


International Technical Support Organization

SG24-4576-00

**IBM PC Server and Novell NetWare
Integration Guide**

December 1995



IBM

**International Technical Support Organization
Raleigh Center**

Take Note!

Before using this information and the product it supports, be sure to read the general information under "Special Notices" on page xv.

First Edition (December 1995)

This edition applies to IBM PC Servers, for use with an OEM operating system.

Order publications through your IBM representative or the IBM branch office serving your locality. Publications are not stocked at the address given below.

An ITSO Technical Bulletin Evaluation Form for reader's feedback appears facing Chapter 1. If the form has been removed, comments may be addressed to:

IBM Corporation, International Technical Support Organization
Dept. HZ8 Building 678
P.O. Box 12195
Research Triangle Park, NC 27709-2195

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

© **Copyright International Business Machines Corporation 1995. 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.

Abstract

This document describes the procedures necessary to successfully implement Novell NetWare on an IBM PC Server platform. It describes the current IBM PC Server line and discusses the technology inside the machines. It outlines step-by-step procedures for installing both NetWare V3.12 and V4.1 using both IBM ServerGuide and the original product media. It has a detailed section on performance tuning. It covers IBM's NetFinity systems management tool, which ships with every IBM PC Server and IBM premium brand PC.

This document is intended for IBM customers, dealers, systems engineers and consultants who are implementing NetWare on an IBM PC Server platform.

A basic knowledge of PCs, file servers, DOS, and NetWare is assumed.

(212 pages)

Contents

Abstract	iii
Special Notices	xv
Preface	xvii
How This Document is Organized	xvii
Related Publications	xviii
International Technical Support Organization Publications	xviii
ITSO Redbooks on the World Wide Web (WWW)	xviii
Acknowledgments	xix
Chapter 1. IBM PC Server Technologies	1
1.1 Processors	1
1.1.1 Clock Rate	1
1.1.2 External Interfaces	2
1.1.3 Processor Types	2
1.2 Multiprocessing	3
1.3 Memory	3
1.3.1 Caches	4
1.3.2 Memory Interleaving	6
1.3.3 Dual Path Buses	7
1.3.4 SynchroStream Technology	8
1.4 Memory Error Detection and Correction	9
1.4.1 Standard (Parity) Memory	9
1.4.2 Error Correcting Code (ECC)	9
1.4.3 Error Correcting Code-Parity Memory (ECC-P)	10
1.4.4 ECC on SIMMs (EOS) Memory	11
1.4.5 Performance Impact	12
1.4.6 Memory Options and Speed	12
1.5 Bus Architectures	12
1.5.1 ISA Bus	13
1.5.2 EISA Bus	13
1.5.3 Micro Channel Bus	13
1.5.4 PCI Bus	15
1.6 Disk Subsystem	16
1.6.1 Hard Disk Interfaces	16
1.6.2 SCSI Technology	17
1.6.3 SCSI Adapters	19
1.6.4 Hard Disk Drives	21
1.6.5 RAID Technology	22
1.6.6 RAID Classifications	23
1.6.7 Recommendations	31
1.7 LAN Subsystem	31
1.7.1 Shared RAM Adapters	31
1.7.2 Bus Master Adapters	32
1.7.3 PeerMaster Technology	33
1.8 Security Features	34
1.8.1 Tamper-Evident Cover	35
1.8.2 Secure I/O Cables	35
1.8.3 Passwords	35
1.8.4 Secure Removable Media	36

1.8.5	Selectable Drive Startup	37
1.8.6	Unattended Start Mode	37
1.9	Systems Management	37
1.9.1	DMI	37
1.9.2	SNMP	39
1.9.3	NetFinity	41
1.9.4	SystemView	45
1.10	Fault Tolerance	47
1.10.1	NetWare SFT III	48
1.11	Uninterruptible Power Supply (UPS)	52
1.11.1	APC PowerChute	53
 Chapter 2. IBM PC Server Family Overview		55
2.1	IBM PC Server Model Specifications	56
2.1.1	IBM PC Server 300	56
2.1.2	IBM PC Server 310	57
2.1.3	IBM PC Server 320 EISA	58
2.1.4	IBM PC Server 320 MCA	59
2.1.5	IBM PC Server 500	60
2.1.6	IBM PC Server 520 EISA	61
2.1.7	IBM PC Server 520 MCA	62
2.1.8	IBM PC Server 720	63
 Chapter 3. Hardware Configuration		65
3.1	The Setup Program	67
3.1.1	Main Menu	67
3.1.2	Advanced Menu	68
3.1.3	Security	71
3.2	EISA Configuration Utility	74
3.3	SCSI Select Utility Program	78
3.4	System Programs	82
3.4.1	Starting From the System Partition	83
3.4.2	Starting From the Reference Diskette	84
3.4.3	Main Menu Options	84
3.4.4	Backup/Restore System Programs Menu	86
3.4.5	Set Configuration Menu	86
3.4.6	Set Features Menu	91
3.4.7	Test the Computer	97
3.4.8	More Utilities Menu	98
3.4.9	Advanced Diagnostic Program	99
3.5	RAID Controller Utility	101
3.5.1	Drive Information	103
3.5.2	Formatting the Disks	104
3.5.3	Defining a Hot-Spare Disk	106
3.5.4	Creating a Disk Array	108
3.5.5	Defining Logical Drives	109
3.5.6	Setting the Write Policy	112
3.5.7	Initializing the Array	112
3.5.8	Backup/Restoring the Configuration	112
 Chapter 4. Novell NetWare Installation		115
4.1	ServerGuide Overview	115
4.2	Starting ServerGuide	115
4.3	Installing NetWare 4.1 with ServerGuide	118
4.4	Installing NetWare 3.12 with Diskettes	127

4.4.1	Hardware Requirements	127
4.4.2	Software Requirements	127
4.4.3	Information Requested at Time of Installation	127
4.4.4	Installation Files	127
4.4.5	Installation Procedure	128
4.5	Installing NetWare 4.1 with the Original CD-ROM	135
4.5.1	Hardware Requirements	135
4.5.2	Software Requirements	135
4.5.3	Installation Procedure	135
4.6	NetFinity Services for NetWare	144
4.6.1	System Requirements	144
4.6.2	Installing NetFinity Services for NetWare	144
4.7	The RAID Administration for NetWare Utility	151
4.7.1	Installing the Utility	151
4.8	Hard Disk Failure Simulation	152
4.8.1	Simulating with a Hot Spare Drive	153
4.8.2	Simulating without a Hot Spare Drive	160
Chapter 5.	Performance Tuning	167
5.1	Hardware Tuning	167
5.1.1	General Performance Characteristics	168
5.2	Performance Analysis Tools	172
5.2.1	DatagLANce	172
5.2.2	NetWare Monitoring Tools	173
5.3	Tuning NetWare	175
5.3.1	Disk Subsystem	176
5.3.2	Network Subsystem	182
5.3.3	System Memory	185
5.3.4	Memory Requirements	187
5.3.5	System Processor	187
Appendix A.	EISA Configuration File	189
Appendix B.	Hardware Compatibility, Device Driver, and Software Patch Information	199
B.1	Finding Compatibility Information on the World Wide Web	199
B.2	Finding Device Drivers on the World Wide Web	200
B.3	Finding Software Patches on the World Wide Web	201
Appendix C.	Configuring DOS CD-ROM Support	203
C.1	Installing CD-ROM Support for PCI Adapters.	203
C.2	Installing CD-ROM Support for Adaptec Adapters	203
C.3	Installing CD-ROM Support for Micro-Channel Adapters	203
List of Abbreviations		205
Index		207

Figures

1.	SMP Shared Secondary Cache	5
2.	SMP with Dedicated Secondary Cache	6
3.	Two-Way Interleaved Memory Banks	7
4.	Dual Path Bus Implementation	8
5.	ECC Memory Operation	10
6.	ECC-P Memory Implementation	11
7.	Micro Channel - Basic Data Transfer (20 MBps)	14
8.	Micro Channel - Data Streaming Transfer (40 MBps)	14
9.	Micro Channel - Data Streaming Transfer (80 MBps)	15
10.	SCSI Disk Interface	17
11.	RAID-0 (Block Interleave Data Striping without Parity)	24
12.	RAID-1 (Disk Mirroring)	25
13.	RAID-1 (Disk Duplexing)	25
14.	RAID-1 Enhanced, Data Strip Mirroring	26
15.	RAID-6,10 - Mirroring of RAID 0 Drives	26
16.	RAID-2 (Bit Interleave Data Striping with Hamming Code)	27
17.	RAID-3 (Bit Interleave Data Striping with Parity Disk)	28
18.	RAID-4 (Block Interleave Data Striping with One Parity Disk)	29
19.	RAID-5 (Block Interleave Data Striping with Skewed Parity)	29
20.	NetFinity Services Folder	41
21.	IBM PC Server Family of Products	55
22.	Hardware Configuration Steps	66
23.	PC Server 320 Setup Program - Main Menu	67
24.	PC Server 320 Setup Program - Advanced Menu	68
25.	PC Server 320 Setup Program - Boot Options Menu	69
26.	PC Server 320 Setup Program - Integrated Peripherals Menu	70
27.	PC Server 320 Setup Program - Security Menu	71
28.	EISA Configuration Utility - Main Panel	74
29.	EISA Configuration Utility - Steps	74
30.	EISA Configuration Utility - Step 2	75
31.	EISA Configuration Utility - Move Confirmation Panel	75
32.	EISA Configuration Utility - Step 3	76
33.	EISA Configuration Utility - Step 4	77
34.	IBM PC Server SCSISelect Utility Program - Main Menu	78
35.	IBM PC Server SCSI Select Utility Program - Host Adapter Settings	78
36.	PC Server 320 SCSI Select Utility Program - SCSI Device Configuration	80
37.	PC Server 320 SCSISelect Utility Program - Advanced Configuration	81
38.	PC Server 320 SCSISelect Utility Program - DASD Information	82
39.	System Programs - Main Menu	84
40.	System Programs - Backup/Restore System Programs Menu	86
41.	System Programs - Set Configuration Menu	87
42.	System Programs - View Configuration Screen	88
43.	Set Configuration - Memory Map	90
44.	Set Configuration - SCSI Device Configuration	91
45.	Set Features Menu	92
46.	Set Passwords and Security Features	93
47.	Set Startup Sequence Screen	96
48.	Set Power-On Features Screen	97
49.	More Utilities Menu	98
50.	Display Revision Level Screen	98
51.	System Error Log Screen	99

52.	Advanced Diagnostic Menu	100
53.	Test Selection Menu	100
54.	RAID Configuration Program - Adapter Selection	102
55.	RAID Configuration Program - Main Menu	102
56.	RAID Configuration Program - Drive Information	103
57.	RAID Configuration Program - Advanced Functions Menu	104
58.	RAID Configuration Program - DASD Formatting	105
59.	RAID Configuration Program - Change RAID parameters	105
60.	RAID Configuration Program - Create/Delete Array Menu	107
61.	RAID Configuration Program - Hot-Spare Disk Definition	107
62.	RAID Configuration Program - Disk Array Creation	108
63.	RAID Configuration Program - Logical Drive Definition	109
64.	RAID Configuration Program - Logical Drive Definition - Array Selection	109
65.	RAID Configuration Program - RAID level Selection	110
66.	RAID Configuration Program - Size Definition	111
67.	RAID Configuration Program - Result	111
68.	RAID Configuration Program - Advanced Functions Menu	113
69.	ServerGuide Language	116
70.	ServerGuide Main Menu	117
71.	Installing NetWare	119
72.	Configuring NetWare	120
73.	Configuring IBM NetFinity	121
74.	Partitioning the Hard Disk	122
75.	Reviewing Configuration	123
76.	Unlocking Programs	124
77.	Installing NetWare Directory Services (NDS)	125
78.	Assigning a Name to a Directory Tree	125
79.	Assigning a Context for the Server	126
80.	Server Context Information.	126
81.	NetWare Installation	129
82.	NetWare V3.12 Installation - Main Menu	130
83.	Create Partition	130
84.	Partition Information	131
85.	Creating a New Volume	131
86.	Volume Status	132
87.	Copy System and Public Files	132
88.	Path for STARTUP.NCF File	133
89.	STARTUP.NCF File	134
90.	AUTOEXEC.NCF File	134
91.	Installation Menu	136
92.	Disk Driver Options	137
93.	Network Driver Options	137
94.	Create Partition	138
95.	Disk Partition Information	138
96.	New Volume Information	139
97.	Optional NetWare Files	139
98.	Install NetWare Directory Services(NDS)	140
99.	Assigning a Name to the Directory Tree	140
100.	Context for the Server	141
101.	Server Context Information	141
102.	Editing STARTUP.NCF File	142
103.	Editing AUTOEXEC.NCF File	142
104.	File Copy Status	143
105.	Other Installation Options	143

