVIEWER.EXE**, part of Windows 2000 Support Tools in the **\SUPPORT\TOOLS** directory of the Windows 2000 CD-ROM) is a Windows tool that displays information about a running process and allows you to stop the process and change the process priority.
- ▶ The System Monitor creates a chart, report or view log data.
- ▶ The Performance Logs and Alerts creates a log or alert.
- ▶ The Performance console is a snap-in for Microsoft Management Console (MMC). The Performance console is used to access the System Monitor and Performance Logs and Alerts tools.

4.4.3 Spotting a bottleneck

There are a number of reasons why to tune a server. Some examples are: you have suspicions that a slow server is slowing down productivity, users are complaining about slow server response, the customer has upgraded a subsystem in their server (for example, they added a second CPU), but they see no improvement in the server throughput.

The simplest way is to run Performance Console in Windows 2000 from the server that is being analyzed. A performance log of the server should be created during its peak time of operation (for example, 9:00 a.m. to 5:00 p.m.). When creating the log, the following objects should be included:

- ▶ Processor
- ▶ System
- ▶ Server work queues
- ▶ Memory
- ▶ Page file
- ▶ Physical disk
- ▶ Redirector
- ▶ Network segment
- ▶ Network interface

Once you determine which subsystem is the bottleneck, you should examine the options for solving the problem.

It is important to remember that the greatest gains are obtained from upgrading a bottleneck's component when the other components in the server have ample "power" left to sustain an elevated level of performance. In other words, ensure the other components in the server are not "latent bottlenecks", working just below the utilization of the bottleneck's component.

Components that are latent bottlenecks will limit improvements realized by any upgrade. In general, components that have average utilization between 60-70% are likely to be latent bottlenecks. If there are latent bottlenecks in a system, then both the primary component causing the bottleneck and the component that has a latent bottleneck must be upgraded to obtain optimal performance.

When attempting to fix a performance problem, remember the following:

- ▶ Take measurements before you upgrade or modify anything so that you can tell if the change had any effect (that is, take baseline measurements).
- ▶ Examine the options that involve reconfiguration of existing hardware, not just those that involve adding new hardware.
- ▶ Once you upgrade a specific subsystem, other latent bottlenecks may appear in other subsystems.

The above is a general approach that you should take when analyzing a server. When trying to find bottlenecks, you should also consider the type of server you are monitoring and what subsystems are potential bottlenecks.

Before making any recommendation/decision:

- ▶ Make sure you understand what is causing the bottleneck.
- ▶ Research your recommendations, and be sure what you are proposing will improve server performance.
- ▶ Know how much the upgrade/reconfiguration will cost.

Note: For complete documentation for performance increases for xSeries, Netfinity, and Windows 2000, refer to *Tuning IBM e(logo)server xSeries Servers for Performance*, SG24-5287 at:

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



Implementing DB2 UDB EE V7 Database Server using iSCSI

This chapter describes the planning, setup, installation, and configuration of DB2 UDB Enterprise Edition 7.1, FixPak 3, using an iSCSI attached disk in a Windows NT 4.0 environment.

Before going into the details of DB2 installation and configuration, a brief review of RDBMS concepts are deemed necessary. The following section will discuss how DB2 related to basic RDBMS concepts.

5.1 Fundamentals of RDBMS

Relational Database Management Systems (RDBMS) share a common set of principles. The purpose of this section is to explain a subset of these principles (that a systems administrator needs to understand) for logical and physical systems planning as well as performance, availability characteristics, and the backup or recovery procedures for data held on a relational database. Please note that although all RDBMS products are based on the same set of principles, not all use the same terminology or structures. For example, the concept of tablespace does not exist on some RDBMSs.

5.1.1 Database

A database presents data as a collection of tables. A table consists of a defined number of columns and any number of rows. A database can include:

- ▶ A data dictionary or a set of system tables that describe the logical and physical structure of the data
- ▶ A configuration file containing the parameter values allocated for the database
- ▶ A recovery log with ongoing transactions and archivable transactions
- ▶ Some RDBMSs use control files as an extension of the data dictionary

5.1.2 Tables

A table consists of data logically arranged in columns and rows. Figure 5-1 shows that tables are assigned to tablespaces and that users interact with tables. Table data is accessed through Structured Query Language, a standardized language for defining and manipulating data in a relational database. The data in the table is logically related, and relationships can be defined between tables.

5.1.3 Data dictionary

It is common for RDBMS to maintain a data dictionary in a set of system tables. They describe the logical and physical structure of the data. They are like any other tables, but are owned by the database administrator or by the database. They are created either when the database is created or when the database administrator runs a set of scripts supplied by the RDBMS. These tables contain information about the definitions of database objects, such as user tables, and indexes, as well as security information and details relating to recovery.

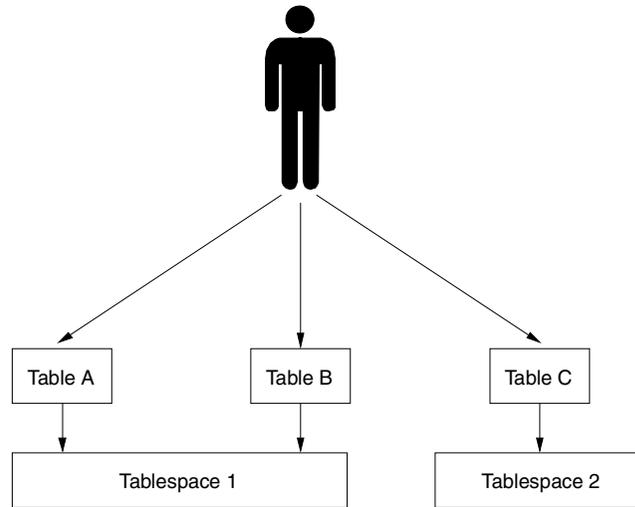


Figure 5-1 Table relationship to tablespaces

5.1.4 Indexes

An index is a set of keys, each pointing to rows in a table. For example, Table A has an index based on the employee numbers in the table. This key value provides a pointer to the rows in the table: Employee Number 19 points to employee KMP. An index allows more efficient access to rows in a table by creating a direct path to the data through pointers. It is possible for the data storage of an index to grow larger than the table to which it refers.

5.1.5 Tablespaces

Tables and indexes are assigned to tablespaces as shown in Figure 5-2. This figure also shows that one or more data files can be allocated to a tablespace, whereas, different tablespaces cannot share the same data files. Some RDBMSs allow several tablespaces to be defined on the same logical volume. Indexes can be assigned to a different tablespace, where their tables reside to improve access speed. Normally, the data dictionary tables reside in their own tablespace.

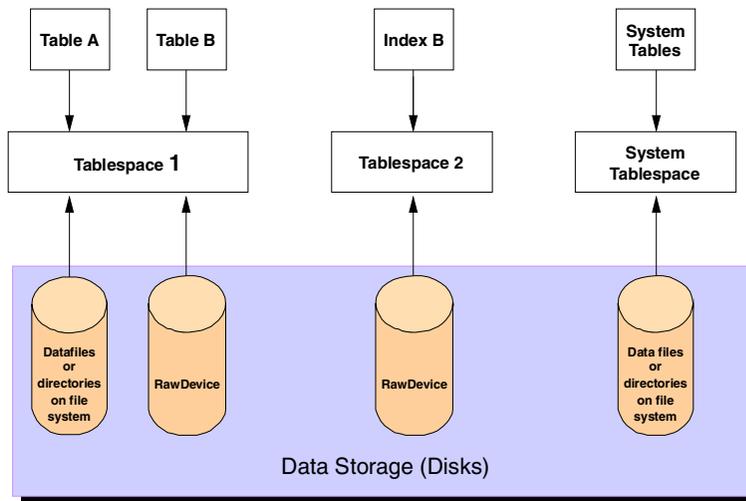


Figure 5-2 Data storage hierarchy

Tablespaces are logical concepts used with RDBMSs. They provide a convenient way of separating the user's view of data from some of the practical considerations associated with storing that data on a disk. For example, a database administrator can make more disk space available to several tables by adding disk space to the appropriate table space, therefore ensuring that tables do not run out of space and that disk space is used efficiently. Furthermore, the tablespace concept means that neither users nor application programs need to be aware of the fact that the database administrator has made more disk space available. Data storage in tablespaces can be implemented using either data files or directories on files systems or raw devices. (For information on file systems and raw devices, please see your operating system documentation.) Tablespaces provide the link between logical views and data storage. Here are some points to note:

- ▶ The data for a table or index may be contained in only one data file.
- ▶ Alternatively, the data for a table or index may be spread over several data files.
- ▶ Each of the data files may contain data for one or more tables in the tablespace.
- ▶ Each data file or directory may reside in a separate file system. The significance of these alternatives is that the only way to back up or recover individual tables is by using the facilities that the RDBMS provides.

Normally, you would back up or restore tablespaces instead of the individual data files of the tablespaces. This ensures that all data storage for a tablespace is backed up consistently with the same timestamp.

You would use tablespace backup instead of full database backup depending on the volatility or importance of data. You have the option of backing up tablespaces which have more update activity more often than tablespaces which have less activity.

The tablespace where the data dictionary tables reside is the most important tablespace. You must ensure that this tablespace is backed up successfully and consistently with the other databases. Corruptions in the data dictionary can cause the database to be unusable.

5.1.6 Partitioning options

RDBMSs may provide partitioning options to handle very large amounts of data. This will allow workload parallels of very large objects, and will allow for the manipulating of subsets of these large objects. You should investigate whether partitioning is used in your database and know how partitioning options are implemented by your RDBMS. This can affect your backup strategy.

5.1.7 Log files

As shown in Figure 5-3, most RDBMSs maintain details of updates to databases in log files. If for some reason, a transaction that makes a change to the database fails to complete successfully, the RDBMS's recovery procedure will use the log file to detect that an update may be only partially complete and undo any changes that the transaction had made to the database.

Some RDBMSs support the use of log files to perform forward recovery. Forward recovery takes advantage of the fact that log files hold details of all changes that have been made to the database, and therefore you do not necessarily need to undo changes, but instead can reapply changes. With forward recovery, the recovery process can:

- ▶ Restore a database to the state it was in at the time the last backup was taken
- ▶ Use the log files to reapply the changes that had been made since the last backup was taken
- ▶ Back out of (undo) any partially completed changes

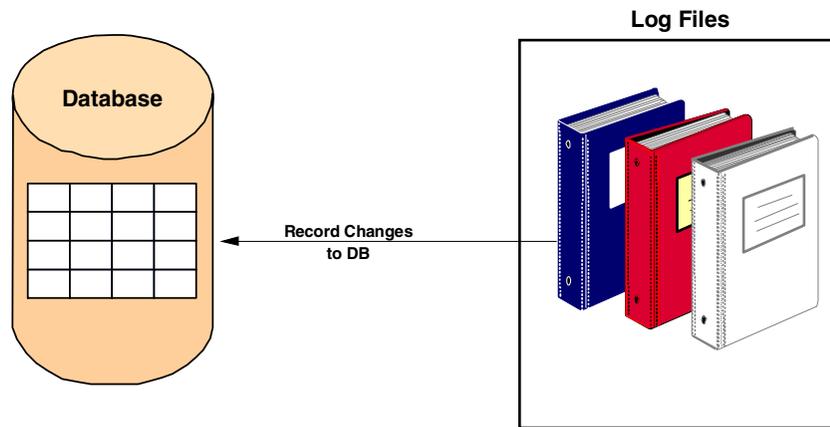


Figure 5-3 Database logfiles

A standard RDBMS concept related to log files is the checkpoint process. All RDBMSs use buffers in memory to hold changes to the database and log files. The purpose of buffers is to improve the operational performance of the RDBMS. However, the use of buffers means that most changes to databases and log files do not get written to disk until some time after the RDBMS has indicated to the user application that the update has been made successfully.

Checkpoints ensure that all database and log file changes held in the RDBMS's buffers are flushed out to disk. This shortens the time it takes to recover a database after a system crash, because the number of redundant log records processed during the recovery is reduced. All RDBMSs support checkpoints and issue them automatically at intervals. It is a basic recommendation that these files are mirrored or duplexed.

5.1.8 Control files

Some RDBMSs maintain control files to hold additional information about the physical structure of the database, such as which physical files are used by each tablespace and which is the current log file. For those RDBMSs which use control files, policies need to be defined for backing up those files. We recommend that these files are mirrored or duplexed.

5.1.9 Configuration parameters

All RDBMSs provide a range of options. Some are set permanently, and others can be modified even when a database is in use (running). Some options allow you to tune the performance of the database; others allow you to specify how you want logging to be implemented, for example. Depending on the RDBMS, you can either change the configuration parameters using database commands or modify them in an initialization file. The configuration parameters may be stored in a file as in the case of an initialization file.

For RDBMSs which use initialization files, a database may have multiple initialization. One reason for having multiple initialization files for a single database might be to optimize performance for different circumstances. For example, you may decide to allocate one set of values when the database is used for batch processing and another set when it is used for online transactions. Although some of the options are set differently for each situation, many will be the same.

Some RDBMSs allow you to specify options that are common to multiple initialization files in configuration files. Instead of repeating all options and their values in each of the initialization files, you can select the configuration file that contains the options that you want to use. You need to define policies for backing up both initialization files and configuration files.

5.2 DB2 UDB EE V7 concepts

The previous section provided a generic overview of databases for system administrators, who are responsible for designing a data held on a relational database. The purpose of this section is to introduce DB2 UDB concepts specific to DB2.

5.2.1 DB2 UDB products

DB2 UDB comes in many flavors. There is a Personal Edition (PE), a Workgroup Edition (WE), an Enterprise Edition (EE), and an Extended Enterprise Edition (EEE). For system administrators, we only need to understand the DB2 UDB EEE product databases that have the partitioning option. That means that one database can be created across one or more nodes or even machines. For the other products, each database created is typically only on one machine.

5.2.2 DB2 UDB EEE partitioning option

Databases created in DB2 UDB EEE can use database partitioning. Each node in a partitioned database supports a subset of the overall database. They are called database partitions. As shown in Figure 5-4, a node will contain a database partition of the database which has its own data, configuration files, and transaction logs.

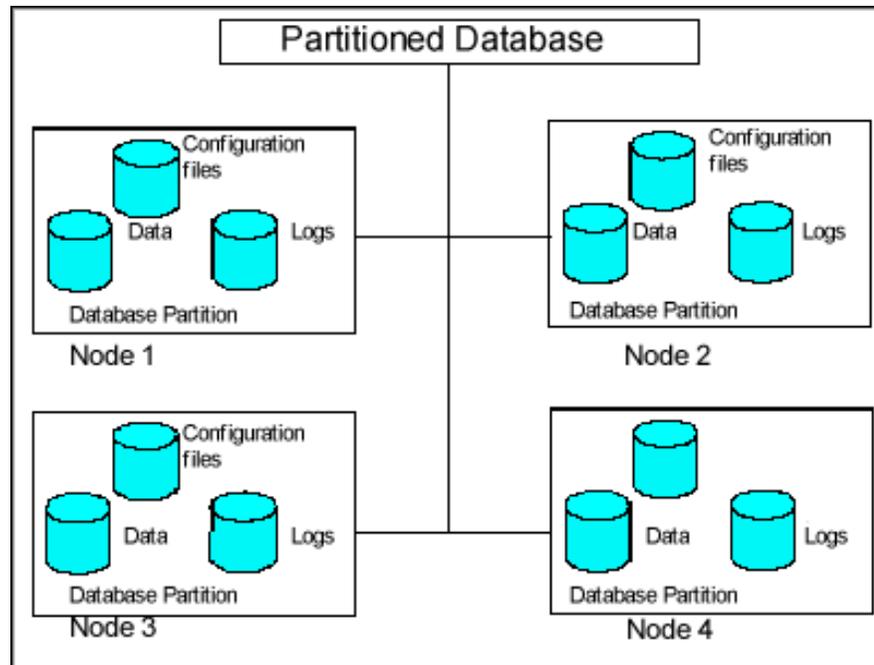


Figure 5-4 DB2 partitioning

The partitioned database can be configured to be flexible so that objects can be created on one or on many nodes depending on their use and on their size. So a small table can be stored on a tablespace that only exists in one node, while a large table can be stored on a tablespace which spans several nodes. Data in a table can, therefore, be distributed on several nodes. Data distribution is done using a hashing algorithm. When there is a data retrieval or update requests on tables which span several nodes, the request is decomposed automatically into subrequests, and executed in parallel among the applicable database partitions.

A database partition fits together with the MPP hardware architecture that is called a shared-nothing architecture, because each has its own data, configuration files, and transaction logs. When using media storage managers, like Tivoli Storage Manager, you should install and configure the media manager on all machines with database partitions. Backup must also be done on each database partition. Recovering to a point in time must be carefully planned with database partitioning.

5.2.3 Instance

An instance (database manager) is an environment for managing data and system resources assigned to it. A machine or system can have more than one instance. Each instance will have its own database manager configuration parameters and security. An instance can have one or more databases. Figure 5-5 shows the hierarchy of DB2 objects in an instance.

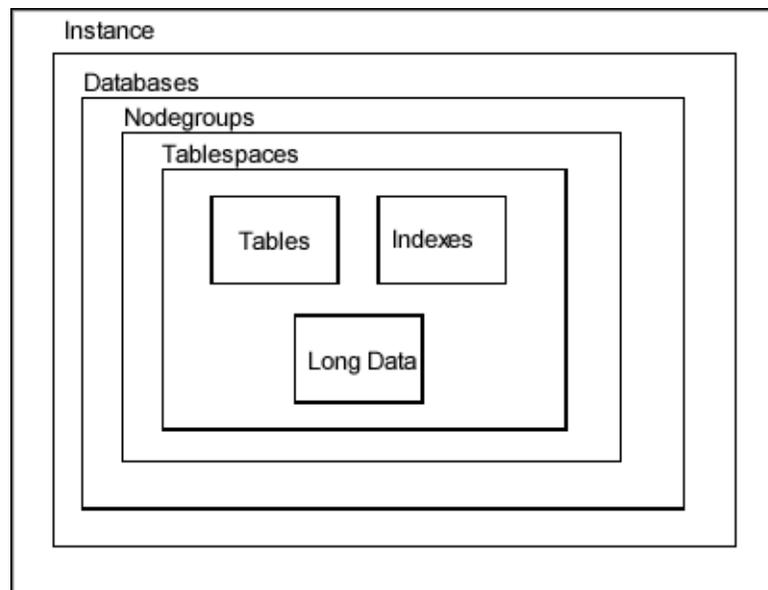


Figure 5-5 DB2 instance hierarchy

5.2.4 Database and database partitions

An instance can have one or more databases. Each database has its own set of system tables (data dictionary). In DB2 UDB EEE, an instance can also have database partitions. A database or database partition has its own data, configuration files, and transaction logs.

Database partitions are initially defined in a file called `db2nodes.cfg`. Each database partition is assigned a node number which represents a node. A node is assigned to a hostname (machine). A node can be physical or logical. Physical nodes are nodes assigned on separate machines. Nodes participating on the same machine are called logical nodes. Logical nodes can be useful when exploiting the symmetrical multiprocessor (SMP) architecture. It is therefore possible to have more than one database partition for a partitioned database residing on the same machine.

You should perform individual backup for each database partition even if the database partitions reside on the same machine, because each database partition is a shared-nothing architecture. You must also have a policy to backup the `db2nodes.cfg` file.

5.2.5 Nodegroups

A nodegroup is a set of one or more database partitions. It is only relevant for DB2 UDB EEE. Tablespaces are created in nodegroups, and tables and indexes are created in tablespaces. The data of a table or index can be distributed across several database partitions if it is defined in a tablespace, and the tablespace is defined in a nodegroup containing more than one database partition. A database partition can be a member of more than one nodegroup as shown in Figure 5-6.

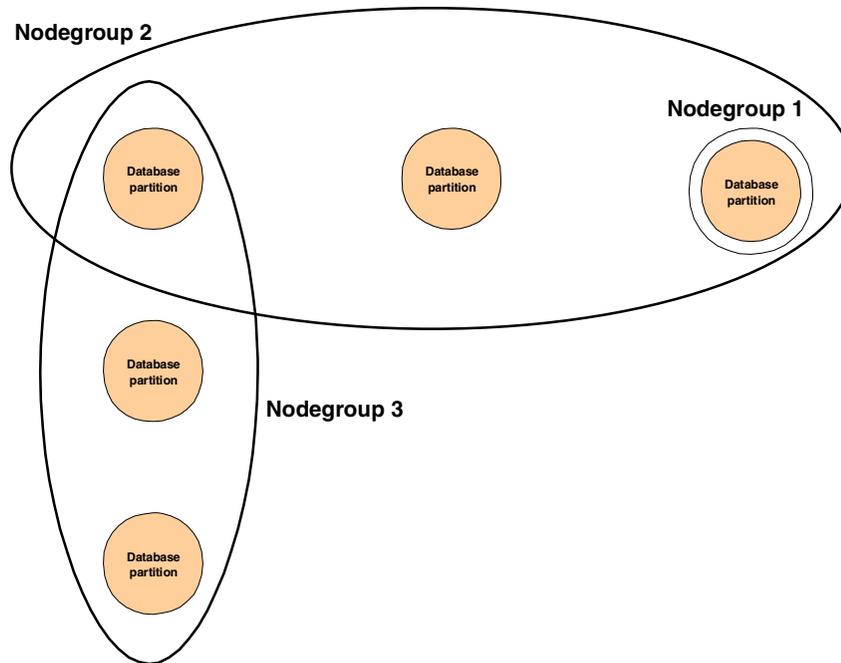


Figure 5-6 Nodegroups

5.2.6 Tablespaces

Tablespaces are created in nodegroups. A tablespace can span one or more physical storage device called containers. A container can be a directory name, a file name or a device name (raw device). Backup can be done on a tablespace level if roll-forward recovery is enabled.

There are two types of tablespaces, System Managed Space and Database Managed Space.

► System Managed Space (SMS)

A SMS tablespace will have directories as containers. The operating system manages the space for the data. Files are created in the directories when the tablespace is created. The size of these files are increased when required. You should not make direct changes to these files, move, or remove them. When tables are created in an SMS tablespace, their indexes and long fields or large objects (LOBs) will be created in the same tablespace.

► **Database Managed Space (DMS)**

A DMS tablespace provides better performance in certain situations. A DMS tablespace will have files and raw devices as containers. The database manager controls the storage space. The space for the files or raw device are pre-allocated, and they cannot increase in size. To increase space for a tablespace, containers can be added. Tables created in DMS tablespaces can have their indexes and long fields or large objects (LOBs) in separate tablespaces. Better performance can be achieved, for example, when indexes are created on tablespaces with faster devices. A good design will have a tablespace to contain tables only, indexes only, or LOBs only data.

A DMS tablespace can be defined as:

- A regular tablespace to store tables and indexes
- A long tablespace to store long fields or LOBs

When using tablespace backup for DMS tablespaces containing tables, consider backing up their corresponding index and LOB tablespaces at the same time.

5.2.7 Recovery logs

Recovery logs or log files are used to recover from applications or system errors. By default, the log files are used in a circular fashion. When the last log file is full, the first log file is reused. With circular logging, you can only do offline database backup. Users cannot use the database while an offline backup is ongoing. Also, you can only use version recovery, that is, recovered from the last available backup. Log files are not applied for version recovery.

Roll-forward recovery can be enabled by setting *logretain* or *userexit* to *on*. When roll-forward recovery is enabled, log files are kept and not reused, so they can be applied when performing roll-forward recovery. You can do online database backup where users can remain connected to the database while the backup is ongoing. You can also do tablespace backup either offline or online.

There are two types of log files:

- Active logs which contain current transaction data needed to do rollback or crash recovery
- Archive logs which contain committed data

Since DB2 does not provide multiplexing of log files, you must mirror the directory or the file system where the active logs are located.

Recovery history file

The recovery history file contains backup summary information that you can use to recover all or part of the database to a point in time. The recovery history file is always backed up during a database or tablespace backup. You have the option to restore only the history file when it is deleted or corrupted.

5.3 Database planning considerations

Planning is one of the most important areas to consider before beginning to implement a database. It is important that the database administrator and other relevant systems administrators work together to anticipate the circumstances in which the database will be utilized by the firm, as well as the resource and configuration requirements. These ideas apply to all types of databases. A fact statement will be presented first.

Important: Data is a strategic asset to the firm. Depending on the type of data kept in the database, the loss of data will result in economic losses to the firm in the form of:

- ▶ Lost customer information resulting in customer dissatisfaction which may have legal consequences
- ▶ Lost information required by governmental entities which may have legal consequences
- ▶ Lost employee information which may have legal consequences
- ▶ Lost employee productivity
- ▶ Lost sales
- ▶ Information loss that is relevant to managing the firm

To assist with these considerations, we will include details of some data recovery situations. We also cover some factors that should be weighed against each other in planning for recovery, for example, type of database, backup windows and relative speed of backup and recovery methods.

5.3.1 Backup requirements

A backup strategy is only one part of your overall data management plan. You must consider how important your data is to the function (or even existence) of your organization. The less time that your organization can function without its data the more important that data is to you. Your system must be designed in such a way as to keep important data available when a failure occurs. However, reliance on backups is not necessarily sufficient. You must also consider these resources:

- ▶ Redundant Array of Inexpensive Disk (RAID) devices
- ▶ Dual access paths
- ▶ Dual I/O controllers
- ▶ Dual power supplies
- ▶ Backup or standby processors
- ▶ Uninterruptable power supplies

None of these on their own can guarantee the availability of your data, but in combination they can reduce the impact of a failure.

Before you can design a physical and logical implementation strategy you need to define the requirements that the strategy must satisfy. Factors that you will need to consider when defining the requirements for your implementation strategy include:

- ▶ Types of events (The categories of incidents that may occur)
- ▶ Speed of recovery (How quickly you need to be able to recover)
- ▶ Backup windows (The periods of time at which backups can be performed)
- ▶ Recovery points (To which points in time you need to be able to recover)
- ▶ Units of recovery (Which other tables and files need to be recovered to the same point in time)
- ▶ Performance Matrix
- ▶ System Availability Matrix

5.4 Creating a database on the iSCSI storage device

This section discusses the creation of a database on different logical disks as implemented on the IBM TotalStorage IP Storage 200i. The server used in this model was an IBM Netfinity 7600. More information can be found at:

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

Basic configuration

The basic configuration of the machine is:

- ▶ IBM Netfinity 7600 Server:
 - 512 MB SDRAM
 - One 700 mhz processor
 - Two 18 GB internal disk attached via ServRAID 4x:
 - Logical drives C: D: and E:
 - Three logical disk drives on the *IBM TotalStorage IP Storage 200i* attached via TCP/IP Ethernet iSCSI:
 - Logical drives F: G: and H:
- ▶ The operating system is Windows NT 4.0 SP6a.
- ▶ The database version is DB2 UDB EE V7.1 with Fixpak 3.

5.4.1 Creating a database

The *DB2 V7.1 Administration Guide* provides detailed information regarding database administration. DB2 documentation may be downloaded at the following Web site:

<http://www.ibm.com/db2>

The IBM TotalStorage IP Storage 200i was used in this project. Example 5-1 shows a representative SQL for the creation of a database on different logical drives attached to a Window NT/2000 server.

Before creating the database, DB2 must be installed on the server. Drive D: was the drive chosen as the drive for the installation of DB2 binaries.

Example 5-1 Sample SQL

```
CONNECT RESET;
CREATE DATABASE ISCSIDB ON H: ALIAS iSCSIdb
USING CODESET IBM-1252 TERRITORY US COLLATE
USING SYSTEM USER TABLESPACE MANAGED BY
DATABASE USING (FILE 'F:\db2\node0000\sql00005\data1' 5120)
EXTENTSIZE 16 PREFETCHSIZE 16 OVERHEAD 14.06 TRANSFERRATE 0.33
CATALOG TABLESPACE MANAGED BY SYSTEM USING ('H:\db2\iSCSIdb\syscat\syscat1')
EXTENTSIZE 8 PREFETCHSIZE 8 OVERHEAD 14.06 TRANSFERRATE 0.33
TEMPORARY TABLESPACE MANAGED BY SYSTEM USING
('G:\db2\node0000\sql00005\temp\temp1')
EXTENTSIZE 32 PREFETCHSIZE 32 OVERHEAD 11.67 TRANSFERRATE 0.31 WITH iSCIS
Database
```

Note: SQL command sequences may be used in a batch file, saved as a script in the DB2 Command Center, or executed from the DB2 Command Line Processor (CLP).

5.5 Implementation of the AutoExNt-Service

At system reboot, various system services are started. To insure that network attached disks are available, the Microsoft Windows AutoExNt-Service is implemented. This package is available in the Microsoft Windows NT 4.0 Resource Kit at:

<http://www.microsoft.com> and/or <http://support.microsoft.com/directory>

The script in Example 5-2 will attempt to connect the iSCSI drives, if it fails to connect it will not start the DB2 Services, and intervention is then required.

Example 5-2 Single attempt startup script

```
net stop "Computer browser" /yes
net stop "server" /yes

:drive_offline
net start "IBMiSCSIcfg"

sleep 10

net start "server"
net start "computer browser"
net start "Distributed File System"

if exist g:\ goto iscsi_yes
Goto ISCSI_Error

    sleep 5

goto drive_offline

:iscsi_yes
net start "DB2 - DB2"
net start "DB2 Governor"
net start "DB2 JDBC Applet Server"
:iscsi_error
```

The script in Example 5-3 will keep attempting to reconnect the G: drive until it is successful. After connection of the G: drive, the required DB2 services will be started. We recommend that this script is implemented at the end of the instigation, after the drive Letter has been selected.

Example 5-3 Multiple attempt startup script

```

net stop "Computer browser" /yes
net stop "server" /yes

:drive_offline
net start "IBMiSCSIcfg"

sleep 10

net start "server"
net start "computer browser"
net start "Distributed File System"

if exist g:\ goto iscsi_yes

net stop "IBMiSCSIcfg" /yes
net stop "computer browser" /yes
net stop "server" /yes

sleep 5

goto drive_offline

:iscsi_yes

net start "DB2 - DB2"
net start "DB2 Governor"
net start "DB2 JDBC Applet Server"

```

Note: For more information on scripts, refer to 2.2.2, "Disk drive addressing" on page 34.

Table 5-1 gives all the affected services and required startup settings.

Table 5-1 AutoExNt-Service

Service	Status @sys boot	Startup
DB2 - DB2	Started	Manual
DB2 Governor		Manual
DB2 JDBC Applet Server	Started	Manual

Service	Status @sys boot	Startup
IBMiSCSIcfg	Started	Manual

Note: The term external disk refers to a disk that is not installed locally on the relevant server and normally SCSI attached via an adapter. The term internal disk refers to a disk attached via TCP/IP Ethernet on the IBM Total Storage iSCSI solution or via the Cisco SN5420 Storage Network Router.

DB2 UDB binaries installed on the IBM 200i external disk

When installing DB2 EE V7, you are given the option to choose which drive to install. In this example, we tried using IBM TotalStorage 200i as the destination disk. This is to test the effect of having the DB2 executables on an external drive. A database was also created on another IBM TotalStorage 200i drive as represented in Example 5-1 on page 169.

To ensure a proper startup procedure, all DB2 UDB EE services must be *manually* started in the **AutoExNt.bat** file.

DB2 UDB binaries installed on server internal disk

In this case, the installation utilized may be the default method of DB2 UDB EE or another *internal disk* may be chosen for the installation of DB2 UDB EE. During the modeling of this method, as long as the **IBMiSCSIcfg** service was manually started, there were no noticeable impacts to DB2 UDB operation. However, there were intermittent messages in the Windows event logs.



Implementing Oracle9i on iSCSI

This chapter contains information about Oracle9i database servers. Oracle improvements were designed to increase availability and supply better performance and easier management. Oracle release notes state, “New Oracle9i database features deliver the performance, scalability, and availability essential to hosted service software made available to anyone anywhere”. This chapter explores features that affect Oracle9i database performance and key aspects of tuning.

A vast amount of information is available in other redbooks regarding the topics covered in this chapter. Therefore to avoid duplication, you will see references to other books and documentation throughout this chapter.

In order to evaluate performance and draw comparisons, Oracle instance was installed on Windows 2000 Advanced Server and Linux. The databases were created locally and on iSCSI disks.

The Oracle9i new architecture includes: a database, application server and a developers suite. This chapter will discuss only the database component, but it is important to recognize that the other components play an important role in performance when used.

The topics that are *not* covered in this book are: real application clusters or RAC (RAC replaces Oracle parallel server), replication, partitioning, data warehousing, security, associated utilities, development tools, and Internet options and tools. Note that there are several new Oracle9i database features associated with the topics mentioned above. For a list of features, see “Other improvements” on page 189.

The topics covered in this chapter are wide, from the Oracle9i database new features, database requirements, configuration, and tools, to performance and tuning.

The purpose of this chapter is to evaluate performance on iSCSI disk. To do this, we need to understand the different database areas that affect performance. We begin by describing the new database architecture and new features to understand the effect that they will have on performance. We ask ourselves certain questions. Which Oracle features are key to performance?

Planning, implementation and configuration of the database are critical to performance. What type of installation and database configuration is best for your business?

Once we have the appropriate components in a database and we expect good performance, How do you measure performance? How do you tune your system and database? And, What do we recommend for better performance? Do you have issues? We will explore all of these questions.

6.1 Database architecture

The Oracle9i *basic* architecture has not changed much since Oracle8i. Oracle9i database logical and physical structures, memory structures, processes and store procedures are basically the same.

In Oracle, you have in one hand, memory and processes, and in the other hand database files, control files and redo logs. Even though you have an Oracle database, you also have an Oracle instance. And, you also have the client software with processes running. In the next sections we will differentiate between these three components in order to better understand the Oracle architecture.

Figure 6-1 is an architectural diagram which shows memory and process structures that manage the Oracle9i database.

The user/business data resides on the physical datafiles. Within datafiles we have logical structures such as segments, tablespaces, tables, indexes, and so on.

Note: For more information about Oracle Database Architecture and Oracle Architecture, in general, please refer to Section 4.2 “Oracle Database Architecture” in the redbook *Database Performance on AIX in DB2 and Oracle Environments*, SG24-5511.

It is important to distinguish between an Oracle instance, an Oracle database and an Oracle client. The following sections describe the differences among these components.

6.1.1 Oracle9i instance

An Oracle9i instance is the group of processes and system area that manages an Oracle9i database; see Figure 6-1. Attached to this instance is a Program Global Area (PGA) which will be discussed later in this section.

An Oracle9i instance has two components: System Global Area (SGA) and background processes.

Here are brief descriptions of these components:

- ▶ **The SGA** is a shared memory region that contains data and control information for one instance. When an instance starts, Oracle allocates the SGA, and the SGA is deallocated when the instance is shutdown. Memory structures created within the SGA are fixed in size.
 - *Memory Structures* — There are two basic memory structures: the system global area (which includes the database buffers, redo log buffers, and the shared pool) and the program global areas. Memory structures are created and used to complete jobs, for example, memory stores program code being executed and data that is shared among users.
 - *Database buffers cache* is a set of database buffers in an instance. Buffer cache contains the most recently used block data, with modified and unmodified blocks. The right sizing of buffer cache and disk I/O reduction improves performance.
 - *Shared pool* contains shared memory constructs such as SQL areas. This area processes the unique SQL statement submitted to a database, and contains a parse tree and an execution plan for that unique statement.
 - *Redo log buffer* logs all changes made to the database. The redo entries stored in the redo log buffers are written to an physical redo log file. These redo files are critical in database recovery.

- ▶ **Background processes** — Every Oracle instance uses several background processes (Figure 6-1). The names of these processes are:
 - DBWn — the *database writer process* is responsible for writing modified blocks from the database buffer cache to the datafiles. An instance could have up to nine DBWn processes running. A modified block, in this case, is a set of modified records that will be written to disk (datafile) by this process.
 - LGWR — the *log writer process* writes the redo log entries to disk. Redo logs performs sequential writes into the online redo logs. In other words this process takes the data records from the redo log buffer and copies them into the physical redo log files.
 - CKPT — the *checkpoint process* takes place in a timely manner. At a checkpoint, a signal to the DBWn process occurs, then the DBWn process will update all the datafiles and control files of the database. This process ensures that at an interval the DBWn process will be signaled to write records.
 - SMON — the *system monitor process* performs crash recovery when a failed instance starts up again. This process also performs regular necessary maintenance.
 - PMON — the *process monitor* performs recovery on a user process and free resources no longer needed.
 - ARCn — this *archiver process* process is active only when the archive mode is enabled. This process is responsible for copying the online redo log files to archival storage. LGWR will automatically start more ARCn processes, up to the maximum of ten, during workload increases.
 - RECO — the *recoverer process* is used to resolve distributed database transactions that are pending due to network or system failure.
 - Dnnn — the *dispatcher processes* are optional background processes, and present only when a multi threaded server configuration is used. At least one dispatcher process is created for every communication protocol in use.
 - LMS — the *lock manager server process* is used for Oracle9i Real Application Clusters.
 - Jnnn — the *job queue processes* are used for batch processing and managed dynamically. These processes not always active.
 - QMNn — the *queue monitor processes* are *optional* background processes that monitor the message queues for Oracle Advanced Queuing.

Program Global Area is a nonshared memory region which contains data and control information for a server process. This memory is created by the Oracle instance when a server process is started. The server process has exclusive access, and is read and written by the instance code on behalf of it. Total PGA memory allocated and attached to an instance is referred to as an aggregate PGA memory allocation.

When a session is connected through a dedicated server, private SQL areas are located in the server processes' PGA. However, if a session is connected through a shared server, part of the private SQL area is kept in the SGA.

6.1.2 Oracle9i database

An Oracle9i database holds the data and has several components: control files, online redo log files, database datafiles, and a parameter file. This set of components is associated with only one database.