106.	NetFinity Network Driver Configuration	145
107.	NetFinity Installation	147
108.	NetFinity Services for NetWare	148
109.	NetFinity Installation - Copying Files	148
110.	Network Driver Configuration	149
111.	Configuration Update	150
112.	NetFinity Installation Complete	150
113.	Raid Administration for NetWare - Main Menu	152
114.	RAID Administration Utility - Main Menu	153
115.	Verifying Array Configuration	154
116.	Using RAID Manager to View Array Configuration	154
117.	Detecting the Disk Failure	155
118.	Disk Failure - NetFinity Alert	155
119.	Disk Failure - NetFinity RAID Service	156
120.	View Last Event Message	156
121.	RAID Administration - Recovery Message	157
122.	NetFinity Recovery Alert	157
123.	Changes in Array Configuration	158
124.	RAID Administration - Replace a Defunct Drive	158
125.	RAID Administration - Verifying the Replacement of a Defunct Drive	159
126.	NetFinity New Hot Spare Drive Alert	159
127.	NetFinity RAID Service - New Hot Spare	160
128.	RAID Administration - Array Configuration	160
129.	NetFinity RAID Service - Verifying Configuration	161
130.	Detecting the Disk Failure	161
131.	NetFinity Alert Log	162
132.	NetFinity RAID Service - Disk Failure	162
133.	Last Event Message	163
134.	RAID Administration Utility - Reviewing Disk Status	163
135.	RAID Administration - Replace a Defunct Drive	164
136.	RAID Administration - Rebuild Progress	164
137.	RAID Administration - Verifying the Rebuild Status	165
138.	NetFinity Alert - New Disk Online	165
139.	LAN Server Controlled Subsystems	167
140.	File Server Performance - General Characteristics	169
141.	Differences in LAN Adapters	171
142.	Differences in Disk Subsystems	172
143.	MONITOR Utility	174
144.	SERVMAN Utility	175
145.	Sample Compatibility Report Showing Ethernet LAN Adapters	200

Tables

1.	ECC Memory Performances	12
2.	Summary of Memory Implementations	12
3.	SCSI Adapters Summary	21
4.	PCI SCSI Adapters Summary	21
5.	Summary of Disks Performance Characteristics	22
6.	RAID Classifications	23
7.	Summary of RAID Performance Characteristics	30
8.	IBM PC Servers 300 Models	56
9.	IBM PC Servers 310 Models	57
10.	IBM PC Servers 320 EISA Models	58
11.	IBM PC Servers 320 MCA Models	59
12.	IBM PC Server 500 Models	60
13.	IBM PC Servers 520 EISA Models	61
14.	IBM PC Servers 520 MCA Models	62
15.	IBM PC Servers 720 Models	63
16.	Host Adapter SCSI Termination Parameter	79
17.	Volume Block Size and Cache Buffer Size Recommendations	176
18.	Default Block Sizes Based on Volume Size	177
19.	NetWare Memory Pools	185

Special Notices

This document is intended for IBM customers, dealers, systems engineers and consultants who are implementing Novell NetWare on an IBM PC Server. The information in this publication is not intended as the specification of any programming interfaces that are provided by IBM.

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, 500 Columbus Avenue, Thornwood, NY 10594 USA.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The information about non-IBM (VENDOR) products in this manual has been supplied by the vendor and IBM assumes no responsibility for its accuracy or completeness. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Any performance data contained in this document was determined in a controlled environment, and therefore, the results that may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data for their specific environment.

Reference to PTF numbers that have not been released through the normal distribution process does not imply general availability. The purpose of including these reference numbers is to alert IBM customers to specific information relative to the implementation of the PTF when it becomes available to each customer according to the normal IBM PTF distribution process.

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

AIX	AIX/6000
AT	DB2/2
DataHub	DatagLANce
EtherStreamer	First Failure Support Technology
IBM	LANStreamer
Micro Channel	NetFinity

NetView
PS/2
Power Series 800
SystemView
VM/ESA

OS/2
Personal System/2
Presentation Manager
Ultimedia

The following terms are trademarks of other companies:

C-bus is a trademark of Corollary, Inc.

PC Direct is a trademark of Ziff Communications Company and is used by IBM Corporation under license.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

Windows is a trademark of Microsoft Corporation.

386	Intel Corporation
486	Intel Corporation
AHA	Adaptec, Incorporated
AppleTalk	Apple Computer, Incorporated
Banyan	Banyan Systems, Incorporated
CA	Computer Associates
DECnet	Digital Equipment Corporation
EtherLink	3COM Corporation
HP	Hewlett-Packard Company
IPX	Novell, Incorporated
Intel	Intel Corporation
Lotus 1-2-3	Lotus Development Corporation
Lotus Notes	Lotus Development Corporation
MS	Microsoft Corporation
Micronics	Micronics Electronics, Incorporated
Microsoft	Microsoft Corporation
Microsoft Excel	Microsoft Corporation
NFS	Sun Microsystems Incorporated
NetWare	Novell, Incorporated
Novell	Novell, Incorporated
OpenView	Hewlett-Packard Company
Pentium	Intel Corporation
Phoenix	Phoenix Technologies, Limited
PowerChute	American Power Conversion
SCO	The Santa Cruz Operation, Incorporated
SCSI	Security Control Systems, Incorporated
SCSISelect	Adaptec, Incorporated
VINES	Banyan Systems, Incorporated
Windows NT	Microsoft Corporation
X/Open	X/Open Company Limited
i386	Intel Corporation
i486	Intel Corporation
i960	Intel Corporation

Other trademarks are trademarks of their respective companies.

Preface

This document describes the procedures necessary to implement Novell NetWare on IBM PC Server platforms. It provides detailed information on installation, configuration, performance tuning, and management of the IBM PC Server in the NetWare environment. It also discusses the features and technologies of the IBM PC Server brand and positions the various models in the brand.

How This Document is Organized

The document is organized as follows:

- Chapter 1, "IBM PC Server Technologies"

This chapter introduces many of the technologies used in the IBM PC Server brand and gives examples of system implementations where they are used.

- Chapter 2, "IBM PC Server Family Overview"

This chapter positions the various models within the IBM PC Server brand and gives specifications for each model.

- Chapter 3, "Hardware Configuration"

This chapter provides a roadmap for configuring the various models of the IBM PC Server line and describes the configuration process in detail.

- Chapter 4, "Novell NetWare Installation"

This chapter gives a step-by-step process for installing both NetWare V3.12 and V4.1 and the NetFinity Manager using both ServerGuide and the original product diskettes and CD-ROM. It also contains an overview of the ServerGuide product. It also covers the RAID administration tools and details a process for simulating and recovering from a DASD failure.

- Chapter 5, "Performance Tuning"

This chapter presents an in-depth discussion of tuning NetWare as it relates to the major hardware subsystems of the file server. It also discusses performance monitoring tools.

- Appendix A, "EISA Configuration File"

This appendix contains a sample report printed from the EISA configuration utility.

- Appendix B, "Hardware Compatibility, Device Driver, and Software Patch Information"

This appendix gives information on where to find the latest compatibility information, device drivers, and code patches in the NetWare environment.

- Appendix C, "Configuring DOS CD-ROM Support"

This appendix gives information on how to configure your IBM PC Server for CD-ROM support in the DOS environment.

Related Publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this document.

- *IBM PC Server 310 System Library*, S52H-3697
- *IBM PC Server 320 System Library for Non-Array Models*, S19H-1175
- *IBM PC Server 320 System Library for Array Models*, S19H-1196
- *IBM PC Server 320 PCI/Micro Channel System Library*, S30H-1778
- *IBM PC Server 520 System Library*, S52H-3695
- *The PC Server 720 System Library*, S30H-1782

International Technical Support Organization Publications

- *Advanced PS/2 Servers Planning and Selection Guide*, GG24-3927
- *NetWare 4.0 from IBM: Directory Services Concepts*, GG24-4078
- *NetWare from IBM: Network Protocols and Standards*, GG24-3890

A complete list of International Technical Support Organization publications, known as redbooks, with a brief description of each, may be found in:

International Technical Support Organization Bibliography of Redbooks, GG24-3070.

To get a catalog of ITSO redbooks, VNET users may type:

```
TOOLS SENDTO WTSCPOK TOOLS REDBOOKS GET REDBOOKS CATALOG
```

A listing of all redbooks, sorted by category, may also be found on MKTTOOLS as ITSOCAT TXT. This package is updated monthly.

How to Order ITSO Redbooks

IBM employees in the USA may order ITSO books and CD-ROMs using PUBORDER. Customers in the USA may order by calling 1-800-879-2755 or by faxing 1-800-445-9269. Most major credit cards are accepted. Outside the USA, customers should contact their local IBM office. Guidance may be obtained by sending a PROFS note to BOOKSHOP at DKIBMVM1 or E-mail to bookshop@dk.ibm.com.

Customers may order hardcopy ITSO books individually or in customized sets, called BOFs, which relate to specific functions of interest. IBM employees and customers may also order ITSO books in online format on CD-ROM collections, which contain redbooks on a variety of products.

ITSO Redbooks on the World Wide Web (WWW)

Internet users may find information about redbooks on the ITSO World Wide Web home page. To access the ITSO Web pages, point your Web browser to the following URL:

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

IBM employees may access LIST3820s of redbooks as well. Point your web browser to the IBM Redbooks home page at the following URL:

<http://w3.itsc.pok.ibm.com/redbooks/redbooks.html>

Acknowledgments

This project was designed and managed by:

Tim Kearby
International Technical Support Organization, Raleigh Center

The authors of this document are:

Wuilbert Martinez Zamora
IBM Mexico

Jean-Paul Simoen
IBM France

Angelo Rimoldi
IBM Italy

Tim Kearby
International Technical Support Organization, Raleigh Center

This publication is the result of a residency conducted at the International Technical Support Organization, Raleigh Center.

Thanks to the following people for the invaluable advice and guidance provided in the production of this document:

Barry Nusbaum, Michael Koerner, Gail Wojton
International Technical Support Organization

Tom Neidhardt, Dave Laubscher, Marc Shelley
IBM PC Server Competency Center, Raleigh

Ted Ross, Ron Abbott
IBM PC Company, Raleigh

Gregg McKnight, Phil Horwitz, Paul Awoseyi
IBM PC Server Performance Laboratory, Raleigh

John Dinwiddie, Alison Farley, Victor Guess, Dottie Gardner-Lamontagne
IBM PC Server Unit, Raleigh

Parts of this document are based on an earlier version of the *NetWare Integration Guide*, which was produced by the IBM European Personal Systems Center in Basingstoke, U.K.

Thanks also to the many people, both within and outside IBM, who provided suggestions and guidance, and who reviewed this document prior to publication.

Chapter 1. IBM PC Server Technologies

IBM PC Servers use a variety of technologies. This chapter introduces many of these technologies and gives examples of system implementations where they are used.

1.1 Processors

The microprocessor is the central processing unit (CPU) of the server. It is the place where most of the control and computing functions occur. All operating system and application program instructions are executed here. Most information passes through it, whether it is a keyboard stroke, data from a disk drive, or information from a communication network.

The processor needs data and instructions for each processing operation that it performs. Data and instructions are loaded from memory into data-storage locations, known as registers, in the processor. Registers are also used to store the data that results from each processing operation, until the data is transferred to memory.

The microprocessor is packaged as an integrated circuit which contains one or more arithmetic logic units (ALUs), a floating point unit, on-board cache, registers for holding instructions and data, and control circuitry.

Note: The ALUs and the floating point unit are often collectively referred to as execution units.

1.1.1 Clock Rate

A fundamental characteristic of all microprocessors is the rate at which they perform operations. This is called the clock rate and is measured in millions of cycles per second or Megahertz (MHz). The maximum clock rate of a microprocessor is determined by how fast the internal logic of the chip can be switched. As silicon fabrication processes are improved, the integrated devices on the chip become smaller and can be switched faster. Thus, the clock speed can be increased.

For example, the Pentium P54C processor in the IBM PC Server 720 operates at a clock speed of 100 MHz. The P54C is based on a fabrication process where transistors on the chip have a channel width of .6 microns (a .6 micron BiCMOS process). The original P5 processor is based on a .8 micron process and could only be clocked at a maximum of 66 MHz.