- ▶ *Control files* are small binary files used to start and operate the database. These files are updated constantly while that database is active (open). These files have to be present and are critical for the database to function properly. Here is some of the information contained in the control files:
 - Database name
 - Timestamp of database creation
 - Names and locations of associated datafiles and online redo log files
 - Tablespace information
 - Datafile offline ranges
 - Log history
 - Archived log information
 - Backup set and backup piece information
 - Backup datafile and redo log information
 - Datafile copy information
 - Current log sequence number
 - Checkpoint information
- ▶ *Redo log files*, when archiving mode on, hold the redo records written from the redo log buffer. The LGWR process actively writes to the current online redo log file. These files are critical, because they are required for instance recovery.
- ▶ *Database datafile* is the physical disk that holds the Oracle logical structure called tablespace. There can be multiple datafiles and tablespaces per database, and several datafiles that expand over one tablespace. These files hold the data. Files sizes vary, but it is recommended that files are not bigger than 2G.
- ▶ *Parameter file* is called the initialization parameter file, because it is read whenever an Oracle instance is started. This parameter file holds database

configuration parameters and values. Certain parameters cannot be changed in this file after the database is created: database block size, database name, database domain name and compatibility parameter.

6.1.3 The Oracle client

The Oracle9i client software is called SQL*Net Version 2.1. In order to establish communication to a database, client software needs to be present and running. The types of information the client software needs to know to establish communication are:

- ▶ Which network protocol to use
- ▶ The location of the host, IP address, or host name
- ▶ The Oracle SID, which is the database name

When we have a user logged to a client application that needs to communicate to a database. A request is sent to SQL*Net, which handles the different internal datatype representations and character set resolutions between the client and host. SQL*Net will then pass the request, via a process called *listener*, to the Transparent Network Substrate layer (TNS). TNS will deal with connectivity issues and host location.

A listener process establishes a communication pathway to Oracle (Figure 6-2). If a user process makes a connection request, the listener determines whether it should use a shared server dispatcher process or a dedicated server process, to establish an appropriate connection. A listener process is an independent process that runs on the host where the Oracle database resides. SQL*Net and its components (listener process and TNS layer) are configured when the Oracle client is installed through Oracle Net Client Assistant.

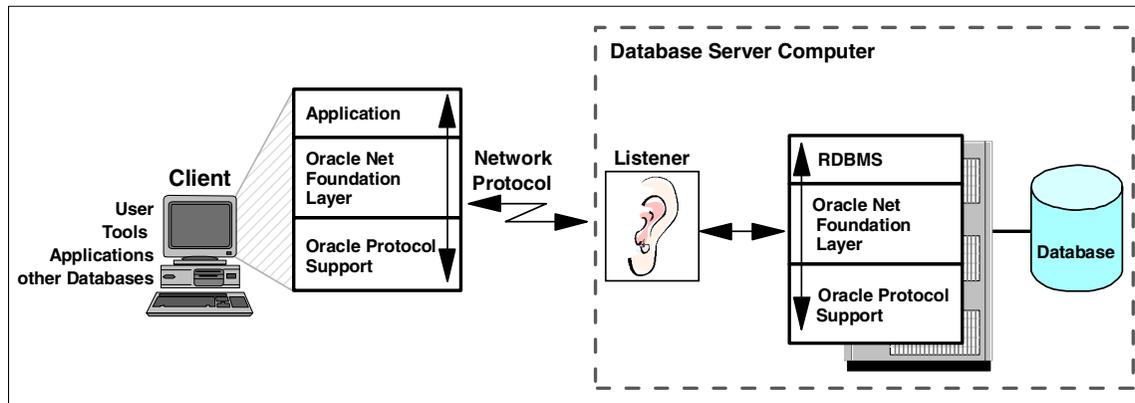


Figure 6-2 Listener process listens for incoming requests

The listener process need to be started in order for the users to reach the Oracle database. Listener components need to be configured with port, communication protocols, database names, database domains and host names (or IP addresses). The listener process could be set up per database or per host.

Note: For more information on Oracle9i database architecture or the listener process, review the *Oracle9i Database Administrator's Guide* or visit the Oracle Web site:

<http://www.oracle.com>

6.2 Oracle9i new features

This section will focus on Oracle's new features and enhancements. Later we discuss how these features and enhancements affect database performance.

Improvements to the Oracle9i databases were also accomplished by new tools and utilities. In combination with the new features, improvements deal with different aspects of performance. Unfortunately because of time constraints, we were not able to implement and test the performance of all features. The testing of performance was accomplished using a set of features on a standard database.

Table 6-1 lists the new and enhanced Oracle9i database features and tools, except for cluster/parallel servers, data warehousing, replication, security and development tools, because they are beyond the scope of this book.

Table 6-1 Oracle9i new features

Area/option	Feature	SE	EE	PE	Description
Content Management	Dynamic Services	Y	Y	Y	Provides Web Service Aggregation and Transformation, Administration, Definition and Registration, and Delivery. ICE 1.0 support.
Content Management	Oracle Database Workspace Manager	Y	Y	Y	Creates heretical workspaces and uses metadata and views to version enabled tables in a long transaction framework.
Content Management	Parallel Text index creation	N	Y	Y	Enables parallel Text index creation on non-partitioned as well as partitioned tables
Content Management	Ultra Search	Y	Y	Y	Provides unified searching across the Oracle database, the Web (HTML), and external data sources.

Area/option	Feature	SE	EE	PE	Description
Database Features	Advanced Queuing	Y	Y	Y	Provides direct support in the database for high performance queuing and messaging operations, integrated with LDAP.
Database Feature	Flashback Query	Y	Y	Y	Allows data to be queried from a point in the past.
Database Feature	Online table reorganization	N	Y	Y	Redefines the logical and physical structures of tables while the user is online
Database Feature	Quiesce database	N	Y	Y	Allows for database maintenance without it forcing a shutdown
Database Feature	Trial Recovery	N	Y	Y	Allows users of the database server to foresee and test recovery
Development	iSQL*Plus	Y	Y	Y	Provides three-tier implementation of SQL*Plus running in a Web browser.
Development	PL/SQL native compilation	Y	Y	Y	Allows PL/SQL byte code to be converted to C and dynamically linked into the Oracle9i database server
Development	User defined aggregates	Y	Y	Y	Allows for user defined aggregation functionality such as group-by on object types and collections to allow efficient aggregation for data cartridges
Development	XML	Y	Y	Y	Stores XML documents in the databases via SQL and renders traditional database data as XML documents.
Integration	Messaging Gateway to IBM MQ Series	N	Y	Y	Integrates the database server with IBM MQ Series.
System Management	Basic Standby Database	Y	Y	Y	Allows the DBA to manually clone a database, and to copy and to apply log files to the standby. Users can reconnect to standby should primary system fail.
System Management	Block-level Media Recovery	N	Y	Y	Restores and recovers individual blocks instead of entire files.

Area/option	Feature	SE	EE	PE	Description
Systems Management	Data Guard	N	Y	Y	Provides for automated physical standby instantiation, configuration, monitoring, switchover, switchback, failover, zero loss mode, and delayed apply support, all through a GUI tool. Can open a standby database as read only but no DML is allowed.
System Management	Fast-start selectable recovery time	N	Y	Y	Provides an init.ora parameter to set how fast the database server should perform crash recovery, guaranteeing a predictable recovery time.
System Management	Global index maintenance during DDL operations	Y	Y	Y	Automatically maintains global indexes when DDL operations are executed against partitioned tables.
System Management	Multiple Block Size	Y	Y	Y	Allows a database to have more than one block size.
System Management	Oracle Managed Files	Y	Y	Y	Automates creation and deletion of physical database components.
System Management	Resumable Space Allocation	Y	Y	Y	Allows execution of large database operations to be suspended and resumed if there is a repairable failure.
System Management	Standby Database GUI	N	Y	Y	Supports easy installation and configuration of Standby and supports monitoring of critical Standby operations.
System Management	Unused index identification	Y	Y	Y	Allows database server to track unused indexes.
* Option	Oracle Data Mining	N	Y	Y	Oracle Data Mining embeds data mining functionality for making classifications, predictions and associations. All model building and scoring functions are accessible through a Java-based API.
* Option	Oracle OLAP	N	Y	Y	Oracle OLAP provides the Oracle OLAP API, which provides support for complex, multidimensional queries

Note: Abbreviations on columns are SE Standard Edition, EE Enterprise Edition, and PE Personal Edition. Also, the table columns' SE, EE and PE values indicate (Y/N) if the features or options are available for installation type. Oracle options are preceded by a *.

Some of these features and utilities will incur a performance benefit and a *cost*, so be careful with the choices offered.

6.2.1 New feature categories

The following sections are divided by categories, which contain further details on relevant new features. If a feature from Table 6-1 is not found in the following sections, it is because a particular feature is outside the scope of this chapter. The new functionality are in areas of disaster recovery, system fault recovery, availability, scalability, performance and manageability.

Availability

When the database is not available, it can decrease performance, and performance is null. The Oracle availability improvements are in three areas: disaster recovery, online data evolution, and precision data repair.

Disaster recovery

A disaster recovery environment is accomplished using four features:

1. Disaster recovery

- **LogMiner and LogMiner Viewer utilities** include log analysis and new graphical user interface. Viewer is a component of Oracle Enterprise Manager.
- **Data Guard, Data Guard Monitoring and Automation** are used for real disaster situations, standby database with automatic switch over. Data Guard requires to set the automatic switching of databases, low administrative intervention and monitoring utility.
- **Zero Data Log Transport component of Data Guard** allows log file updates that are synchronously written directly from the primary database to the physical identical standby database.
- **Delay mode** is the time lag that protects the standby database from applying corrupted or erroneous data coming from primary database. Delay mode delays the application of archive redo log when they arrive at the standby database.

2. *System fault recovery*

- **Fast-start fault recovery** allows a database administrator to specify the expected mean time to recover (MTTR) a single Oracle instance.
- **Fast instance freeze and resume** allows the system to take a diagnostic snapshot of the entire system at the time of failure. This allows the database administrator to start the database and continue processing. Then diagnostic analysis will be offline.

3. *Human error safeguard*

These types of errors are common to databases usage, and three features enable the correction of errors.

- **Flashback query** allows a user to find old data error, by specifying day and time of data update, then issue a standard query on the data as it appeared on the specified day and time. This will enable error correction through data updates.
- **Log analysis** — LogMiner covered in “Disaster recovery” on page 183 allows the database administrator to examine all updates on the database, including all data manipulation, definition and administrative commands.
- **Resumable space allocation** allows the database administrator to suspend an operation/process that cannot be complete because of a lack of resources (example disk space). Once the resource issue is resolved, the operation/process is allowed to resume from the point of interruption.

4. *Planned downtime*

Certain maintenance operations can be executed online instead of waiting for a schedule database downtime.

- **Online schema changes** enables the redefinition of a table structure while table continues to be available to users.
- **Online table and index reorganization** is used to reorganize and redefine tables while objects are available to users and applications. Table and index reorganization increase data availability, reduces downtime and disk fragmentation while improving data access performance. All these while data continues to be fully available on a normal business day.
- **Dynamic database parameters** affect the system memory area (SGA of the database) and can now be reset online. This could, depending on the values set, increase performance. Changes can be accomplished with an ALTER SYSTEM statement, and will take effect immediately. For a change to stay permanently, you will have to change the parameters file.

Online data evolution

Earlier in “Planned downtime” on page 184, we discussed online reorganization and redefinition of tables and indexes. This new architecture is discussed in more detail, here.

- ▶ *New online architecture* provides the following capabilities:
 - As we mentioned earlier, physical attributes of a database object can be changed online, such as: table move to new location, table partitioning, and conversion from one type of table organization (heap-organized) to another (index-organized).
 - Logical attributes can be changed, such as: column names, types, and sizes. Columns can be added, deleted or merged. The only exception is that primary key column(s) modifications are not allowed.
 - Indexes can be created and analyzed online at the same time.
 - Online repair of invalid physical guesses by fixing the physical guess component of logical ROWID(s) stored in secondary indexes.
- ▶ *Quiesce database* is used to perform operations that demand no active transactions.
- ▶ *Buffer cache and share pool* can be resized dynamically. For more information see Memory management – on page 187.
- ▶ *Full object* (such as, tables, indexes) access while the analyze validate command is performed to validate structures.

Precision database repair

Precision database repair is a set of features that prevents and improves the handling of log corruption, reducing the risk of extended downtimes due to failures.

- ▶ *Rapid crash recovery* uses a two-pass recovery algorithm that ensures that only the blocks that need to be processed are read from and written to the datafiles (database data files). Also crash recovery time is controlled by a time-based mean time to recover the MTTR parameter to set the time limit.
- ▶ *Failed state diagnostics* is covered in Fast-start fault recovery – on page 184.

Scalability

Scalability could decrease performance as more users and database systems grow. If database usage increases or the database grows disproportional, the database and system will come to a halt, affecting performance directly.

Oracle9i scalability improvements have two areas: scalable session state management and fine-grained, automatic resource management.

- ▶ *Scalable session state management* allows hosting of more users and applications on the same or larger hardware platforms by improving:
 - Shared memory capabilities
 - Networking
 - Shared server to substantially reduce the overhead required for each user
- ▶ *Fine-grained, automatic resource management* has three new features that enhance resource management:
 - **Granular control over resources** allows maximum active sessions control and query execution time estimation. For example, an administrator is able to limit the number of concurrently active sessions per consumer group. When the limit is reached, other requests will be queued and processed once the active sessions finish.
 - **Automatic consumer group switching** allows automatic conditional switching to other consumer groups of a long running session. For example, an OLTP consumer group switched (once it reaches certain conditional criteria) to a more process intensive consumer group for batch processing.
 - **Undo pool quota** limits the amount of rollback data generated per resource consumer group, and prevents transactions from consuming rollback space and impacting system operation.

Performance

Performance is a concern in most database systems. Oracle9i performance improvements focus in three areas: native compilation and PL/SQL optimization management, latch contention, and network and distributed database performance. Here we discuss the enhanced Oracle9i performance features:

- ▶ *Native compilation and PL/SQL optimization management* allows for native compilation and improves PL/SQL optimizations. Cost-based optimizer with better optimized plans, less resource usage and more memory, results in performance increase.
- ▶ *Latch contention* — Reduction or elimination of latches in several areas improved performance in active systems: I/O improvements are: self-tuning, direct I/O, prefetching, skip/scan indexes row resource operations, and improve performance on OLTP environments.
- ▶ *Network and distributed database performance* enhancements are:
 - Database-to-database communication using ICO
 - Network interface optimization
 - Virtual circuit I/O

- Unified event/wait model
- Distributed query optimization

Manageability

The objective of Oracle9i System Management is to make the database server inherently self tuning and managing. Several features enable the database administrator to delegate some of the day-to-day administrative tasks to the server.

- ▶ *Database self-management* — The implementation of an autonomous and self-managing database server is accomplished by:
 - **Self-managing undo** allows the database to take care of undo block contention, consistent read retention, and space utilization. These features require that an administrator allocate their undo space in a single undo tablespace. Rollback segment planning and tuning is eliminated.
 - **Memory management** — Dynamic resizing of buffer cache and shared pool. It also includes a buffer cache size advice mechanism that predicts the performance of running with different sizes for the buffer cache.
- ▶ *Streamlined operational management of the database* — These features are simple database administration improvements:
 - **Persistent initialization parameters** allow certain parameters to be changed online by the administrator, OEM, or internal self-tuning to maintain values after a database is shutdown. This feature also allows the database to start from a remote machine without a need to have a copy of the initialization file. Some parameters persist across database shutdowns.
 - **Oracle database configuration assistant** creates existing database templates from which databases can be generated. This tool also provides sample schemas that can be used as needed.
 - **Oracle-managed files** manage files by automating the routine tasks of creation and deletion of database files. Administrators are still required to plan and administer space.
 - **Corresponding sub-cache configuration** allows databases with multiple block sizes, and allows administrators:
 - To configure corresponding sub-cache within each alternative block size
 - To locate objects in tablespaces of appropriate block size to maximize I/O performance
 - To transport tablespaces between different databases, for example, OLTP environment to a data warehousing environment

- **Mean time to recover (MTTR) specification** was discussed in Fast-start fault recovery – on page 184.
- **SQL queryable through V\$SQL_PLAN View** allows the administrator to see the actual plan use to execute a query statement.
- ▶ *Enhanced recovery manager (RMAN) performance* provides several features:
 - **Persistent RMAN configuration** allows for one-time backup configuration that applies to any session. Permanent settings could be created for:
 - Automatic channels
 - Channel parallelism
 - Retention policies
 - Backup options
 - Auxiliary filenames
 - **Recovery window** — Recovery manager reduces the administrator’s time and effort on backup tasks, such as:
 - Backup expiration policy control
 - Obsolete backup automatic marking
 - Identification of archive logs no longer required for point in time database recovery
 - **Block media recovery (RMAN)** allows media recovery on individual blocks in a datafile while the datafile is online.
 - Other manageability enhancements are:
 - Control file auto backup
 - Batch termination of online backup mode
 - Perform trial media recovery
 - Multiple conversion pairs for the *_FILE_NAME_CONVERT parameters
- ▶ *Easier administration using Oracle Enterprise Manager (OEM)* — These features are:
 - **Advice-incorporated tools**
 - This OEM feature provides advice and recommendations for appropriate Oracle configurations.
 - This feature monitors performance and quickly resolves problems, with appropriate reporting and alerts.

SQL and PL/SQL improvements

A database’s basic components have improved in this Oracle version. How SQL and PL/SQL perform is relevant to the overall performance of the database. These improvements increased the performance within Oracle9i database:

- ▶ Extracting of database definitions, using `dbms_metadata` mechanism. The results are in XML or SQL DDL formats.
- ▶ SQL compilation: PL/SQL immediately supports all SQL syntax changes to embedded SQL.
- ▶ SQL parallel query mechanism extended for store procedures written in 3GL languages. Stored procedures can return incrementally to the calling SQL statement.
- ▶ Support for ANSI-style *CASE* statements.
- ▶ Native compilation provides better support for compute-intensive processing.
- ▶ SQL and PL/SQL engines are more tightly integrated improving performance.
- ▶ Reduced overhead from calling PL/SQL procedures from SQL.
- ▶ Improved data conversion between raw and numeric types.

Other improvements

Here are other improvements in the Oracle9i database:

- ▶ *Index-organized tables* — The advancements include:
 - Bitmap indexes and parallel DML on index-organized tables
 - Hash method partitioning
 - B-tree index on UROWID columns
- ▶ *Bitmap Join Index* is an index structure which spans multiple tables improving performance on joined tables.
- ▶ *Materialized View Mechanisms*
 - Query-rewrite enables a single materialized view to address a broader class of queries.
 - Incremental refreshes for a wider variety of materialized views.
- ▶ *Automatic Memory Tuning* allows memory-intensive queries sufficient memory to run while memory-light queries will be given a lesser amount of memory. This increases overall query performance.
- ▶ *Multiple Inserts* — There are two SQL capabilities to increase data load scalability while reducing complexity:
 - Multiple inserts, single SQL statement to insert into multiple tables, more efficient than separate SQL statements.
 - Merge, conditional incremental refresh that requires two tasks: new records will be inserted and existing records will be updated simultaneously by a single SQL statement.

Notes:

- ▶ Oracle Net8 for Oracle9i is SQL*Net. Service manager (svrmgrl) and CONNECT INTERNAL are not supported in Oracle9i. These features were replaced within SQL*Plus. New syntax is discussed in 6.5, “Hints and tips” on page 283.
- ▶ For more information on Oracle9i database new features, review *Oracle9i Database Release Notes*, *Oracle9i Database Administrator’s Guide* or visit the Oracle Web site:
<http://www.oracle.com>
- ▶ For more information on Oracle clustering, review the Redpaper: *Implementing Oracle9i RAC with Linux on IBM @server xSeries Servers*, REDP0410.

6.3 Database planning

Database planning should be given careful consideration. Performance will improve by choosing the appropriate combination of:

- ▶ Oracle9i installation type
- ▶ Oracle9i database type
- ▶ Hardware and communication protocols
- ▶ Applications and accessing tools

Note:

- ▶ Hardware and communication protocols are discussed in other sections of this book.
- ▶ Applications and accessing tools are business-specific and therefore will not be covered in this book.

In this section we will discuss Oracle9i installation and database planning for a small database environment. A Standard Edition installation and a general purpose database is created.

To plan for an Oracle9i installation and database configuration type, you should first answer some questions. This will ensure your installation and database choices match your business needs:

- ▶ What type of data will be stored in the database? Is the data critical to the business?
- ▶ What is the database size expected? How much data? How much will the data will grow in 3 months, 6 months, a year?

- ▶ How many users will be accessing the database? How many users can your system handle?
- ▶ What application and tools will be used to access the data?
- ▶ What are the hours in which the database needs to be available? When will you perform maintenance on the database? When will you schedule backups?
- ▶ What is the life expectancy of this database? Will you migrate historical data, and when? Will you migrate to a large system in a couple years?
- ▶ Will batching and the loading of data be performed on the database during heavy transaction periods?
- ▶ Once the database grows, do you have resources to fine tune the database? When should you revisit performance issues?

The answers to these questions accomplish two things. It helps you think about the planning issues, and gives us an opportunity to explain the planning for our database. Please keep in mind, this is only a guideline for our planning. Here are the questions revisited, and our responses.

What type of data will be stored in the database? Is the data critical to the business?

You have critical data, when your business *cannot* function properly without the data. Your data should be available as needed in a timely manner. Your *backup and recovery* process should be in place. Backup and recovery are covered in another chapter of this book.

What is the database size expected? How much data? How much will your data grows in 3 months, 6 months, a year?

For the system discussed in this book, we recommend a database size of less than 4 GB for a small database, with growth under 20% per year.

How many users will be accessing the database? How many users can your system handle?

For this system, the database should have under 75 concurrent users.

What application and tools will be used to access the data?

The applications and tools used to access the database should be '*performance friendly*'. Be aware that most of the applications and tools today in the market place are *not* performance friendly.

Performance friendly means software that includes performance benefits in its design.

What are the hours of availability for the database? When is maintenance performed on the database? When are database backups planned for?

A database which is required to be available 24x7 and /or contains critical data should:

- Have planned maintenance time to avoid an unplanned maintenance crisis
- Have planned database changes (to database structures, data loads, and so on)
- Have a scheduled backup on a regular basis
- Have healthy backup and recovery process in place. By healthy we mean, that the process has been tested successfully and data can be recovered anytime

What is the life expectancy of this database? Will you migrate historical data, and when? Will you migrate to a large system in a couple years?

If life expectancy of a database is more than two or three years and database growth is 20% per year, you should plan to migrate historical data or migrate to a larger system before you experience problems. You need to review your needs, growth rates and performance at least once a year — for a small, slow growing system.

Will batching and the loading of data be performed on the database during heavy transaction periods?

Planning for heavy transaction periods would reduce critical problems. Some of the issues are disk space, performance degradation and resource contention. Planning gives time for the users and database administrator to prepare for possible problems. For example, we have observed that after heavy loading of data, a database's performance degrades to a point of complete inactivity for users. With proper planning, performance degradation is unnecessary.

As the database grows, do you have the resources to fine tune the database? When should you revisit performance issues?

Database configurations do reach performance thresholds, and that is when a database will need a performance adjustment. It is not easy to guess when you will have a problem in your hands, but being aware of data growth rates, performance issues, configuration parameter (and how they affect performance), optimizer options, Oracle trace facility, applications which are known to degrade performance, and understanding the contents of the Oracle StatsPack report, gives you a better opportunity to surmise the Oracle.

While Oracle9i provides several new tools and features, we will take advantage only of the most basic components that this new version provides. We chose a small database, an Oracle9i Standard Edition installation, and a general purpose database with a basic database configuration. For Windows 2000, Oracle 9.0.1.1.1 will be installed; however, only the Oracle9i Version 9.0.1 for Linux is available at this time.

What are the hardware and software configurations necessary for our environments: Windows 2000, Linux, Windows 2000 with iSCSI disks, and Linux with iSCSI disks? Next, we will discuss the requirements for these four environments.

6.3.1 Requirements for the environments

We have chosen these environments for our performance tuning and testing. The reason for choosing these four environments is to draw performance comparisons among them. The following sections describe the requirements for these environments:

- ▶ Windows 2000
- ▶ Linux
- ▶ Windows 2000 with iSCSI disks
- ▶ Linux with iSCSI disks

Oracle9i requirements on Windows 2000

Here is a list of installation requirements on Windows 2000. The database will be installed on a local drive. Follow this list to avoid implementation problems:

- ▶ Windows 2000 Professional, Windows 2000 Server, Windows 2000 Advanced Server, and Windows 2000 Datacenter.
- ▶ Windows 2000 Service Pack, certified with 1 or higher.
- ▶ Minimal processor for the Standard Edition installation is Pentium 166 or Pentium 200. If you are installing more than just the basic components, you will need a better processor.
- ▶ RAM of 128 MB (256 MB recommended).
- ▶ NTFS file system:
 - Oracle software local home drive requires 2.7 GB disk (see “Virtual memory”, below).
 - Database local drive. Size is determined by your database needs.
 - System local drive 140 MB.
- ▶ Temporary space 400 MB.
- ▶ Virtual memory:

- Initial size 200 MB.
- Maximum size 400 MB.
- ▶ Video 256 colors.
- ▶ Cd-ROM drive.
- ▶ Oracle 9.0.1.1.1 database installation software for Standard Edition, either on CD or it can be downloaded from the Oracle Web site — requires 2.7 GB of disk space (see “Virtual memory”, above).
- ▶ Oracle9i client installation for Windows 2000. The administrative option will require 650 MB on the Oracle home drive. The run time version requires 486 MB of hard disk space. Client software comes on a CD or can be downloaded from the Oracle Web site.
- ▶ TCP/IP protocol.

These requirements are minimal, and you may need to upgrade (increasing temporary space, better processor, and so on) as other applications and tools are added, the data grows, or more users log on to the host.

Oracle9i requirements on Linux

Here is a list of installation requirements on Linux. The database will be installed on a local disk. Follow this list to avoid implementation problems:

- ▶ Minimum memory of 512 MB to install Oracle9i Server.
- ▶ Minimum memory of 128 MB to install the Oracle9i Client.
- ▶ Disk space for Oracle9i instance installation (database software) 2.5 GB.
- ▶ Disk space for database (depends on database size required). This is where the data will reside, the physical datafiles. The seed database will take 1 GB.
- ▶ The swap space required should be equal to twice the amount of RAM or at least 400 MB, whichever is greater.
- ▶ The processor that is recommended for the Standard Edition installation is Pentium II 233 MHz.
- ▶ Temporary disk space required for Oracle Universal Installer of 400 MB.
- ▶ A CD-ROM drive which is able to read CD-ROM disks, ISO 9660 format with RockRidge extensions.
- ▶ Oracle9i is certified on SuSe 7.1, Kernel 2.4.4, and operating system library GNU Lib C 2.2. Here is a list of other software required:
 - Any X server supported by UNIX operating system vendor.
 - Hummingbird eXceed using native window manager.
 - Try to run xclock to determine if X windows system is working properly.
- ▶ Executable required in user bin directory.

- ▶ Operating system latest patches, both kernel and library patches.
- ▶ Kernel parameters set to appropriate values for:
 - shmsys:shminfo_shmmax
 - shmsys:shminfo_shmmin
 - shmsys:shminfo_shmseg
 - shmsys:shminfo_shmmni
 - semsys:seminfo_semmns
 - semsys:seminfo_semmni
 - semsys:seminfo_semmsl
- ▶ Oracle9i components require Blackdown JDK/JRE 1.1.8_v3.
- ▶ Oracle 9.0.1. database installation software for Standard Edition, either on CD or it can be downloaded from the Oracle Web site.
- ▶ Oracle9i client for Linux.
- ▶ TCP/IP protocol.
- ▶ X windows server installed or X windows emulator, for Oracle installation and tools usage (for example, Database Configuration Assistant).

Oracle9i requirements on Windows 2000 and iSCSI disks

The installation requirements will be the same as for Windows 2000 (see “Oracle9i requirements on Windows 2000” on page 193), except that the database files (datafiles) will reside on the iSCSI disk. The requirements for this installation are slightly different than from Windows 2000. Here are the details:

- ▶ The iSCSI client initiator is required; see 2.3.8, “Client initiator setup procedure” on page 53.
- ▶ The requirements are the same as for the “Oracle9i requirements on Windows 2000” on page 193, except for the NTFS file system. The database location is on the iSCSI disk, and the requirements are:
 - Oracle software local home drive requires 2.7 GB disk.
 - Database *iSCSI* drive. Size is determined by your database needs.
 - System local drive 140 MB.

Once these requirements are satisfied, the next step is to plan for the Oracle installation by selecting the installation type and database configuration type.

Oracle9i requirements on Linux and iSCSI disks

The installation requirements will be the same as for Linux (see “Oracle9i requirements on Linux” on page 194), except the database files (datafiles) will reside on the iSCSI disk. The requirements for this installation are slightly different than from Linux. Here are the details:

1. The iSCSI client initiator is required; see the *IBM iSCSI Client for Linux Installation and Configuration Instruction Manual* (included in the IP Storage 200i box).
2. The requirements are the same as for “Oracle9i requirements on Linux” on page 194, except for the disk space for the database depends on the database size required. The database location is on the iSCSI disk:
 - Database *iSCSI* disk. Size is determined by your database needs.

Once these requirements are satisfied, the next step is to plan for the Oracle installation by selecting the installation type and database configuration type.

6.3.2 Oracle9i installation and database types

There are different types of Oracle installations and well as different types of databases to choose from. For example, Oracle9i Database Enterprise Edition will allow you to select different types of installations, while installing the software.

Oracle Universal Installer (OUI) is the software that allows you to install Oracle software. For Oracle9i Database Enterprise Edition downloaded software, the installation options are:

- ▶ Enterprise Edition, used for large companies and large complex systems.
- ▶ Standard Edition, used for smaller companies and basic configurations.
- ▶ Custom Edition, used for complex database systems with specific needs.

There is a Personal Edition that was not an option for this Oracle9i Database Enterprise Edition. The Personal Edition is a desktop database used mostly for development. This edition includes replication and distributed features that enable easy deployment to enterprise environments.

Note: For information regarding installations, please refer to Oracle's documentation or Web site:

<http://www.oracle.com>

Another option of the Oracle Universal Installer is the database configuration type. The installer will create a database for you at the end of the installation process. There are four configuration types of databases to choose from:

- ▶ **General purpose** — used for variety of database tasks, from simple transactions to complex queries. This is a general purpose database for relatively small systems.
- ▶ **Transaction processing** — used for online transaction processing, such as billing (Internet commerce sites). Processes small amounts of data and with simple operations, such as: read, write and changes to database tables.

- ▶ **Data warehousing** — used for large volumes of data with complex queries. Response time, accuracy and availability are a must for this type of database. Also used for Decision Support System (DSS).
- ▶ **Customized** — used for specialized database with selected Oracle components. This database requires configuration knowledge and time to configure database during installation.

The installer will also give you the option to install software only, which means that the installer will not run the configurations tools. No database creation or configuration, and no client configuration. These tasks will have to be performed manually.

When the Database Configuration Assistant is used to create and configure a database, the initialization parameters are set to certain values. At a glance, only the basic configuration parameters are different among the database types (general purpose, transaction processing and data warehousing), which are:

- ▶ Database cache size
- ▶ Java pool size
- ▶ Large pool size

As we discussed earlier in this chapter, we chose a Standard Edition installation for a general purpose database.

6.4 Database implementation and configuration

In this section we discuss the details for an Oracle9i software and client installation, database creation and configuration, and client configuration.

Note: Oracle Client Net8 for Oracle9i is SQL*Net V2.1.

The following sections will describe the installation and configuration steps for these different environments:

- ▶ Section 6.4.1, “Windows 2000 installation and configuration” on page 197
- ▶ Section 6.4.2, “Linux installation and configuration” on page 220
- ▶ Section 6.4.3, “Database creation on iSCSI disks” on page 242

6.4.1 Windows 2000 installation and configuration

Oracle installations and configurations require a minimum set of hardware and software levels. Make sure you comply with the installation requirements before proceeding. Ensure that you also fulfill the pre-installation requirements before continuing with the Oracle9i installation and configuration.

- ▶ “Oracle pre-installation on Windows 2000” on page 198
- ▶ “Oracle installation on Windows 2000” on page 199
- ▶ “Oracle post-installation on Windows 2000” on page 216

Oracle pre-installation on Windows 2000

For Windows 2000, check that you meet the pre-installation requirements. Because this is a basic installation of the Standard Edition, there are only a few pre-installations tasks:

1. Verify that all of the requirements are installed and available.
2. Make sure space is allocated for an Oracle instance, client software, and database creation. Review the section “Oracle9i requirements on Windows 2000” on page 193.
3. Create a user who is a member of the administrator’s group (used for Oracle installed).
4. Verify that network protocols TCP/IP is available by pinging on a command prompt:

```
ping localhost
```

You will get an output similar to Example 6-1.