The clock rate of the external components can be different from the rate at which the processor is clocked internally. Clock doubling is a technique used in the Intel DX2 and DX4 class processors to clock the processor internally faster than the external logic components. For example, the 486DX2 at 66/33 MHz clocks the processor internally at 66 MHz, while clocking the external logic components at 33 MHz. This is an efficient systems design technique when faster external logic components are not available or are prohibitively expensive.

One might think that the faster the clock speed, the faster the performance of the system. This is not always the case. The speed of the other system components, such as main memory, can also have a dramatic effect on

performance. (Please see 1.3, "Memory" on page 3 for a discussion of memory speeds and system performance.) The point is that you cannot compare system performance by simply looking at the speed at which the processor is running. A 90 MHz machine with a set of matched components can out perform a 100 MHz machine which is running with slow memory. IBM PC Servers are always optimized to incorporate these factors and they always deliver a balanced design.

1.1.2 External Interfaces

The processor data interface, or data bus, is the data connection between the processor and external logic components. The Pentium family of processors utilizes a 64-bit data bus, which means that they are capable of reading in 8 bytes of data in one memory cycle from processor main memory. The Intel 486 has a data bus of only 32-bits, which limits its memory cycles to 4 bytes of data per cycle.

The width of the processor address interface, or address bus, determines the amount of physical memory the processor can address. A processor with a 24-bit address bus, such as the i286 class of processors, can address a maximum of 16 megabytes (MB) of physical memory. Starting with the i386 class of processors, the address bus was increased to 32 bits, which correlates to 4 gigabyte (GB) of addressability.

1.1.3 Processor Types

IBM currently uses two processors in the PC Server line:

- 80486DX2

The 80486DX2 has a 32-bit address bus and 32-bit data bus. It utilizes clock doubling to run at 50/25 MHz or 66/33 MHz. It is software compatible with all previous Intel processors. The 80486DX2 has an internal two-way set associative 8KB cache.

- Pentium

The Pentium has a 32-bit address bus and 64-bit data bus. It has internal split data and instruction caches of 8KB each. The instruction cache is a write-through cache and the data cache is a write-back design. The Pentium microprocessor is a two-issue superscalar machine. This means that there are two integer execution units (ALUs) in addition to the on-board floating point unit. The superscalar architecture is one of the key techniques used to improve performance over that of the previous generation i486 class processors. Intel was able to achieve this design while maintaining compatibility with applications written for the Intel i368/i486 family of processors.

Note: A superscalar architecture is one where the microprocessor has multiple execution units, which allow it to perform multiple operations during the same clock cycle.

1.2 Multiprocessing

Multiprocessing uses two or more processors in a system to increase throughput. Multiprocessing yields high performance for CPU intensive applications such as database and client/server applications.

There are two types of multiprocessing:

- Asymmetric Multiprocessing
- Symmetric Multiprocessing

Asymmetric Multiprocessing: In asymmetric multiprocessing the program tasks (or threads) are strictly divided by type between processors and each processor has its own memory address space. These features make asymmetric multiprocessing difficult to implement.

Symmetric Multiprocessing (SMP): Symmetric multiprocessing means that any processor has access to all system resources including memory and I/O devices. Threads are divided evenly between processors regardless of type. A process is never forced to execute on a particular processor.

Symmetric multiprocessing is easier to implement in network operating systems (NOSs) and is the method used most often in operating systems that support multiprocessing. It is the technology currently used by OS/2 SMP, Banyan Vines, SCO UNIX, Windows NT, and UnixWare 2.0.

The IBM PC Server 320, 520, and 720 support SMP. The PC Server 320 and 520 support two-way SMP via an additional Pentium processor in a socket on the planar board. The 720 supports two-to-six way SMP via additional processor complexes.

1.3 Memory

The system design of PC servers (in fact all microprocessor-based systems) is centered around the basic memory access operation. System designers must always *tune* this operation to be as fast as possible in order to achieve the highest possible performance.

Processor architectures always allow a certain number of clock cycles in order to read or write information to system memory. If the system design allows this to be completed in the given number of clock cycles, then this is called a zero wait state design.

If for some reason the operation does not complete in the given number of clocks, the processor must *wait* by inserting extra *states* into the basic operation. These are called *wait states* and are always an integer multiple of clock cycles. The challenge is that as each new generation of processors is clocked faster, it becomes more expensive to incorporate memory devices that have access times allowing zero wait designs. For example, state of the art Dynamic Random Access Memory, or DRAM, has a typical access time of about 60 nanoseconds (ns). A 60 ns DRAM is not fast enough to permit a zero wait state design with a Pentium class processor. Static RAM, or SRAM, has an access time of less than 10 ns. A 10 ns SRAM design would allow for zero waits at current processor speeds but would be prohibitively expensive to implement as main memory. A basic trade-off that all system designers must face is simply that as the access time goes down, the price goes up.

The key is to achieve a balanced design where the speed of the processor is matched to that of the external components. IBM engineers achieve a balanced design by using several techniques to reduce the *effective* access time of main system memory:

- Cache
- Interleaving
- Dual path buses
- SynchroStream technology

1.3.1 Caches

Research has shown that when a system uses data, it will be likely to use it again. As previously discussed, the faster the access to this data occurs, the faster the overall machine will operate. Caches are memory buffers that act as temporary storage places for instructions and data obtained from slower, main memory. They use static RAM and are much faster than the dynamic RAM used for system memory (typically five to ten times faster). However, SRAM is more expensive and requires more power, which is why it is not used for all memory.

Caches reduce the number of clock cycles required for a memory access since they are implemented with fast SRAMs. Whenever the processor must perform external memory read accesses, the cache controller always pre-fetches extra bytes and loads them into the cache. When the processor needs the next piece of data, it is likely that it is already in the cache. If so, processor performance is enhanced, if not, the penalty is minimal.

Caches are cost-effective because they are relatively small as compared to the amount of main memory.

There are several levels of cache implemented in IBM PC servers. The cache incorporated into the main system processor is known as Level 1 (L1) cache. The Intel 486 incorporates a single 8KB cache. The Intel Pentium family has two 8KB caches, one for instructions and one for data. Access to these on-board caches are very fast and consume only a fraction of the time required to access memory locations external to the chip.

The second level of cache, called second-level cache or L2 cache, provides additional high speed memory to the L1 cache. If the processor cannot find what it needs in the processor cache (a first-level *cache miss*), it then looks in the additional cache memory. If it finds the code or data there (a second-level *cache hit*) the processor will use it and continue. If the data is in neither of the caches, an access to main memory must occur.

L2 caches are standard in all IBM PC server models.

With all types of caching, more is not always better. Depending on the system, the optimum size of Level 2 Cache is usually 128KB to 512KB.

L2 Caches can be of two types:

- **Write-Through Cache**

Read operations are issued from the cache but write operations are sent directly to the standard memory. Performance improvements are obtained only for read operations.

- **Write-Back Cache**

Write operations are also performed on the cache. Transfer to standard memory is done if:

- Memory is needed in the cache for another operation
- Modified data in the cache is needed for another application

The third level of cache or L3 cache is sometimes referred to as a victim cache. This cache is a highly customized cache used to store recently evicted L2 cache entries. It is a smaller cache usually less than 256 bytes. An L3 cache is implemented in the IBM PC Server 720 SMP system.

1.3.1.1 SMP Caching

Within SMP designs, there are two ways in which a cache is handled:

- Shared cache
- Dedicated cache

Shared Cache: Sharing a single L2 cache among processors is the least expensive SMP design. However, the performance gains associated with a shared cache are not as great as with a dedicated cache. With the shared secondary cache design, adding a second processor can provide as much as a 30% performance improvement. Additional processors provide very little incremental gain. If too many processors are added, the system will even run slower due to memory bus bottlenecks caused by processor contention for access to system memory.

The IBM PC server 320 supports SMP with a shared cache.

Figure 1 shows SMP with shared secondary cache.

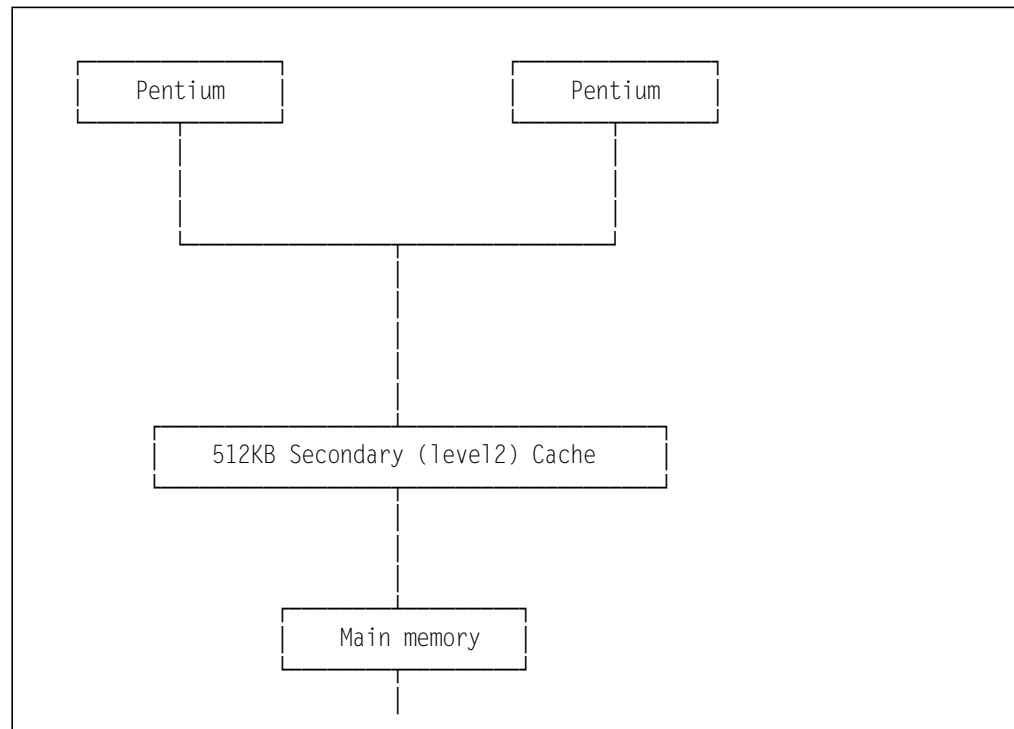


Figure 1. SMP Shared Secondary Cache

Dedicated Cache: This SMP design supports a dedicated L2 cache for each processor. This allows more cache hits than a shared L2 cache. Adding a second processor using a dedicated L2 cache can improve performance as much as 80%. With current technology, adding even more processors can further increase performance in an almost linear fashion up to the point where the addition of more processors does not increase performance and can actually decrease performance due to excessive overhead.

The IBM PC Server 720 implements SMP with dedicated caches.

Figure 2 shows SMP with dedicated secondary cache.

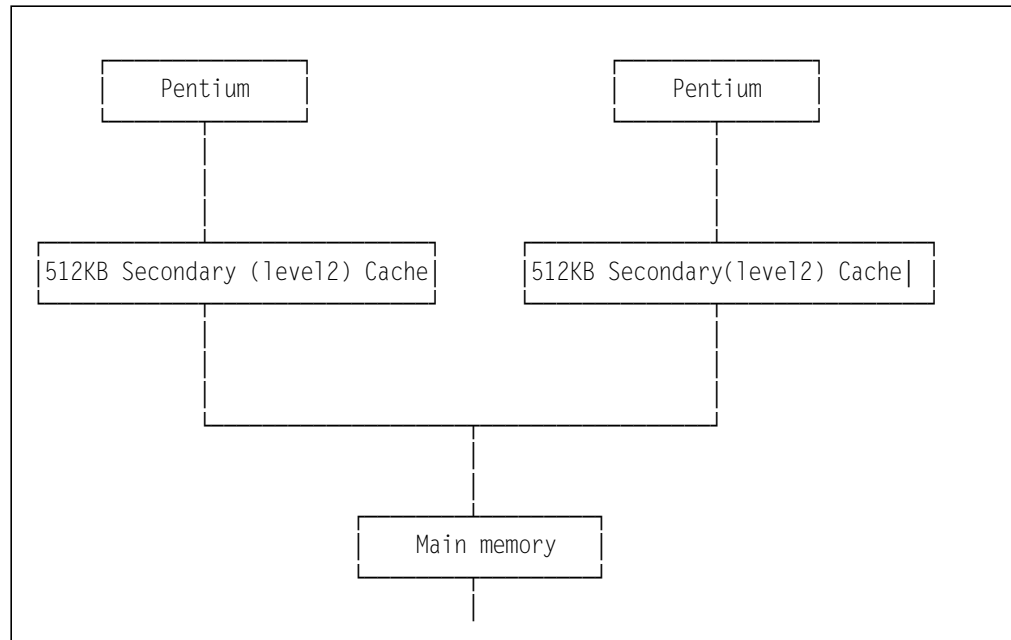


Figure 2. SMP with Dedicated Secondary Cache

Dedicated caches are also more complicated to manage. Care needs to be taken to ensure that a processor needing data always gets the latest copy of that data. If this data happens to reside in another processor's cache, then the two caches must be brought into sync with one another.

The cache controllers maintain this *coherency* by communicating with one another using a special protocol called MESI, which stands for *Modified, Exclusive, Shared, or Invalid*. These refer to tags that are maintained for each line of cache, and indicate the state of each line.

The implementation of MESI in the IBM PC server 720 supports two sets of tags for each cache line, which allows for faster cache operation than when only one set of tags is provided.

1.3.2 Memory Interleaving

Another technique used to reduce effective memory access time is interleaving. This technique greatly increases memory bandwidth when access to memory is sequential such as in program instruction fetches.

In interleaved systems, memory is currently organized in either two or four banks. Figure 3 on page 7 shows a two-way interleaved memory implementation.

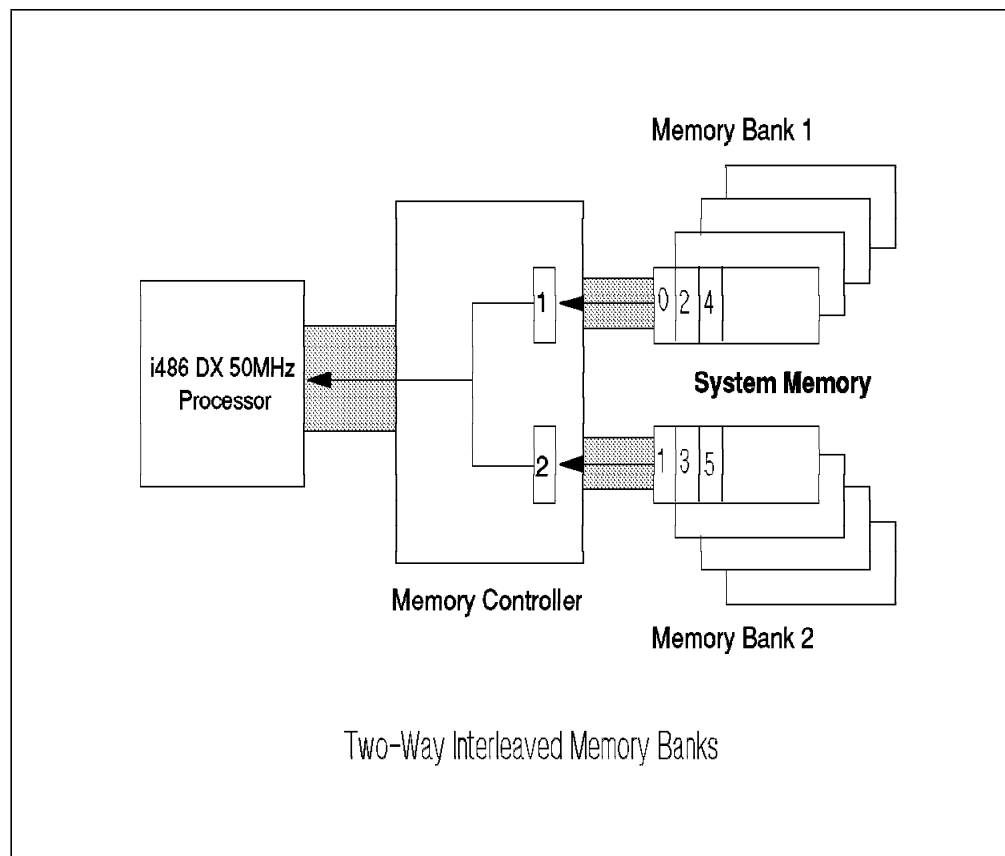


Figure 3. Two-Way Interleaved Memory Banks

Memory accesses are overlapped so that as the controller is reading/writing from bank 1, the address of the next word is presented to bank 2. This gives bank 2 a head start on the required access time. Similarly, when bank 2 is being read, bank 1 is fetching/storing the next word.

The PC server 500 uses a two-way interleaved memory. In systems implementing two-way interleaved memory, additional memory must be added in pairs of single in-line memory modules (SIMMs) operating at the same speed (matched SIMMs).

The PC server 720 uses a four-way interleaved memory with a word length of 64 bits. In this system, in order to interleave using 32-bit SIMMs, it is necessary to add memory in matched sets of eight SIMMs each.

1.3.3 Dual Path Buses

A dual path bus allows both the processor and a bus master to access system memory simultaneously. Figure 4 on page 8 shows a dual path bus implementation.

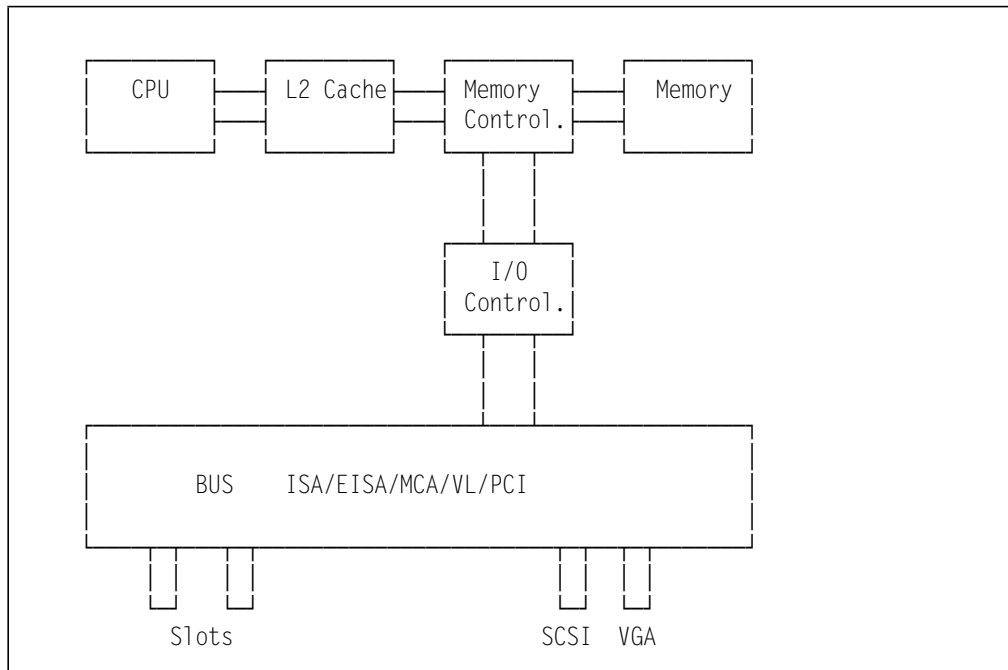


Figure 4. Dual Path Bus Implementation

Without a dual path bus, there is often contention for system resources such as main memory. When contention between the processor and a bus master occurs, one has to wait for the other to finish its memory cycle before it can proceed. Thus, fast devices like processors have to wait for much slower I/O devices, slowing down the performance of the entire system to the speed of the slowest device. This is very costly to the overall system performance.

1.3.4 SynchroStream Technology

SynchroStream is an extension of the dual bus path technique. The SynchroStream controller synchronizes the operation of fast and slow devices and streams data to these devices to ensure that all devices work at their optimum levels of performance.

It works much like a cache controller in that it pre-fetches extra data on each access to memory and buffers this data in anticipation of the next request. When the device requests the data, the IBM SynchroStream controller provides it quickly from the buffer and the device continues working. It does not have to wait for a normal memory access cycle.

When devices are writing data into memory, the IBM SynchroStream controller again buffers the data, and writes it to memory after the bus cycle is complete.

Since devices are not moving data to and from memory directly, but to the SynchroStream controller, each device has its own logical path to memory. The devices do not have to wait for other, slower devices.

1.4 Memory Error Detection and Correction

IBM PC servers implement four different memory systems:

- Standard (parity) memory
- Error Correcting Code-Parity
- Error Correcting Code (ECC) memory
- ECC Memory on SIMMs (EOS) Memory

1.4.1 Standard (Parity) Memory

Parity memory is standard IBM memory with 32 bits of data space and 4 bits of parity information (one check bit/byte of data). The 4 bits of parity information are able to tell you an error has occurred but do not have enough information to locate which bit is in error. In the event of a parity error, the system generates a non-maskable interrupt (NMI) which halts the system. Double bit errors are undetected with parity memory.

Standard memory is implemented in the PC Servers 300 and 320 as well as in the majority of the IBM desktops (for example the IBM PC 300, IBM PC 700, and PC Power Series 800).

1.4.2 Error Correcting Code (ECC)

The requirements for system memory in PC servers has increased dramatically over the past few years. Several reasons include the availability of 32 bit operating systems and the caching of hard disk data on file servers.

As system memory is increased, the possibility for memory errors increase. Thus, protection against system memory failures becomes increasingly important. Traditionally, systems which implement only parity memory halt on single-bit errors, and fail to detect double-bit errors entirely. Clearly, as memory is increased, better techniques are required.

To combat this problem, the IBM PC servers employ schemes to detect and correct memory errors. These schemes are called Error Correcting Code (or sometimes Error Checking and Correcting but more commonly just ECC). ECC can detect and correct single bit-errors, detect double-bit errors, and detect some triple-bit errors.

ECC works like parity by generating extra check bits with the data as it is stored in memory. However, while parity uses only 1 check bit per byte of data, ECC uses 7 check bits for a 32-bit word and 8 bits for a 64-bit word. These extra check bits along with a special hardware algorithm allow for single-bit errors to be detected and corrected in real time as the data is read from memory.

Figure 5 on page 10 shows how the ECC circuits operate. The data is scanned as it is written to memory. This scan generates a unique 7-bit pattern which represents the data stored. This pattern is then stored in the 7-bit check space.




Figure 5. ECC Memory Operation

As the data is read from memory, the ECC circuit again performs a scan and compares the resulting pattern to the pattern which was stored in the check bits. If a single-bit error has occurred (the most common form of error), the scan will always detect it, automatically correct it and record its occurrence. In this case, system operation will not be affected.

The scan will also detect all double-bit errors, though they are much less common. With double-bit errors, the ECC unit will detect the error and record its occurrence in NVRAM; it will then halt the system to avoid data corruption. The data in NVRAM can then be used to isolate the defective component.

In order to implement an ECC memory system, you need an ECC memory controller and ECC SIMMs. ECC SIMMs differ from standard memory SIMMs in that they have additional storage space to hold the check bits.

The IBM PC Servers 500 and 720 have ECC circuitry and provide support for ECC memory SIMMs to give protection against memory errors.

1.4.3 Error Correcting Code-Parity Memory (ECC-P)

Previous IBM servers such as the IBM Server 85 were able to use standard memory to implement what is known as ECC-P. ECC-P takes advantage of the fact that a 64-bit word needs 8 bits of parity in order to detect single-bit errors (one bit/byte of data). Since it is also possible to use an ECC algorithm on 64 bits of data with 8 check bits, IBM designed a memory controller which implements the ECC algorithm using the standard memory SIMMs.

Figure 6 on page 11 shows the implementation of ECC-P. When ECC-P is enabled via the reference diskette, the controller reads/writes two 32-bit words and 8 bits of check information to standard parity memory. Since 8 check bits are available on a 64-bit word, the system is able to correct single-bit errors and detect double-bit errors just like ECC memory.