Example 6-1 Ping output

```
Pinging a78rp1lg.almaden.ibm.com [127.0.0.1] with 32 bytes of data:
Reply from 127.0.0.1: bytes=32 time<10ms TTL=128
Ping statistics for 127.0.0.1:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

Once the pre-installation tasks are completed, you are ready to install Oracle9i.

Note: Oracle installer will create the Oracle home directory and database file directory. The paths for these directories will be the Oracle standard path default values. During the installation, you have the option to change directory names and paths.

Oracle installation on Windows 2000

The Oracle9i installation is performed using the Oracle Universal Installer (OUI) Version 2.0.1. There are several components installed with the basic Oracle software, such as replication, queueing, clustering, Oracle Enterprise Manager, Oracle Net Assistant, to name a few. Some of these components are used, and some are not. However, these components are part of the installation process unless you specify otherwise.

For the Oracle9i installation on Windows 2000, you downloaded three zip files, from the Oracle Web site or have three CD-ROMs. Follow these steps to begin the installation:

1. Log on as administrator or user belonging to an administrative group.
2. Create four directories:
 - OracleInstall
 - Under this directory to create three directories: Disk1, Disk2, and Disk3.
3. If you downloaded the files, extract files from zip files onto Disk1, Disk2, and Disk3. One zip file per directory.

Or, if you have the CD-ROM, copy one CD_ROM per directory onto Disk1, Disk2, and Disk3.
4. Close as many open applications as you have running.
5. At command prompt run **Disk1\setup.exe**.

The Windows 2000 Oracle Universal Installer (OUI) Welcome window will appear (Figure 6-3).

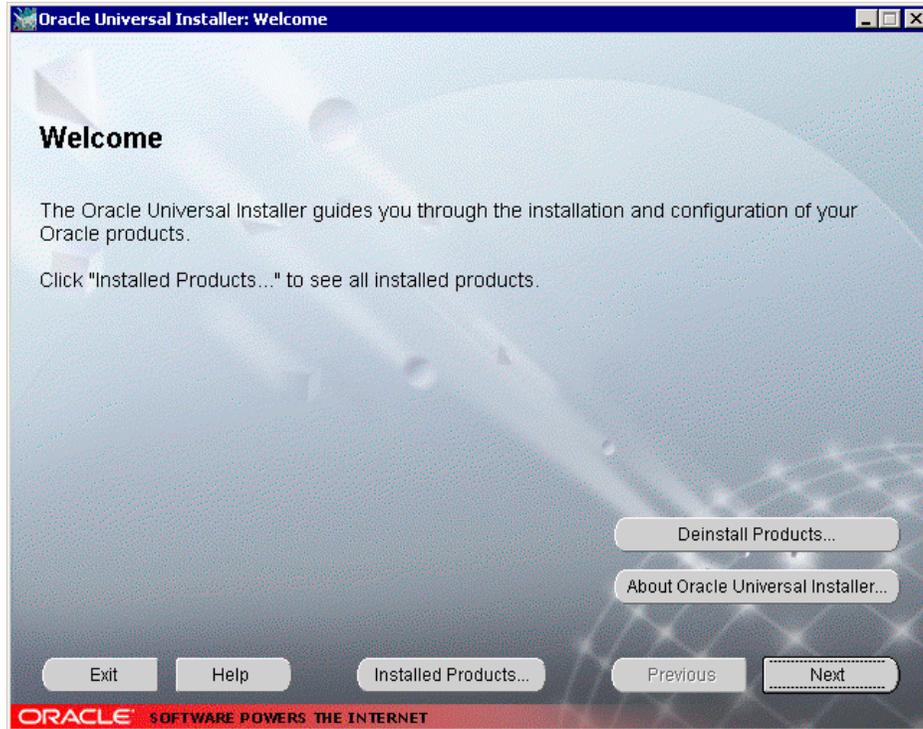


Figure 6-3 Oracle Universal Installer welcome window

To start the installation, click **Next**, which is always located at the bottom right corner.

The OUI File Locations window opens (Figure 6-4).

Note: If your system has Oracle products installed, you have the option to view a list of installed products, by clicking the Installed Products button. Or you can deinstall products, by clicking the Deinstall Products button.

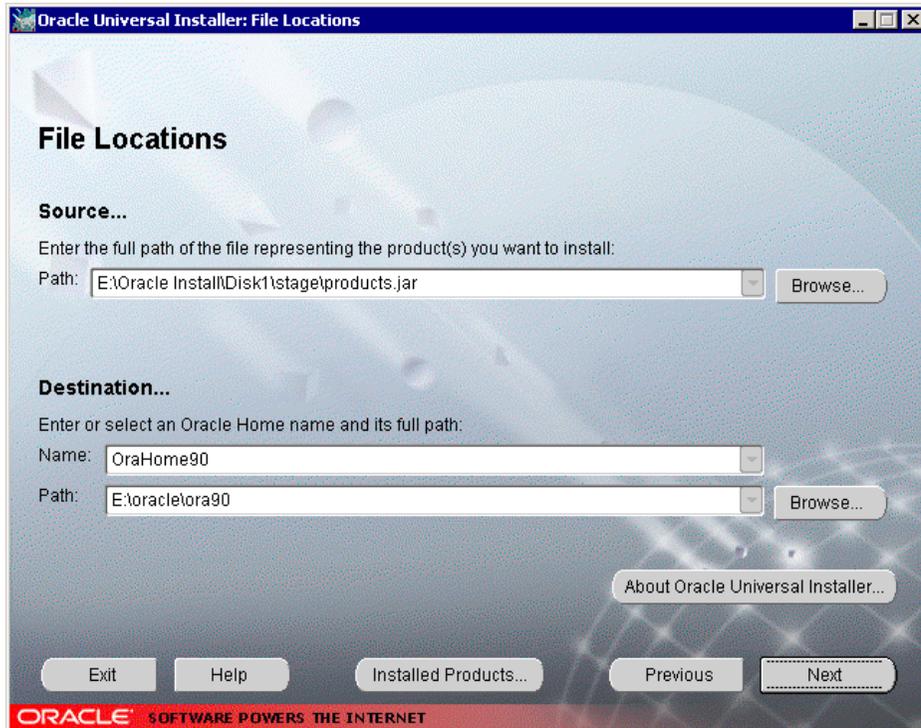


Figure 6-4 OUI File locations window

The options for this window are:

- ▶ Standard default values (that are displayed) *or* choose different paths:
 - The source path is one of the directories created for software installation in Step 2. on page 199, `\OracleInstall\Disk1`, where Oracle installation files reside. The source path is also the full path to where you copied or extracted the Oracle files (from CD-ROM or zip files). Most likely you will not have to change this option.
 - Or, you can choose your own destination names and paths. The destination is the full path where the Oracle software will be installed and run from. The current settings and defaults are Oracle standard destination full paths and names for the Oracle home. You could choose a different path and name to fit your own company standards. We chose OUI default settings.

Tip: You have the option to browse for both source and destination paths.

To continue, click **Next**. The OUI Available Products window opens (Figure 6-5).

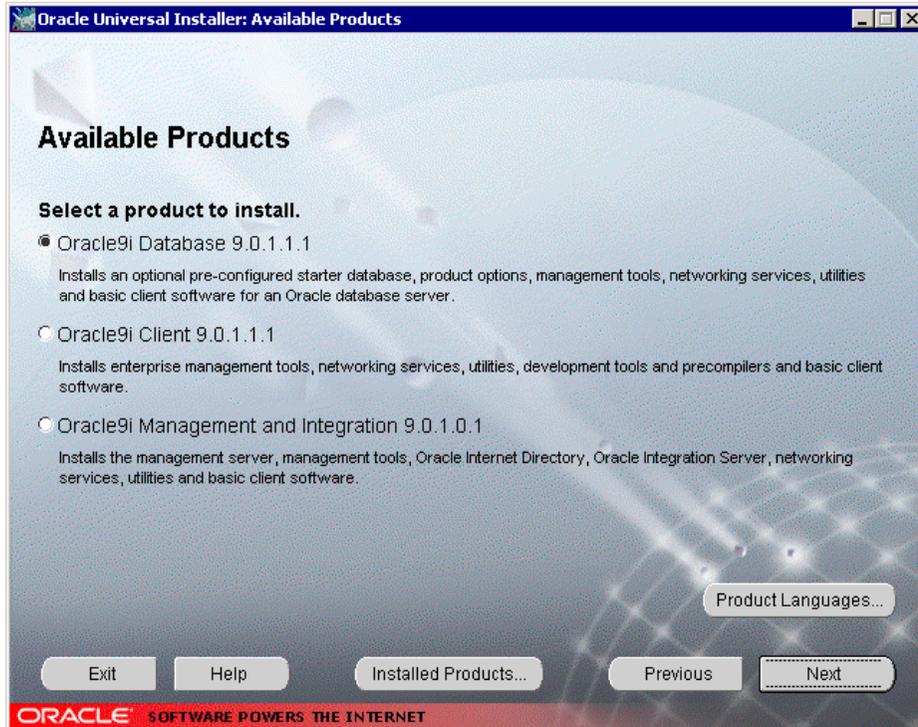


Figure 6-5 Available products window

The Product Language button is available to you. You could continue and allow the defaults, or click this option to view and/or change defaults.

Review your options in this window. Select the product you want to install, and then click **Next** to continue.

Note: We selected the Oracle9i Database 9.0.1.1.1 product. This option will include the Oracle9i software, the client installation and configuration, and database creation and configuration.

The OUI Installation types window opens (Figure 6-6).

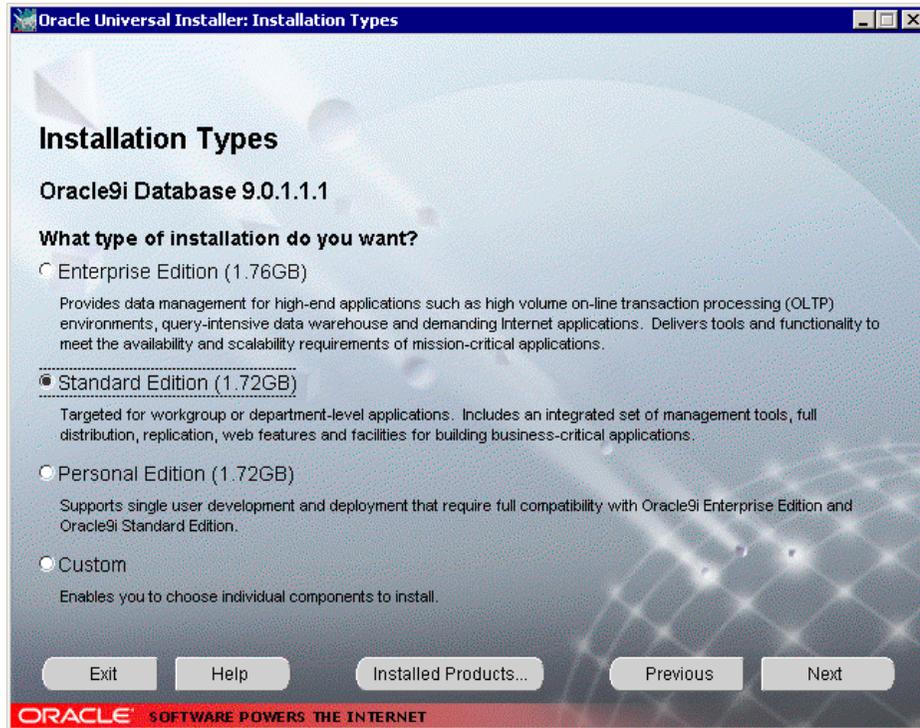


Figure 6-6 Installation types window

Review your options in this window. Select the installation type you want, and then click **Next** to continue.

Remember that during our planning we choose a small, general purpose database configuration, for a Standard Edition installation. We are setting up a departmental or small company database, with a limited amount of concurrent users.

Note: These installation types were discussed in section 6.3.2, “Oracle9i installation and database types” on page 196.

The OUI Database Configuration window opens (Figure 6-7).

Figure 6-7 Database configuration window

Review your options in this window. Select the database configuration type, and then click **Next** to continue.

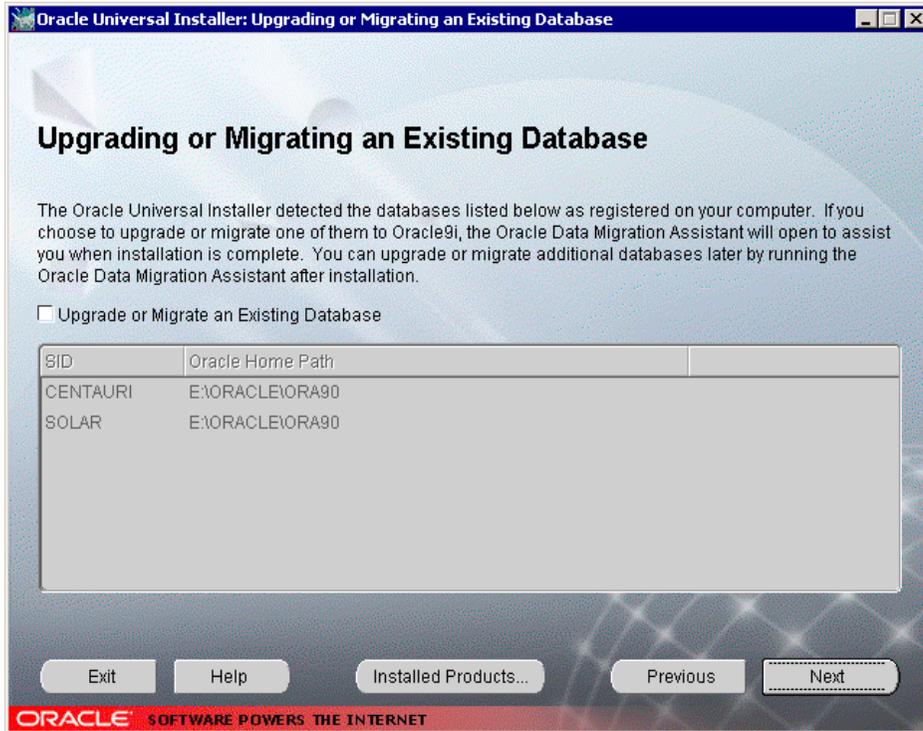


Figure 6-8 Upgrading and migrating an existing database window

Review your options in this window. If you want to upgrade or migrate a database, you should consult the Oracle documentation on upgrading and migrating databases to Oracle9i. Once you know the issues regarding upgrades and migrations, you can proceed.

Note: This book does not include information regarding migrations and upgrades.

In this window do *not* click on any button except for the **Next** button. Make sure that *Upgrade or Migrate an Existing Database* is *not* checked. Click **Next** to continue.

The OUI Database Identification window opens (Figure 6-9).

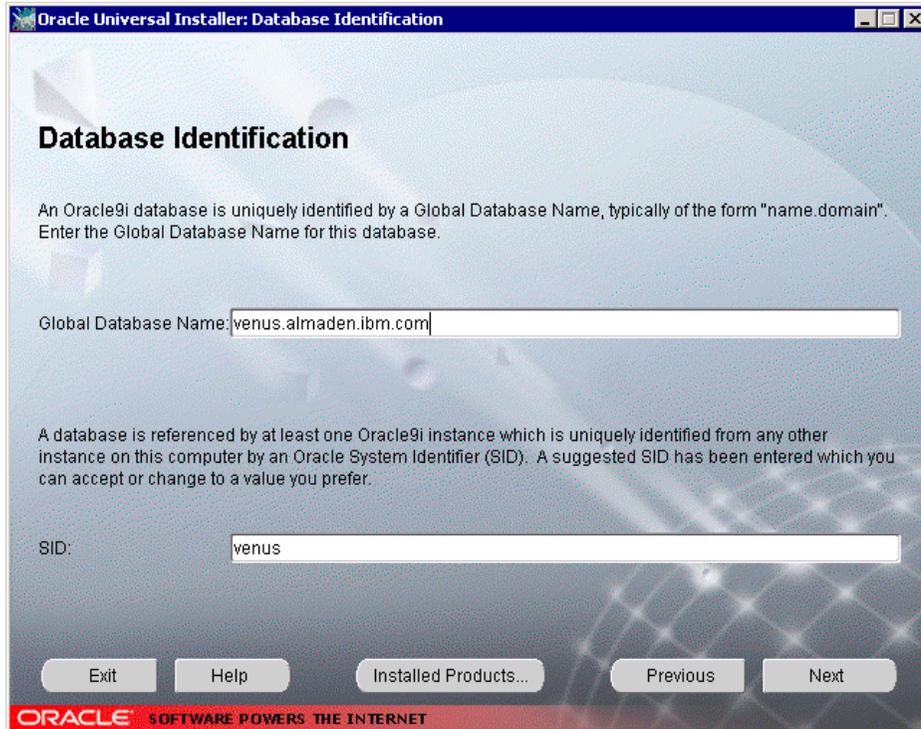


Figure 6-9 Database identification window

Global database domain is the combination of the database name and the network domain where the database exists. This combination is a unique name that will identify the database in the network, or as a stand alone database system.

For example, in Figure 6-9, the database name is 'venus' and the domain is 'almaden.ibm.com'. Therefore, the global database name is 'venus.almaden.ibm.com'.

The Oracle System Identifier, SID, is another name for the database name. The database name should always be unique. We chose 'venus' as database name. If a database name is *not* showing in this window, enter your database name and global database domain, and then click **Next** to continue.

The OUI Database File Location window opens (Figure 6-10).

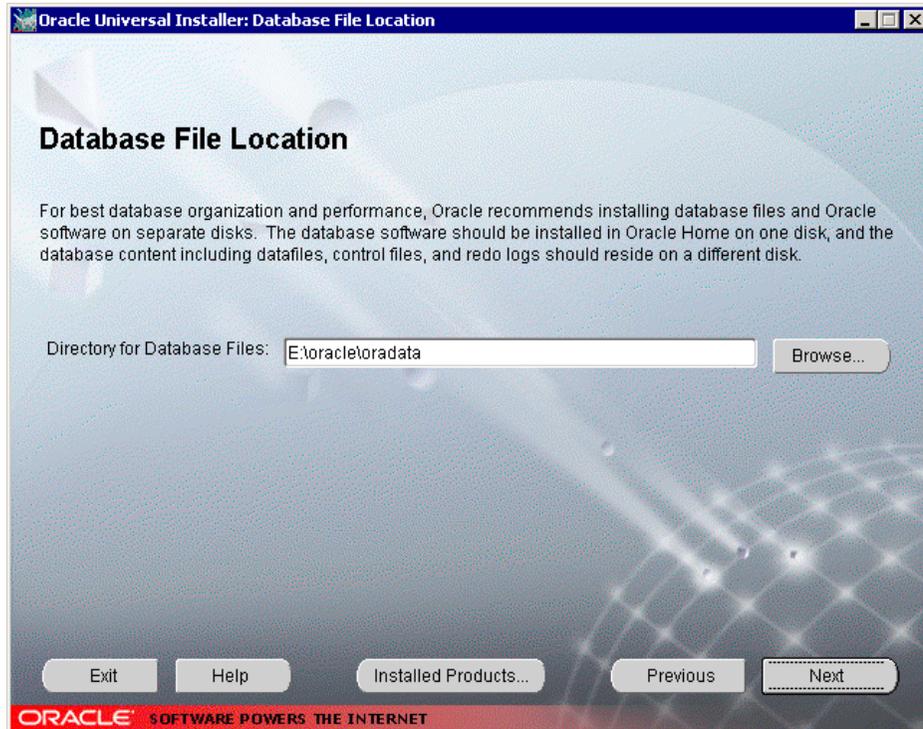


Figure 6-10 Database file location window

The directory for database files defaults to the standard Oracle directory name and directory structure. We chose the default value in this window. The installer will create the directory (make sure there is enough space for database files, that is, 1 GB for the seed database).

Tip: You could use the browse button to search other locations, or enter your database directory.

To continue, click **Next**.

The OUI Database Character Set window opens (Figure 6-11).

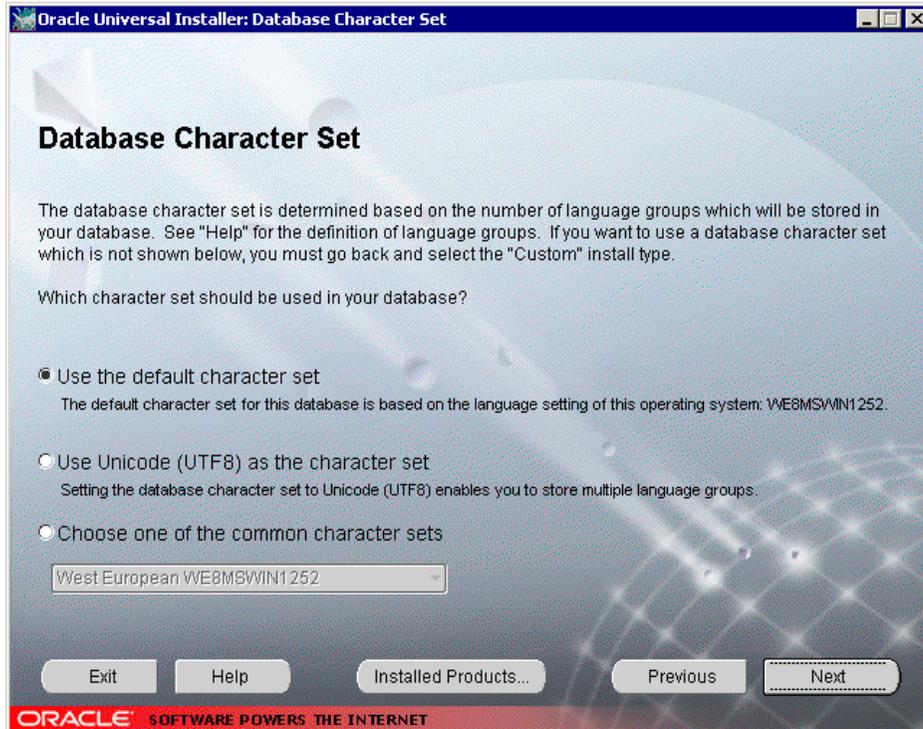


Figure 6-11 Database character set window

Review your options in this window. Select the character set you want, and then click **Next** to continue.

Note: We selected the default character set. Please review the Oracle9i documentation on character set options before selecting another option.

The OUI Summary window opens (Figure 6-12).

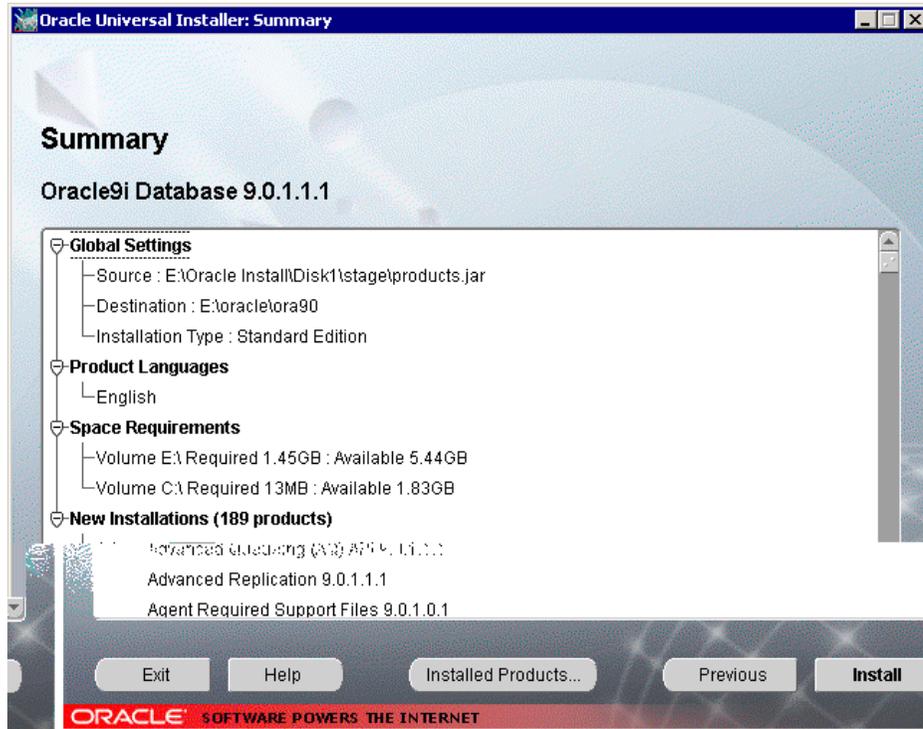


Figure 6-12 Summary window

The summary in Figure 6-12 provides information on all Oracle components to be installed, as well as all the installation options chosen so far.

To continue, click **Next**.

The OUI Install window opens (Figure 6-13).

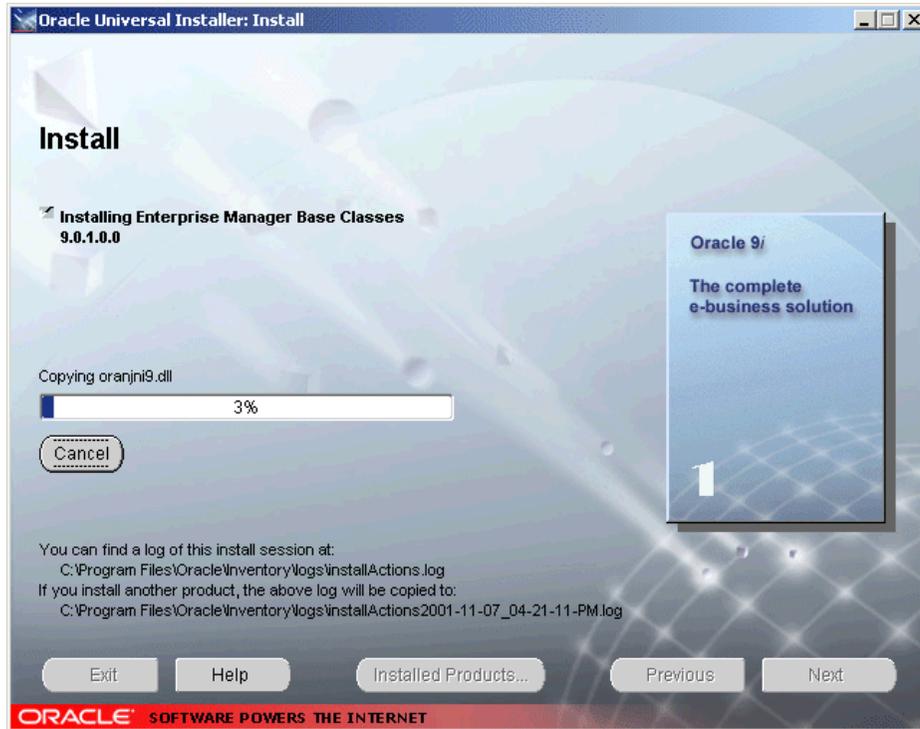


Figure 6-13 Install window

This installation process will take some time to complete. The Oracle9i software for the instance, database and client are being installed. Notice that the rectangular image on the center right side of screen will increase in number, from 1 up to 7, while the software installs.

When it finishes, click **Next** to continue.

The OUI Configuration Tools window opens (Figure 6-14).

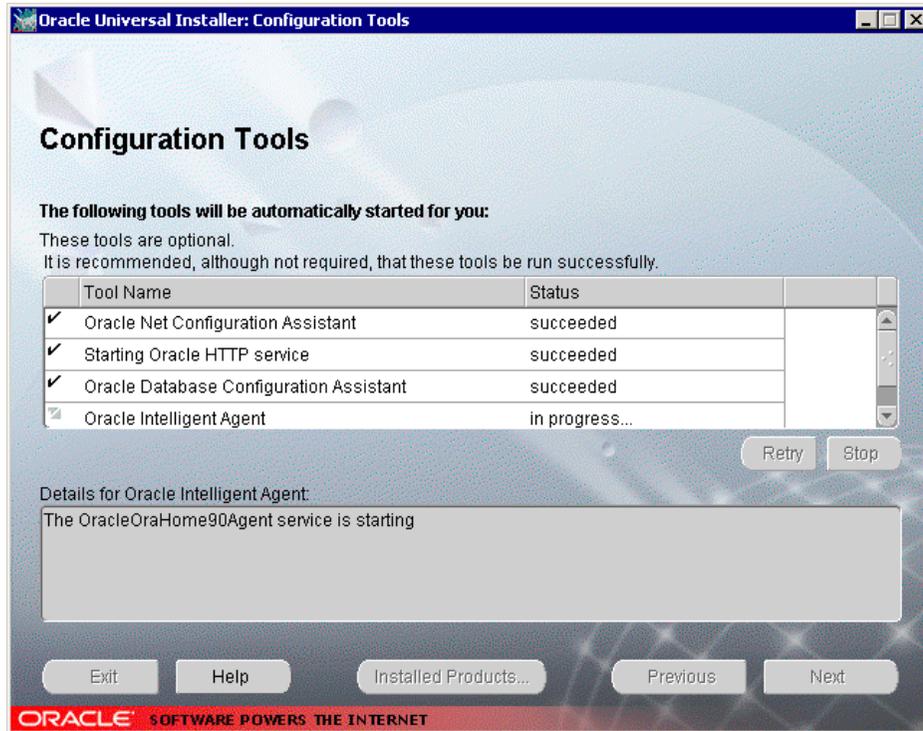


Figure 6-14 Configuration tools window

Once the Oracle installation has completed, the installer automatically will start configuration tools, which creates and configures the database and configures and starts the Oracle client.

Configuration tools will run top to bottom, and as they complete successfully, a permanent check mark will appear:

- ▶ Oracle net assistance
- ▶ Starting Oracle HTTP service
- ▶ Oracle database configuration assistance
- ▶ Oracle intelligent agent

Notice the status column on the window indicates success, failure, or in progress. Also note the result area across the bottom of the screen, where you can see details on the execution of a particular tool.

Tip: Do not open any other window on top of this window while the configuration tools are running. The window may not refresh back to the original screen. You will not know when the configuration completes or if you encountered errors. It will appear as if the configuration tools hang.

While the OUI Configuration Tools is running, the Oracle Database Configuration Assistant window opens (Figure 6-15).

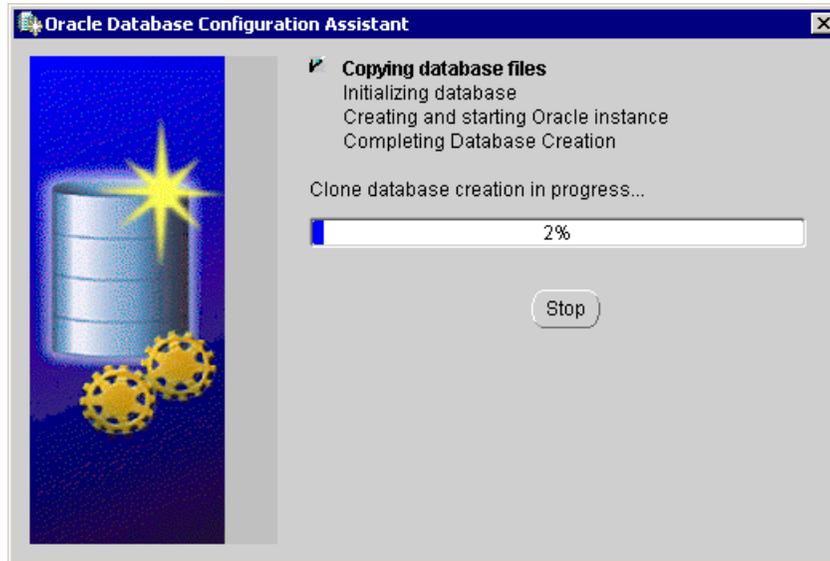


Figure 6-15 Configuration assistant window

The Oracle Database Configuration Assistant (ODCA) begins creating the database. This window does not require user interaction (unless you want to stop this tool). This tool accomplishes four tasks:

- ▶ Copying database files
- ▶ Initializing database
- ▶ Creating and stating the Oracle instance
- ▶ Completing database creation

Notice that as the steps are completed a permanent check mark is displayed.

Another ODCA window opens (Figure 6-16). This window provides database information.

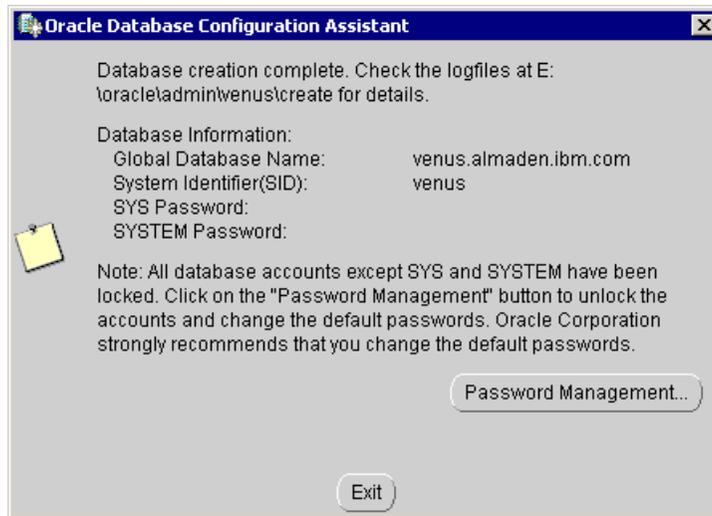


Figure 6-16 ODCA database information window

This window provides database information, and allows password changes on newly created accounts. Some important details to know are:

1. The SYS and SYSTEM password were removed on this window. Your own database SYS and SYSTEM password will be displayed. Take note of them.
2. Other database accounts are part of the seed database, 'venus'. Except for SYS and SYSTEM accounts, all other database accounts are locked.
3. To unlock accounts, click on the Password Management button. The Password Management window opens (Figure 6-17).
4. To continue and avoid unlocking passwords, click **Exit**.

The OUI End of Installation window opens (Figure 6-18).

Note: For security reasons, this Figure 6-16 was modified to remove the SYS and SYSTEM password displayed.

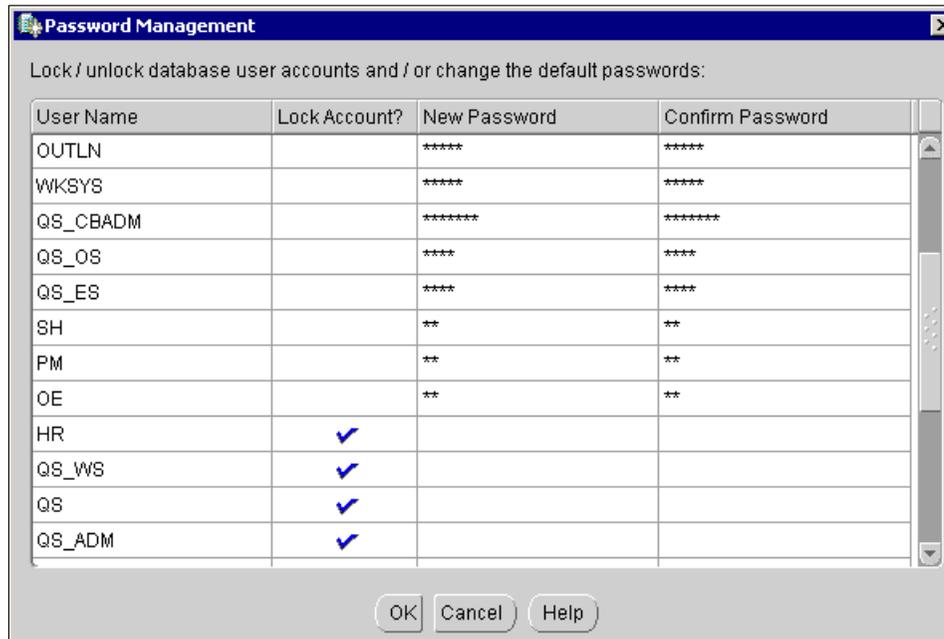


Figure 6-17 Password management window

In this window you are allowed to unlock accounts.

- ▶ Locked accounts have a blue check mark on the lock account column.
- ▶ Unlocked accounts do not have a blue check mark and have '*****' on the new password and confirm password columns. To unlock accounts:
 - Click the blue check mark for a particular account (blue check mark disappears).
 - Enter the same password in new password and confirm password columns. Enter different passwords per user account name.
 - When finished unlocking and entering new passwords, click **OK**.

Important: If you do not unlock accounts at this time, and you need to use them later, you will have to unlock them manually using SQLplus.

Figure 6-16 will appear again, click **Exit**. Figure 6-14 on page 211 will displayed again until all configuration tools finish.

Once the configuration tools complete successfully, the OUI End of Installation window opens (Figure 6-18).

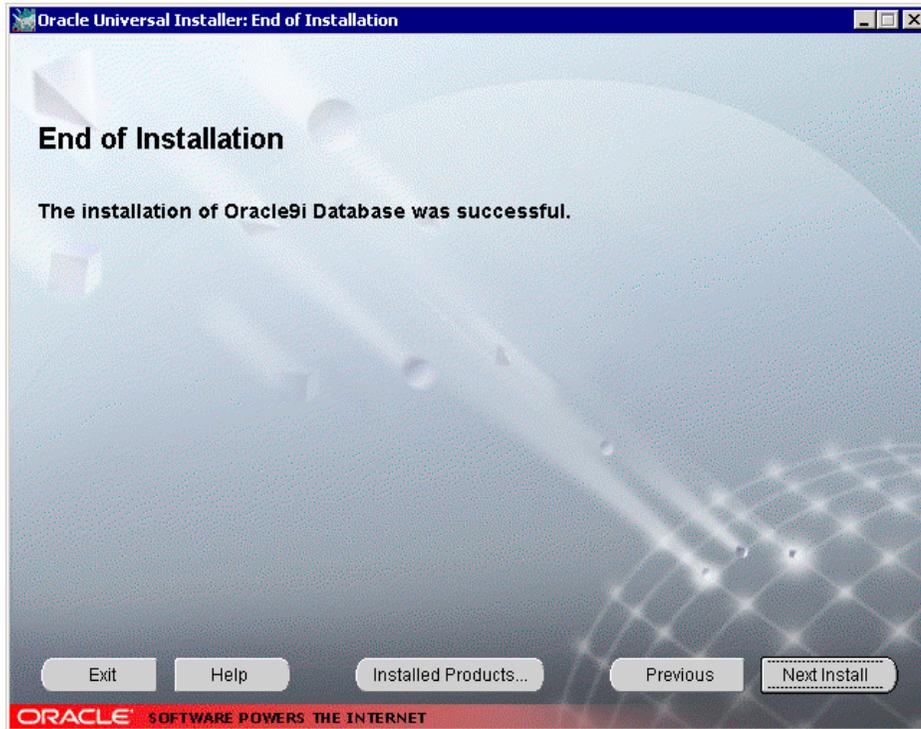


Figure 6-18 End of installation window

If you are satisfied with installation and have no need to install other Oracle9i products, click **Exit**.

The Exit confirmation window opens (Figure 6-19).

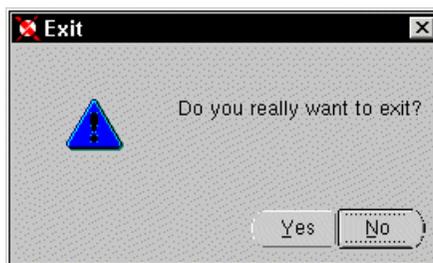


Figure 6-19 Linux OUI exit window

The exit window will open to confirm the exit from installation software. Click **Yes**, otherwise click No, if you are not finished.

Congratulations! You have installed and configured the Oracle instance, client, and database. Your database is ready on Windows 2000!

Tip: If you need to deinstall the Oracle9i products, follow the command menu: **Start>Control Panel> Administrative Tools>Services** and stop any Oracle process running. On the OUI End of Installation window, click the Installed Products button. A window with a list of all installed products will display. Notice the buttons at that bottom of the window. One button is for removing installed products. Click the button and wait until products are deinstalled. Before you install again, make sure you remove directories created by Oracle for the previous installation.

Oracle post-installation on Windows 2000

Once the installation has completed, here are the steps to make sure the Oracle instance, client, and database are running properly:

1. Verify that the Oracle processes are running properly. On your desktop go to **Start->Control Panel-> Administrative Tools->Services**

You will see the Services window (Figure 6-20). Look for the Oracle processes.

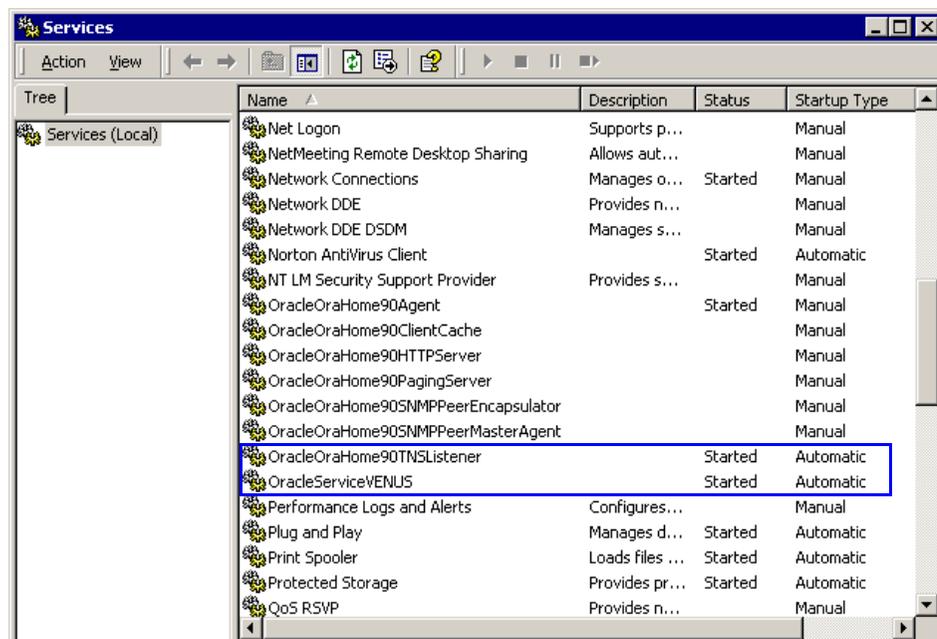


Figure 6-20 Services processes window

Ensure that Oracle processes are present, some processes have different startup types than others (manual or automatic), and some processes have a 'Started' status. For your database to function properly, ensure that the database and the listener have a 'Started' status.

- OracleOraHome90TNSListener — listener process is started.
- OracleServiceVenus — database is started and open.

Notes:

- ▶ More processes could be present, if you have chosen a different installation and database configuration type.
- ▶ For more information about the processes in Figure 6-20, consult the Oracle documentation, or visit the Oracle Web site (<http://www.oracle.com>).

2. Verify that you can access the database and that the database is available on your desktop. Select on your desktop:

Start->Programs->Oracle - OraHome->Application Development->SQL Plus

A log on window will display (Figure 6-21).

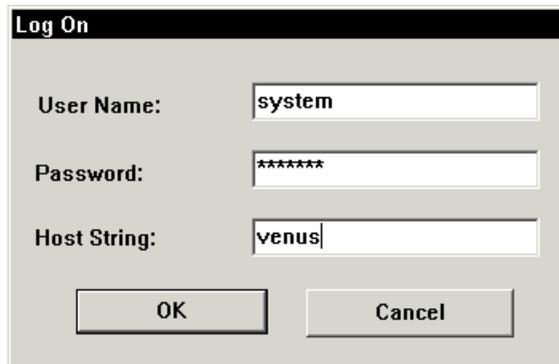


Figure 6-21 SQL Plus log on window

Enter the following information:

- *User name* System
- *Password* The password was given to you in Figure 6-16 on page 213, during the installation.
- *Host string* is the name of your database. Our example database name was 'venus' (see Figure 6-16 on page 213).

Click **OK**, and the Oracle SQL*Plus window opens (Figure 6-22).

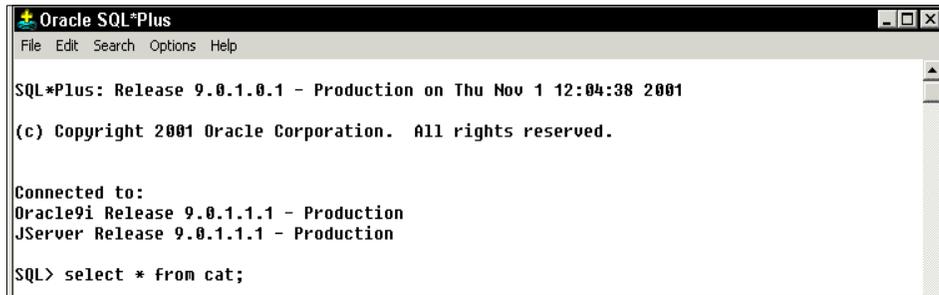


Figure 6-22 Oracle SQL*Plus window

To test the database, enter this SQL statement at SQL prompt:

select * from cat;

The statement will return a result set, that looks like Figure 6-23.

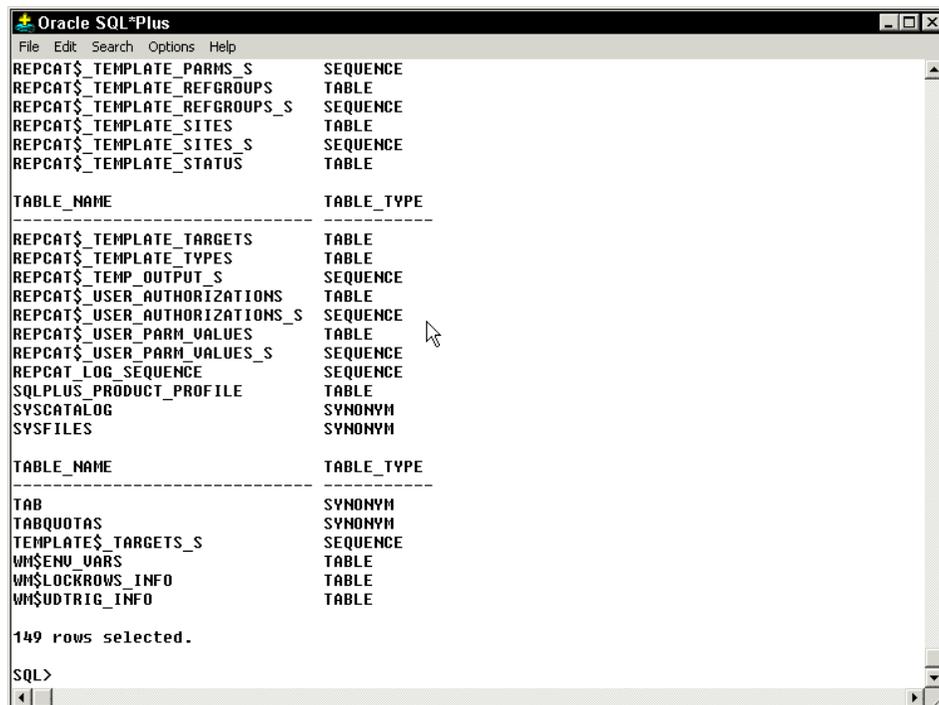


Figure 6-23 Oracle SQL*Plus result set window

You should get a result set like the one in Figure 6-23. If you did, then run the utlrp.sql script in Step 3, below.

If you encounter an error, you will have to investigate and resolve the issue.

Tip: The following step is *only* for an Oracle9i installation on Window 2000.

3. Run the utlrp.sql script, at the SQL prompt:

```
SQL> @oracle_base\ oracle_home\rdbms\admin\utlrp.sql
```

Tip: The utlrp.sql script is used to recompile all existing invalid PL/SQL modules in a database. See the fix 9.0.1X Patch set at:

<http://metalink.oracle.com>

Then, search for the doc id:note:156751.1.

4. Check if the database is in *archive log mode*:

- a. Start SQL*Plus > **sqlplus /nolog**

- b. Connect to the Oracle9i database with the SYSDBA username:

```
SQL> connect / as sysdba
```

- c. If the database is open, shut it down:

```
SQL> shutdown immediate
```

- d. Mount the database:

```
SQL> startup mount
```

- e. Enter the following command:

```
SQL> archive log list
```

The following output indicates the database is not in archive mode:

<i>Database log mode</i>	No Archive Mode
<i>Automatic archival</i>	Disabled
<i>Archive destination</i>	%RDBMS%\
<i>Oldest online log sequence</i>	34
<i>Current log sequence</i>	37

5. Change the database archive mode to *archive log*:

- a. After following the connect steps, and starting database mount command, then enter: **SQL> alter database archive log;**

- b. Enter the following command to verify the change:

```
SQL> archive log list
```

The following output indicates the database is now in archive mode:

<i>Database log mode</i>	Archive Mode
--------------------------	--------------

<i>Automatic archival</i>	Disabled
<i>Archive destination</i>	%RDBMS%\
<i>Oldest online log sequence</i>	34
<i>Current log sequence</i>	37

c. Open the database with:

```
SQL> alter database open;
```

6. The following initialization parameters need to be set within init.ora file:

a. Set *log_archive_format* format to:

```
log_archive_format = log archive format
```

b. The *log_archive_dest_n* parameters have null string values when *log_archive_dest* is set to an operating system-specific default value on instance startup, and is not explicitly defined by user. Set the log archive destination:

```
log_archive_dest_n = filespec
```

Where, filespec is the archive destination.

c. The *log_archive_start* indicates whether archiving should be automatic or manual when the instance starts up. it should be set to true:

```
log_archive_start = true
```

The three initialization parameters look like the following statements within the init.ora file:

```
log_archive_start = true
```

```
log_archive_dest_1 = %oracle_home%\database\archive
```

```
log_archive_dest_2 = %oracle_home%\database\archive2
```

```
log_archive_format = "%oracle_sid%t%ts%s.arc"
```

If all post-installation steps were successful, you are ready to measure performance on this database system.

However, to compare performance, we need another system — Linux.

6.4.2 Linux installation and configuration

Linux Red Hat Version 7.2 will be used for this installation. Please keep in mind that not all of the UNIX options are available on Linux.

Oracle installations and configurations require a minimum set of hardware and software levels. Make sure you comply with the installation requirements before proceeding. Ensure that you also fulfill the pre-installation requirements before continuing with the Oracle9i installation and configuration.

- ▶ “Oracle pre-installation on Linux” on page 221
- ▶ “Oracle installation on Linux” on page 223
- ▶ “Oracle post-installation on Linux” on page 241

Oracle pre-installation on Linux

1. Create the Oracle database administrator group, OSDBA.
2. Create the Oracle account called ‘oracle’. This account is required to be a member of the OSDBA.
3. Create mount points for Oracle software and database. The database may require more than one mount point. To reduce disk I/O you will want to spread the datafiles across several disks. Another option is to mirror or strip disks.
 - Oracle software standard file system:
 - Oracle home - /u01/app/oracle/product/9.0.1
 - Oracle base - /u01/app/oracle
 - Oracle database file system:
 - /oracle/oradata
4. Add the following Linux kernel parameters with values, to the /proc/sys/kernel file system:

- set shmsys:shminfo_shmmax=4294967295
- set shmsys:shminfo_shmmin=1
- set shmsys:shminfo_shmmni=100
- set semsys:seminfo_semmns=1000
- set semsys:seminfo_semmni=100
- set semsys:seminfo_semmsl=250

The following parameters are not supported by Linux Red Hat 7.2:

- shmsys:shminfo_semvmx
- shmsys:shminfo_shmmin
- shmsys:shminfo_shmseg

5. If you added or changed any of the parameters in Step 4, above, reboot the system for the changes to take effect.
6. Set the environment variables in .login or .profile for the Oracle account:

Oracle Environment:

- ORACLE_HOME=/u01/app/oracle/product/9.0.1; export ORACLE_HOME
- ORACLE_BASE=/u01/app/oracle; export ORACLE_BASE
- ORACLE_SID=marte; export ORACLE_SID
- ORACLE_TERM=xterm; export ORACLE_TERM

- TNS_ADMIN=/export/home/oracle/config/9.0.1; export TNS_ADMIN
 - ORA_NLS33=\$ORACLE_HOME/ocommon/nls/admin/data
 - export ORA_NLS33
 - LD_LIBRARY_PATH=\$ORACLE_HOME/lib:/lib:/usr/lib:/usr/local/lib:
 - export LD_LIBRARY_PATH
- # Set shell search paths:**
- PATH=/bin:/usr/bin:/usr/sbin:/etc:/usr/ccs/bin:/bin:
 - PATH=\$PATH:\$ORACLE_HOME/bin; export PATH
- # CLASSPATH must include the following JRE locations:**
- CLASSPATH=\$ORACLE_HOME/JRE:\$ORACLE_HOME/jlib:\$ORACLE_HOME/rdbms/jlib
 - CLASSPATH=\$CLASSPATH:\$ORACLE_HOME/network/jlib
7. Oracle user's .profile and .login files should have values for:
- Umask 022
 - Display variable for the Oracle Universal Installer to display:
DISPLAY=9.1.39.51:0.0
- Where, 9.1.39.51 is desktop IP address from where you will install.
8. Temporary disk space of 400 MB is required during the Oracle installation. Make sure the required space is available.
9. Before attempting to run the Oracle Universal Installer, verify that you can successfully run the following command:
- ```
% /usr/openwin/bin/xclock
```
10. You will have to make sure that file ownership and permissions are given for all of the Oracle file systems created. For example:
- For /u01/app/oracle/product/9.0.1, all file system levels are owned by the Oracle user and they belong to the OSDBA group. Changing ownership on file systems will allow the Oracle user to perform tasks required for installation and proper function of the database. To change ownership of a file structure issue the command:  
**chown oracle:osdba u01**
  - Other permission is required for /usr/local/include. For permission, enter the commands:
    - **chmod 755 usr**
    - **chmod 755 local**
    - **chmod 755 include**

**Tip:** The oratab file in Linux is located in the /etc file structure.

Once the pre-installations tasks are completed, you are ready to install Oracle9i.

## Oracle installation on Linux

The Oracle9i installation is performed using the Oracle Universal Installer (OUI) version 2.0.1. There are several components installed with the basic Oracle software, such as replication, queueing, clustering, Oracle Enterprise Manager, Oracle Net Assistant, to name a few. Some of these components are used, and some are not. However these components are part of the installation process unless you specify otherwise.

For Oracle9i installation on Linux, you downloaded three gzip files, from Oracle Web site or have three Oracle9i CD-ROMs. Follow the steps to begin the installation:

1. Log on as the Oracle user. If you log on as the root user, you will get an error message.
2. Create four directories:
  - Install.
  - Under this directory created three directories: Disk1, Disk2, and Disk3.
3. If you downloaded the files, extract files from gzip files onto Disk1, Disk2, and Disk3. One gzip file per directory — follow these steps:
  - Copy gzip 1, 2, 3 files into directories Disk1, Disk2, Disk3 — one file per directory.
  - Issue the following commands for the three files within the Disk directories:  

```
gunzip linux9i_Disk1.cpio.gz
cpio -idmv < linux9i_Disk1.cpio
```
4. Or, if you have the CD-ROMs, at command prompt run **Disk1\setup.exe**.

The Linux Oracle Universal Installer (OUI) Welcome window will display (Figure 6-24).

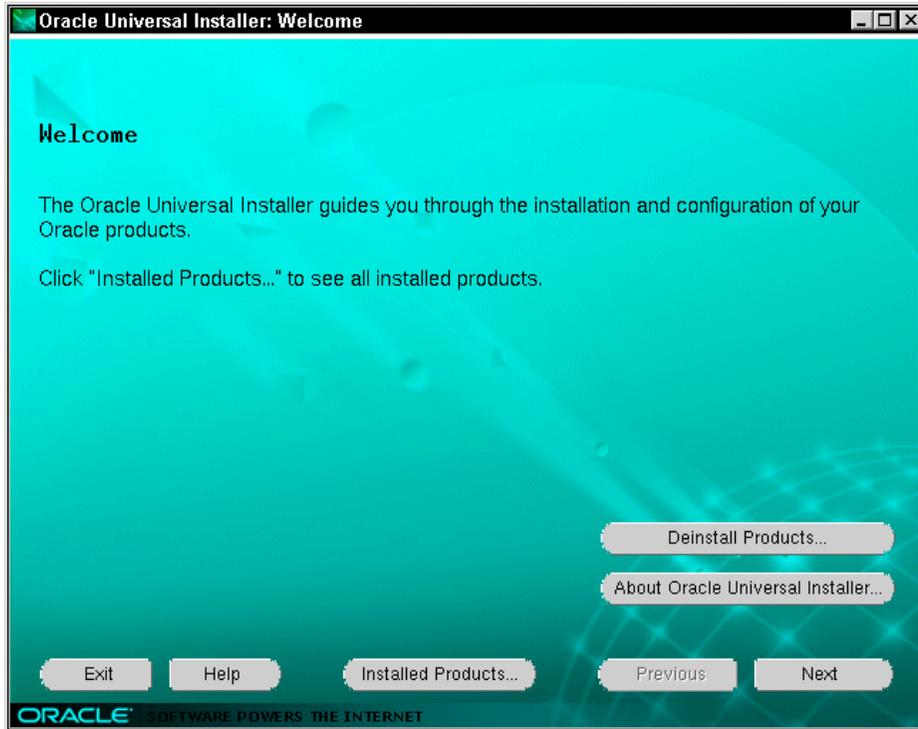


Figure 6-24 Linux OUI welcome window

To start the installation, click **Next**, which is always located at the bottom right corner.

The OUI File Locations window opens (Figure 6-25).

**Note:** If your system has Oracle products installed, you have the option to view a list of installed products, by clicking the Installed Products button. Or, you can deinstall products, by clicking the Deinstall Products button.

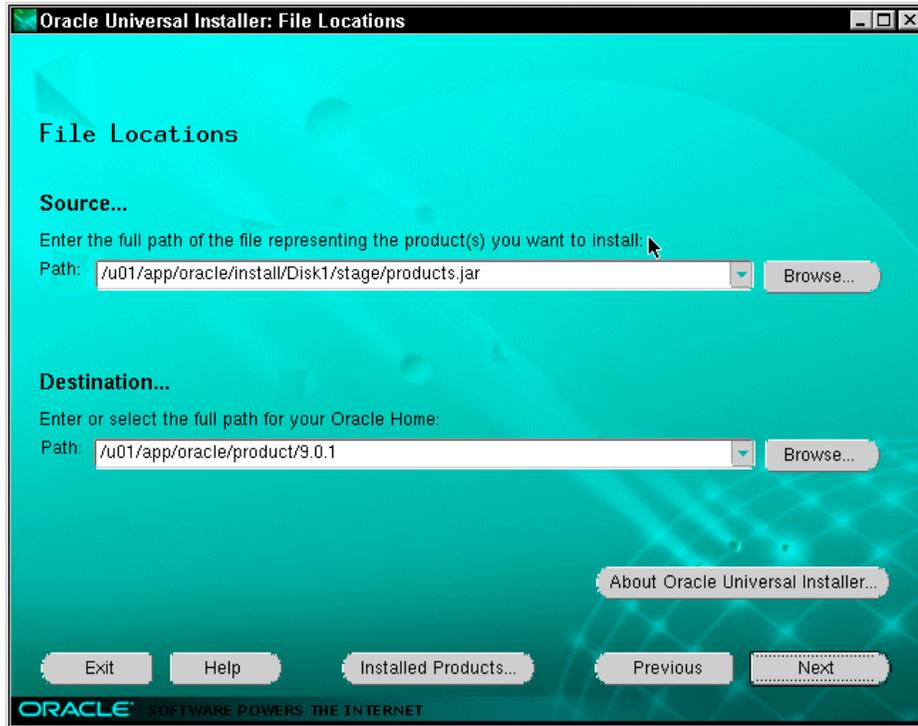


Figure 6-25 Linux OUI file locations window

The options for this window are:

- ▶ Standard default values (that are displayed) *or* choose different paths:
  - The source path is one of the directories created for software installation in Step 2. on page 223, /u01/app/oracle/install/Disk1, where installation files reside. Source path is also the full path to where you mounted from or extracted the Oracle files (from CD-ROM or gzip files). Most likely you will not have to change this option.
  - Or, you can choose your own destination names and paths. The destination is the full path where the Oracle software will be installed and run from. The current settings and defaults are Oracle standard destination full paths and names for the Oracle home. You could choose a different path and name to fit your own company standards and mount points. We choose OUI default settings.

**Tip:** You have the option to browse for both source and destination paths.

To continue, click **Next**. The OUI Available Products window opens (Figure 6-26).



Figure 6-26 Linux OUI available products window

The Product Language button is available to you. You could continue and allow defaults, or click on this option to view and/or change defaults.

Review your options in this window. Select the product you want to install, and then click **Next** to continue.

**Note:** We selected the Oracle9i Database 9.0.1.0.0 product. This option will include the Oracle9i software, the client installation and configuration, and database creation and configuration.

The OUI Installation Types window opens (Figure 6-27).

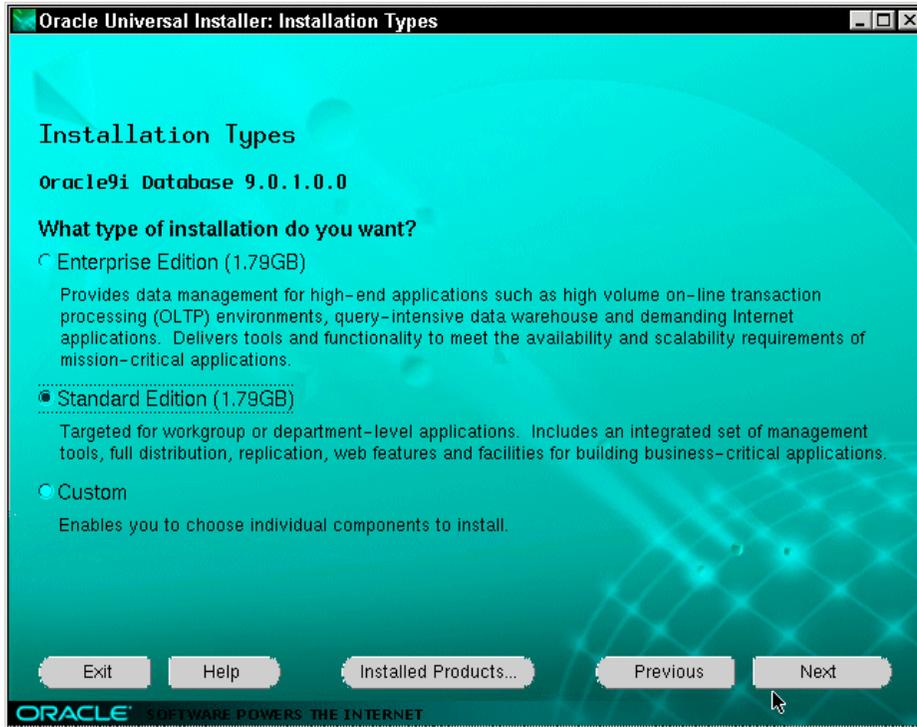


Figure 6-27 Linux OUI installation types window

Review your options in this window. Select the installation type you want, and then click **Next** to continue.

**Note:** We selected the Standard Edition type. These installation types were discussed in section 6.3.2, “Oracle9i installation and database types” on page 196.

The Database Configuration window opens (Figure 6-28).

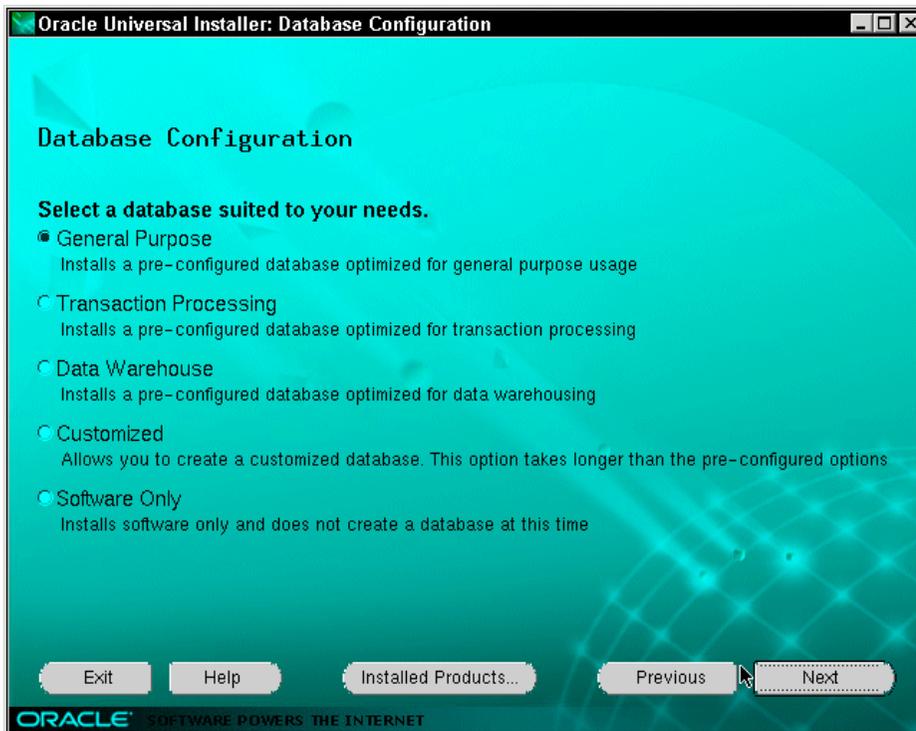


Figure 6-28 Linux OUI database configuration window

Review your options in this window. Select the database configuration type, and then click **Next** to continue.

**Note:** We selected the General Purpose configuration. We discussed the database configuration types in section 6.3.2, “Oracle9i installation and database types” on page 196.

If the OUI Upgrading or Migrating an Existing Database window opens, it is because the OUI detects previously installed databases. See Figure 6-8 on page 205 for details.

The Privileged Operating System Groups window opens (Figure 6-29).



Figure 6-29 Linux OUI privileged operating system groups window

Here are a few things to notice:

- ▶ The group(s) shown for database administrator and operator group are the same.
- ▶ The group shown was created in Step 1. on page 221.
- ▶ If you did not create the group, either this window will not display or no default group will be given.

**Be careful:** If you enter a group, that was not specifically created for Oracle, you may get unpredictable results.

Allow or keep the default group or enter your own group, and then click **Next** to continue.

The Database Identification window opens (Figure 6-30).

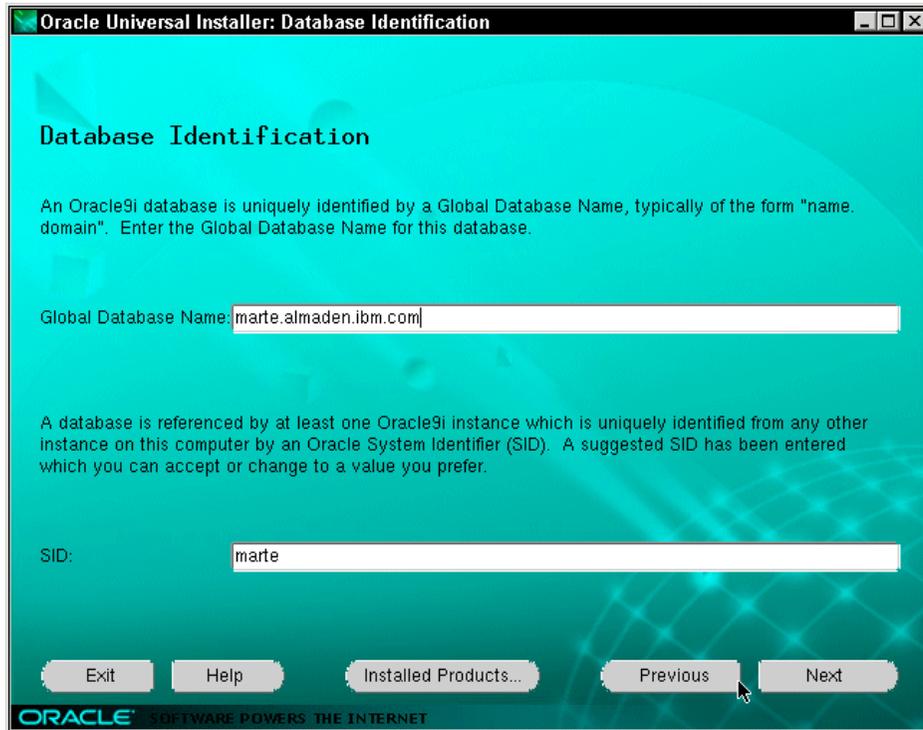


Figure 6-30 Linux OUI database identification window

The global database domain is the combination of the database name and the network domain where the database exists. This combination is a unique name that will identify the database in the network, or as a stand alone database system.

For example, in Figure 6-30, the database name is 'marte' and the domain is 'almaden.ibm.com'. Therefore, the global database name is 'marte.almaden.ibm.com'.

The Oracle System Identifier, SID, is another name for the database name. The database name should always be unique. We chose 'marte' as database name. If a database name is *not* showing in this window, enter your database name and global database domain, and then click **Next** to continue.

**Note:** The database file location window is not displayed. The seed database will be created, by default, under the \$ORACLE\_BASE file structure; see /admin and /oradata for database administration and datafiles, respectively.

The OUI Database Character Set window opens (Figure 6-31).

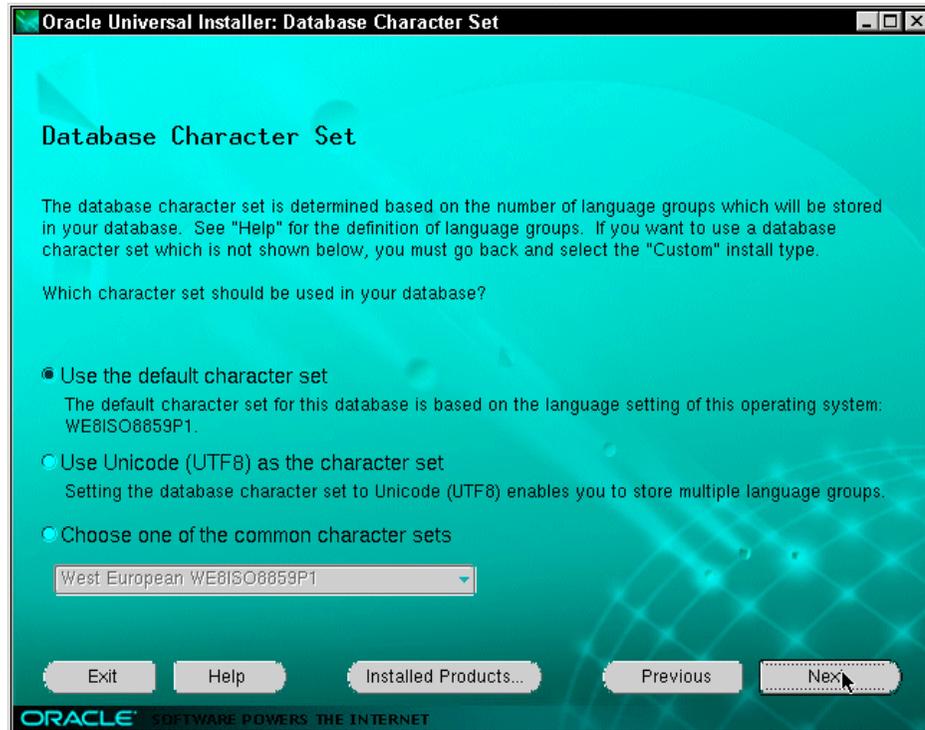


Figure 6-31 Linux database character set window

Review your options in this window. Select the character set you want, and then click **Next** to continue.

**Note:** We selected the default character set. Please review the Oracle9i documentation on character set options before selecting another option.

The OUI Choose JDK Home Directory window opens (Figure 6-32).



Figure 6-32 Linux OUI choose JDK home directory window

If the JDK home entry is not displayed, you could browse for it or enter the location. This product, JDK, is part of the software required.

To continue, click **Next**. The OUI Summary window opens (Figure 6-33).

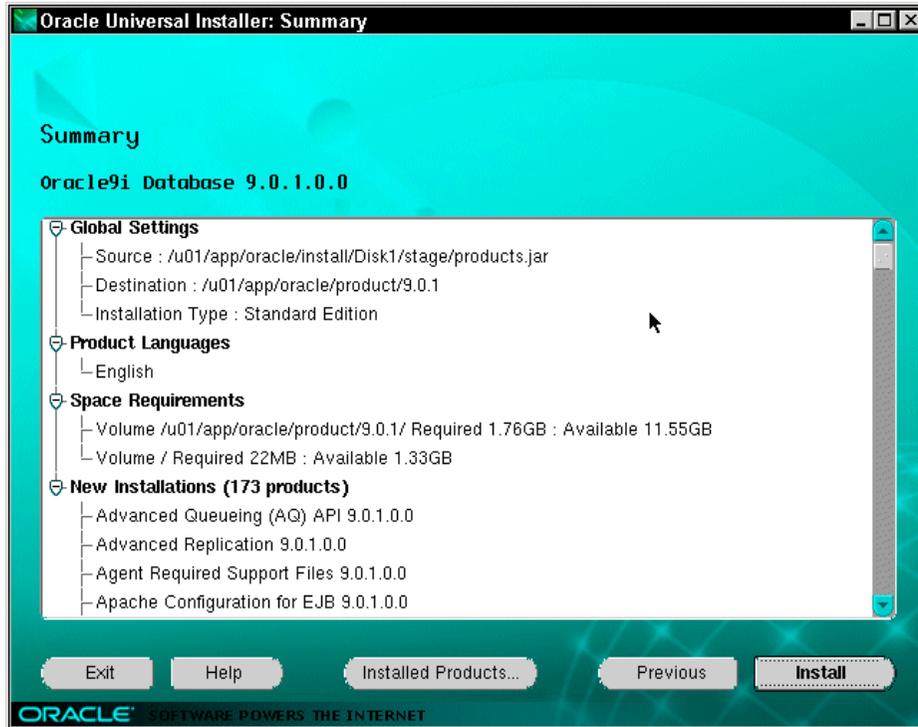


Figure 6-33 Linux OUI summary window

The summary in Figure 6-33 provides information on all Oracle components to be installed, as well as all the installation options chosen so far.

To continue, click **Next**. The OUI Install window opens (Figure 6-34).

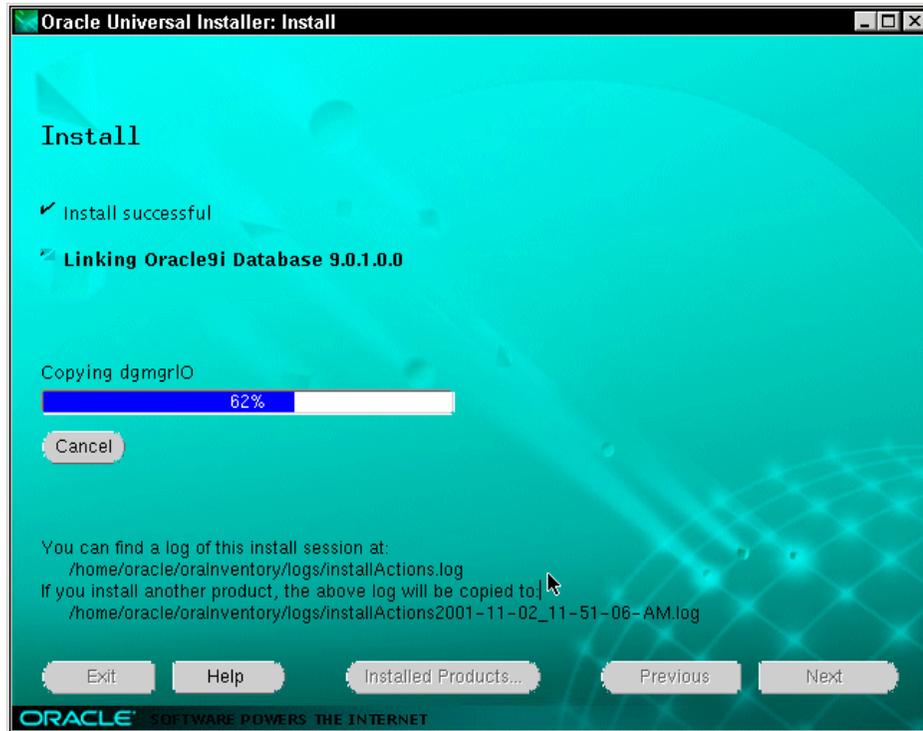


Figure 6-34 Linux OUI install window

This window will show the progress of the Oracle software installation. Notice that at first this process will install software (shown with a solid check mark), then it will link the Oracle binary files.

You may encounter some problems in this process. You could have permission errors (Linux) or script problems (Oracle). Oracle has documented these problems:

- ▶ Installation Fails With Makefile Exception
- ▶ Error invoking target install of makefile  
/u01/app/oracle/product/901/plsql/lib/ins\_plsql.mk

If you encounter these errors, they will appear in an Error window (Figure 6-35).



Figure 6-35 Linux OUI error window

To fix the errors:

- ▶ For *Permission errors*, you may get several of them, until you:
  - Open up the directory to allow access, and give the `/usr/local/include/` non-Oracle directory the correct permissions.
  - In our installation we gave full directory access. After permissions are granted, click the Retry button.
- ▶ For *Oracle script errors*, follow the steps to correct error:
  - a. Open a new console window on the host as the installation user.
  - b. Set up the Oracle environment (that is, `ORACLE_HOME`).
  - c. Make a backup of the `$ORACLE_HOME/genclntsh` script.
  - d. Remove the `${LD_SELF_CONTAINED}` flag from the `ld` command in the 'Create Library' section.
  - e. Run the `genclntsh` script.
  - f. Click the Retry button in the Error window.

The next window, Oracle fixed script, displays part of the `genclntsh` script, where the `${LD_SELF_CONTAINED}` flag was removed. The window also contains the command to run the script, at the bottom (Figure 6-36):

```
./genclntsh
```

```

Command Prompt - telnet 9.138.169
OLIBS="${COLIBS} ${LNLSTRIL}"
OLIBS="${COLIBS} ${LUNLSTRIL}"

#SYSLIBS='cat ${ORACLE_HOME}/lib/sysliblist`" -lao -lm -lc"
SYSLIBS='cat ${ORACLE_HOME}/lib/sysliblist`" -ldl -lm -lc"

Create library
/bin/rm -f ${LIB_DIR}/${CLNT_LIB} ${MAPFILE}
${LD} ${LD_RUNTIME} ${LD_OPT} ${LD_OUT} \
 ${DEF_ON} ${NAUTAB} ${NAETAB} ${NAEDHS} ${SYMS_SQL} \
 ${SYMS_PLSQL} ${SYMS_OCIC} ${SYMS_LDAP} ${SYMS_NETWORK} ${OLIBS} \
 ${SYSLIBS} ${USRLIBS} > ${MAPFILE}

Create ".so" link to library in a non-development environment only
and generate the client static library
if [-z "${SRCHOME}"]
then
 if [-s ${LIB_DIR}/${CLNT_LIB}]
 then
 "gencntsh" 199L, 5230C written
 oracle@wisla bin1$./gencntsh

Created /u01/app/oracle/product/9.0.1/lib/libcintst9.a
oracle@wisla bin1$
oracle@wisla bin1$

```

Figure 6-36 Linux Oracle fixed script window

**Note:** For more information on these errors, search on the Oracle metalink Web site for: [ins\\_plsql.mk](#), [ins\\_precomp.mk](#), or installation fails.

Once errors are fixed, and you have clicked the Retry button, the installation will continue.

The Setup Privileges window opens (Figure 6-37).

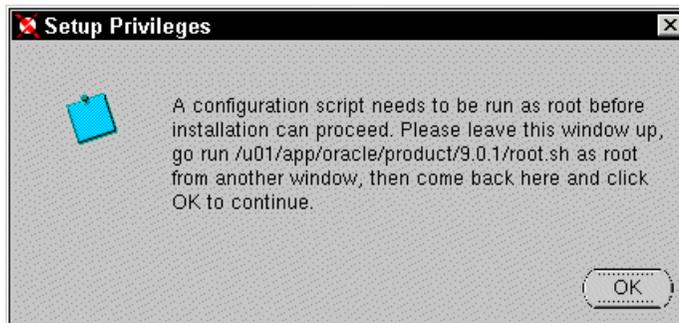


Figure 6-37 Linux OUI setup privileges window

Follow the instructions to run root.sh script, then click **OK** to continue.

The next window, root.sh execution, shows the root.sh execution (Figure 6-38). You may not get the overwrite questions if this is your first installation.