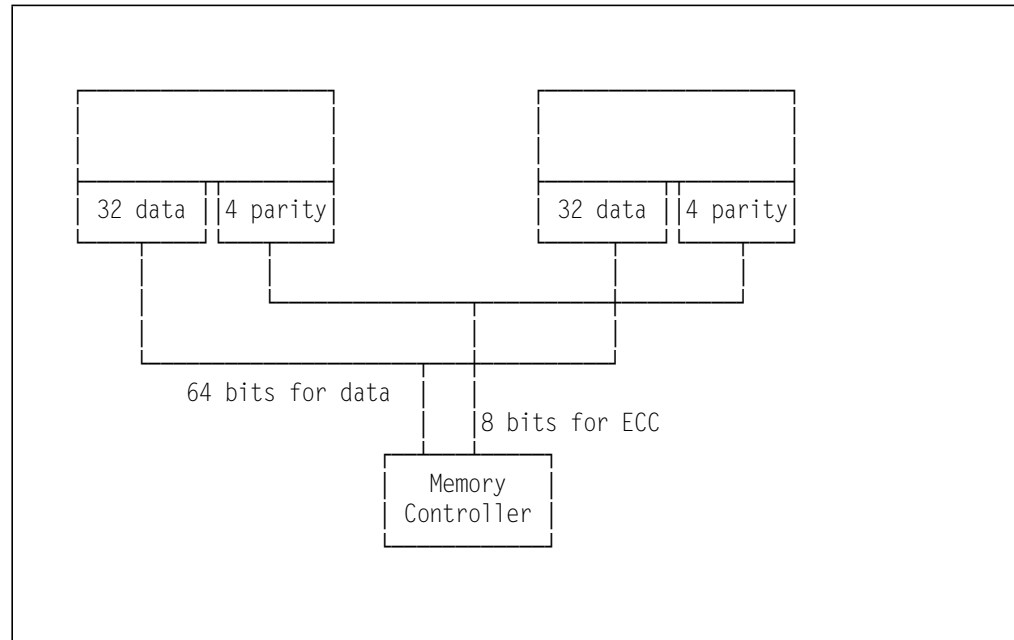


Figure 6. ECC-P Memory Implementation

While ECC-P uses standard non-expensive memory, it needs a specific memory controller that is able to read/write the two memory blocks and check and generate the check bits. Also, the additional logic necessary to implement the ECC circuitry make it slightly slower than true ECC memory. Since the price difference between a standard memory SIMM and an ECC SIMM has narrowed, IBM no longer implements ECC-P.

1.4.4 ECC on SIMMs (EOS) Memory

A server that supports one hundred or more users can justify the additional cost necessary to implement ECC on the system. It is harder to justify this cost for smaller configurations. It would be desirable for a customer to be able to upgrade his system at a reasonable cost to take advantage of ECC memory as his business grows.

The problem is that the ECC and ECC-P techniques previously described use special memory controllers imbedded on the planar board which contain the ECC circuits. It is impossible to upgrade a system employing parity memory (with a parity memory controller) to ECC even if we upgrade the parity memory SIMMs to ECC memory SIMMs.

To answer this problem, IBM has introduced a new type of memory SIMM which has the ECC logic integrated on the SIMM. These are called ECC on SIMMs or EOS memory SIMMs. With these SIMMs, the memory error is detected and corrected directly on the SIMM before the data gets to the memory controller. This solution allows a standard memory controller to be used on the planar board and allows the customer to upgrade a server to support error checking memory.

1.4.5 Performance Impact

As previously discussed, systems which employ ECC memory have slightly longer memory access times depending on where the checking is done. It should be stressed that this affects only the access time of external system memory, not L1 or L2 caches. Table 1 shows the performance impacts as a percentage of system memory access times of the different ECC memory solutions.

Again, these numbers represent only the impact to accessing external memory. They do not represent the impact to overall system performance which is harder to measure but will be substantially less.

Table 1. ECC Memory Performances

	SIMM	Memory Controller	Impact to Access Time	Systems where implemented
ECC	X	X	3%	PC Servers 500 and 720
ECC-P		X	14%	No more (Mod 85)
EOS	X		None	Option for PC Servers 300, 320 Standard for PC Servers 520

1.4.6 Memory Options and Speed

The following memory options are available from IBM:

- 4MB, 8MB, 16MB, 32MB 70 ns Standard (Parity) Memory SIMMs
- 4MB, 8MB, 16MB, 32MB 70 ns ECC Memory SIMMs
- 4MB, 8MB, 16MB, 32MB 60 ns ECC Memory SIMMs
- 4MB, 8MB, 16MB, 32MB 70 ns EOS Memory SIMMs

Table 2 shows the options used by each PC server.

Table 2. Summary of Memory Implementations

PS/2 Model	70 ns Standard	70 ns ECC-P	70 ns ECC	60 ns ECC	70 ns EOS
PC Server 300/310/320	X				OPT
PC Server 500			X		
PC Server 520					X
PC Server 720				X	

1.5 Bus Architectures

There are a number of bus architectures implemented in IBM PC servers:

- ISA
- EISA
- MCA
- PCI

1.5.1 ISA Bus

The Industry Standard Architecture (ISA) is not really an architecture at all but a defacto standard based on the original IBM PC/AT bus design. The main characteristics of the ISA bus include a 16-bit data bus and a 24-bit address bus. The bus speed is limited to 8 MHz and it did not allow for DMA and bus masters in its original form. It does not support automatically configuring adapters and resolving resource conflicts among adapters nor does it allow for sharing of interrupts. Nonetheless, it was an extremely successful design and even with these disadvantages, it is estimated that the ISA bus is in 70% of the PCs manufactured today.

1.5.2 EISA Bus

The Extended Industry Standard Bus Architecture (EISA) is a 32-bit superset of the ISA bus providing improved functionality and greater data rates while maintaining backward compatibility with the many ISA products already available.

The main advancements of the EISA bus are 32-bit addressing and 16-bit data transfer. It supports DMA and bus master devices. It is synchronized by an 8.33 MHz clock and can achieve data transfer of up to 33 MBps. A bus arbitration scheme is also provided which allows efficient sharing of multiple EISA bus devices. EISA systems can also automatically configure adapters.

1.5.3 Micro Channel Bus

The Micro Channel Architecture (MCA) was introduced by IBM in 1987. Micro Channel is an improvement over ISA in all of the areas discussed in the previous section on EISA. In addition, it supports data streaming which is an important performance feature of the MCA architecture.

1.5.3.1 Data Streaming

The data streaming transfer offers considerably improved I/O performance. In order to understand data streaming transfers we need to see how data is transferred between Micro Channel bus master adapters and memory.

The standard method of transfer across the Micro Channel is known as basic data transfer. In order to transfer a block of data in basic data transfer mode, an address is generated on the address bus to specify where the data should be stored; then the data is put on the data bus.

This process is repeated until the entire block of data has been transferred. Figure 7 on page 14 shows basic data transfer in operation. Basic data transfer on the Micro Channel runs at 20 MBps (each cycle takes 200 nanoseconds, and 32 bits or 4 bytes of data are transferred at a time).

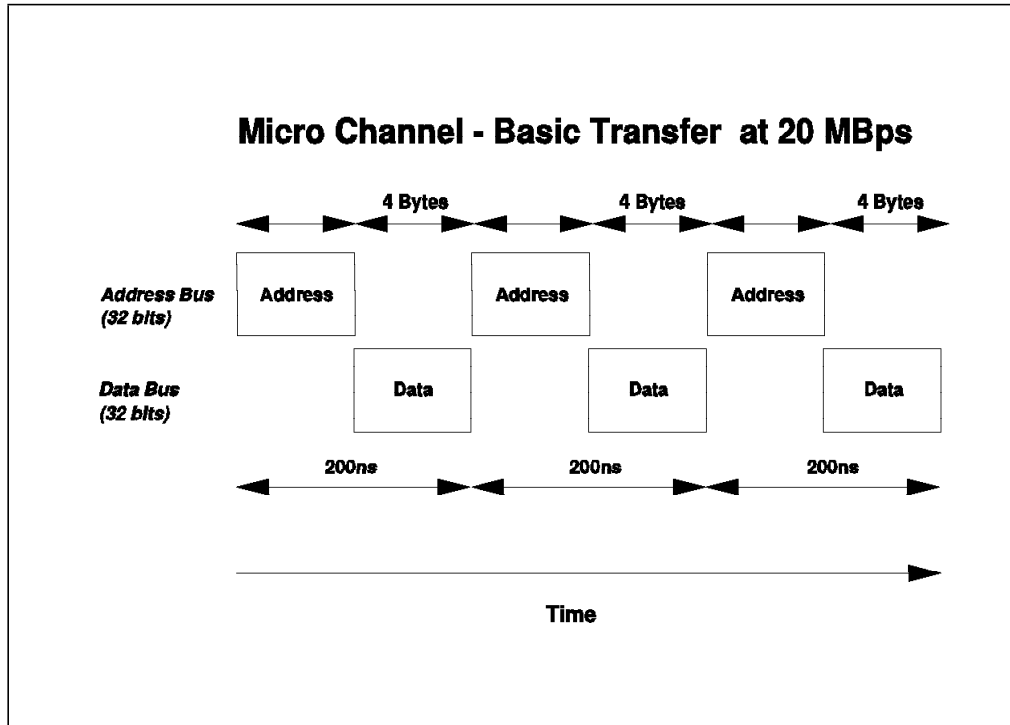


Figure 7. Micro Channel - Basic Data Transfer (20 MBps)

However, in many cases, blocks transferred to and from memory are stored in sequential addresses, so repeatedly sending the address for each 4 bytes is unnecessary. With data streaming transfers, the initial address is sent, and then the blocks of data are sent; it is then assumed that the data requests are sequential. Figure 8 shows 40 MBps data streaming in operation.

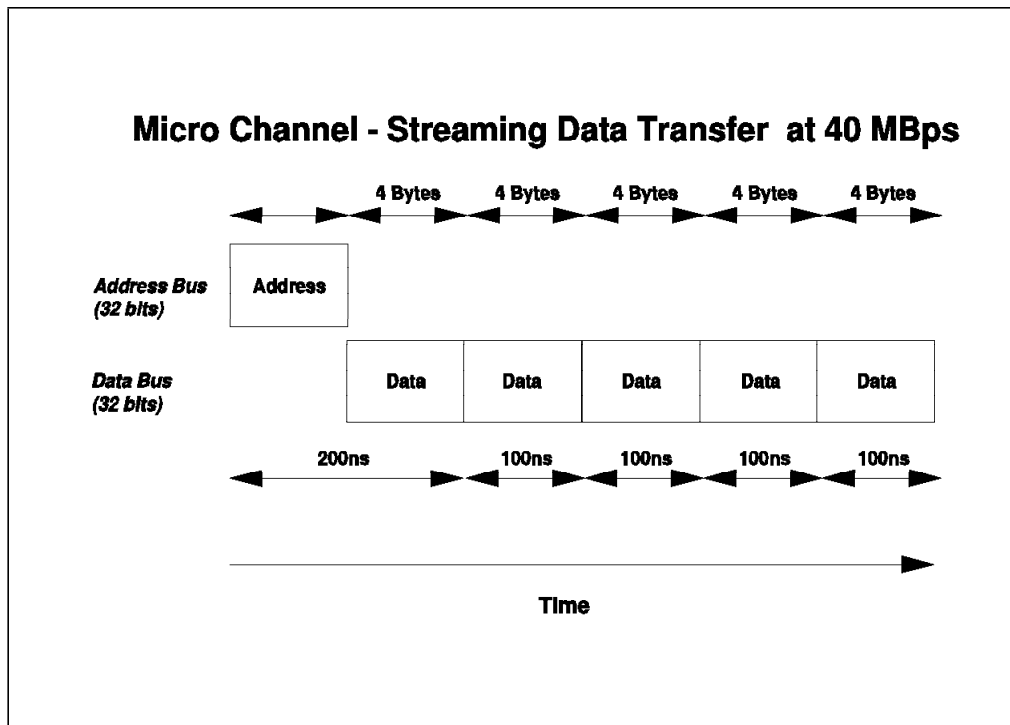


Figure 8. Micro Channel - Data Streaming Transfer (40 MBps)

The Micro Channel supports another mode of data streaming whereby the address bus can also be used to transfer data. This is depicted in Figure 9 on page 15.

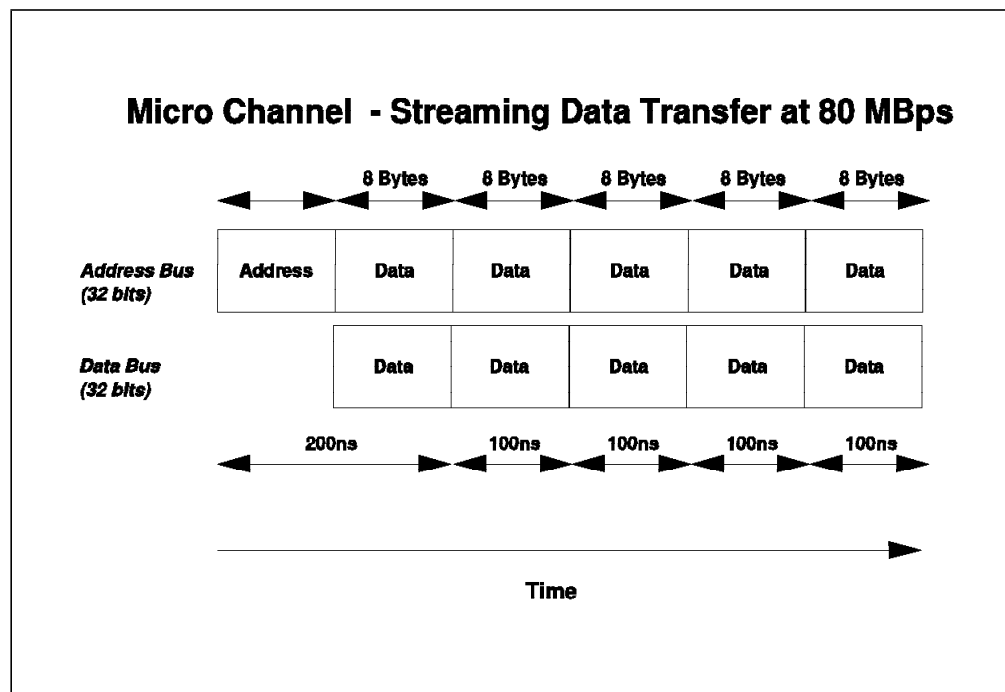


Figure 9. Micro Channel - Data Streaming Transfer (80 MBps)

As can be seen from this figure, in this mode, after the initial address is presented during the first bus cycle, the address bus is then multiplexed to carry an additional 32 bits of data. This results in an effective data transfer rate of 80 MBps.

Data streaming, as well as improving the data transfer rate, also provides a more efficient use of the Micro Channel. Since MCA operations complete in a shorter amount of time, the overall throughput of the system is increased.

Data streaming is useful for any adapters that perform block transfers across the Micro Channel such as the IBM SCSI-2 Fast/Wide Streaming RAID Adapter/A.

MCA is implemented in some models of the IBM PC Server 300 and 500 lines and in all models of the PC Server 720.

1.5.4 PCI Bus

In the later part of 1992, Intel, IBM and a number of other companies worked together to define a new local component bus which was designed to provide a number of new features and work with a wide range of new processors. The result was the Peripheral Component Interconnect (PCI) bus. The PCI bus was designed to provide the Pentium processor with all the bandwidth it needed and to provide for more powerful processors in the future. It was also designed for use in multiprocessing environments.

The PCI bus was designed to work with a number of buses including Micro Channel, ISA and EISA buses. It was designed to provide a local bus, more tightly integrated with the processor, to provide more bandwidth to I/O devices such as LAN adapters and DISK controllers, which require more bandwidth than

is available with previous bus architectures. In order to optimize performance, the PCI architecture strictly limits the number of loads (hence the number of adapters) on the bus. It therefore needs an I/O expansion bus to handle the more routine I/O devices.

The bus has 32 or 64 bits of address and data, is processor independent and is capable of speeds over 50 MHz. 8-bit and 16-bit devices are not supported. The 64-bit data bus width in combination with clock speeds over 50 MHz can result in data transfer of several hundred megabytes per second. In addition to memory space and I/O space, the bus includes a third address space to support automatic resource allocation and configuration of system and adapter boards.

Unique features of the PCI include parity on all bus lines and control lines. The parity is not optional as in other architectures, but is required. All PCI bus masters must support data streaming to memory devices.

1.6 Disk Subsystem

The disk subsystem is a critical element of server design. In this section we examine the controllers, the devices, and the interfaces between them. We will specifically address SCSI technology and also examine RAID technology in some detail.

1.6.1 Hard Disk Interfaces

The disk interface specifies the physical, electrical, and logical connections between the controller and the Direct Access Storage Devices (DASD). There have been four main interfaces developed thus far. Each possesses different characteristics and performance levels. The interfaces are:

1. **ST506** - This interface was the original standard for microcomputers. It has a data transfer rate of 5 million bits per second (Mbps) between the controller and the DASD Device. It is a serial rather than a parallel interface. This interface is classified as a device level interface because the device itself has no logic to interpret commands. Functions such as formatting, head selection, and error detection are directed by the controller which is housed in an adapter card. A device level interface requires specific adapters and device drivers for each different type of device.
2. **Enhanced Small Device Interface (ESDI)** - This is an enhanced version of the ST506 interface. It provides a 10 Mbps data transfer rate (15 Mbps in some implementations). ESDI devices were the first to use a type of data encoding called Run Length Limited (RLL) which results in denser storage and faster data transfer than the older modified frequency modulation (MFM) technique. However, it is still a device level, serial interface.
3. **Integrated Drive Electronics (IDE)** - This is a bus level interface meaning that the device controller is built into the device itself. The IDE interface was designed for the low cost PC market segment. The interface is flexible and has been enhanced over time. The latest enhancements include caching at the adapter level, a CD-ROM interface, and an extension of the maximum disk storage which was previously limited to 500 MB. However, most IDE implementations still limit the maximum number of hard disks per interface to two. This limitation makes IDE more applicable for desktop systems.

4. **Small Computer System Interface (SCSI)** - The SCSI interface is a high speed parallel interface that transfers eight bits at a time rather than one bit at a time for the ST506 and ESDI serial interfaces. Thus data transfer rates for SCSI are measured in megabytes versus megabits and are considerably faster than those of the serial interfaces. SCSI is also a bus level interface which makes it very flexible. Since the commands are interpreted by the device and not the SCSI host bus adapter, new devices (with new commands) can be implemented and used with standard SCSI adapters. The device driver then interacts with the device via the new commands. An example of this would be a CD-ROM device sharing the same adapter as a hard disk drive. Figure 10 shows a SCSI subsystem with a host bus adapter attached to an integrated controller and hard disk.




Figure 10. SCSI Disk Interface

The SCSI flexibility and high performance make it very suitable for the server environment. In fact, SCSI is the most widely used disk subsystem technology in advanced servers today. All the current IBM PC Servers except for a few at the low end use this technology. For these reasons, we will take a closer look at this interface.

1.6.2 SCSI Technology

As previously discussed, SCSI is a bus level interface through which computers may communicate with a large number of devices of different types connected to the system unit via a SCSI controller and daisy-chained cable. The attached devices include such peripherals as fixed disks, CD-ROMs, printers, plotters, and scanners. The SCSI controller may be in the form of an adapter or integrated on the planar board.

There are several terms and concepts used in discussing SCSI technology that require definition.

- **SCSI-I and SCSI-II:**

SCSI is a standard defined by the American National Standards Institute (ANSI). The original SCSI-I standard is defined in ANSI standard X3.131-1986.

It defines an 8-bit interface with a data transfer rate of 5 MBps. SCSI-II is the second SCSI standard and is defined in ANSI standard X3T9.2/375R REV10K. It defines extensions to SCSI-I which allow for 16 and 32-bit devices, a 10 MBps transfer rate, and other enhancements discussed below.

- **Common Command Set**

The SCSI standard defines a set of commands which must be interpreted by all devices that attach to a SCSI bus. This is called the common command set. Unique devices may implement their own commands, which can be sent by a device driver and interpreted by the device. The advantage of this architecture is that the SCSI adapter does not have to change when new devices with new capabilities are introduced.

- **Tagged Command Queuing (TCQ)**

TCQ is a SCSI-II enhancement. It increases performance in DASD intensive server environments. With SCSI-I systems, only two commands could be sent to a fixed disk; the disk would store one while operating on the other. With TCQ it is possible to send multiple commands to the fixed disk and the disk stores the commands and executes each command in the sequence that gives optimal performance.

- **Scatter/Gather**

Scatter/Gather allows devices to transfer data to and from non-contiguous or *scattered* areas of system memory and on-board cache independently of the CPU. This, again, increases CPU overlap. The Scatter/Gather feature allows for high performance, even in systems that have fragmented memory buffers.

- **Fast/Wide Devices and Controllers**

Fast refers to the doubling of the data transfer rate from the SCSI-I 5 MBps to 10 MBps. *Wide* is used in reference to the width of the SCSI parallel bus between the adapter and the device. Wide generically means wider than the original 8-bit path defined in SCSI-I. Its use is currently limited to mean 16-bits as 32-bit implementations are not currently available. With a 16-bit path, the data rate is double that of an 8-bit device. *Fast/Wide* refers to adapters and devices which implement both the fast and wide interfaces defined above. A fast/wide device has a maximum data transfer rate of 20 MBps.

Note

Wide refers to the width of the bus between the SCSI adapter and the disk drive or other SCSI device. Do not get this confused with the width of the host bus interface (for example, a 32-bit MCA or PCI adapter).

- **Disconnect/Reconnect**

Some commands take a relatively long time to complete (for example a seek command which takes roughly 11 ms). With this feature, the controller can disconnect from the bus while the device is positioning the heads (seeking). Then, when the seek is complete and data is ready to be transferred, the device can arbitrate for the bus and then reconnect with the controller to transfer the data. This allows a much more efficient use of the available

SCSI bus bandwidth. If the controller held onto the bus while waiting for the device to seek, then the other devices would be locked out. This is also sometimes referred to as overlapped operations or multi-threaded I/O on the SCSI bus. This feature is very important in multitasking environments.

- **Synchronous versus Asynchronous**

An asynchronous device must acknowledge each byte as it comes from the controller. Synchronous devices may transfer data in bursts and the acknowledgments happen after the fact. The latter is much faster than the former and most newer devices support this mode of operation. The adapters negotiate with devices on the SCSI bus to ensure that the mode and data transfer rates are acceptable to both the host adapter and the devices. This process prevents data from being lost and ensures that data transmission is error free.

1.6.3 SCSI Adapters

The SCSI adapter provides the interface between the host bus (for example Micro Channel or PCI) and the SCSI bus. The SCSI adapters that IBM has developed are:

- ***IBM Personal System/2 Micro Channel SCSI Adapter***

This adapter is a 16-bit Micro Channel bus master adapter adhering to the SCSI-I interface. It is capable of an 8.3 MBps burst data transfer rate on the Micro Channel. It uses a 16-bit data path and can use a 24- or a 32-bit address on the Micro Channel. It can be installed in either a 16- or 32-bit MCA slot, but if the system has more than 16MB of memory, it must be put in a 32-bit slot due to the limitations of 24-bit addressing in a 16-bit slot.

The bus master capability of this SCSI adapter optimizes data flow from each SCSI device configured to the system. This capability can provide performance benefits in applications where multitasking or high-speed data flow is essential. It allows the processor to be off-loaded from many of the input/output activities common to DASD transfers. This adapter also conforms to the Subsystem Control Block (SCB) architecture for Micro Channel bus masters.

- ***IBM Personal System/2 Micro Channel SCSI Adapter with Cache***

This adapter provides a superset of the features of the PS/2 Micro Channel SCSI Adapter. It is a 32-bit Micro Channel bus master adapter containing a 512KB cache buffer. The cache is used to buffer data between system memory and the device, which permits higher efficiency on both the Micro Channel and the SCSI buses. It has a burst data transfer rate on the Micro Channel of 16.6 MBps. This adapter is recommended where improved data transfer rates and multiple SCSI devices are required and system memory is constrained.

- ***IBM SCSI-2 Fast/Wide Adapter/A***

This adapter is a high performance SCSI-II adapter and a 32-bit Micro Channel bus master adapter capable of streaming data at 40 MBps. It has dual SCSI-II fast and wide channels (one 20 MBps internal and one 20 MBps external). It supports devices using either asynchronous, synchronous, or fast synchronous (10 MBps) SCSI data transfer rates. It also supports

standard 8-bit SCSI devices. Up to seven SCSI physical devices may be attached to this adapter.

This adapter has a dedicated 80C186 local processor on board, which allows it to implement advanced features such as TCQ.

The dual bus design of the adapter prevents access to internal DASD from the external port. It also allows the maximum cable length to be calculated individually for each bus. This allows for additional capability externally.

- **IBM SCSI-2 Fast/Wide Streaming-RAID Adapter/A**

This adapter has the same performance advantages of the IBM SCSI-2 Fast/Wide Streaming Adapter/A plus a RAID array controller. This feature offers the additional data protection security inherent in RAID configurations. Also, the adapter microcode is optimized for database and video server environments.

Two independent SCSI buses are available for internal and external array configurations, further enhancing performance and fault tolerant configurations. The dual bus of the adapter allows for a maximum connection of up to 14 drives, seven on each individual bus. One bus cannot support internal and external devices simultaneously.