```

Command Prompt - telnet 9.1.38.169

[root@wisla 9.0.11# ./root.sh
Running Oracle? root.sh script...
The following environment variables are set as:
 ORACLE_OWNER= oracle
 ORACLE_HOME= /u01/app/oracle/product/9.0.1

Enter the full pathname of the local bin directory: [/usr/local/bin]:
The file "dbhome" already exists in /usr/local/bin. Overwrite it? (y/n)
[In]: y
 Copying dbhome to /usr/local/bin ...
The file "oraenv" already exists in /usr/local/bin. Overwrite it? (y/n)
[In]: y
 Copying oraenv to /usr/local/bin ...
The file "coraenv" already exists in /usr/local/bin. Overwrite it? (y/n)
[In]: y
 Copying coraenv to /usr/local/bin ...

Adding entry to /etc/oratab file...
Entries will be added to the /etc/oratab file as needed by
Database Configuration Assistant when a database is created
Finished running generic part of root.sh script.
Now product-specific root actions will be performed.
[root@wisla 9.0.11#

```

Figure 6-38 Linux OUI root.sh execution window

The Configuration Tools window opens (Figure 6-39).

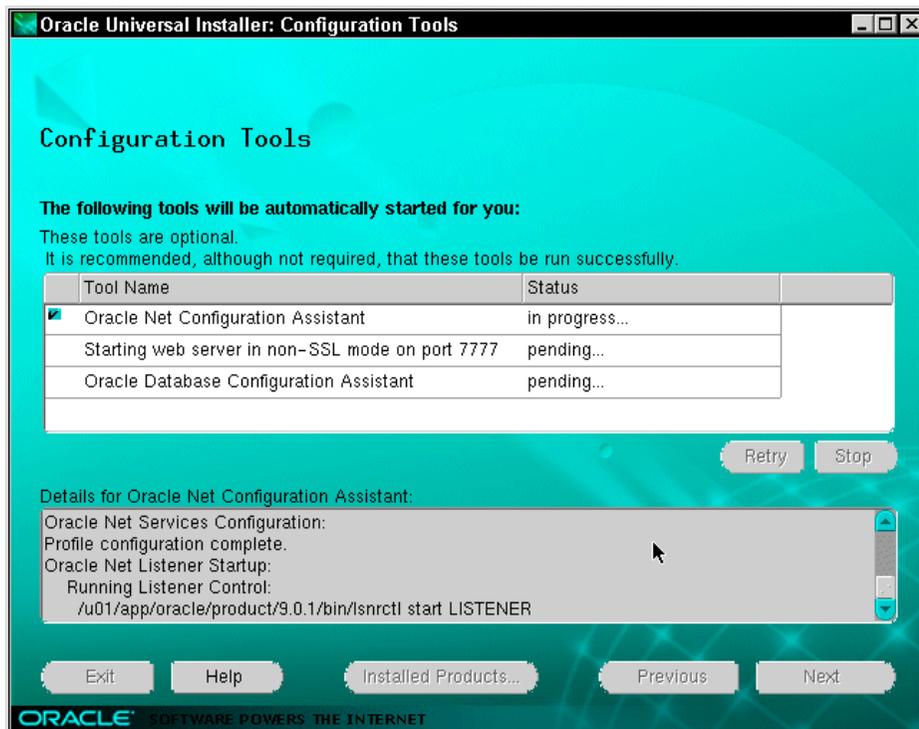


Figure 6-39 Linux OUI configuration tools window

Once Oracle installation has completed, the installer automatically starts the configuration tools, which creates and configures the database, and configures and starts the Oracle client.

The configuration tools will run top to bottom, and as they complete successfully a permanent check mark will appear:

- ▶ Oracle net assistance
- ▶ Starting web server
- ▶ Oracle database configuration assistance

Notice the status column on the window indicates success, failure, or in progress. Also note the result area across the bottom of the screen, where you can see details on execution of a particular tool.

**Tip:** Do not open any other window on top of this window while the configuration tools are running. The window may not refresh back to original screen. You will not know when the configuration completes or if you encountered errors. It will appear as if the configuration tools hang.

While the Oracle database configuration assistance is running, another Oracle window opens (Figure 6-40).

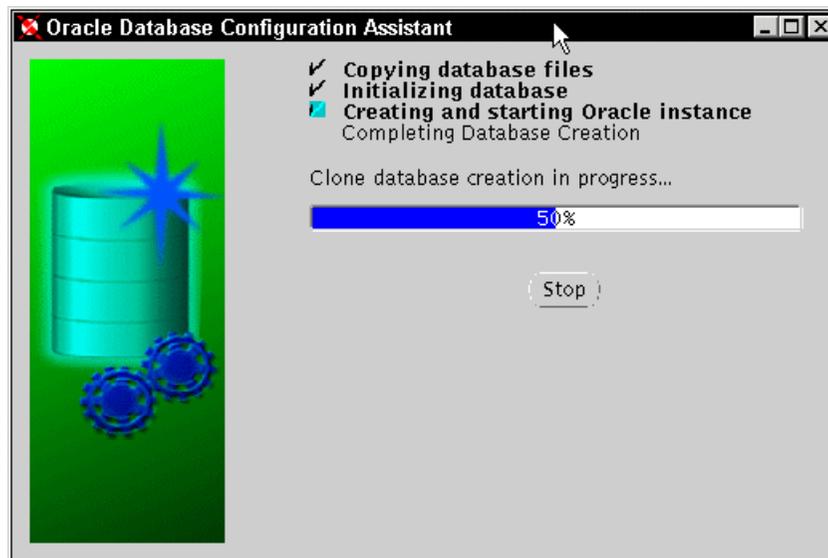


Figure 6-40 Linux database configuration assistant window

The Oracle Database Configuration Assistant (ODCA) begins creating the database. This window does not require user interaction (unless you want to stop this tool). This tool performs four tasks:

- ▶ Copying database files
- ▶ Initializing database
- ▶ Creating and starting the Oracle instance
- ▶ Completing database creation

Notice that as the steps are completed a permanent check mark is displayed.

Another ODCA window opens (Figure 6-41). This window provides database information.

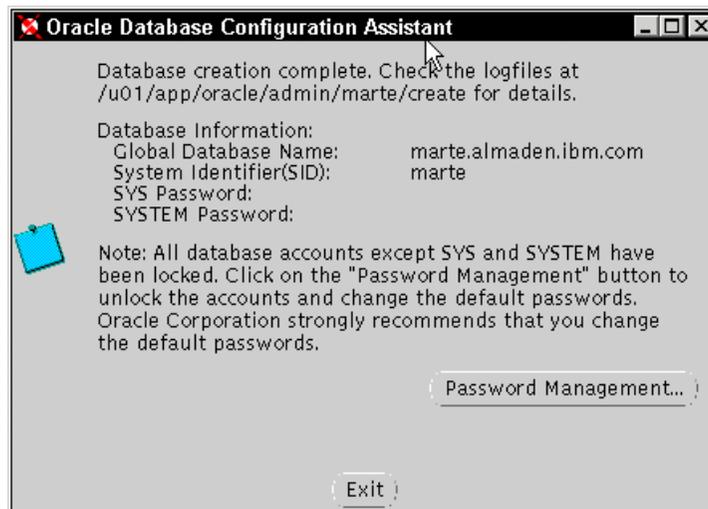


Figure 6-41 Linux ODCA database information window

This window provides database information, and allows password changes on newly created accounts. Some important details to know:

1. The SYS and SYSTEM passwords were removed on this window. Your own database SYS and SYSTEM password will be displayed. Take note of them.
2. Other database accounts are part of the seed database, 'marte'. Except for sys and system accounts, all other database accounts are locked.
3. To unlock the accounts, click the Password Management button. The Password Management window opens.
4. To continue and avoid unlocking passwords, click **Exit**.

The OUI End of Installation window opens (Figure 6-42).

**Note:** For security reasons this window was modified to remove the SYS and SYSTEM password displayed.

If you choose to unlock accounts at this time, see Figure 6-16 on page 213 and follow its details for instructions on how to unlock accounts. The same procedure applies for Linux and Windows 2000 installations.

Figure 6-39 on page 237 will be displayed again until all of the configuration tools finish.

Once the configuration tools complete successfully, the OUI End of Installation window opens (Figure 6-42).

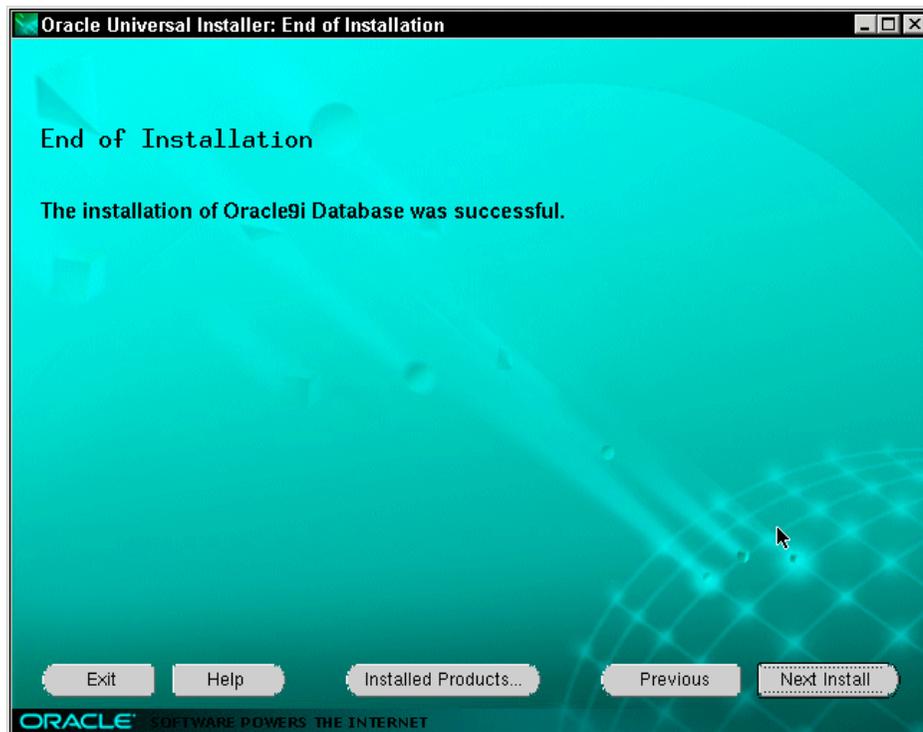


Figure 6-42 Linux OUI end installation window

If you are satisfied with installation and you have no need to install other Oracle9i products, click **Exit**.

The Exit confirmation window opens (Figure 6-43).

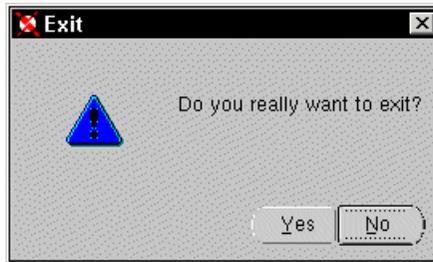


Figure 6-43 Linux OUI exit window

If you are satisfied with installation and you have no need to install other Oracle9i products, click **Yes**.

Congratulations! You have just configured the Oracle instance and client, and created and configured a database. Your database is ready on Linux!

**Tip:** If you need to deinstall the Oracle9i products, make sure that you stop the Oracle processes that are running. On OUI End of Installation window, click the Installed Products button. A window with a list of all installed products will display. Notice the buttons at that bottom of screen. One button is for removing installed products. Click the button and wait until the products are deinstalled. Before you install again, make sure you remove the directories created by Oracle for the previous installation.

## Oracle post-installation on Linux

On a basic installation, only these post-installation requirements follow:

- ▶ For configuring shared server parameters, the shared server requires a higher setting for the *shared\_pool\_size* initialization parameter. If you installed your server with Oracle Universal Installer, then the value of the *shared\_pool\_size* parameter is set automatically by the Oracle Database Configuration Assistant. However, if you created a database manually, increase the *shared\_pool\_size* parameter in the *initsid.ora* file. Typically, you should add 1 KB for each concurrent user.
- ▶ For configuring Oracle Net, the basic configuration of Oracle Net is done by Oracle Net Configuration Assistant, when it is started by OUI during installation. Complete the configuration with the following steps:
  - a. Log in as root and reserve a port for the Oracle listener process, by editing the */etc/services* file.
  - b. Add an entry:

```
listener_name 1521/tcp #Oracle listener process
```

**Tip:** The default port assigned during installation is 1521 for TCP/IP protocol. If you changed to a different port number, make sure that this is reflected in the entry on the etc/services file.

- c. Check the status of the listener following the installation, using this command:

```
$ lsnrctl status listener_name
```

Or,

```
$ lsnrctl status
```

**Tip:** After executing the status command, look for:

```
'Instance database_name, status READY, has # handler(s) for this service...The Command completed successfully'.
```

This will indicate that listener is running properly.

- d. If the listener is not running, start it by using this command:

```
$ lsnrctl start listener_name
```

Or,

```
$ lsnrctl start
```

- e. Start SQL\*Plus to test the connection to the server:

```
$ sqlplus username/password@mart
```

If you can successfully connect to the server with SQL\*Plus, you have established network connectivity over TCP/IP. You could also see the testing performed on the Windows 2000 post-implementation in Step 2. on page 217.

- Check the database archivelog mode; see Step 4. on page 219.

**Note:** For more information regarding post-implementation functions, review Chapter 4 “Post-Installation” in the *Oracle9i Installation Guide Release 1 (9.0.1) for UNIX Systems*.

### 6.4.3 Database creation on iSCSI disks

Here we provide instructions for the database creation on iSCSI for Windows 2000 and Linux, as well as a remote Oracle9i client implementation.

The steps are covered in these sections:

- ▶ “Database implementation on iSCSI disks” on page 243
  - “Windows 2000 iSCSI database” on page 243
  - “Linux iSCSI database” on page 258

For *database implementation on iSCSI disks*, the database files, which contain the user/business data, will reside on iSCSI disks.

The Oracle9i instance *resides locally* on either Windows 2000 or Linux operating systems.

- ▶ “Remote Oracle9i client implementation” on page 261
  - “Install remote Oracle9i client” on page 261
  - “Configure remote database on client” on page 271

*Remote Oracle9i client implementation* is on a Windows 2000 desktop. It is important to point out that “Install remote Oracle9i client” on page 261 includes the configuration of a remote client. Configuration is part of the installation process.

However, if you install the client, and, at a later date, you want to configure a remote database, then follow “Configure remote database on client” on page 271.

The term *remote database* is used to refer to a database located on an iSCSI disk.

## **Database implementation on iSCSI disks**

The only requirement for the database creation on iSCSI disks is the ability to see the iSCSI unit from Windows 2000 or Linux. You have installed the drivers and point to the iSCSI unit on Windows 2000, or mounted a point for the iSCSI unit on Linux. You have access to the iSCSI unit.

To create a remote database, we will use one of the tools that comes with the Oracle database installation tools, the Database Configuration Assistant (DBCA).

### ***Windows 2000 iSCSI database***

To start the DBCA, on Windows 2000 go to:

**Start->Programs->Oracle - OraHome90->Configuration and Migration Tools->Database Configuration Assistant**

The DBCA Welcome window opens (Figure 6-44).



Figure 6-44 DBCA welcome window

To begin the database creation, click **Next**.

The DBCA Step 1 of 9: Operations window opens (Figure 6-45).

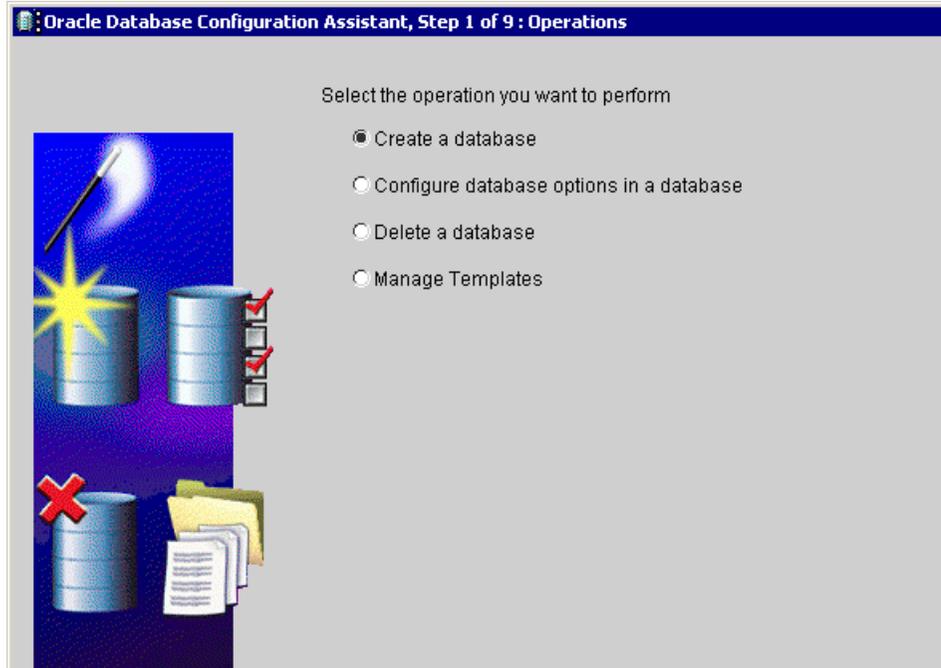


Figure 6-45 DBCA operations window

We are interested in database creation, so we select the *Create a database* option. Click **Next** to continue.

Notice the other options which could be useful for other tasks.

The DBCA Step 2 of 8: Database Templates window opens (Figure 6-46).

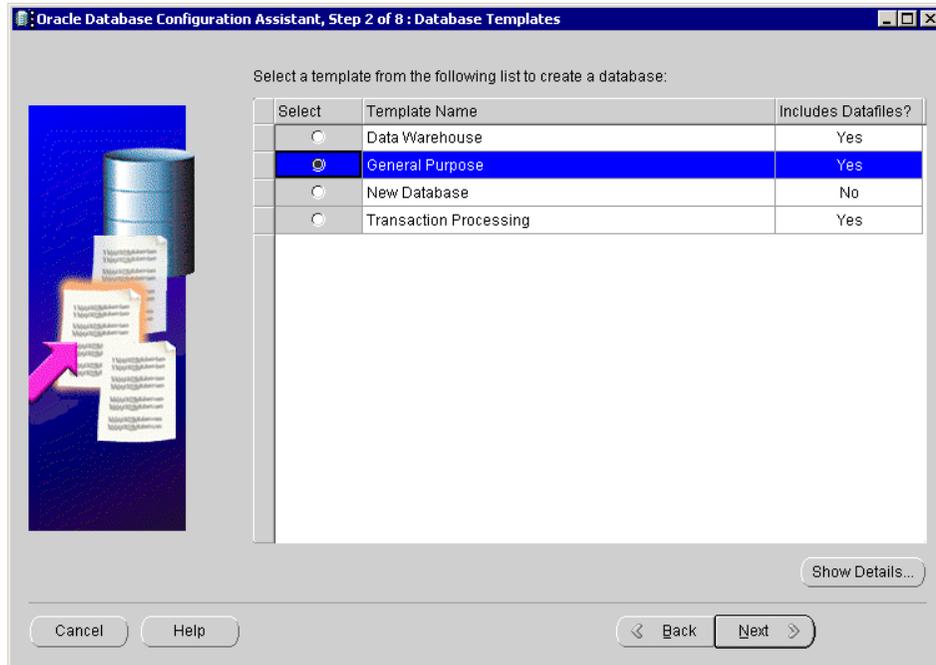


Figure 6-46 DBCA database templates window

Among the choices, we select the general purpose database configuration. The same database configuration is chosen, throughout this chapter.

Select General Purpose, and the row will highlight. Click **Next** to continue.

Notice that you can view detailed configurations by clicking on Show Details button.

The DBCA Step 3 of 8: Database Identification window opens (Figure 6-47).

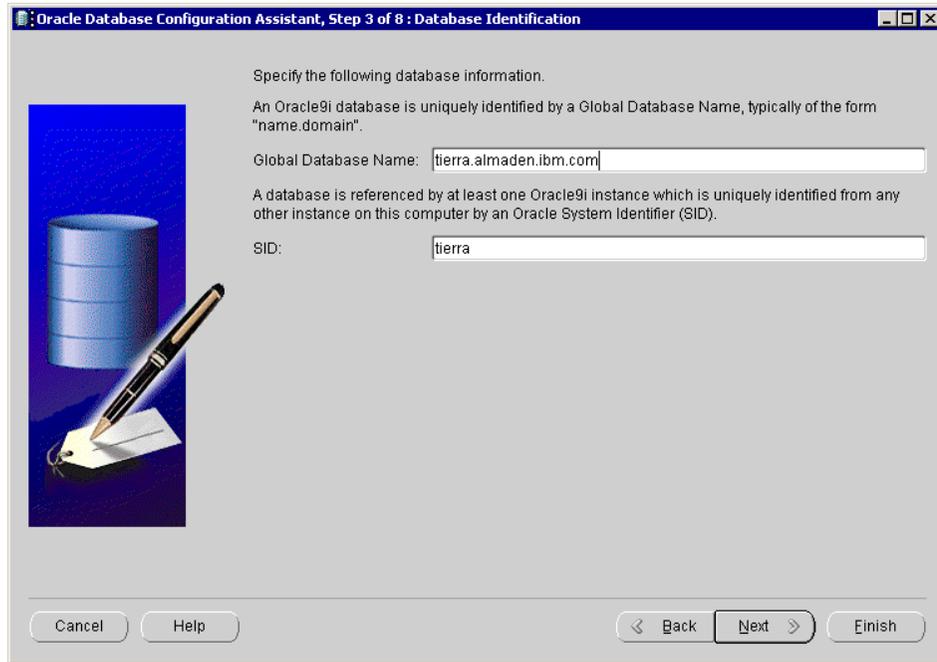


Figure 6-47 Database identification window

The global database domain is the combination of the database name and the network domain where the database exists. This combination is a unique name that will identify the database in the network, or as a stand alone database system.

For example, the database name is 'tierra' and the domain is 'almaden.ibm.com'. Therefore, the global database name is 'tierra.almaden.ibm.com'.

The Oracle System Identifier, SID, is another name for database name. The database name should always be unique. We chose 'tierra' as a database name. Enter your database name and global database domain, and then click **Next** to continue.

The DBCA Step 4 of 7: Database Connection Options window opens (Figure 6-48).

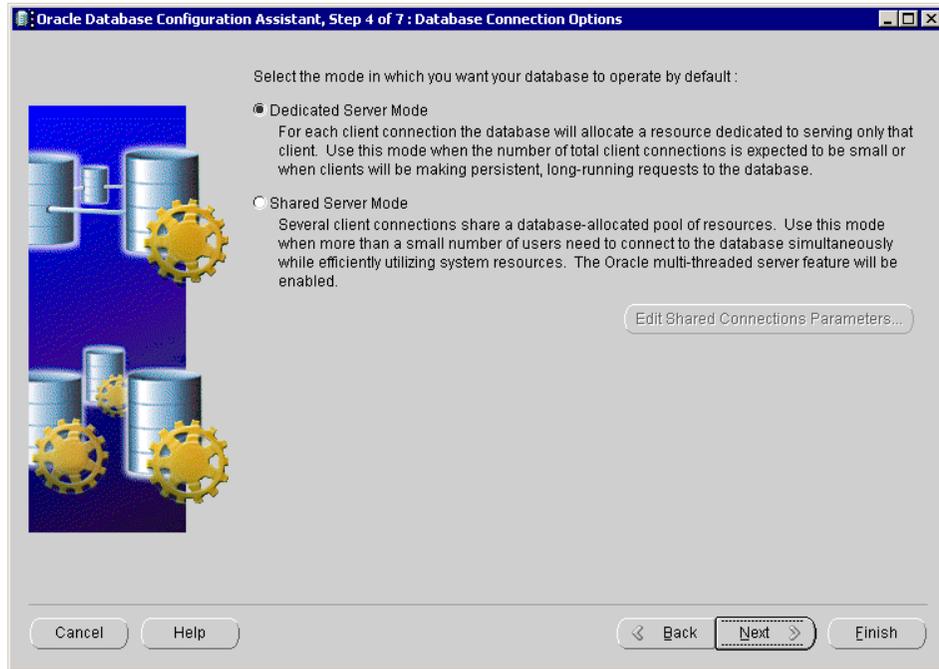


Figure 6-48 Database connection options window

We chose the Dedicated Server Mode, because we are setting up a small environment.

Select the Dedicated Server Mode button. Click **Next** to continue.

The DBCA Step 5 of 7: Initialization Parameters window opens (Figure 6-49).

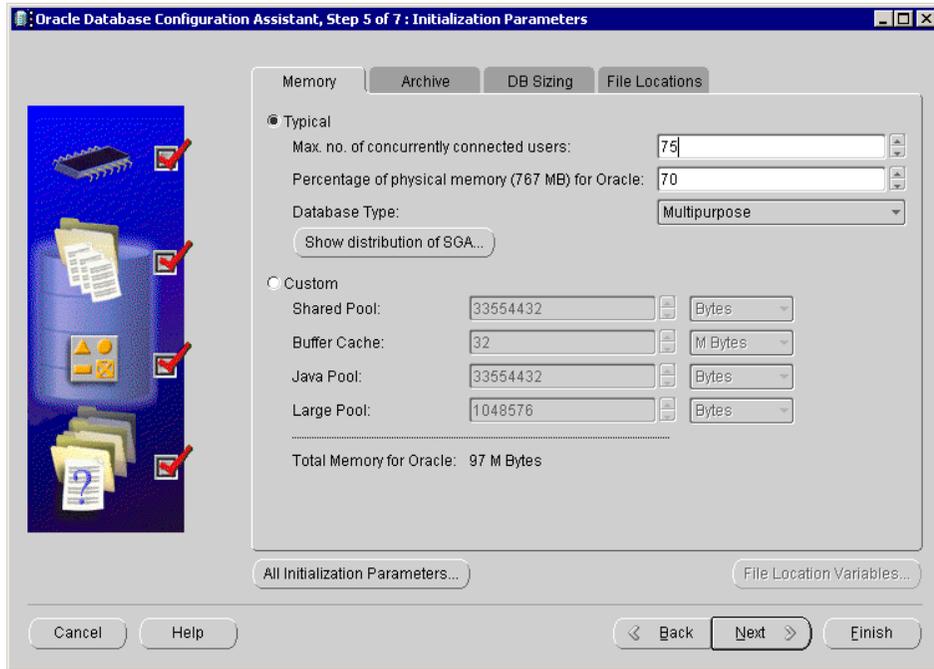


Figure 6-49 Initialization parameters window

Four tabs are displayed in this window. We will modify some of the default parameters in the Memory, Archive, DB sizing, and File Location tabs.

On the Memory tab, we select a basic memory option. Click the Typical button. To continue, click the Archive tab to select archiving options, or click **Next** (to skip the tabs).

Here are a few things to notice:

- ▶ You can change the value of the maximum number of concurrent users.
- ▶ You can change the value of the percentage of physical memory for Oracle.
- ▶ You can view the distribution of System Global Area (SGA).
- ▶ You can view all of the initialization parameters.
- ▶ You can *not* view the file location variables.

You can view these by clicking the appropriate button. You can change the values by entering data in the value boxes.

If you skip the tab selections, go to Figure 6-53, “Database storage window” on page 253.

The DBCA Step 5 of 7: Initialization Parameters Archive tab shows (Figure 6-50).

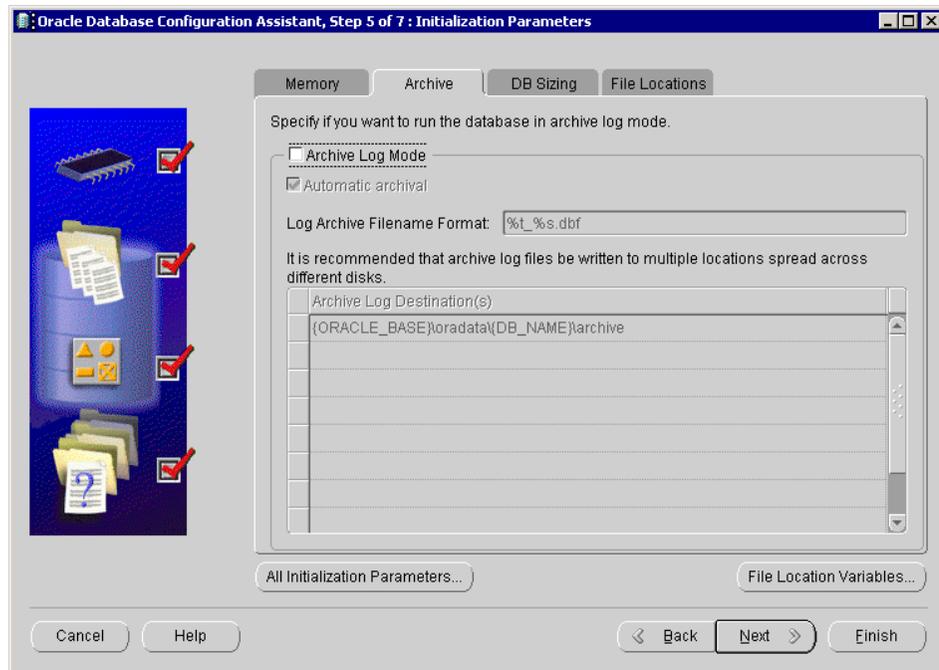


Figure 6-50 Initialization parameters archive tab window

We chose the Automatic Archival *only*. Select Automatic Archival. To continue, click the DB Sizing tab, or click **Next** (to skip the tabs).

Notice the default values displayed for:

- ▶ Log archive file format
- ▶ Archive log destinations

In this tab you are able to view other items by clicking the appropriate button:

- ▶ All initialization parameters
- ▶ File location variables

If you select both the Archive Log mode and Archive Automatic options, you will most likely encounter a series of errors:

- ▶ ORA-00439: feature not enabled: Managed standby
- ▶ ORA-01034: Oracle not available
- ▶ ORA-27101: shared memory realm does not exist

**Attention:** To view the error windows, see Figure 6-60 through Figure 6-62 on page 260. Merely click Abort for these errors, and continue with the instructions for Figure 6-50. Only choose Automatic Archival.

If you skip the tab selections, go to Figure 6-53, “Database storage window” on page 253.

The DBCA Step 5 of 7: Initialization Parameters DB Sizing tab shows (Figure 6-51).

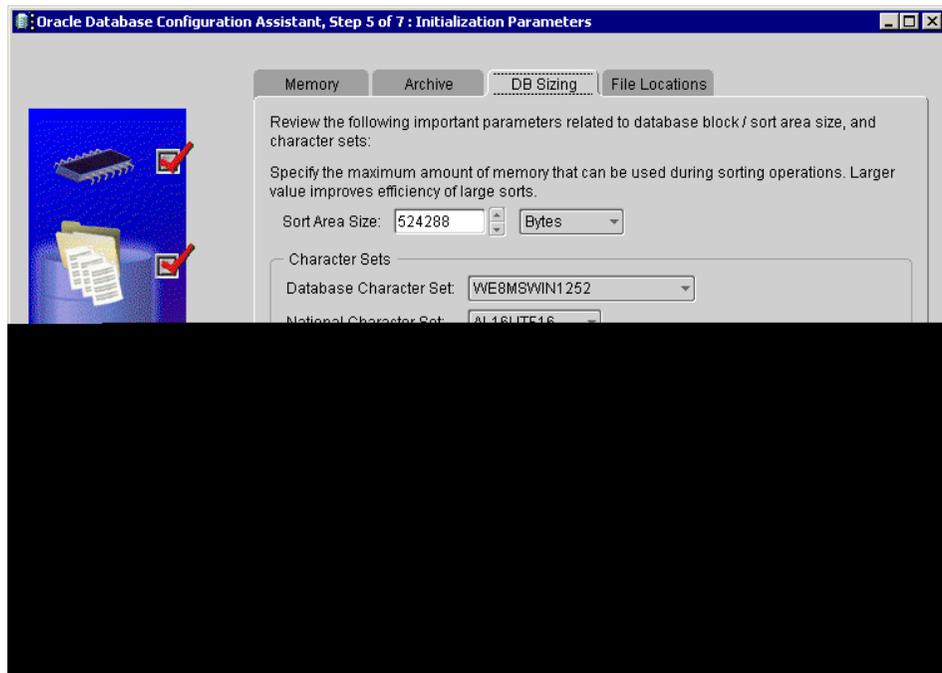


Figure 6-51 Initialization parameters db sizing tab window

This tab is not about database sizing. It is about choosing the following items:

- ▶ Sort area size
- ▶ Database character set
- ▶ National character set

The default values are displayed. We chose default values for the character set. Change them to fit your requirements.

For the sort area size, discretion should be used based on sorting average operations.

To continue, click the File Location tab, or click **Next**.

If you skip the tab selections, go to Figure 6-53, “Database storage window” on page 253.

The DBCA Step 5 of 7: Initialization Parameters File Location tab shows (Figure 6-52).

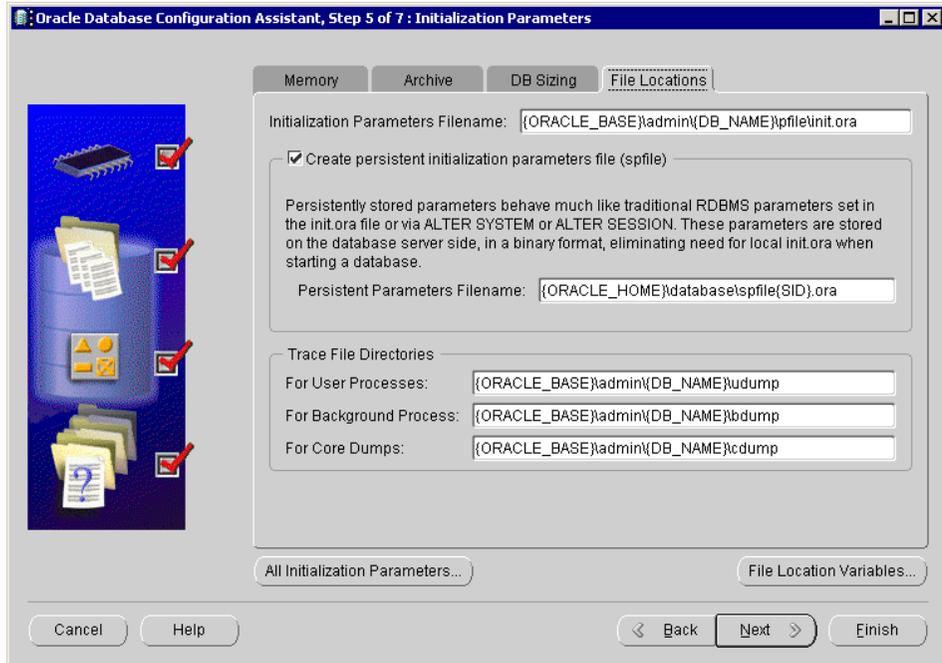


Figure 6-52 Initialization parameters file location tab window

We chose default file locations within the Oracle home directories. Change the file locations to fit your requirements. Click **Next** to continue.

**Note:** Here is a description of the *Create persistent initialization parameters file* button.

- ▶ **Persistent initialization parameters** allow certain parameters to be changed online by the administrator, OEM, or internal self-tuning to maintain values after a database is shutdown. This feature also allows the database to start from a remote machine without a need to have a copy of the initialization file. Some parameters persist across database shutdowns.

The DBCA Step 6 of 7: Database Storage window opens (Figure 6-53).

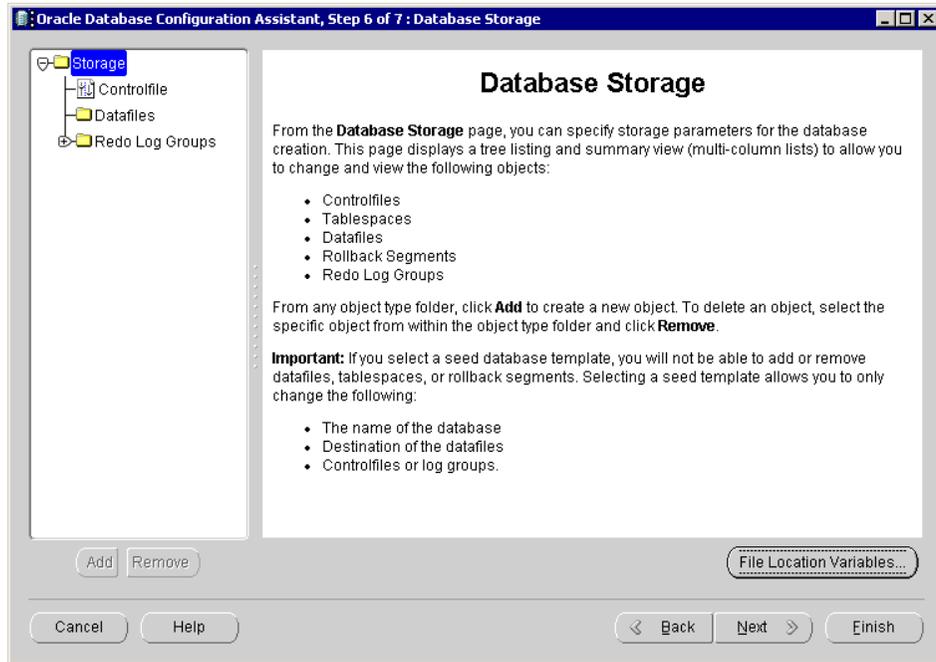


Figure 6-53 Database storage window

In this window you can review the Redo Log Group location, but more importantly you can change the database datafiles location. It is in this window that we change the data files to *reside* on the iSCSI disk.

Click on the datafiles folder, the database storage area changes to display database file names and file directory (Figure 6-54).

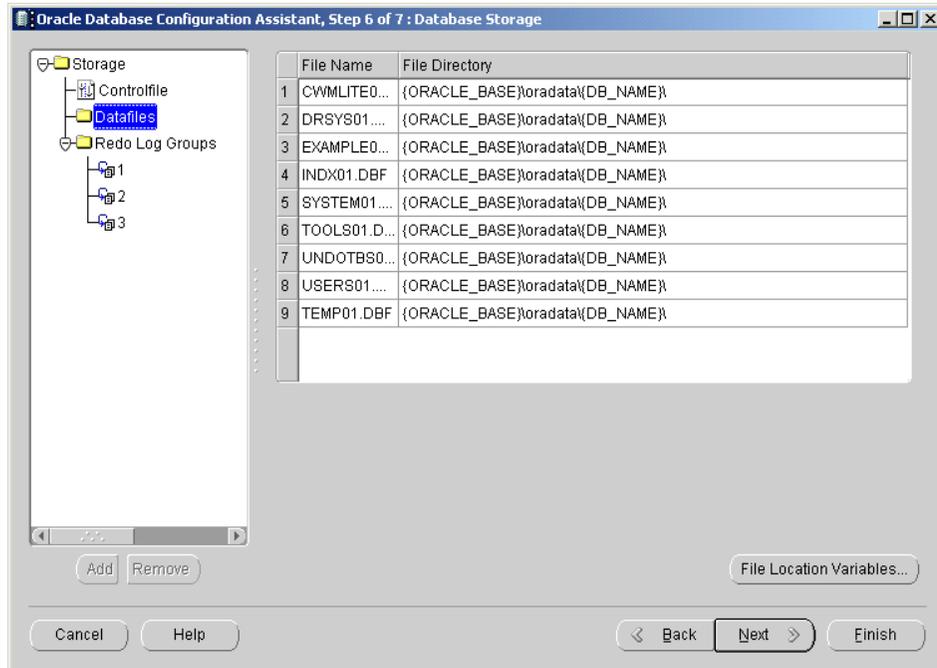


Figure 6-54 File name and file directory window

Place the cursor in the file directory column. Highlight '{ORACLE\_BASE}' on a row (one row at a time), change it to the iSCSI drive. (In our case the drive is G, see Figure 6-55).

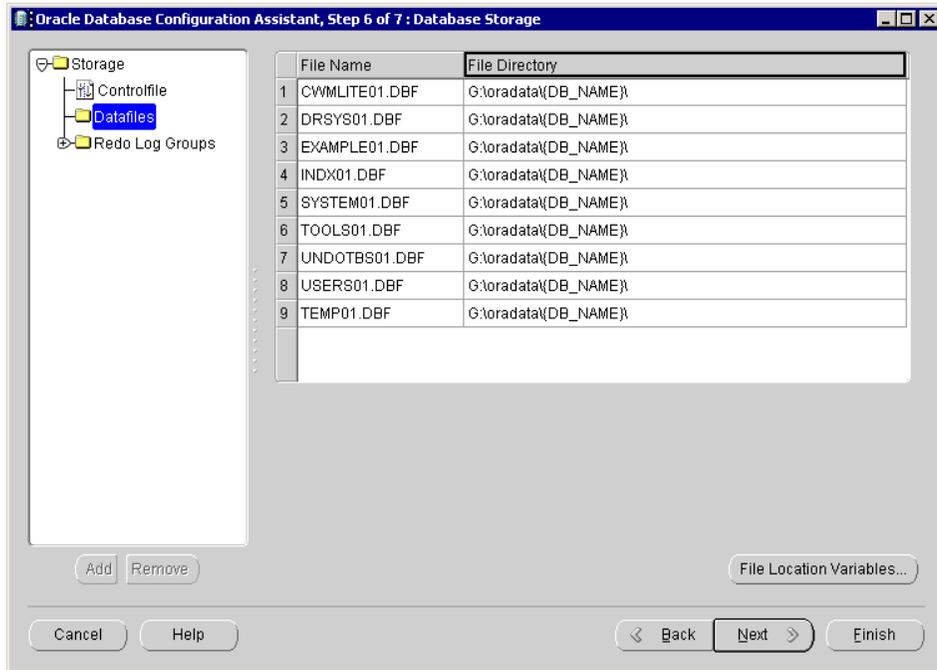


Figure 6-55 Change to iSCSI drive window

Notice that all datafiles will reside on the iSCSI disk. To continue, click **Next**.

The Summary window opens (Figure 6-56).



Figure 6-56 Summary window

You can scroll in this window to view all of the settings before creating the database. Notice that you can also save this information in html format.

To continue, click **OK**.

The DBCA Step 7 of 7: Creation Options window opens (Figure 6-57).

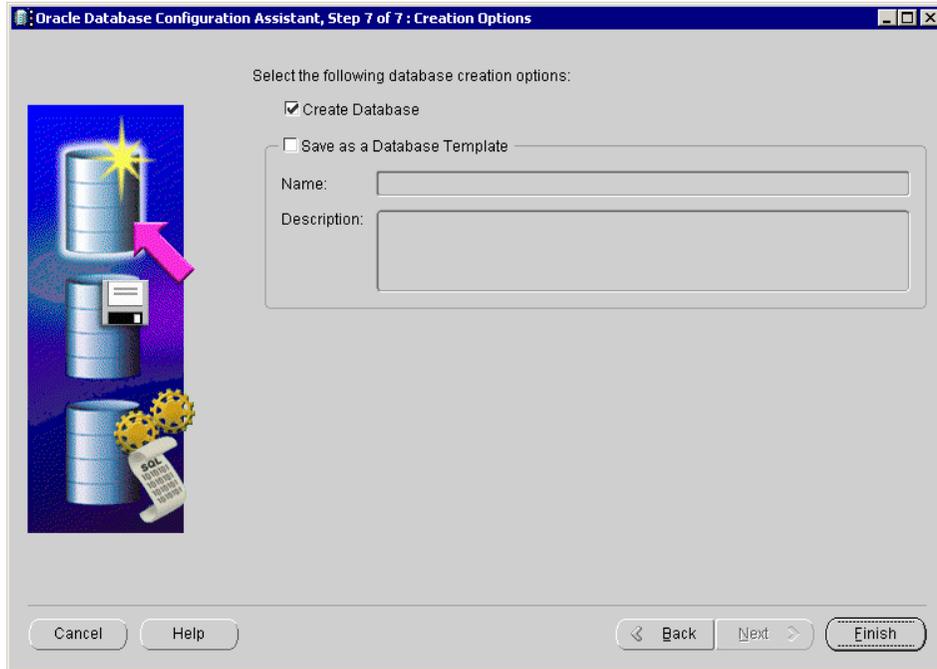


Figure 6-57 Creation options windows

We chose to create a database. Select Create Database. Click **Finish**.

Another DBCA window opens to create the database (Figure 6-58).

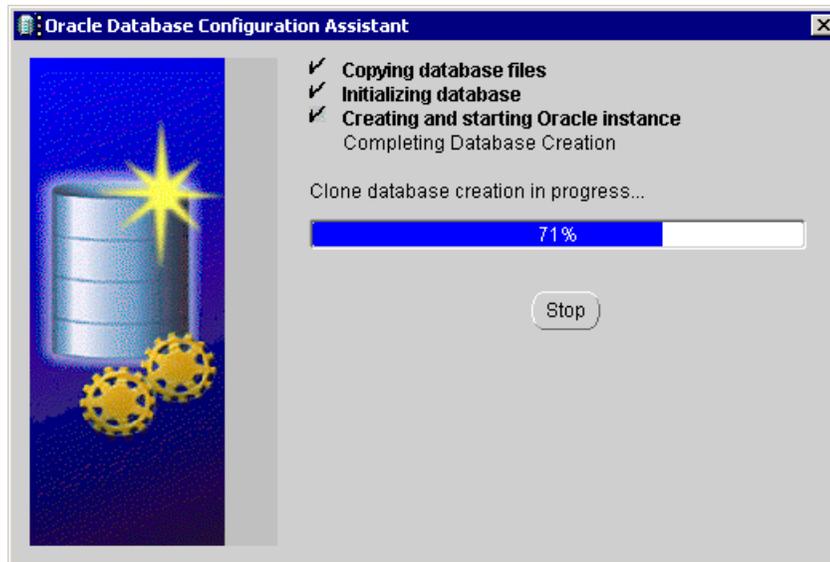


Figure 6-58 Create the database window

This window performs tasks to completion without your intervention, unless you want to cancel the process. To cancel the process, click **Stop**.

When the database creation finishes, another window opens with password and database information. The window allows you to unlock accounts and/or review database information; see Figure 6-16 on page 213 for details. When you click the Exit button on that window, you complete the database creation on the iSCSI disk. All the windows will close.

**Tip:** During the database creation, command windows will open. Minimize them to make sure you can see the progress. Do not open other windows on top of this window.

After installing Oracle on Linux, the next step you'll probably take is to install an Oracle client. The steps are found in "Oracle post-installation on Windows 2000" on page 216.

### **Linux iSCSI database**

To start the Database Configuration Assistant (DBCA), at Linux prompt, and run an X term window emulator, enter:

```
> dbca
```

Follow the instructions in “Windows 2000 iSCSI database” on page 243. The windows in this section will display green colors instead of blue colors. Follow the steps for Figure 6-44 on page 244 through to Figure 6-55 on page 255. There are a few changes, as noted below:

- ▶ Note for Figure 6-47, our Linux iSCSI database is called ‘*lunas*’.
- ▶ Note for Figure 6-55, the changes to the datafile file directory for the Linux database are as follows (see Figure 6-59):

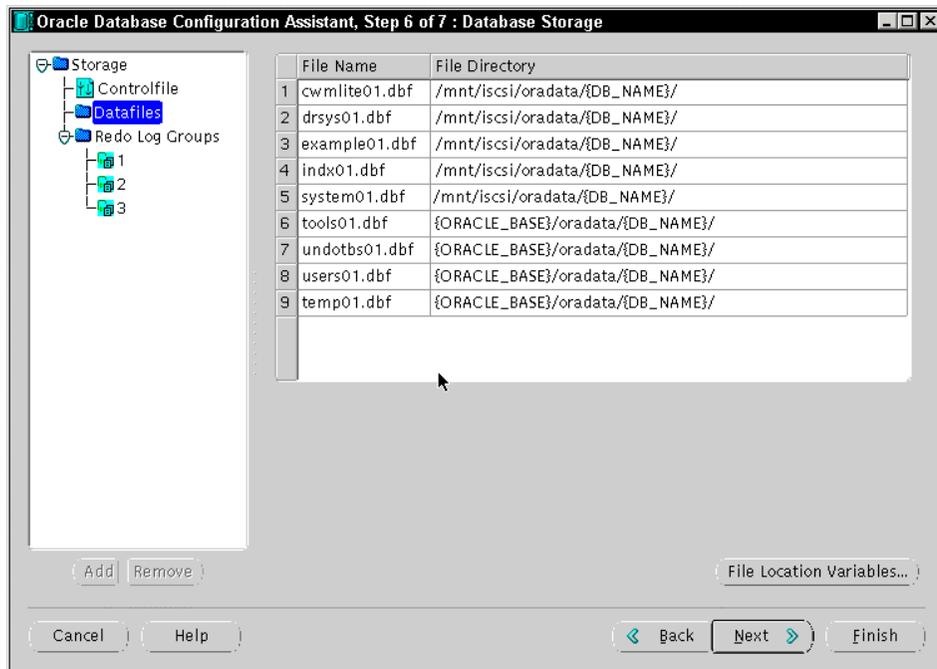


Figure 6-59 Linux database storage

Notice that for Linux database on iSCSI disks, the mount point and file structure is `/mnt/iscsi`. The `{ORACLE_BASE}` part of the file directory column, for every row, is changed to `/mnt/iscsi`. Make sure all rows are changed that way.

Now continue with the instructions for Figure 6-56 on page 256 through Figure 6-58 on page 258.

- ▶ Note that for Figure 6-58, while this window is running and if you selected both automatic archival and archivelog mode, you will encounter the following errors (see Figure 6-60 to Figure 6-62):

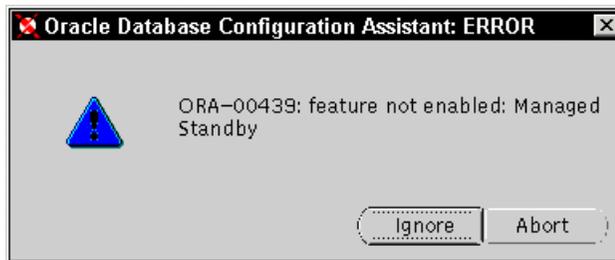


Figure 6-60 Managed standby error window

Our selections did not explicitly include standby options. Click **Abort**.

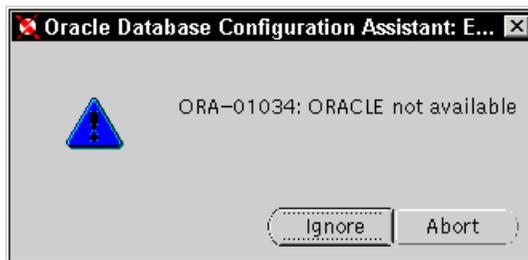


Figure 6-61 Oracle not available error window

Click **Abort** as many times as you get this error.

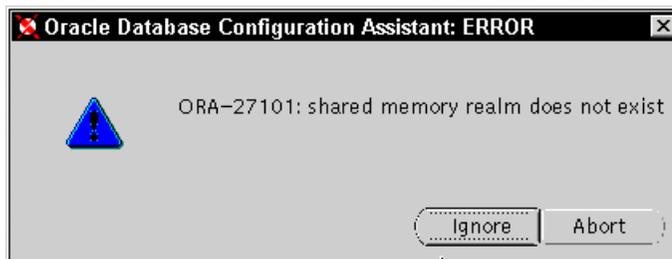


Figure 6-62 Shared memory error window

Click **Abort** as many times as you get this error.

**Tip:** These are Oracle errors. Follow the instructions for Figure 6-50 on page 250 to avoid them. Choose only the automatic archival in that step.

**Note:** The Oracle Web site Metalink has documented these errors. Go to the Web site and search for ora - ##### error.

Follow the post-installation steps in “Oracle post-installation on Linux” on page 241.

## **Remote Oracle9i client implementation**

To access databases remotely, the Oracle9i client software must be installed and configured. We cover the installation and configuration of Oracle9i in the following sections. This installation is for desktops, which do not have Oracle9i client installed already. **Note:** If you have client software installed on your desktop, skip this section and go directly to “Configure remote database on client” on page 271.

### ***Install remote Oracle9i client***

This section includes both the installation and configuration of the Oracle9i client.

To install Oracle9i client, you need to download the software from the Oracle Web site or have a CD-ROM which contains the Oracle9i client.

- ▶ If you have downloaded the client software, you need to unzip the files (on Windows 2000). Open a command prompt window, go to the client/Disk1 directory, and run the **setup.exe** program.
- ▶ If you have a CD-ROM, load it in the CD-ROM drive. Go to the client/Disk1 directory, and double click on the setup icon.

The Oracle Universal Installer (OUI) Welcome window opens (Figure 6-63).

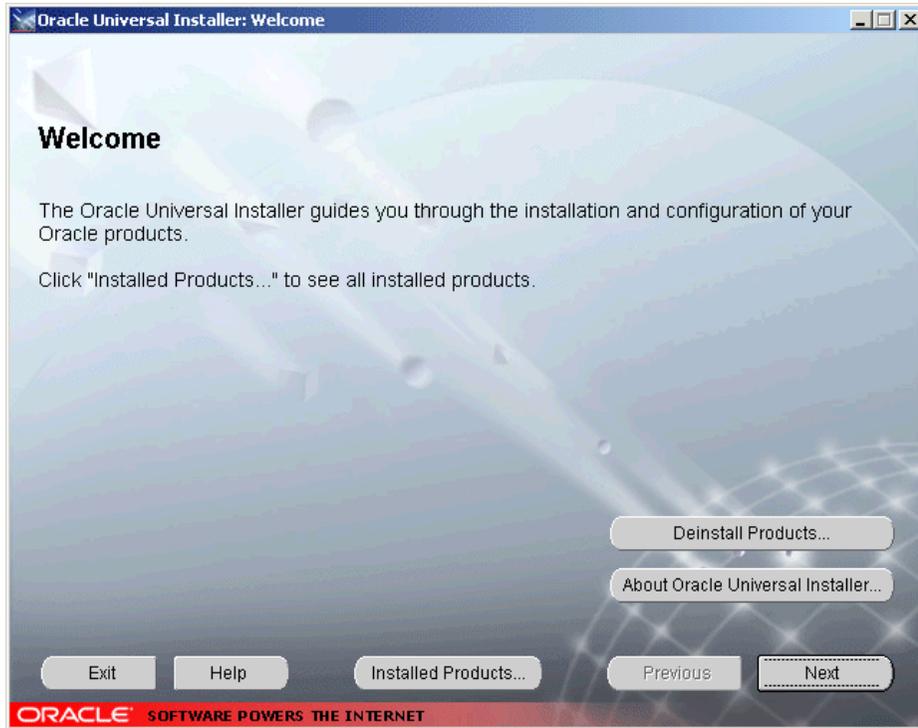


Figure 6-63 OUI client welcome window

To begin the installation, click **Next**.

The OUI Client File Locations window opens (Figure 6-64).

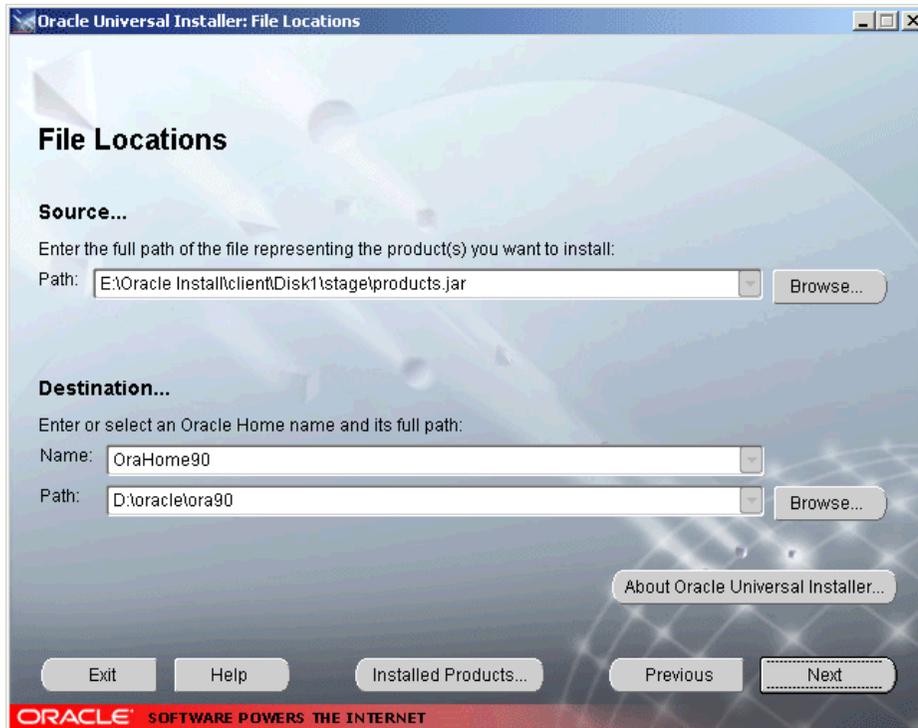


Figure 6-64 OUI client file locations window

The options for this window are:

- ▶ Standard default values (that are displayed) *or* choose different paths:
  - The source path is one of the directories created for software installation in Step 2. on page 199, `\OracleInstall\Disk1`, where Oracle installation files reside. The source path is also the full path to where you copied or extracted the Oracle files (from CD-ROM or zip files). Most likely you will not have to change this option.
  - Or, you can choose your own destination names and paths. The destination is the full path where the Oracle software will be installed and run from. The current settings and defaults are Oracle standard destination full paths and names for the Oracle home. You could choose a different path and name to fit your own company standards. We chose OUI default settings.

**Tip:** You have the option to browse for both source and destination paths.

To continue, click **Next**. The OUI Client Installation types window opens (Figure 6-65).

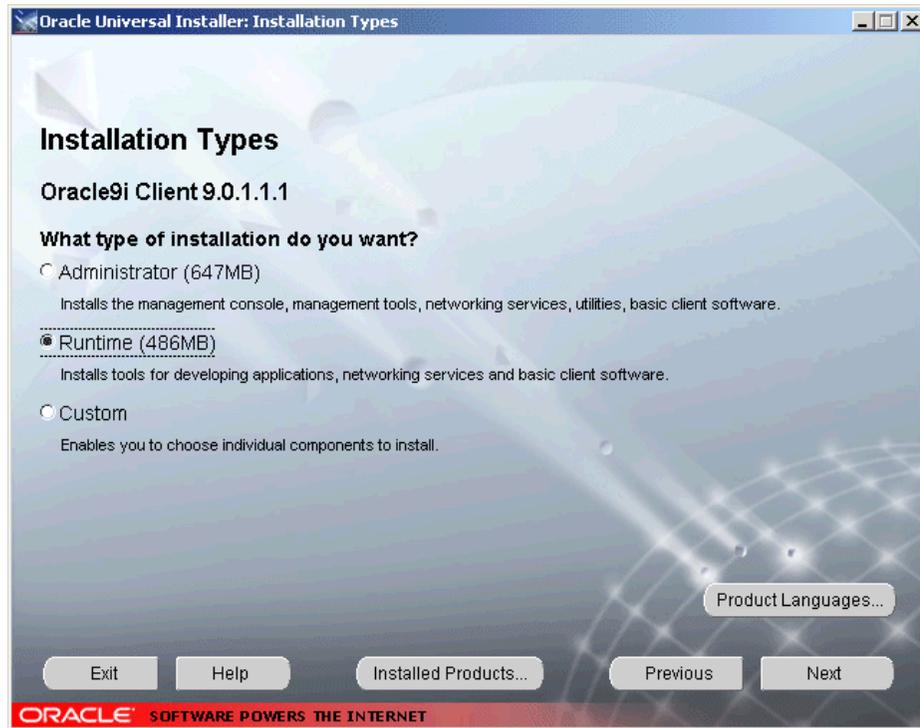


Figure 6-65 OUI client installation types window

In this window, we chose the Runtime option, for remote database access. These options are available:

- ▶ The Administrator option gives you administrative tools and privileges to manage the client software. This option is installed with the Oracle database software.
- ▶ The Runtime option gives you basic components to use the Oracle client.
- ▶ The Custom option allows you to configure the client in more detail.

Review your options in this window. Select the installation type you want, and then click **Next** to continue.

The OUI Client Summary window opens (Figure 6-66).

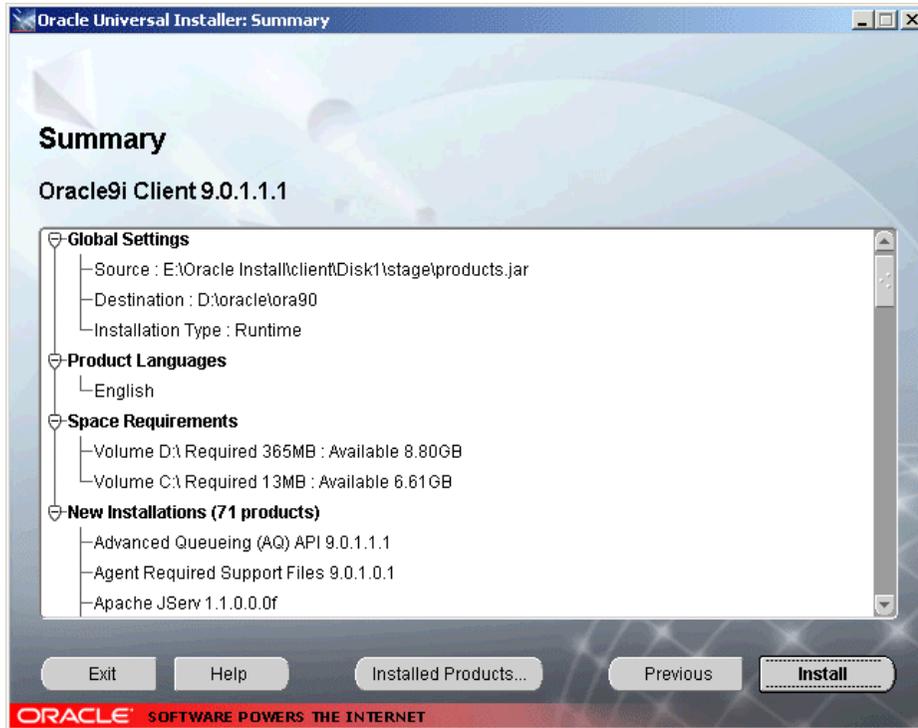


Figure 6-66 OUI client summary window

This summary provides information on all Oracle client components to be installed, as well as all the installation options chosen so far.

To continue, click **Next**. The OUI Client Install window opens (Figure 6-67).

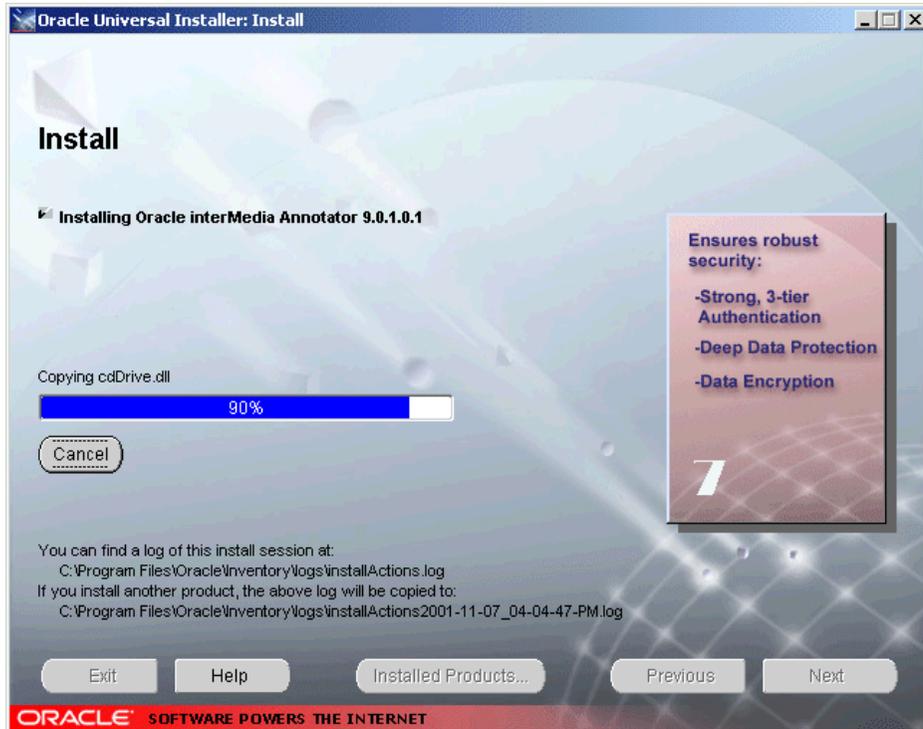


Figure 6-67 OUI client install window

This installation process will take some time to complete. The Oracle9i client software is being installed. Notice that the rectangular image on the center right side of window will increase in numbers, from 1 up to 7, while the software installs.

When it finishes, click **Next** to continue.

The OUI Client Configuration Tools window opens (Figure 6-68).

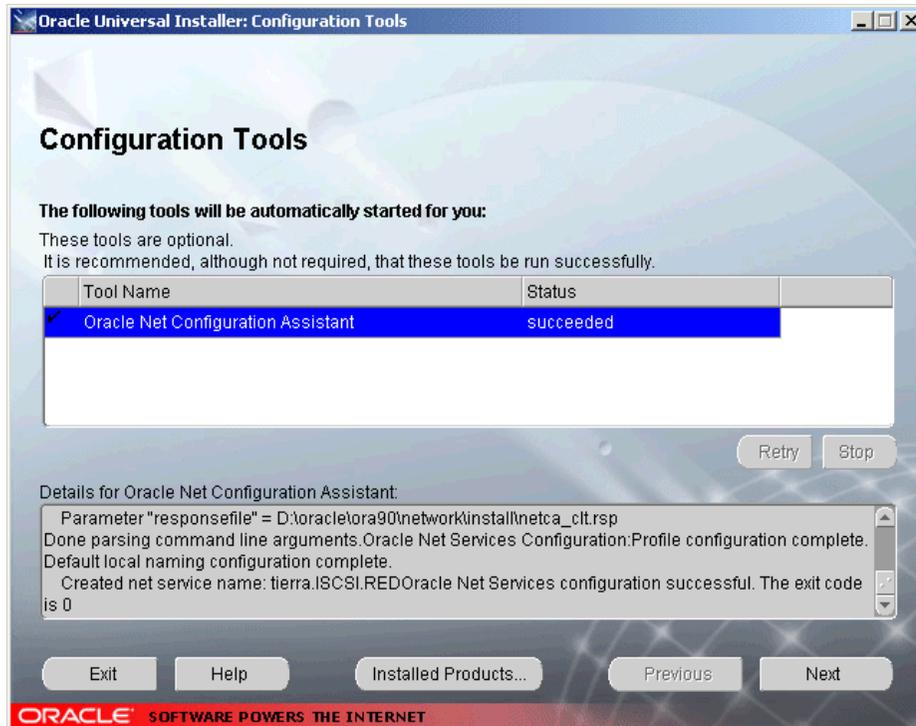


Figure 6-68 OUI client configuration tools window

Once the client installation has completed, the installer automatically starts the configuration tools to configure and to start the Oracle client.

The Oracle Net Configuration Assistant is running, when it shows a solid check mark.

Notice the status column on the window indicates success, failure, or in progress. Also note the result area across the bottom of the screen, where you can see details on the execution of a particular tool.

**Tip:** Do not open any other window on top of this window while the configuration tools are running. The window may not refresh back to the original screen. You will not know when the configuration completes or if you encountered errors. It will appear as if the configuration tools hang.

While the OUI configuration tools are running, the Oracle Net Configuration Assistant (ONCA) Welcome window opens (Figure 6-69).

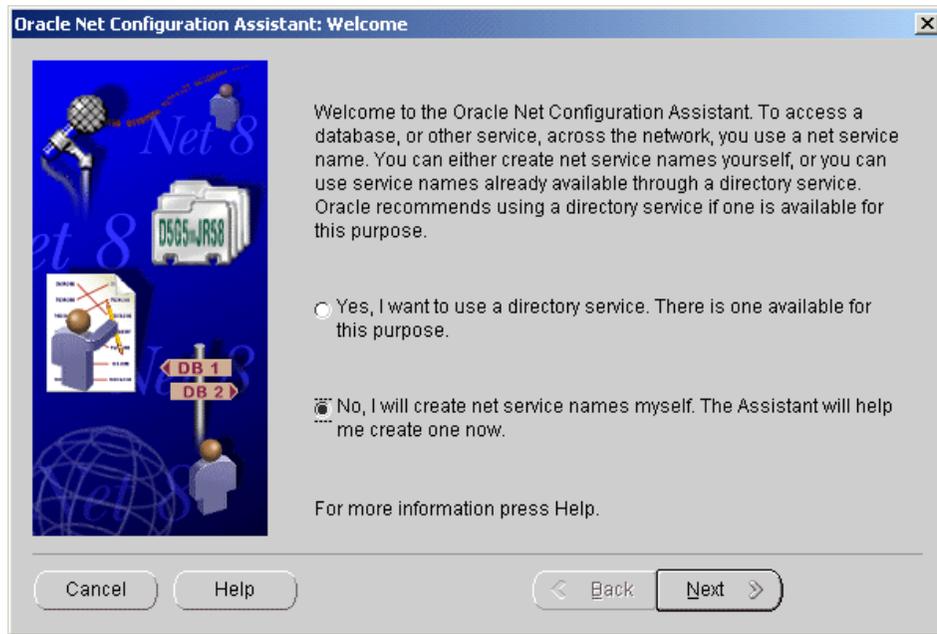


Figure 6-69 ONCA client welcome window

This window gives you two options to choose from:

- ▶ Yes, to use the service names using the directory service
- ▶ No, to create the desired net service names

We chose 'No' to create the net service name, since this is a basic client installation.

Make your selection, and then click **Next** to continue.

The ONCA Net Service Name Configuration (NSNC), Database Version window opens (Figure 6-70).

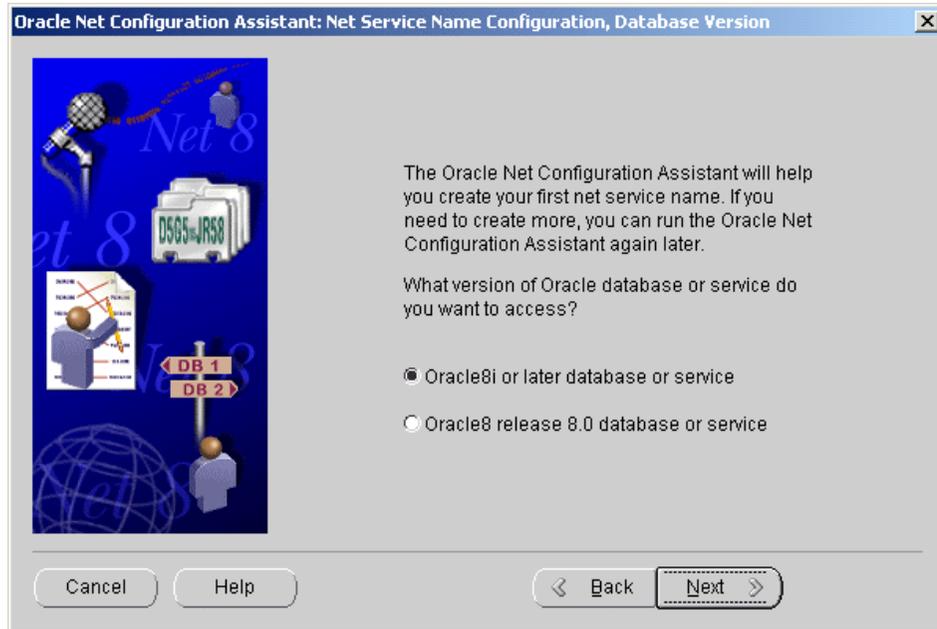


Figure 6-70 NSNC database version window

We chose the latest Oracle version (Oracle9i).

Select your option, and then click **Next** to continue.

The OUI Net Configuration Assistant window opens (Figure 6-71).

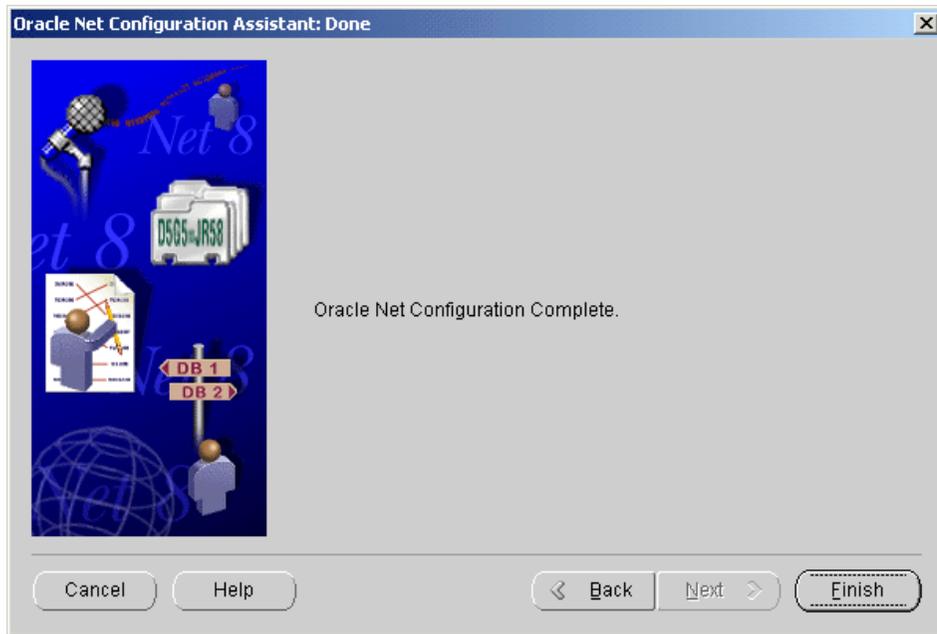


Figure 6-71 ONCA complete window

You just completed the client configuration for your remote database. Click **Finish** to continue.

The OUI End of Installation window opens (Figure 6-72).

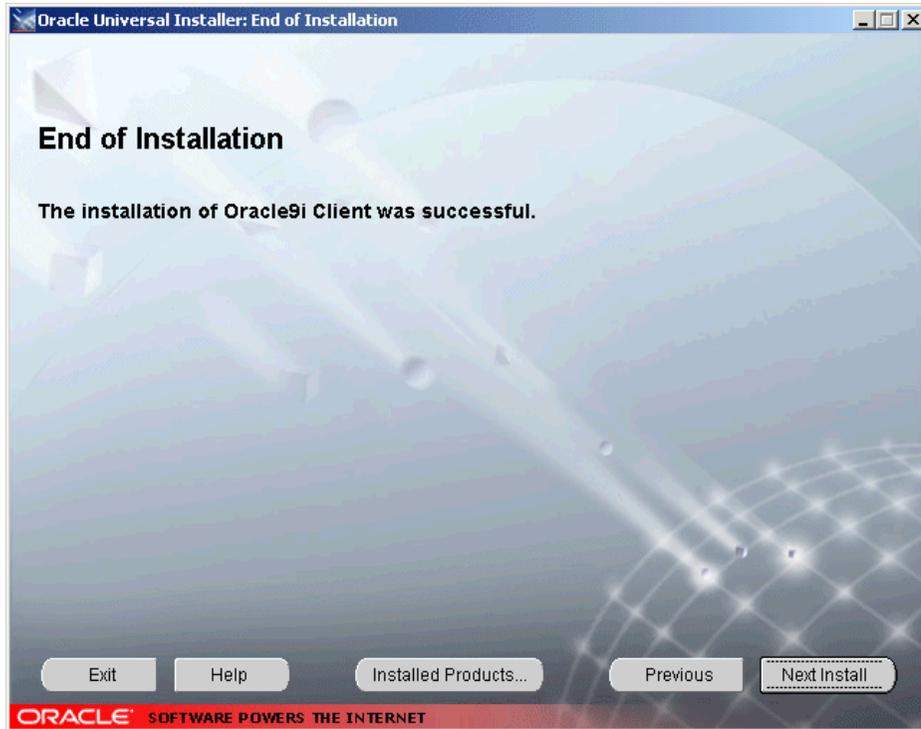


Figure 6-72 OUI end of installation window

Congratulations! You just completed the client installation and configuration for your remote database. You can now remotely access your database.

### ***Configure remote database on client***

From a Windows 2000 desktop, where the Oracle9i client is installed, open the Oracle9i Net Configuration Assistant by selecting:

**Start->Programs->Oracle - OraHome90->Configuration and Migration Tools->Net Configuration Assistant**

The Oracle Net Configuration Assistant (ONCA) Welcome window opens (Figure 6-73).

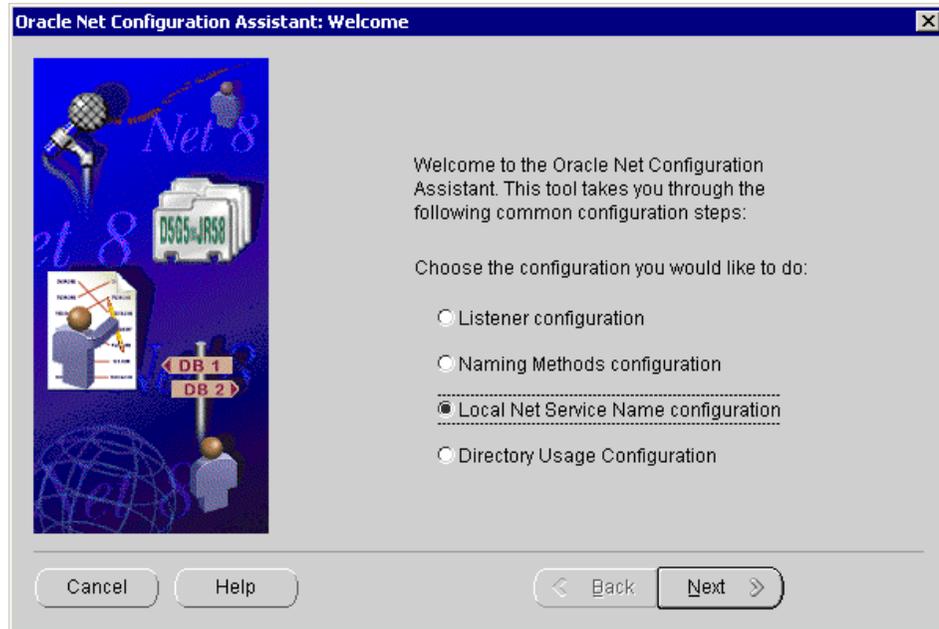


Figure 6-73 ONCA welcome window

This window gives you several options. The option to configure a remote database on the local network is *Local Net Service Name configuration*. Select this option, and then click **Next** to continue.

The ONCA Net Service Name Configuration window opens (Figure 6-74).

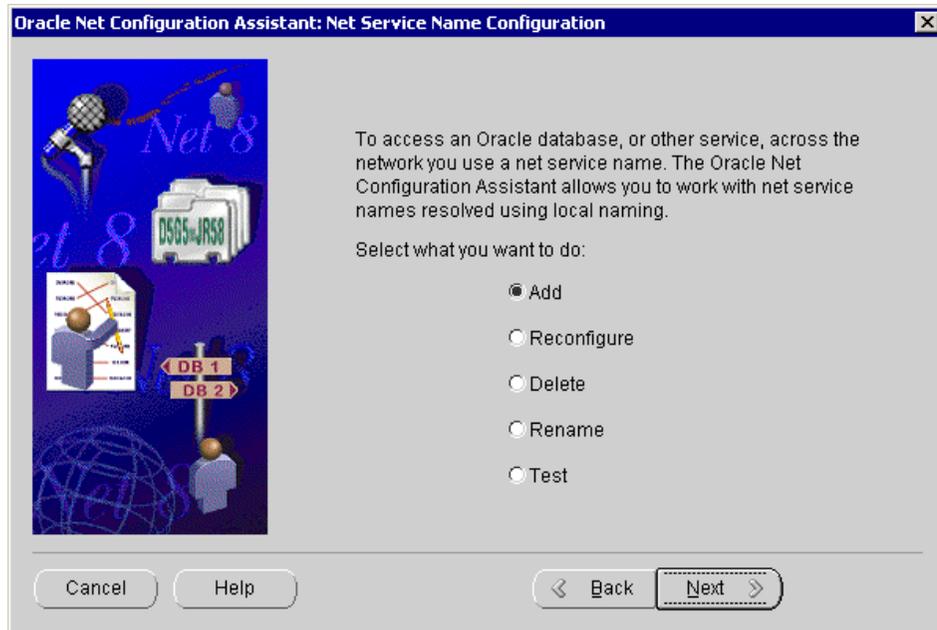


Figure 6-74 ONCA net service name configuration window

This window allows you to select a task. We want to configure a new database to the client software. To add a new database, click **Add**. Click **Next** to continue.

The Net Service Name Configuration (NSNC), Database Version window opens (Figure 6-75).

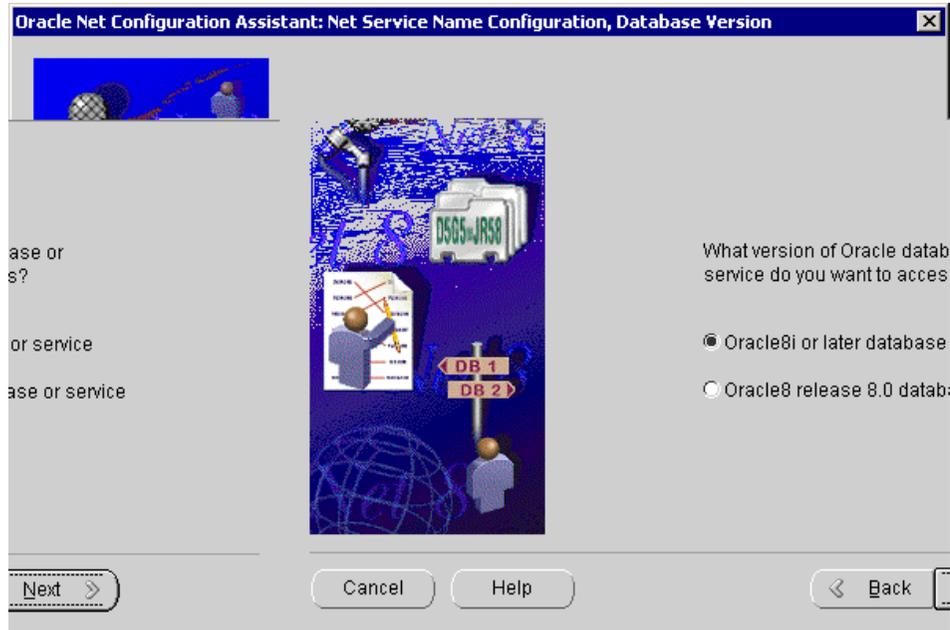


Figure 6-75 NSNC database version window

We chose latest Oracle version (Oracle9i).

Select your option, and then click **Next** to continue.

The NSNC, Service Name window opens (Figure 6-76).



Figure 6-76 NSNC service name window

Enter the service name, and click **Next** to continue. In our case, it is the global database name. The global database name is composed of database name and domain name.

Our database name is 'tierra', and the domain name is 'almaden.ibm.com'.

The NSNC, Select Protocols window opens (Figure 6-77).

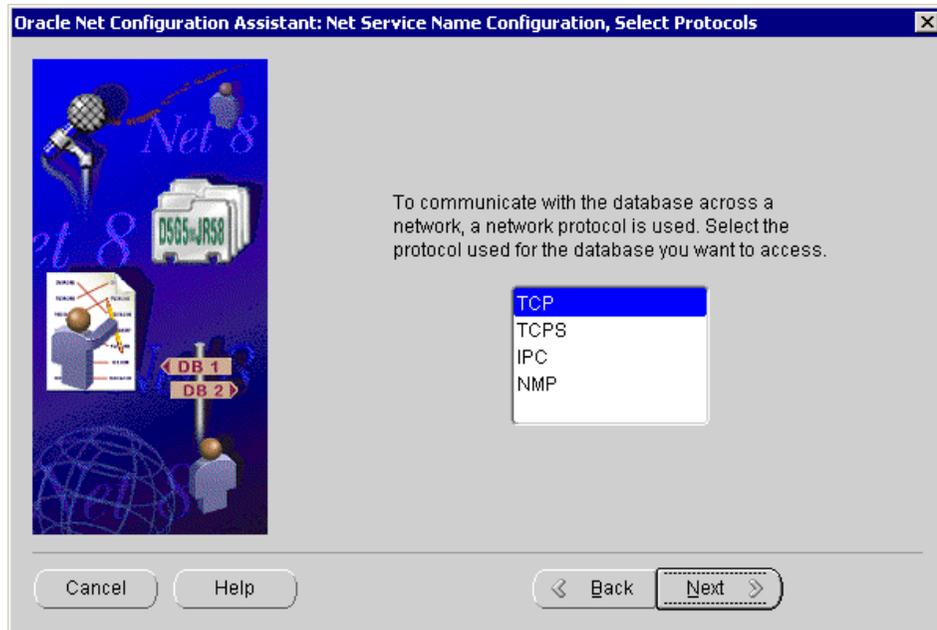


Figure 6-77 NSNC select protocols window

Our network protocol is TCP/IP. We selected TCP by clicking on desired protocol.

Select your protocol, and click **Next** to continue.

The NSNC, TCP/IP protocol window opens (Figure 6-78).

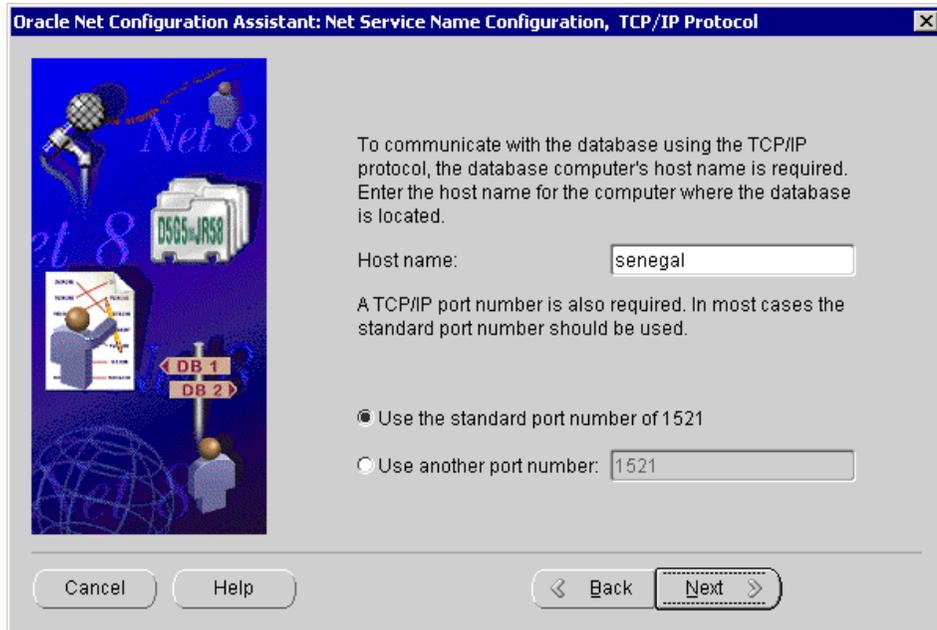


Figure 6-78 NSNC TCP/IP protocol window

Enter the host name where the Oracle software resides.

In our case it is *'senegal'*. We used the default standard port number 1521.

Select your port number. You can select another port number as long as is not used by another application. Click **Next** to continue.

The NSNC, Test window opens (Figure 6-79).

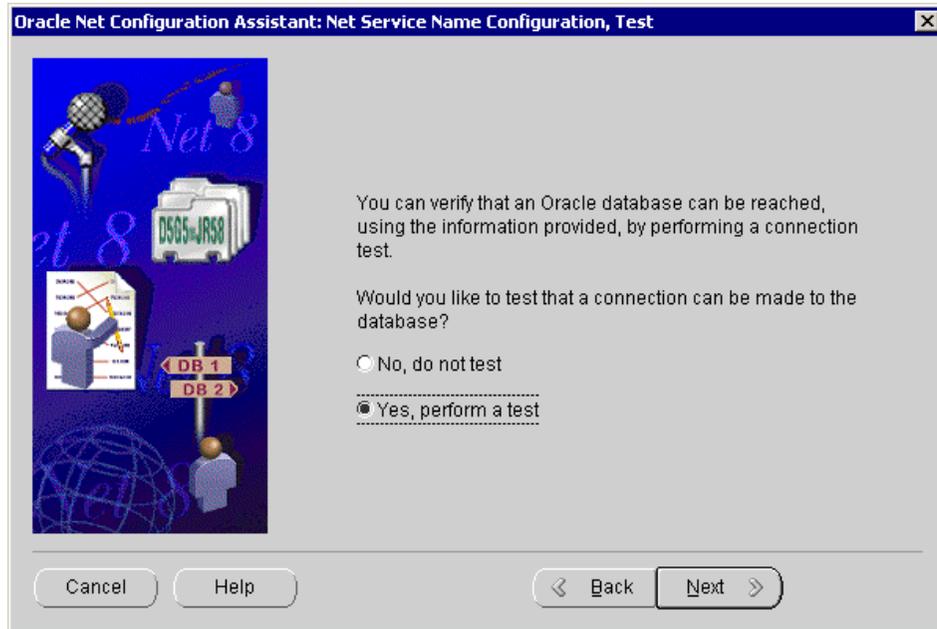


Figure 6-79 NSNC test window

Once the database is configured on the Oracle client, it is a good idea to test your configuration. We chose to test for connectivity by selecting Yes.

Make your choice, and click **Next** to continue.

The NSNC, Connecting window opens (Figure 6-80).

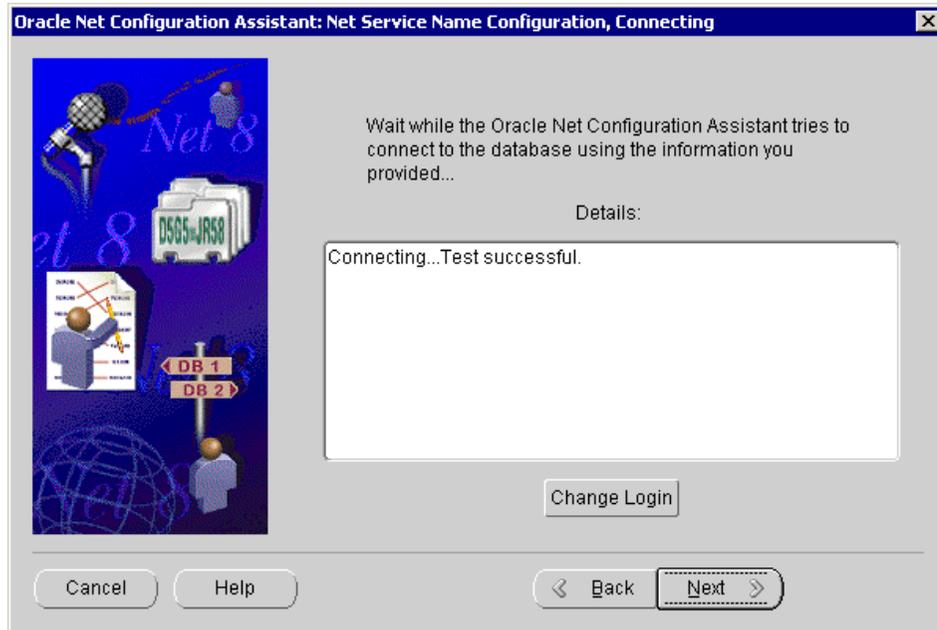


Figure 6-80 NSNC connecting window

Our connectivity to the configured database tested successfully.

For another option, you could change your login by clicking the Change Login button. You could also go back to review you configuration choices and retest. To go back, make consecutive clicks on Back button, until you find the configuration items you want to review.

Click **Next** to continue.

The NSNC, Net Service Name window opens (Figure 6-81).

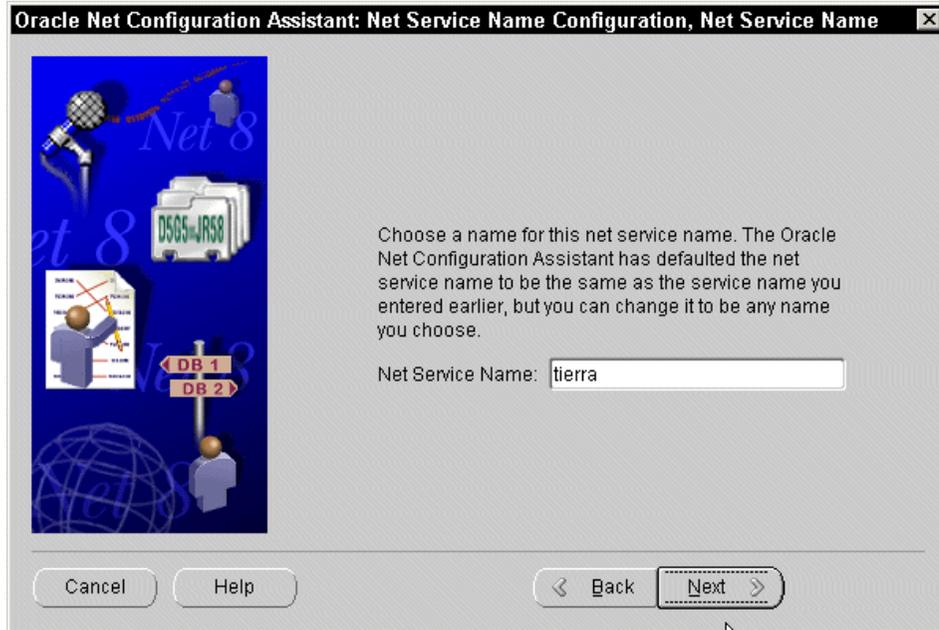


Figure 6-81 NSNC net service name window

The net service name on the client software is the database name. The default used from previous windows is 'tierra'. If this is not correct, enter the correct database name, and then click **Next** to continue.

The NSNC, Another Service Name? window opens select (Figure 6-82).

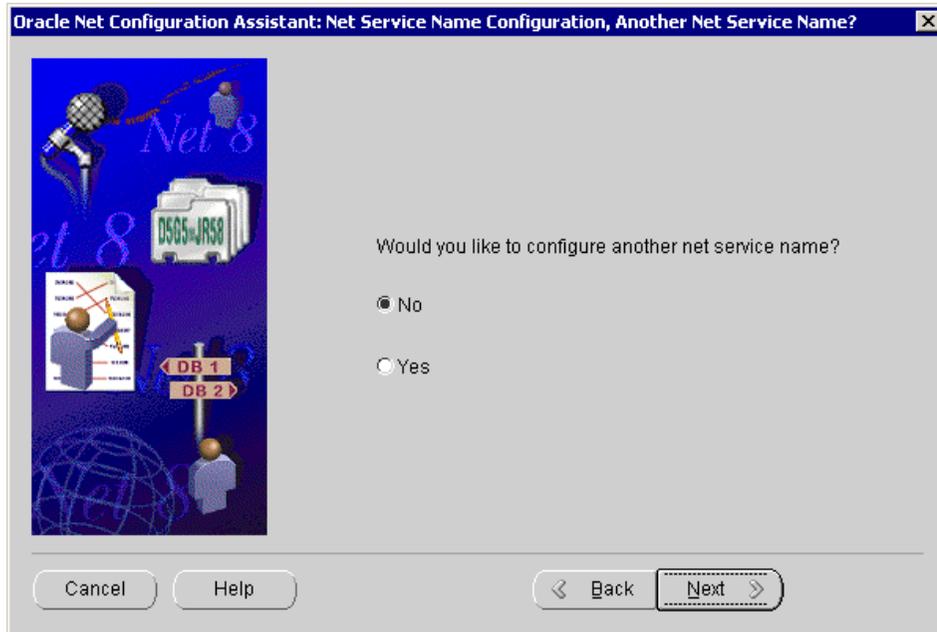


Figure 6-82 NSNC another service name window

We chose not to configure another database.

Select No, and then click **Next** to continue.

The NSNC, Complete window opens (Figure 6-83).

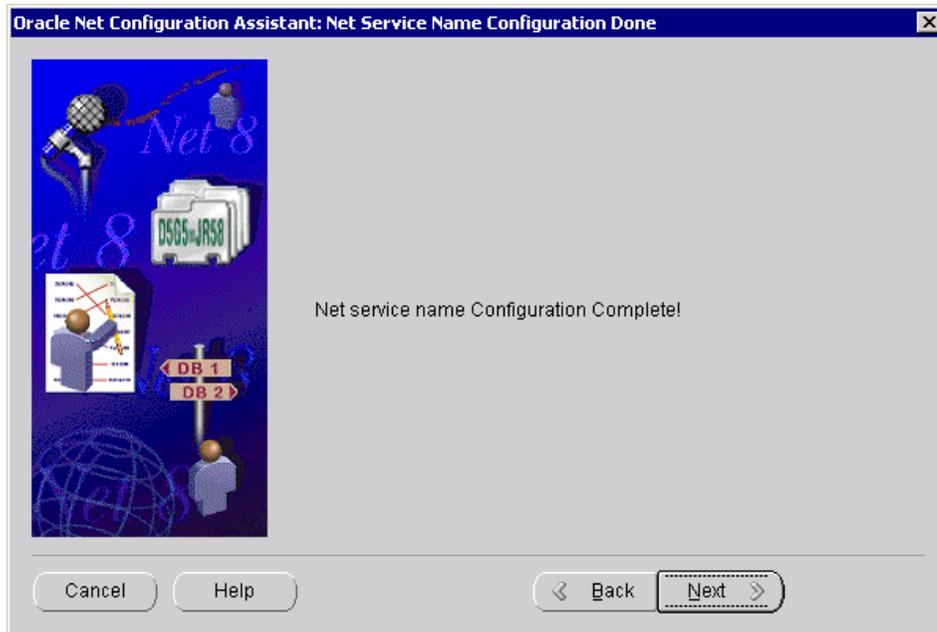


Figure 6-83 NSNC complete window

Congratulations! You just completed a net service name configuration on your remote database. You can now remotely access your database.

Click **Next** to continue.

The ONCA Welcome window opens again (Figure 6-84).

**Tip:** If after successful database client configuration and testing, access to database is not possible, even though your database was started properly, then stop and start the listener process. See 6.5, “Hints and tips” on page 283.

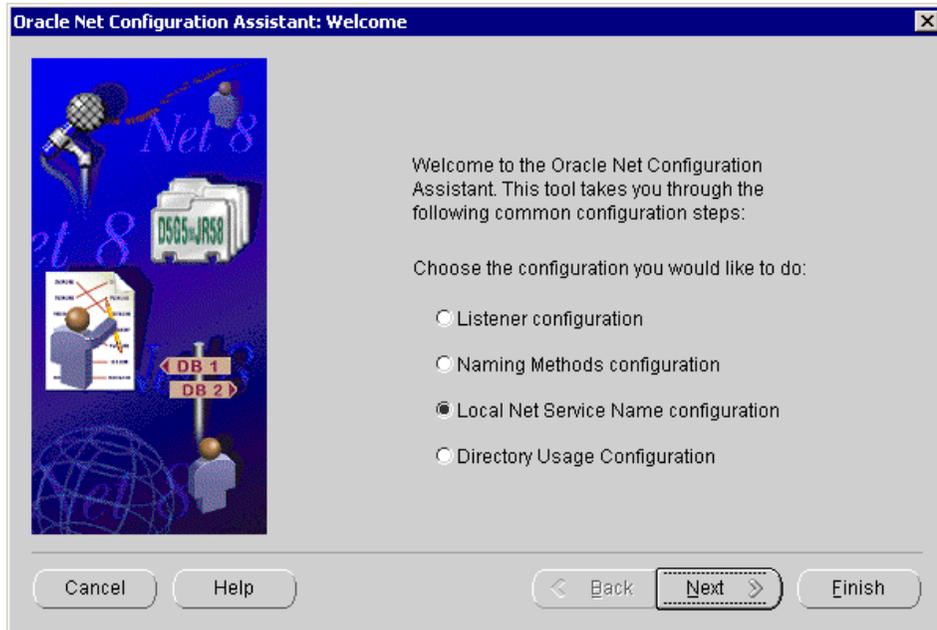


Figure 6-84 ONCA welcome window

You completed the configuration, click **Finish** to exit.

Once you have completed the configuration, the only other task is to ensure that your remote hosts and protocol are defined on your desktop files:

- ▶ C:\WINNT\system32\drivers\etc\services  
Add a new entry for the host with IP address: **9.1.38.166 senegal**
- ▶ C:\WINNT\system32\drivers\etc\hosts  
Add a new entry for the Oracle listener port number and protocol:  
**oralsnr 1521/tcp #oracle port and protocol**

## 6.5 Hints and tips

Here are some recommendations, hints and tips. Some of the information was included in the steps which you've already completed, and it is repeated here for convenience. Other information here is helpful for you to know. Please review it carefully.

### **Shutdown and start database and listener**

Make sure ORACLE\_SID is set properly. See the next three list items for Windows or Linux. Otherwise, at the command line, use these commands:

```
>lsnrctl stop (stop listener process)
```

Or,

```
>lsnrctl start (start listener process)
```

```
>sqlplus "/ as sysdba" (connect to database)
```

```
SQL>shutdown immediate (shutdown database)
```

Or,

```
SQL>startup (start database)
```

1. Every time you configure the Oracle9i client to add a database, stop and start the listener process.
2. For Linux, ensure the \$ORACLE\_HOME/bin/oracle executable has the appropriate permissions. It should be rwsr-s--x. If it's not, then shutdown the database and change permissions. Then start the database, to change permissions, and enter at the prompt:  

```
chmod 6751 $oracle_home/bin/oracle
```
3. For Windows 2000 or Linux, set the ORACLE\_SID variable to the desired database name before accessing the database.
  - Windows 2000 — change registry
  - Linux — export variable with appropriate settings

### **Recommendations**

- ▶ For database migration to a new Oracle9i version, consult the Oracle migration documentation.
- ▶ While Oracle9i net assistant is running, we recommend that you do *not* open other windows on top of it. The window will not refresh in some instances and you will not know when the installation process finishes or if it hangs.

### **Hints and tips**

- ▶ For Linux, to invoke the Database Configuration Assistant (you must have an X-term window or emulator); enter this at the Linux prompt:  

```
> dbca
```
- ▶ For Linux, to invoke the Net Configuration Assistant, (you must have an X-term window or emulator); enter this at the Linux prompt:  

```
> netca
```

- ▶ SVRMGRL, server manager, and the CONNECT INTERNAL command was desupported in Oracle9i. The new syntax to access desupported functionality, is by using SQLplus:

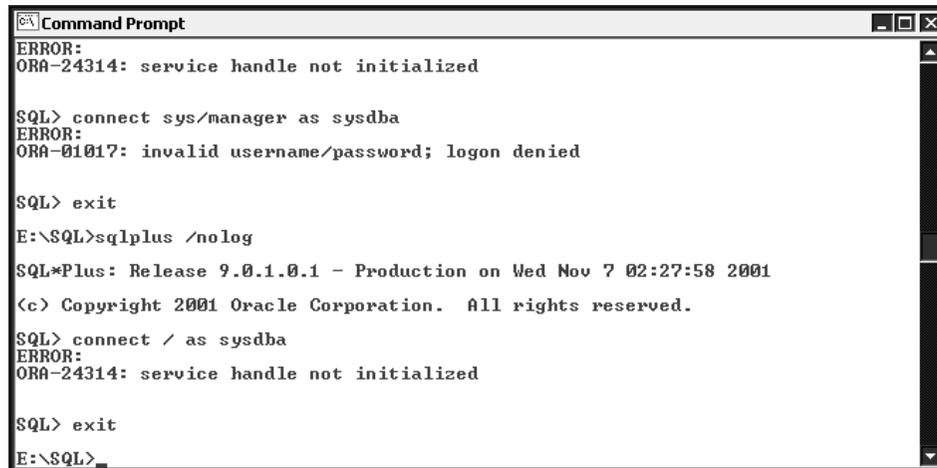
```
sqlplus "/ as sysdba" or
sqlplus /nolog with following command
connect / as sysdba or
connect <username/password> as sysdba
```

When the **connect** command is used consecutive times, an error will occur in both UNIX and Windows environments. The Oracle Metalink Web site has documented these errors for UNIX, mostly for Oracle 8 versions.

This error number is *ora-24314 service handle not initialized* (Figure 6-85). On Windows 2000, the only way to get around this error was by rebooting the system. We were not able to see the SQLplus process running, therefore could not end the process, as recommended by Metalink UNIX documentation on this error. As of the time of writing this redbook, no Metalink documentation for Windows 2000 was found for this error. For updated information, please check:

<http://metalink.oracle.com>

On Linux, we tried to reproduce this same error, but were unsuccessful. Here is the command error window.

A screenshot of a Windows Command Prompt window titled "Command Prompt". The window shows the following text:

```
ERROR:
ORA-24314: service handle not initialized

SQL> connect sys/manager as sysdba
ERROR:
ORA-01017: invalid username/password; logon denied

SQL> exit
E:\SQL>sqlplus /nolog
SQL*Plus: Release 9.0.1.0.1 - Production on Wed Nov 7 02:27:58 2001
(c) Copyright 2001 Oracle Corporation. All rights reserved.

SQL> connect / as sysdba
ERROR:
ORA-24314: service handle not initialized

SQL> exit
E:\SQL>
```

Figure 6-85 Command prompt error window

- ▶ In Oracle9i database there are several, desupported and deprecated parameters. Here are the lists:

#### **Desupported Initialization Parameters in Oracle9i:**

- always\_anti\_join

- always\_semi\_join
- job\_queue\_interval
- optimizer\_percent\_parallel
- hash\_multiblock\_io\_count
- db\_block\_lru\_latches
- db\_block\_max\_dirty\_target
- sort\_multiblock\_read\_count
- db\_file\_direct\_io\_count
- gc\_defer\_time
- gc\_releasable\_locks
- gc\_rollback\_locks
- lm\_locks
- lm\_ress
- use\_indirect\_data\_buffers

#### **Deprecated Initialization Parameters in Oracle9i:**

- rollback\_segments
  - fast\_start\_io\_target
  - transactions\_per\_rollback\_segment
  - log\_checkpoint\_interval
  - db\_block\_buffers
  - buffer\_pool\_keep
  - buffer\_pool\_recycle
- To view the configuration parameters for existing databases, perform one or more of the following statements:
- `select * from v$parameters;` (shows the current value for the session performing the query)
  - `select * from v$system_parameter;` (shows the instance-wide value for the parameter)
  - `select * from v$option;` (options that are installed with the Oracle server)

## **6.6 Configuring startup scripts for Oracle**

This section provides a complete script for starting Oracle safely, after the iSCSI HDD has been attached as a local drive.

In Table 6-2 are all the affected services that need to be changed to start Oracle safely.

Table 6-2 Affected services that need to be changed

| Service name                                 | Startup type |
|----------------------------------------------|--------------|
| AutoExNt                                     | Automatic    |
| Computer browser                             | Manual       |
| server                                       | Manual       |
| Distributed File System                      | Manual       |
| OracleOraHome90TNSLis<br>tener               | Manual       |
| OracleServiceTIERRA1<br>(Required data base) | Manual       |
| OracleMTSRecoveryServi<br>ce                 | Manual       |

The script in Example 6-2 will attempt to connect the iSCSI drives. If it fails to connect, it will not start the Domino Services, and intervention is then required.

Example 6-2 Attempt to connect the iSCSI drives

---

```
net stop "Computer browser" /yes
net stop "server" /yes

:drive_offline
net start "IBMiSCSIcfg"

sleep 10

net start "server"
net start "computer browser"
net start "Distributed File System"

if exist g:\ goto iscsi_yes
Goto ISCSI_Error

 sleep 5

goto drive_offline

:iscsi_yes

net start "OracleOraHome90TNSListener"
net start "OracleServiceTIERRA1"
net start "OracleMTSRecoveryService"
```

---

The Script in Example 6-3 will keep attempting to reconnect the G: drive until it is successful. After connection of the G: drive, the required Oracle services will be started. We recommend that this script is implemented at the end of the instigation, after the drive Letter has been selected.

*Example 6-3 Attempting to reconnect*

---