- **IBM SCSI-2 Fast PCI Adapter**

This adapter features a RISC processor which reduces the time the adapter needs to process SCSI commands. It also supports system DMA which reduces the CPU overhead by transferring data into the system memory directly.

It contains two SCSI-II ports, one internal and one external, and supports up to seven SCSI-II fast devices at a 10 MBps data transfer rate. It implements advanced features such as:

- Multi-thread I/O
- Scatter/Gather
- Tagged Command Queueing
- Synchronous and asynchronous Fast SCSI modes

- **IBM SCSI-2 Fast/Wide PCI Adapter**

In addition to the features supported with the IBM SCSI-2 Fast PCI Adapter characteristics, the IBM SCSI-2 Fast/Wide Adapter provides a 20 MBps data transfer rate.

It contains three SCSI-II ports:

- One 50-pin 8-bit internal connector
- One 68-pin 16-bit wide internal connector
- One 68-pin 16-bit wide external connector

It can support up to 15 units, on two of the three ports. If you connect external devices to the adapter, you can attach internal SCSI devices to one or the other, *but not to both* internal SCSI connectors.

- **IBM SCSI-2 Fast/Wide RAID PCI Adapter**

In addition to the features supported with the IBM SCSI-2 F/W PCI Adapter, the IBM SCSI-2 F/W RAID adapter provides a RAID controller. Please reference 1.6.5, "RAID Technology" on page 22 for a discussion on RAID.

1.6.3.1 Summary

The following tables summarize the features of the IBM SCSI adapters.

Attribute	SCSI Adapter with no Cache	Enhanced SCSI Adapter with Cache	SCSI-2 Fast/Wide Adapter/A	SCSI-2 Fast/Wide Streaming RAID Adapter/A
SCSI Bus Width	16-bit	32-bit	32-bit	32-bit
SCSI Data Transfer Rate	5 MBps	5 MBps	20 MBps	20 MBps
Micro Channel Data Transfer Rate	8.3 MBps	16.6 MBps	20 MBps	40/80 MBps
Parity	Optional	Optional	Yes	Yes
Tagged Command Queueing (TCQ)	N/A	N/A	Yes	Yes
Systems where implemented			500	500/720
<p>Note: Fast and Wide are data transfer methods as defined in SCSI-II. The 720 will support 80 MBps data streaming with the SCSI-2 Fast/Wide Streaming RAID Adapter/A N/A = not available</p>				

Attribute	PCI SCSI-2 Fast Adapter	PCI SCSI-2 Fast/Wide Adapter	PCI SCSI-2 Fast/Wide RAID Adapter
SCSI Bus Width	32-bit	32-bit	32-bit
SCSI Data Transfer Rate	10 MBps	20 MBps	20 MBps
Parity	Yes	Yes	Yes
Tagged Command Queueing (TCQ)	Yes	Yes	Yes
Systems where implemented	PC Server 300/310	PC Server 320/520	PC Server 320/520
<p>Note: Fast and Wide are data transfer methods as defined in SCSI-II</p>			

1.6.4 Hard Disk Drives

Ultimately, the hard disk drive is the component that has the most effect on subsystem performance. The following specs should be considered when evaluating hard disks in order to optimize performance:

- Average access time
- Maximum transfer rate
- On-board cache size

Average access time is one of the standard indicators of hard drive performance. This is the amount of time required for the drive to deliver data after the computer sends a read request. It is composed of two factors, the seek time and the rotational delay. The seek time is the time necessary to position the heads

to the desired cylinder of the disk. The latency is the amount of time it takes for the disk to rotate to the proper sector on that cylinder.

It should be noted that two disks of the same physical size, for example 3.5-inch disks, will differ in their access times with the larger capacity disk having a better access time. This is due to the fact that the distance between cylinders is shorter on the larger disk and, therefore, seek time is reduced. This is the primary reason that disk access times have been reduced as capacities have been increased.

Maximum transfer rate is the rate at which the device can deliver data back to the SCSI adapter. It mainly depends on the processor/DMA controller integrated on the device but can be no more than the SCSI maximum data transfer rate, for example 20 MBps for a SCSI-II Fast/Wide interface.

Caching is important for the same reason it is important on other subsystems; namely, it speeds up the time it takes to perform routine operations. For high performance, the drive should be able to provide write caching. With write caching, the drive signals the completion of the write immediately after it receives the data but before the data is actually written to disk. The system then continues to do other work while the hard disk is actually writing the data. Performance is significantly better because subsequent disk operations can be overlapped with this cached write operation.

The following table summarizes the specifications on current IBM PC Server hard disks:

Table 5. Summary of Disks Performance Characteristics

Disk	Average Seek Time	Burst Transfer Rate	Sustained Transfer Rate	Average Latency	MTBF (K hours)
1GB Fast	8.6 ms	10 MBps	3.7-4.5 MBps	5.56 ms	500
2GB Fast	9.5 ms	10 MBps	3.7-4.5 MBps	5.56 ms	500
1.08GB Fast	10.5 ms	10 MBps	3.2-4.0 MBps	5.56 ms	500
1.12GB F/W	6.9 ms	20 MBps	5.5-7.4 MBps	4.17 ms	1000
2.25GB F/W	7.5 ms	20 MBps	5.5-7.4 MBps	4.17 ms	1000
4.51GB F/W	8.0 ms	20 MBps	5.5-7.4 MBps	4.17 ms	1000

1.6.5 RAID Technology

Several factors have contributed to the growing popularity of disk arrays:

- Performance

The capacity of single large disks has grown rapidly, but the performance improvements have been modest, when compared to the advances made in the other subsystems that make up a computer system. The reason for this is that disks are mechanical devices, affected by delays in positioning read/write heads (seeks) and the rotation time of the media (Latency).

- Reliability

Disks are often among the least reliable components of the computer systems, yet the failure of a disk can result in the unrecoverable loss of vital business data, or at the very least a need to restore from tape with consequent delays.

- Cost

It is cheaper to provide a given storage capacity and a given performance level with several small disks connected together than with a single disk.

There is nothing unusual about connecting several disks to a computer to increase the amount of storage. Mainframes and minicomputers have always had banks of disks. It becomes a disk array when several disks are connected and accessed by the disk controller in a predetermined pattern designed to optimize performance and/or reliability.

Disk arrays seem to have been invented independently by a variety of groups, but it was the Computer Architecture Group at the University of California, Berkeley who invented the term RAID. RAID stands for *Redundant Array of Inexpensive Disks* and provides a method of classifying the different ways of using multiple disks to increase availability and performance.

1.6.6 RAID Classifications

The original RAID classification described five levels of RAID (RAID-1 through 5). RAID-0 (data-striping) and RAID-1 Enhanced (data stripe mirroring) have been added since the original levels were defined. RAID-0 is not a pure RAID type, since it does not provide any redundancy.

Different designs of arrays perform optimally in different environments. The two main environments are those where high transfer rates are very important, and those where a high I/O rate is needed, that is, applications requesting short length random records.

Table 6 shows the RAID array classifications, and is followed by brief descriptions of their designs and capabilities.

<i>Table 6. RAID Classifications</i>	
RAID Level	Description
RAID-0	Block Interleave Data Striping without Parity
RAID-1	Disk Mirroring/Duplexing
RAID-1 (Enhanced)	Data Stripe Mirroring
RAID-2	Bit Interleave Data Striping with Hamming Code
RAID-3	Bit Interleave Data Striping with Parity Disk
RAID-4	Block Interleave Data Striping with one Parity Disk
RAID-5	Block Interleave Data Striping with Skewed Parity

1.6.6.1 RAID-0 - Block Interleave Data Striping without Parity

Striping of data across multiple disk drives without parity protection is a disk data organization technique sometimes employed to maximize DASD subsystem performance (for example, Novell NetWare's *data scatter* option).

An additional benefit of this data organization is *drive spanning*. With data striped across multiple drives in an array, the logical drive size is the sum of the individual drive capacities. The maximum file size may be limited by the operating system.




Figure 11. RAID-0 (Block Interleave Data Striping without Parity)

Data striping improves the performance with large files since reads/writes are overlapped across all disks. However, reliability is decreased as the failure of one disk will result in a complete failure of the disk subsystem according to the formula:

$$\text{Mean Time to Failure} = \frac{\text{Mean Time to Failure of a single disk}}{\text{Number of Disks in the array}}$$

1.6.6.2 RAID-1 - Disk Mirroring/Duplexing

This approach keeps two complete copies of all data. Whenever the system makes an update to a disk, it duplicates that update to a second disk, thus mirroring the original. Either disk can fail, and the data is still accessible. Additionally, because there are two disks, a read request can be satisfied from either device, thus leading to improved performance and throughput. Some implementations optimize this by keeping the two disks 180 degrees out of phase with each other, thus minimizing latency.

However, mirroring is an expensive way of providing protection against data loss, because it doubles the amount of disk storage needed (as only 50% of the installed disk capacity is available for data storage).




Figure 12. RAID-1 (Disk Mirroring)

Disk mirroring involves duplicating the data from one disk onto a second disk using a single controller.

Disk duplexing is the same as mirroring in all respects, except that the disks are attached to separate controllers. The server can now tolerate the loss of one disk controller, as well as or instead of a disk, without loss of the disk subsystem's availability or the customer's data. Since each disk is attached to a separate controller, performance and throughput may be further improved.

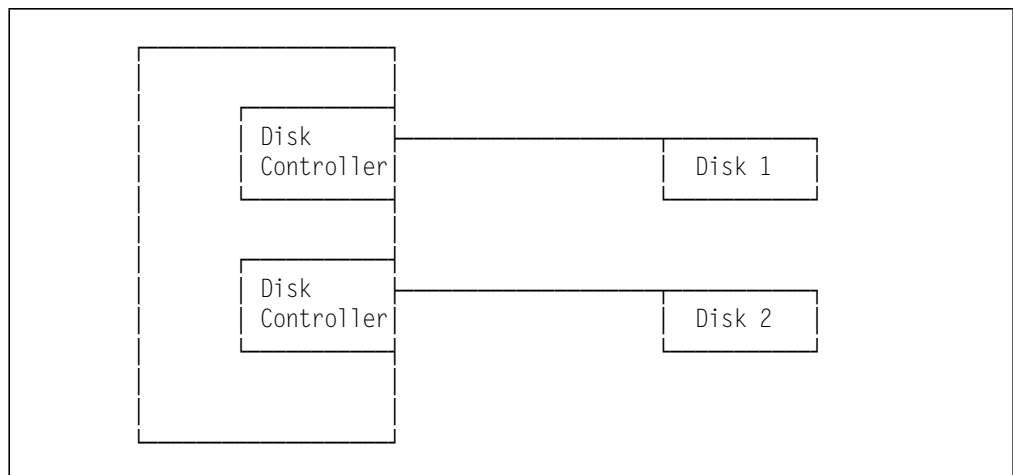


Figure 13. RAID-1 (Disk Duplexing)

1.6.6.3 RAID-1 Enhanced - Data Strip Mirroring

RAID level 1 supported by the IBM PC Server array models provides an enhanced feature for disk mirroring that stripes data and copies of the data across all the drives of the array. The first stripe is the data stripe; the second stripe is the mirror (copy) of the first data stripe, but, it is shifted over one drive. Because the data is mirrored, the capacity of the logical drive, when assigned to RAID 1 Enhanced, is 50 percent of the physical capacity of the hard disk drives in the array.




Figure 14. RAID-1 Enhanced, Data Strip Mirroring

Some vendors have also implemented another slight variation of RAID-1 and refer to it as RAID-10 since it combines features of RAID-1 and RAID-0. Others refer to this technique as RAID-6, which is the next available RAID level.

As shown in Figure 15, this solution consists of mirroring a striped (RAID-0) configuration. In this example, a RAID-0 configuration consisting of 2 drives (drive 1 and 2) is mirrored to drives 3 and 4.




Figure 15. RAID-6,10 - Mirroring of RAID 0 Drives

Performance and capacity are similar to RAID 1. 50% of total disk capacity is usable. However, this solution always uses an *even* number of disks. RAID-1 enhanced can use an *odd* number of disks. RAID-1 also can be mixed with others RAID levels.

1.6.6.4 RAID-2 - Bit Interleave Data Striping with Hamming Code

This type of array design is another form of data striping: it spreads the data across the disks one bit or one byte at a time in parallel. This is called bit (or byte) interleaving.

Thus, if there were five disks in the array, a sector on the first drive will contain bits 0 and 5, and so on of the data block; the same sector of the second drive will contain bits 1 and 6, and so on as shown in Figure 16.