```
net stop "Computer browser" /yes
net stop "server" /yes

:drive_offline
net start "IBMiSCSIcfg"

sleep 10

net start "server"
net start "computer browser"
net start "Distributed File System"

if exist g:\ goto iscsi_yes

net stop "IBMiSCSIcfg" /yes
net stop "computer browser" /yes
net stop "server" /yes

sleep 5

goto drive_offline

:iscsi_yes

net start "OracleOraHome90TNSListener"
net start "OracleServiceTIERRA1"
net start "OracleMTSRecoveryService"
```

---

**Note:** For more information on scripts refer to Example 2-1 on page 35.



## SAP R/3 considerations for iSCSI

The intent of this chapter is to discuss using the IBM TotalStorage IP Storage 200i device and the Cisco SN5420 IP Storage Router as the I/O subsystem for an IBM @server xSeries Server with Microsoft Windows NT/2000.

Previously in Chapter 5, we discussed DB2 UDB EE V7 implementation specifics. In this chapter, we present specific issues regarding DB2 in an SAP R/3 implementation. In this chapter we provide a “proof of concept” for installation and operational viability only.

Performance/Stress testing, Clustering and/or High Availability models were not the object of this implementation. Additionally, backup and recovery scenarios using Tivoli Storage Manager (TSM) or other third party backup software tools were not the object of this effort. However, we should mention that DB2 backup administration tools and SAP R/3 backup administration tools were used successfully to local iSCSI disk and locally attached tape drives.

It is not the intent of this chapter to reiterate the installation of DB2 UDB or SAP R/3. There are several IBM Redbooks regarding SAP R/3 and DB2 UDB on xSeries Servers and Microsoft Windows NT/2000. Refer to IBM Redbooks or appropriate vendor documentation for these activities at the following Web sites:

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

<http://www.ibm.com/db2>

## 7.1 A general model of SAP R/3

Figure 7-1 is a graphical representation of the SAP R/3 application product. Note that a thorough review of the firm's business requirements should be performed before implementing this product. Additionally, review the prerequisites and considerations from SAP AG, the relevant hardware vendor, the software operating system vendor, and the database software vendor.

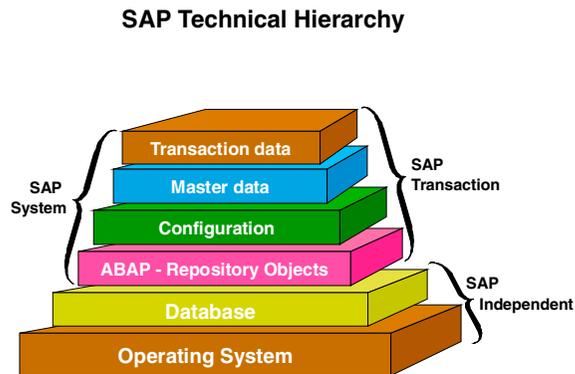


Figure 7-1 SAP technical hierarchy

As is depicted in this figure, the SAP R/3 application product is constructed in such a manner as to support various database and operating system products. SAP AG regularly supplies updated information regarding release levels and supported database, operating systems, and hardware products — contact SAP AG for details or refer to their Web site for the latest information on SAP R/3 applications:

<http://www.sap.com>

**Attention:** SAP AG recommends that only SAP R/3 Certified technicians install the SAP R/3 application.

This redbook is based on an installation of SAP R/3 4.0b and is appropriate to other versions of SAP R/3 4.xx.

The SAP R/3 model allows for three basic types of installations:

- ▶ Single tier
- ▶ Two tier
- ▶ Three tier

Refer to SAP R/3 documentation at <http://www.sap.com> for additional information regarding SAP R/3 implementation specifics.

In this implementation, a two tier model was used. The database instance and central instance application server was installed on an IBM Netfinity 7600 server and the client instances used were IBM Thinkpads. No other SAP application servers were in use.

## 7.2 SAP R/3 certification requirements (iXOS)

This section discusses iXOS certification requirements. This section is available in other redbooks, regarding SAP R/3 implementations. It is presented again here for informational purposes.

iXOS Software AG (iXOS), in close cooperation with SAP, has developed a server hardware certification process for the purpose of investigating the performance and stability of SAP R/3 on Microsoft Windows NT platforms. This certification is essential for all hardware manufacturers. Since 1993, the R/3 NT Competency Center (R3 NTC) of iXOS has exclusively performed the certification process, as an independent assessor, for a large number of Microsoft Windows NT platform hardware vendors.

**Note:** SAP supports the operation of certified R/3 systems only.

### 7.2.1 Certification categories

There are five different hardware certification categories as developed by iXOS in conjunction with SAP. The hardware certification testing is executed by iXOS:

- ▶ **Initial certification:** The first (and only the first) Microsoft Windows NT server offering from a hardware vendor must undergo the initial certification. A very detailed test sequence is performed.
- ▶ **Ongoing certification:** All server offerings for R/3 on Microsoft Windows NT currently offered to the market are subject to the process of the ongoing certification. Twice a year, the ongoing certification tests are repeated on each server offering. The ongoing certification is executed to validate the operation of the system in conjunction with a new release of Microsoft Windows NT or an upgrade of hardware or firmware by the hardware vendor.
- ▶ **Ongoing controller certification:** This level of certification allows for a set of reduced tests to be performed to certify I/O controllers offered by hardware vendors. Once an I/O controller is certified by iXOS, it is certified for use in all the hardware vendor's Microsoft Windows NT server products that have been certified by iXOS.

- ▶ **zSeries and S/390 certification:** This certification requires reduced tests to be performed against an already certified NT platform that is to have access to a DB2 database on an S/390. For this certification category, it is the connectivity type that is certified. Each connectivity type (for example, FDDI, Fast Ethernet, ESCON) must be certified once per hardware vendor.
- ▶ **Outgoing certification:** Hardware platforms no longer being sold for R/3 on Microsoft Windows NT, but still used by customers in a productive environment are subject to an outgoing certification.

To enter the SAP R/3 on Microsoft Windows NT market, a hardware vendor must secure an initial certification for its server platform. The first server offering and all subsequent server offerings that a hardware vendor may supply to the market are subjected to the ongoing certification, until they are no longer offered for SAP R/3 on Microsoft Windows NT.

## 7.2.2 Hardware components

According to the iXOS certification process a hardware platform consists of three different types of components:

- ▶ **Critical components**
  - The chip set: A platform is defined by its chip set and corresponding chipset extensions that enable the data transfer between processor, memory, and I/O. Changing the chip set requires that a Microsoft Windows NT server system must undergo an ongoing certification.
  - The I/O controller: A particular I/O controller must be certified once for each hardware vendor. If an I/O controller has been successfully tested with one Microsoft Windows NT server offering, supplied by a hardware vendor, it is certified with all other iXOS certified Microsoft Windows NT server offerings supplied by the vendor (if supported by the vendor).

The same applies to the certification of S/390 connectivity types.

- ▶ **Peripheral components**

The hardware vendor is obligated to provide a list of all peripheral components associated with the server system to be used in support of SAP R/3 on Microsoft Windows NT. The hardware vendor guarantees function and support of the components listed. If any of the peripheral components are replaced the list is to be updated by the hardware vendor and no new certification is necessary.

Peripheral components are:

- Hard disks
- Memory

- Network adapter
  - Backup device
- **Non-critical components**

All components that are not defined as critical components or peripheral components, are defined as non-critical components. Changing non-critical components does not affect the certification of the platform. Non-critical components are, for example:

- Monitor
- Graphic adapter
- Mouse
- Others

At the time of writing this redbook, the IBM TotalStorage IP Storage 200i and the Cisco SN5420 IP Storage Router I/O storage subsystems have not been certified by iXOS.

For more information, refer to the iXOS Web site at:

<http://www.ixos.com>

**Note:** SAP R/3 certification rules do not require the certification of I/O subsystems. However, iXOS offers, to vendors of I/O subsystems, special tests that validate the stability and measure the I/O performance of the storage solution. At the time of writing this redbook, the iSCSI I/O subsystems have not been certified by iXOS.

## 7.3 Use of Windows AutoExNt-Service and iSCSI

For an SAP R/3 system on Microsoft Windows NT/2000, several system services will need to be available at system startup and operation. With the iSCSI attached disk, it is necessary to use Microsoft Windows AutoExNt-Services. These services are available with Microsoft Windows 2000 and are also supported on Microsoft Windows NT.

### 7.3.1 Windows AutoExNt-Service for SAP R/3 and DB2 UDB

The rationale for the use of Microsoft Windows AutoExNt-Service is to insure disk availability, before the database application and the SAP R/3 application need access to the disk — for the services listed in the table below that are typical for an SAP R/3 system with DB2 UDB.

Table 7-1 presents various systems services.

Table 7-1 AutoExNt-Service

| Service                | Status @sys boot | Startup   |
|------------------------|------------------|-----------|
| DB2 - DB2              |                  | Manual    |
| DB2 - DB2C11           | Started          | Automatic |
| DB2 - DB2DAS00         | Started          | Automatic |
| DB2 Governor           |                  | Manual    |
| DB2 JDBC Applet Server |                  | Manual    |
| IBMiSCSIcftg           | Started          | Manual    |
| SAPC11_00              | Started          | Automatic |
| SAPOSCOL               | Started          | Automatic |
| SAProuter              |                  | Manual    |

**Note:** The SAP R/3 application product allows for SAProuter to be started in the SAP R/3 System *startup profile* and this is recommended. Reference the SAP R/3 documentation for details on how this is implemented. However, it may also be started by Microsoft Windows AutoExNt-Service.

The installation procedure of DB2 UDB creates the appropriate services for DB2. During the installation of SAP R/3, these services are modified as required for use by SAP R/3 and will differ slightly from a DB2 UDB implementation only.

During installation planning, the logical and physical layout of the disk should be done only after the firm's business requirements are known and understood for the given system and application use. The business requirements include, but are not limited to:

- ▶ System availability
- ▶ Backup and recovery requirements
- ▶ Disaster recovery requirements
- ▶ Performance expectations of the end users, for example, SAP R/3 dialog response time
- ▶ Operational and system administrative requirements

## 7.3.2 Implementation of the AutoExNt-Service

At system reboot, various system services are started. To insure that network attached disks are available, the Microsoft Windows AutoExNt-Service is implemented. This package is available in the Microsoft Windows NT 4.0 Resource Kit, at:

<http://www.microsoft.com>  
<http://support.microsoft.com/directory>

The Microsoft Windows AutoExNt-Service is made up of four files:

- ▶ **instexnt.exe**
- ▶ **sleep.exe**
- ▶ **autoexnt.exe**
- ▶ **AutoExNt.bat**

The procedure is relatively simple:

1. Place the first three files in the \Winnt\System32 directory of the server.
2. Execute **instexnt.exe install** to install service.
3. Use **edit.com** or other ASCII editor to create the **AutoExNt.bat** command file.

Example 7-1 is a script that will attempt to connect the iSCSI drives, if it fails to connect it will not start the SAP Services, and intervention is then required.

*Example 7-1 Single attempt startup script*

---

```
net stop "Computer browser" /yes
net stop "server" /yes

:drive_online
net start "IBMiSCSIcfg"

sleep 10

net start "server"
net start "computer browser"
net start "Distributed File System"

if exist g:\ goto iscsi_yes
Goto ISCSI_Error

 sleep 5

goto drive_online

:iscsi_yes
net start "DB2 - DB2"
net start "SAProuter"
```

```
:iscsi_error
```

---

The script in Example 7-2 will keep attempting to reconnect the G: drive until it is successful. After connection of the G: drive, the required SAP services will be started. We recommend that this script is implemented at the end of the instigation, after the drive letter has been selected.

*Example 7-2 Attempting to reconnect*

---

```
net stop "Computer browser" /yes
net stop "server" /yes

:drive_online
net start "IBMiSCSIcfg"

sleep 10

net start "server"
net start "computer browser"
net start "Distributed File System"

if exist g:\ goto iscsi_yes

net stop "IBMiSCSIcfg" /yes
net stop "computer browser" /yes
net stop "server" /yes

sleep 5

goto drive_online

:iscsi_yes

net start "DB2 - DB2"
net start "SAProuter"
```

---

**Note:** For more information on scripts refer to 2.2.2, "Disk drive addressing" on page 34.

## 7.4 Information about disk layouts

Throughout this section, the term, *Windows disk n*, means the *disk n* that we observe in the Windows NT/2000 Disk Administrator window. In a database or clustered configuration this is normally a RAID disk array consisting of two or more physical disks or a subdivision of a RAID disk array (so-called logical drive).

**Note:** This model may lend itself well for clustering environments, when or if clustering becomes a future direction of the firm and the IBM 200i.

The local/internal disks contain the Windows NT/2000 operating system (production and emergency) and the page files and software needed independently of the database log and data files. The recommended data security for the local disks should be achieved with hardware RAID-1. Additionally, if desired and as used in many shops, RAID-1 could be implemented for the database log files on the Tower Model 200 and RAID-5 for the database data files.

The following shared disks, implemented in RAID-1 or RAID-5, are needed:

- ▶ R/3 files (R/3 binaries, SAP control files, etc.)
- ▶ DB2 or database binaries
- ▶ Active database log files
- ▶ Archived database log files
- ▶ Database data files

### 7.4.1 Disk considerations for DB2 and SAP R/3

**Note:** Microsoft clustering is not covered, however some differences are noted. At this time, clustering and/or high availability models are not supported by the iSCSI solutions presented in this redbook.

1. All partitions are NTFS partitions.

**Notes:** The terms “local” or “internal” drives are used interchangeably. Essentially the intent expressed is a drive that is internal to the xSeries Server and not a disk that is internal or in a rack drawer of the Tower Model 100 or Rack Model 200. The IBM Total Storage iSCSI 200i solution drives are represented to the Windows NT/2000 operating system as ‘local’ drives. The term “shared” drive implies a drive that is part of the IBM Total Storage iSCSI 200i solution.

2. Windows NT disk 0 and 1 must have at least 4.5 GB of available space.
3. For R/3 4.xx, SAP “requires” a page file with a size at least four times the amount of physical memory. In a cluster environment, SAP requires a page file with a size at least five times the amount of physical memory. With 3 GB of memory you need 15 GB for page files. In this configuration you need to use 9.1 GB or even 18 GB disks for Windows NT disk 0 and 1.

4. All software that has to be run locally on a node, regardless of the accessibility of any shared disk, and has to be installed on a local disk. Besides the operating system, this is:
  - DB2 UDB software in \SQLLIB
  - Backup software, for example, ADSM client, on local drive
  - The SAP OS collector and SAP binaries on a local drive
  - Paging file space on C: and/or D: drives
5. We highly recommend that you install the emergency/backup for Windows NT/2000. This Windows NT/2000 product is used for backups and the recovery of the production Windows NT.
  - To correct problems with the boot partition you should create a bootable diskette and boot into the emergency Windows NT/2000 and correct the error from there.
  - To correct problems with drivers and the registry in the production Windows NT/2000, you boot into the emergency Windows NT and correct the error from there.

### **SAP R/3 files**

1. The main directory for SAP R/3 is \usr\sap. This directory contains the R/3 executable, R/3 profiles, instance-specific directories, and normally the central transport directory.
2. The shared resource is “SAPMNT” of type “file share”, part of the SAP application server file group, which points to the directory \usr\sap.

Note: In an MSCS environment you cannot store any database related files on the same disk, because this disk is a requirement for the R/3 Application Server.

### **DB2 UDB database files**

1. The database software, DB2 UDB Enterprise Edition, has to be installed locally on the node. The standard directory is \SQLLIB. The databases <SID> and <SID>adm need also to be installed on a shared drive.
2. The database data files, the active database logs, and the archived database logs have to be stored on different disks to avoid that data is lost or at least the amount of data lost is minimized in the case of a RAID array failure.
3. The disk holding the active database logs, is the disk with the largest number of synchronous disk writes. Try to connect this disk to a channel or adapter with low other load. If possible, separate this disk at least from the disk which stores the archived database logs, because every time the “DB2 User Exit”

process is started, a log file is copied from the active database logs to the archive database logs while a new log file data to the active log disk.

**Note:** In a cluster configuration, it is not possible to store the archived database logs on the same drive as the R/3 binaries. The “DB2 user exit” process is always started on the machine running the database instance. If the R/3 server is running on the other node, the “DB2 user exit” process has no access to the “R/3 files” drive. Therefore, you need an additional disk, which is also part of the database cluster group.

4. When a database log file becomes inactive, the RDBMS calls the DB2 UserExit (DB2UE). The DB2UE, in turn, copies the active database log file to the archive database file directory and then deletes the original database log file. Every write access requires that besides the real data, the new parity be written back to disk in synchronous mode. If you use a different controller type in conjunction with RAID-5, make sure that it has a write back cache and that this cache can be activated.
5. The maximum number of RAID-5 arrays to be configured for the database data should not exceed six for a new installation, because the R3 SETUP offers only the distribution of “sapdata1” up to “sapdata6” to different Windows NT disks. You can add additional drives later as needed.

## 7.4.2 Recommended disk layout and configuration

For performance reasons, the database data files should be distributed over six Windows NT disks. One possible optimal distribution of database data on different Windows NT disks is shown in Table 7-2.

**Attention:** Thought should be given to the volume label names. This will aid you with administrative duties and troubleshooting should the need arise. An example volume label name would include information regarding the location of the disk, for example, server or storage device, and general function of the array, for example, sap, db2, etc.

**Note:** Refer to the SAP R/3 documentation for filesystem sizes, because this will depend on the specific version of SAP R/3 in use.

Table 7-2 Filesystem layouts for SAP R/3 on Microsoft Windows NT/2000

| Drive letter | Disk number | Filesystem                                                       | Description                                                           | Raid level | Location    |
|--------------|-------------|------------------------------------------------------------------|-----------------------------------------------------------------------|------------|-------------|
| C            | 0           | \winnt                                                           | Windows OS                                                            | RAID-1     | Local       |
| D            | 1           | paging                                                           | paging                                                                | RAID-1     | Local       |
| E            | 2           | \usr\sap<sid><br>\usr\sap\trans<br>\sqlib<br>\other applications | SAP binaries<br>DB2 binaries<br>Backup software<br>other applications | RAID-1     | Local       |
| F            | 3           | \db2\<sid>\log_dir                                               | DB2 active logs                                                       | RAID-1     | 200i device |
| G            | 4           | \db2\<sid>\log_archive<br>\db2\<sid>\sapreorg                    | DB2 archive logs<br>DB2 reorg space                                   | RAID-1     | 200i device |
| H            | 5           | \db2\<sid>\sapdata1                                              | Tablespaces                                                           | RAID-5     | 200i device |
| I            | 6           | \db2\<sid>\sapdata2                                              | Tablespaces                                                           | RAID-5     | 200i device |
| J            | 7           | \db2\<sid>\sapdata3                                              | Tablespaces                                                           | RAID-5     | 200i device |
| K            | 8           | \db2\<sid>\sapdata4                                              | Tablespaces                                                           | RAID-5     | 200i device |
| L            | 9           | \db2\<sid>\sapdata5                                              | Tablespaces                                                           | RAID-5     | 200i device |
| M            | 10          | \db2\<sid>\sapdata6                                              | Tablespaces                                                           | RAID-5     | 200i device |

**Important:** The default installation of SAP R/3 with DB2 UDB mixes data tablespaces and index tablespaces as well as DB2 system tablespaces in the same directories. On-going performance monitoring may call for the separation of data, index, and system tablespaces to their own disk arrays.

## 7.5 SAPRouter configuration and recommendations

SAP OSS Note 30289 provides a thorough explanation of SAPRouter's use, purpose, and features. Because the IBM TotalStorage IP Storage 200i uses transport of SCSI storage over IP (*SAN over IP*) and delivers this storage over Ethernet, a special mention of this important SAP R/3 function and the **saprountab** file is discussed in this section.

It is desirable to isolate the IP I/O traffic from the IP communications traffic used by SAP R/3. This may be accomplished by several different methods. An additional step is recommended here, specifically in the **saproustab** definition file. SAProuter may be used to either permit or deny access to the SAP system.

In the example below, 9.24.105.202 is the host IP address, and 9.24.105.101 is the IP host address of the data port on the IBM TotalStorage IP Storage 200i. The specific entries will be determined by the network model in use. The general recommendation is to isolate the data traffic from the user traffic by use of switching and/or subnetting. A sample **saproustab** is shown in Example 7-3.

*Example 7-3 Sample saproustab*

---

```
this is a sample saproustab:

D host3
iSCSI data path
D 9.24.105.101 9.24.105.202
P 9.24.*.* 9.24
P 9.24.105.202 localhost 3299

deny all routes from host3
deny all routes from 9.24.105.101 to 9.24.105.202 (iSCSI storage)
permit all routes to/from addresses matching 9.24
permit all routes from host to local host

```

---

## 7.6 Networking recommendations

The `\winnt\system32\drivers\etc\hosts` file is the primary networking file in use (Example 7-4).

*Example 7-4 Hosts file*

---

```
Contents of hosts file
Following are the contents of hosts file used during the installation:
hosts file for R/3 DB2 UDB cluster installation
127.0.0.1 localhost
9.24.105.202 nf7600 SAPTRANSHOST
9.24.105.101 iSCSIhosta #data port
10.0.0.2 iSCSIhostb #maintenance port
```

---

## 7.7 Cisco SN5420 IP Storage Router and IBM ESS

The IBM Enterprise Storage Server (ESS) allows for RAID-5 LUNS and/or JBOD disk configurations (just a bunch of disks), as of the time of writing this redbook. RAID-1 is not generally supported. The IBM ESS disk ranks can be configured as RAID 5 or non-RAID arrays (just a bunch of disks). Therefore RAID-1 configurations are not generally possible. A determination will need to be made regarding this issue. If RAID-1 is the preferred model, then an internal disk may be the best choice. If RAID-5 arrays are acceptable, then logical external drives may be configured on the IBM ESS.



## Data migration considerations for DB2

In this appendix we provide an overview of migration from DB2 on one platform to DB2 on a different platform. In the case of migration from other vendor's RDBMS to DB2, we suggest you use application specific tools as necessary. When you move data between DB2s on different platforms, use **db2move** and **db2look** tools. Always consult with the vendors that may be using an application that is integrated with DB2. Examples of such vendors would be R/3 and Siebel. In many cases, these other vendors may have developed specific tools pertaining to database migrations.

We describe how to copy database objects (for example, tables, and tablespaces) between DB2 on the same or a different operating system. For this purpose you can use db2move (database movement) and db2look (DB2 statistics and DDL extraction) tools.

Procedures described in the following sections are applicable to different versions of DB2 on operating systems where DB2 UDB for UNIX, Windows, and OS/2 are available.

## Moving database objects

In this section we describe some examples of moving database objects using DB2 tools.

### DB2MOVE EXPORT and DB2MOVE IMPORT

If you only need to copy tables between two DB2s, or if you want to import tables into a new database (the other option is to LOAD), follow the steps described below. This approach is used if your tables are not already created. Import is generally slower than load, so consider using **db2move load**, described in subsequent pages.

1. Log on to the platform from which you want to export tables as a user, with the required authorization for export.
2. Use the **db2move** tool to export all tables and their contents:

```
db2move <database-alias> export
```

To export a subset of tables, use the options described in *DB2 UDB Command Reference*.

You will get the following files generated by EXPORT:

- **tabnnn.ixf**: The exported PC/IXF file of a specific table that is needed for creation of tables.
  - **db2move.lst**: The list of all files can be used as input to the db2move IMPORT or LOAD action.
  - **tabnnn.msg files**: The export message file of the corresponding table.
  - **export.out**: The summarized result of the export action.
  - **tabnnnc.yyy**: The exported LOB files of a specific table when your tables contain LOBs.
  - **system.msg**: The message file containing system messages when your table contains LOBs.
3. Transfer db2move.lst and all tabnnn.ixf files (also tabnnnc.yyy, if your tables contain LOBs) to your target server by using, for example, FTP. You do not need to copy tabnnn.msg files, because they include only messages produced by export.
  4. Log on to the target operating system as a user, with the required authorization for import.
  5. Create a new database:  

```
create db <database-alias>
```

6. Use db2move to import tables into an existing database:

```
db2move <database-alias> import
```

## DB2MOVE EXPORT, DB2LOOK, DB2MOVE LOAD

If you need to load the contents of the tables into a database, or if you want to copy objects other than just tables, you have to use db2look and db2move. In order to use the db2move load command, the target tables must already exist.

Use this approach if you have a larger amount of data, because load is faster than import. For creating tables and other database objects, use the output created from db2look.

1. Log on as a user, with the required authorization for export.
2. As in the previous example, use db2move to export all tables and their contents:

```
db2move <database-alias> export
```

3. Use db2look to get DDL statements for creation of all required objects:

```
db2look -d <database-alias> -a -e -l -x -o <outfile-name>
```

You should have the information about all objects created by all users in the outfile:

- DDL statement for database objects such as tables, views, automatic summary tables, indexes, triggers, user defined functions, etc.
  - Authorization DDL (for example, GRANT statement)
  - DDL statements for user-defined tablespaces, user-defined bufferpools, and user-defined nodegroups
4. Transfer <file name>.sql, db2move.lst, and all tabnnn.ixf files (as well as tabnnnc.yyy, if your tables contain LOBs) to your target operating system by using, for example, FTP. Be careful to upload the <file name>.sql file in text mode, not binary mode. You do not need to copy tabnnn.msg files, because they include only messages produced by export.
  5. Log on onto the target operating system as a user, with the required authorization for load.
  6. Create a new database:

```
create db <database-alias>
```
  7. Edit <outfile-name>.sql and correct the CONNECT TO <database name> statement to match your new database name.

8. Run the <outfile-name>.sql file containing DDL to create all database objects. The file is created by db2look (see Step 3):

```
db2 -tvf <outfile-name>.sql
```

9. Use db2move to load tables into an existing database:

```
db2move <database-alias> load
```

**Note:** User-defined functions are not copied by using db2move and db2look. You have to reinstall UDFs manually.

# Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

## IBM Redbooks

For information on ordering these publications, see “How to get IBM Redbooks” on page 308.

- ▶ *Tuning IBM e(logo)server xSeries Servers for Performance*, SG24-5287
- ▶ *Using LDAP for Directory Integration: A Look at IBM SecureWay Directory, Active Directory, and Domino*, SG24-6163
- ▶ *Getting the Most From Your Domino Directory*, SG24-5986
- ▶ *LDAP Implementation Cookbook*, SG24-5110
- ▶ *A Roadmap for Deploying Domino in the Organization*, SG24-5617
- ▶ *Database Performance on AIX in DB2 and Oracle Environments*, SG24-5511
- ▶ *Implementing SAP R/3 4.5B Using Microsoft Cluster Server on IBM Netfinity Servers*, SG24-5170
- ▶ *Optimizing IBM Netfinity Servers for SAP R/3 and Windows NT*, SG24-5219
- ▶ *DB2 UDB V7.1 Performance Tuning Guide*, SG24-6012
- ▶ *IP Storage Networking: IBM NAS and iSCSI Solutions - Update*, SG24-6240

## Other resources

These publications are also relevant as further information sources:

- ▶ *Microsoft Exchange 2000 Server Resource Kit*. Microsoft Press, September 2000. ISBN: 0735610177
- ▶ *Microsoft Windows NT Server Resource Kit: For Windows NT Server Version 4.0*. Microsoft Press, November 1996. ISBN: 1572313447

## Referenced Web sites

These Web sites are also relevant as further information sources:

- ▶ <http://www.snia.org/English/Overview>  
Storage Networking Industry Association

- ▶ <http://www.storage.ibm.com>  
IBM Storage
- ▶ <http://java.sun.com>  
Source for Java Technology by Sun Systems
- ▶ <http://kernel.org>  
Linux Kernel Archives
- ▶ <http://www.cisco.com/cgi-bin/tablebuild.pl/sn5420>  
Cisco SN5420 product page
- ▶ <http://www.lotus.com/domino>  
Domino product page
- ▶ <http://www.microsoft.com/exchange>  
MS Exchange
- ▶ <http://www.microsoft.com/servers/downloads>  
Microsoft support download page
- ▶ <http://www.pc.ibm.com/us/eserver/xseries/index.html>  
IBM Netfinity support page
- ▶ <http://www.ibm.com/db2>  
IBM DB2 UDB product page
- ▶ <http://support.microsoft.com/directory>  
Microsoft Product Support Services
- ▶ <http://www.sap.com>  
SAP AG Official Web site
- ▶ <http://www.ixos.com>  
SAP R/3 Certification
- ▶ <http://www.oracle.com>  
Oracle home page
- ▶ <http://metalink.oracle.com>  
Oracle MetaLink

## How to get IBM Redbooks

Search for additional Redbooks or redpieces, view, download, or order hardcopy from the Redbooks Web site:

[ibm.com/redbooks](http://ibm.com/redbooks)

Also download additional materials (code samples or diskette/CD-ROM images) from this Redbooks site.

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.

## **IBM Redbooks collections**

Redbooks are also available on CD-ROMs. Click the CD-ROMs button on the Redbooks Web site for information about all the CD-ROMs offered, as well as updates and formats.



# Special notices

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 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.

The following terms are trademarks of other companies:

Tivoli, Manage. Anything. Anywhere., The Power To Manage., Anything. Anywhere., TME, NetView, Cross-Site, Tivoli Ready, Tivoli Certified, Planet Tivoli, and Tivoli Enterprise are trademarks or registered trademarks of Tivoli Systems Inc., an IBM company, in the United States, other countries, or both. In Denmark, Tivoli is a trademark licensed from Kjøbenhavns Sommer - Tivoli A/S.

C-bus is a trademark of Corollary, Inc. in the United States and/or other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and/or other countries.

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

PC Direct is a trademark of Ziff Communications Company in the United States and/or other countries and is used by IBM Corporation under license.

ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel Corporation in the United States and/or other countries.

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

SET, SET Secure Electronic Transaction, and the SET Logo are trademarks owned by SET Secure Electronic Transaction LLC.

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

# Index

## Symbols

/etc/services file 241

## Numerics

200i 22

5420 Storage Router 63

## A

Access Control Lists (ACLs) 67

active database logs 298

Active Directory 118

Active Directory Connector 118

active logs 166

ADC 114

AIX 93

archive 249

archive logs 166

archived database logs 298

archiver process 177

ARCn 177

ARP (Address Resolution Protocol) 26

assign client 51

authorization DDL 305

AutoExNt 34, 35, 131, 170

automatic memory tuning 189

automatic resource management 186

## B

backup 159

bandwidth 150

basic disks 59

binaries on external 172

binaries on internal 172

bitmap join index 189

Block I/O 11

Block I/O versus File I/O 11

bottleneck 152

buffer cache 176, 185

## C

cache 24

certification categories 291

certification requirements 291

checkpoint process 160, 177

checkpoints 160

circular logging 166

Cisco 21, 63

Cisco SN5420 293, 302

I/O subsystem 289

implementation 106

introduction 64

technical specifications 67

Cisco SN5420 IP Storage Router 302

Cisco SN5420 Storage Router 63

CKPT 177

client initiator 53

clustering 71, 289, 297

command line interface (CLI) 67, 69

Competency Center 291

console port 73

control files 160, 178

control port 42

corruptions 159

CPU 113, 150, 151

creating a database 108, 169

## D

data dictionary 156, 163

Data Guard 183

data migration 303

data port 32, 46

database 156

Database Configuration Assistant (DBCA) 243, 258

database datafile 178

database managed space 165, 166

database self-management 187

database writer process 177

DB2 UDB 155, 293

DB2 UDB concepts 161

DB2 UDB database files 298

DB2 UDB products 161

DB2 UserExit (DB2UE) 299

db2look 303, 305

- db2move 303, 304
- db2move load 304
- db2nodes.cfg 164
- DBCA 243, 246
- DBWn 177
- DDL 182
- DDL extraction 303
- deinstall 200
- disaster recovery 183
- disk 150
- disk administrator 296
- disk layout 299
- disk management 59
- dispatcher processes 177
- DML 182
- DMS 166
- DMS tablespace 166
- DNS 114, 135
- domain 247
- domainprep 122
- Domino Administrator 141
- Domino Server setup 139
- duplex SC connector 69
- duplexed 160
- DWORD 127
- dynamic disks 59

## E

- e-mail user account 129
- enhanced recovery manager 188
- Enterprise Edition (EE) 161
- enterprise mail 111
- ESCON 292
- ESCON Director 107
- Ethernet 12, 28
- Ethernet network management access 70
- Exchange 2000 111
- EXP 22
- Extended Enterprise Edition (EEE) 161

## F

- failed state diagnostics 185
- Fast Ethernet 29, 292
- Fibre Channel (FC) 3
- Fibre Channel infrastructure 67
- Fibre Channel port 69
- Fibre Channel SAN 10, 12
- File I/O 11

- FixPak 155
- forestprep 120
- full object 185

## G

- Gigabit 28, 69
- Gigabit Ethernet 11, 25, 65
- global database domain 206, 247
- graphical user interface 91
- GUI 70, 91

## H

- HA 65, 70
- HA Ethernet 69
- high availability 289
- high availability (HA) interface 81
- high availability option 107
- hotfix Q302895 93
- HP UX 93
- HTML 180
- HTTP 117, 211
- human error safeguard 184
- HyperTerminal 57

## I

- I/O Inbound/Outbound 301
- IBM Enterprise Storage Server (ESS) 107, 302
- IBM iSCSI client configuration 56
- IBMiSCSIcfg 35
- IETF 3
- ifconfig
  - IP address under Linux 34
- IIS 115, 135, 137
- index 157
- index-organized tables 189
- initial start up screen 74
- initialization file 161
- initiators 19, 53
- instance 163
- instext 34, 131
- Integrated Advanced System Management Processor 25
- Internet Engineering Task Force (IETF) 2, 3
- Internet Information Services 117
- Internet SCSI (iSCSI) 3
- IP 65
- IP address assignment

- two methods 26
- IP Internet protocol 301
- IP Storage 200i
  - departmental model
    - 4125 Model 210 24
  - technical specifications
    - model 110 24
    - model 210 24
  - workgroup model
    - 4125 Model 110 24
- IP Storage 200i optional features
  - 1 GB 133 MHz ECC SDRAM RDIMM (#3404) 29
  - 10/100/1000 Gigabit Ethernet Copper Adapter (#3303) 29
  - 512 MB 133 MHz ECC SDRAM RDIMM (#3403) 29
  - 5600 Tower-To-Rack Kit (#3601) 29
  - Gigabit Ethernet SX Adapter (#3302) 28
- ipconfig
  - IP address under Windows 34
- iSCSI 3, 21, 84, 128, 132, 243
- iSCSI driver installation 97
- iSCSI login 15
- iSCSI PDU 15
- iSCSI protocol data unit (iSCSI PDU) 15
- iSCSI SAN 11
- iSCSI targets 86
- ISO 9660 194
- iXOS 291

## J

- JDK 231
- job queue processes 177

## L

- LAN 145
- LanmanServer 34, 35
- large objects (LOBs) 165, 166
- latch contention 186
- LDAP 135, 181
- level 2 cache 24
- LGWR 177, 178
- Light Path Diagnostics 25
- Linux 21, 93, 220
- Linux configuration 220
- Linux installation 220
- listener 179

- listener process for incoming requests 179
- LMS 177
- LOB Large Object Data 304
- Local Net Service Name configuration 272
- localhost 117, 128
- lock manager server process 177
- log writer process 177
- log\_archive\_dest 220
- log\_archive\_dest\_n 220
- log\_archive\_format 220
- log\_archive\_start 220
- logfiles 160
- logical drive 296
- logical nodes 164
- logical unit numbers (LUNs) 67
- LogMiner 184
- logretain 166
- LoopID 86
- Lotus Domino 111, 137

## M

- machine 164
- mail storage group 128
- mailbox store 128
- mailboxes 151
- maintenance port 34
- management console 32
- management GUI
  - Internet Explorer 33
  - Netscape 32
- management ports 69
- mask 74
- materialized view mechanisms 189
- McDATA Enterprise Fibre Channel Director 107
- mean time to recover 188
- memory 25, 150, 152, 249
- memory structures 176
- MGMT 70
- Microsoft management console 134
- MMC 134
- MPP 163
- MSCS Microsoft Cluster Server 298
- MTTR 184, 188
- multiple inserts 189
- multiplexing 166

## N

- NAS 3, 5, 6

Net Service Name Configuration (NSNC) 273  
Netfinity 28  
Netware 93  
network 25, 113, 150, 152  
network and distributed database performance 186  
networking 301  
NNTP 114, 115, 117, 137  
node 164  
nodegroups 164, 165  
NTFS 116, 137

## O

ODCA 212  
OEM 188  
OLAP 182  
OLTP 186  
ONCA Net Service Name Configuration 272  
ONCA Net Service Name Configuration (NSNC)  
268  
online data evolution 185  
Oracle 173  
Oracle advanced queuing 177  
Oracle data mining 182  
Oracle Database Configuration Assistant (ODCA)  
212, 239  
Oracle Enterprise Manager (OEM) 188, 223  
Oracle installation 199, 223  
Oracle Net Assistant 223  
Oracle Net Configuration Assistant 241, 267  
Oracle Net Configuration Assistant (ONCA) 268,  
271  
Oracle parallel server 174  
Oracle post-installation 216, 241  
Oracle pre-installation 198, 221  
Oracle script errors 235  
Oracle system identifier 230  
Oracle Universal Installer (OUI) 196, 199, 223, 261  
Oracle9i 173  
Oracle9i architecture 175  
OS Operating System 300  
OSDBA 221  
OSDBA group 222  
OSS 300  
OUI 202, 226  
OUI Configuration Tools 212  
OUI Net Configuration Assistant 269  
Outlook Express 131

## P

page file 152  
parameter file 178  
partitioned database 162  
password 67  
PCI 22  
performance matrix 168  
peripheral component storage 22  
permission errors 235  
persistent initialization parameters 252  
Personal Edition (PE) 161  
PGA 176  
physical nodes 164  
PL/SQL optimization management 186  
PMON 151, 177  
point in time 163  
POP3 128, 131  
process monitor 177  
processor 24, 152  
Program Global Area 178  
proxy server 32  
PVIEWER 152

## Q

quality of service (QoS) 108  
query 156  
queue monitor processes 177  
quiesce database 185

## R

RAC 174  
RAID 24, 113, 168, 296  
RAM 113  
rapid crash recovery 185  
raw device 165  
RDBMS 155, 160  
RDBMS concepts 156  
RDBMS Relational Database Management System  
299  
real application clusters 174  
RECO 177  
recoverer process 177  
recovery history file 167  
recovery logs 156, 166  
recovery points 168  
recovery process 159  
Redbooks Web site 308  
Contact us xx

redo log files 178  
redo log group 253  
redo logs 174, 176  
register 146  
registration 147  
Relational Database Management Systems (RDBMS) 156  
remote database 243  
RJ-45 70  
RMAN 188  
roll-forward recovery 166  
router 65  
ROWID 185  
RS-232 69

## S

S/390 292  
SAN 300  
SAP 290  
SAP R/3 files 298  
SAProuter 301  
saprountab 300, 301  
scalability 66  
scalable session state management 186  
scripts 131, 148, 170, 234  
SCSI Small Computer Systems Interface 289  
SDRAM 22  
server audience 140  
ServeRAID 24, 25, 26, 27  
setup.exe 53  
SGA 249  
shared pool 176  
shared\_pool\_size 241  
shutdown script 134  
SID 179, 206, 247  
Siebel 303  
Simple Network Management Protocol (SNMP) 67  
sleep 34  
SMON 177  
SMP 164  
SMS 165  
SMS tablespace 165  
SMTP 117, 128, 131  
SNIA 3  
SNMP 70  
Solaris 93  
SQL 178  
SQL\*Net 179

SQL\*Plus 181, 242  
startup profile 294  
Storage Networking Industry Association (SNIA) 3  
Storage Wide Area Networks (SWAN) 10  
streamlined operational management 187  
symmetrical multiprocessor 164  
system availability matrix 168  
system errors 166  
system fault recovery 184  
System Global Area (SGA) 176, 249  
system identifier 206  
system managed space 165  
system monitor process 177  
system tablespaces 300

## T

tables 156  
tablespaces 157, 158  
TCO 67  
TCP/IP 115, 137  
telnet 117, 128  
telnet session 117  
terminal emulator 69  
Tivoli Netview 10  
Tivoli Storage Manager (TSM) 163, 289  
TNS 179  
TNS layer 179  
TotalStorage 22  
troubleshooting 39, 61  
TSM 163  
tuning  
    aspects 173

## U

UDF user defined functions 306  
unique name 206  
unshielded twisted pair (UTP) 70  
userexit 166  
UTP 70

## V

VLUN 42, 47, 48

## W

Windows 2000 93  
Windows NT 93  
workgroup 161

World Wide Node Number (WWNN) 86  
World Wide Port Name (WWPN) 86

## **X**

xSeries 28

## **Z**

Zero Data Log Transport 183

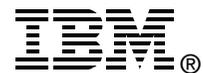


# Using iSCSI: Solutions' Planning and Implementation

(0.5" spine)  
0.475" <-> 0.875"  
250 <-> 459 pages







# Using iSCSI Solutions' Planning and Implementation

**Plan and implement  
for Oracle9i, DB2  
UDB, SAP R/3,  
Domino, and  
Exchange**

**Install and configure  
IPStorage 200i and  
Cisco SN5420**

**Learn the inner  
workings of iSCSI  
protocol**

iSCSI is a proposed industry-standard that allows SCSI block I/O protocols (commands, sequences and attributes) to be sent over a network using the popular TCP/IP protocol. This book covers the concepts of the iSCSI protocol, how it works, how it got started, where iSCSI is right now and where it is going as a standard.

IBM TotalStorage IP Storage 200i and Cisco SN5420 Storage Router are the first products to exploit this technology. These products use different approaches. Cisco SN5420 uses your existing SAN storage, while IBM IP Storage 200i comes integrated with its own storage. This book details the installation and configuration of the IBM TotalStorage IP Storage 200i and the Cisco SN5420 Storage Router.

This IBM Redbook is intended for IBM Professionals, Customers, and Business Partners who are tasked to install and implement an iSCSI storage network. This book will help you understand what considerations to make when using iSCSI. We also document some suggestions, recommendations, hints and tips to help your in iSCSI implementation.

This book covers the installation steps involved for applications, such as Oracle9i for Windows 2000 and Linux, Microsoft Exchange 2000, Lotus Domino for Windows 2000, SAP R/3 and DB2 UDB, when using iSCSI storage. It also covers useful scripts that can be used when using iSCSI storage with these applications to ensure proper startup of the applications.

## **INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION**

### **BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE**

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

**For more information:  
[ibm.com/redbooks](http://ibm.com/redbooks)**