RAID-2 improves on the 50% disk overhead in RAID-1 but still provides redundancy by using the Hamming code. This is the same algorithm used in ECC memory. The check bits can be generated on a nibble (4 bits), a byte (8 bits), a half word (16 bits) or a word (32 bits) basis but the technique works most efficiently with 32-bit words. Just like with ECC, it takes 7 check bits to implement the Hamming code on 32 bits of data.

Generating check bits by byte is probably the process used most frequently. For example, if data were grouped into bytes, 11 drives in total would be required, 8 for data and 3 for the check bits. An 8-3 configuration reduces the overhead to 27%.

Note: For clarity, the Hamming Code drives are not shown in Figure 16.

An array of this design will perform optimally when large data transfers are being performed. The host will see the array as one logical drive. The data transfer rate, however, will be the product of the number of drives in the array and the transfer rate of the individual drives.

This design is unable to handle multiple, simultaneous small requests for data, unlike the previous design; so, it is unlikely to satisfy the requirements for a transaction processing system that needs a high transaction rate.




Figure 16. RAID-2 (Bit Interleave Data Striping with Hamming Code)

1.6.6.5 RAID-3 - Bit Interleave Data Striping with Parity Disk

The use of additional disks to redundantly encode customer's data and guard against loss is referred to as check sum, disk parity or error correction code (ECC). The principle is the same as memory parity, where the data is guarded against the loss of a single bit.

Figure 17 shows an example of RAID-3. Four of the disks hold data, and can be accessed independently by the processor, while the fifth is hidden from the processor and stores the parity of the other four. Writing data to any of the disks (1, 2, 3 or 4) causes the parity to be recomputed and written to disk 5. If any of the data disks subsequently fail, the data can still be accessed by using the information from the other data disks along with the parity disk which is used to reconstruct the data on the failed disk.

Since the files are held on individually addressable disks, this design offers a high I/O rate. Compared to a single disk of similar capacity, this array has more actuators for the same amount of storage. These actuators will work in parallel, as opposed to the sequential operation of the single actuator, thus reducing average access times.

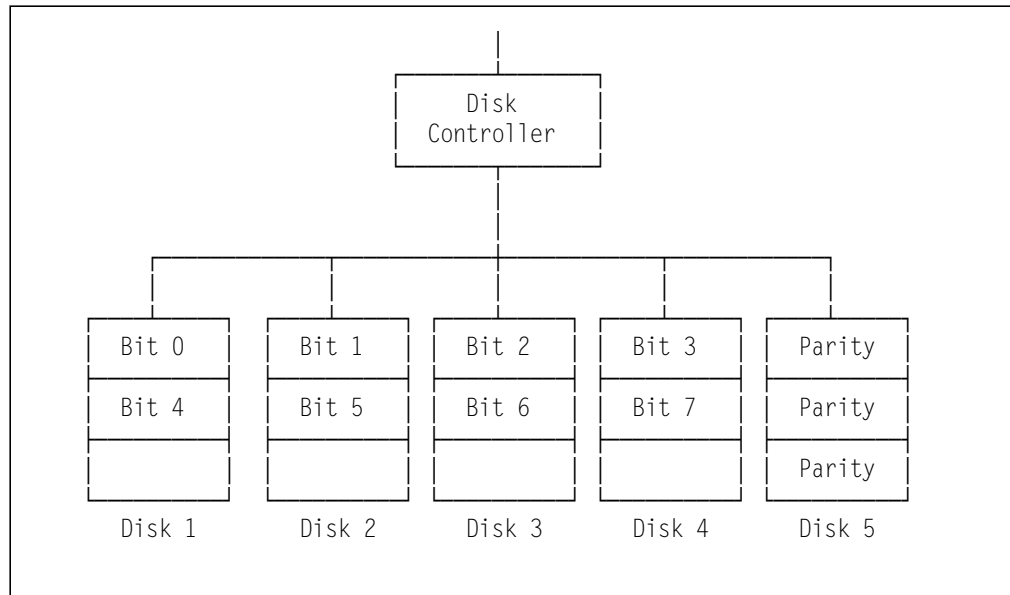


Figure 17. RAID-3 (Bit Interleave Data Striping with Parity Disk)

Multiple disks are used with the data scattered across them. One disk is used for parity checking for increased fault tolerance.

1.6.6.6 RAID-4 - Block Interleave Data Striping with One Parity Disk

The performance of bit-interleaved arrays in a transaction processing environment, where small records are being simultaneously read and written, is very poor. This can be compensated for by altering the striping technique, such that files are striped in block sizes that correspond to the record size being read. This will vary in different environments. Super-computer type applications may require a block size of 64KB, while 4KB will suffice for most DOS applications.

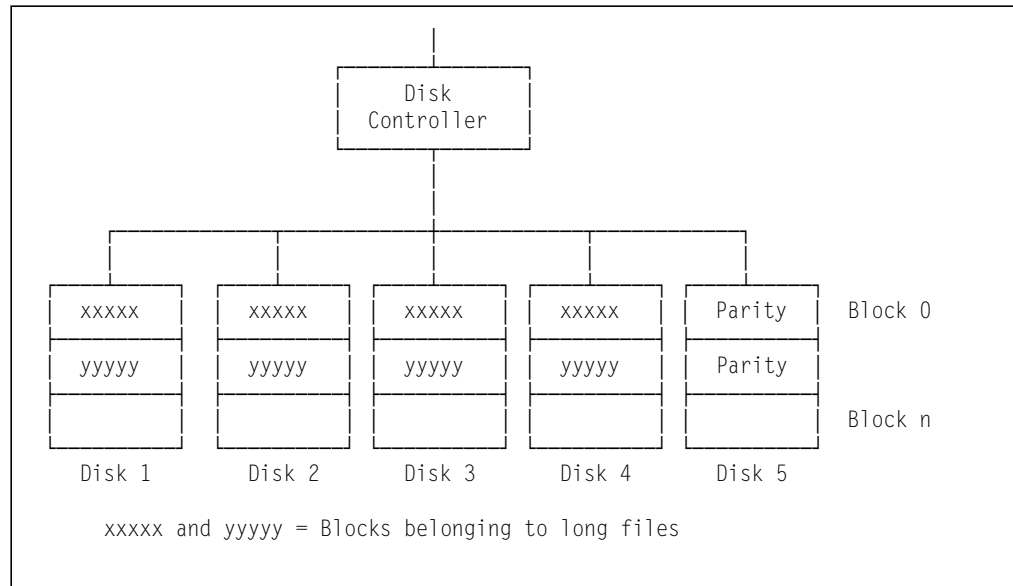


Figure 18. RAID-4 (Block Interleave Data Striping with One Parity Disk)

1.6.6.7 RAID-5 - Block Interleave Data Striping with Skewed Parity

The RAID-5 design tries to overcome all of these problems. Data is striped across the disks to ensure maximum read performance when accessing large files, and having the data striped in blocks improves the array's performance in a transaction processing environment. Parity information is stored on the array to guard against data loss. Skewing is used to remove the bottleneck that is created by storing all the parity information on a single drive.




Figure 19. RAID-5 (Block Interleave Data Striping with Skewed Parity)

1.6.6.8 Summary of RAID Performance Characteristics

RAID-0: Block Interleave Data Striping without parity

- Fastest data-rate performance
- Allows seek and drive latency to be performed in parallel
- Significantly outperforms single large disk

RAID-1: Disk Mirroring/Disk Duplexing and Data Strip mirroring (RAID-1, Enhanced)

- Fast and reliable, but requires 100% disk space overhead
- Data copied to each set of drives
- No performance degradation with a single disk failure
- RAID-1 enhanced provides mirroring with an odd number of drives

RAID-2: Bit Interleave Data Striping with Hamming Code

- Very fast for sequential applications, such as graphics modelling
- Almost never used with PC-based systems

RAID-3: Bit Interleave Data Striping with Parity

- Access to all drives to retrieve one record
- Best for large sequential reads
- Very poor for random transactions
- Poor for any write operations
- Faster than a single drive, but much slower than RAID-0 or RAID-1 in random environments

RAID-4: Block Interleave Data Striping with one Parity Disk

- Best for large sequential I/O
- Very poor write performance
- Faster than a single drive, but usually much slower than RAID-0 or RAID-1

RAID-5: Block Interleave Data Striping with Skewed Parity

- Best for random transactions
- Poor for large sequential reads if request is larger than block size
- Better write performance than RAID-3 and RAID-4
- Block size is key to performance, must be larger than typical request size
- Performance degrades in recovery mode (when a single drive has failed)

Table 7. Summary of RAID Performance Characteristics

RAID Level	Capacity	Large Transfers	High I/O Rate	Data Availability
Single Disk	Fixed (100%)	Good	Good	1
RAID-0	Excellent	Very Good	Very Good	Poor 2
RAID-1	Moderate (50%)	Good	Good	Good
RAID-2	Very Good	Good	Poor	Good
RAID-3	Very Good	Very Good	Poor	Good
RAID-4	Very Good	Very Good	Poor	Good
RAID-5	Very Good	Very Good	Good	Good

Note:

1 The MTBF (mean time before failure) for single disks can range from 10,000 to 1,000,000 hours.

2 Availability = MTBF of one disk divided by the number of disks in the array.

1.6.7 Recommendations

- Use IDE on smaller systems

IDE actually outperforms SCSI on systems where only one or two devices are attached. Several models of the IBM PC Server 300 and 320 lines implement IDE as an integrated controller on the planar board. This is more than adequate if no more than a couple of hard disks will be used.

- Distribute the workload on large systems

Research has shown that a single 66 MHz Pentium processor doing database transactions needs as many as 6-10 drives to optimize system performance. Therefore, do not determine the number of drives you need by simply adding up your total storage requirements and dividing this by the capacity of your drives. Instead, distribute the disk intensive workload from a single physical disk drive to multiple disk drives and use the striping features of RAID technology.

1.7 LAN Subsystem

The LAN adapter is another important component in the file server design. While there are many different types of LAN adapters, for file servers they fall into two main categories: bus master and shared RAM (non-bus master). The following discussion centers on the benefits of using bus master LAN adapters, although for small, lightly loaded LANs, non-bus master LAN adapters are quite adequate.

1.7.1 Shared RAM Adapters

Shared RAM adapters derive their name from the fact that they carry on-board RAM that is shared with the system processor. The memory on the adapter card is mapped into a reserved block of system address space known as the upper memory block (UMB) area. The UMB area is reserved for I/O adapters and is addressed between the addresses of 640KB and 1MB. The server processor can access this memory in the adapter in the same manner in which it accesses system memory.

Shared RAM can be 8, 16, 32, or 64KB in size depending on which adapter is used and how it is configured. Adapter cards with 64KB support RAM paging which allows the system to view the 64KB of memory on the card in four 16KB pages. This scenario only requires 16KB of contiguous system memory instead of the 64KB required when not using RAM paging. All IBM NetBIOS products support RAM paging.

The starting address of the shared RAM area is determined by the adapter device driver, switch settings, or in the case of an EISA or MCA adapter, via the setup utility or the reference diskette, respectively.

The main disadvantage of shared RAM architecture is that any data movement between the shared RAM area and system memory must be done under direct control of the system's CPU. This movement of data to and from the shared RAM must be done because applications cannot operate on data while it resides in the shared RAM area. To compound matters, MOVE instructions from/to the shared RAM are much slower than the same MOVE instruction from/to the

system memory because they occur across an I/O expansion bus. This means that when shared RAM adapters are involved, the CPU spends a significant amount of time doing the primitive task of moving data from point A to point B.

On lightly loaded servers providing traditional productivity applications such as word-processing, spreadsheets, and print sharing, this is not really a problem. But, for applications such as database or for more heavily loaded file servers, this can be a major source of performance degradation.

The IBM Token Ring Network 16/4 Adapters I and II for MCA and ISA are examples of shared RAM adapters.

1.7.2 Bus Master Adapters

Bus master adapters utilize on-board Direct Memory Access (DMA) controllers to transfer data directly between the adapter and the system memory without involving the system processor. The primary advantage of this architecture is that it frees up the system processor to perform other tasks, which is especially important in the server environment.

The IBM 16/4 Token-Ring bus master adapter/A: This adapter was the first generation of bus master LAN adapters from IBM. It employed the 64KB on-board adapter memory as a frame buffer which was used to assemble frames before they were sent to the server or sent from the server to the network. The time elasticity provided by this buffer allowed the token-ring chip set to complete its processing and forwarding of the frame before the frame was lost; this is a condition known as overrun (receive) or underrun (transmit).

This adapter was a 16-bit Micro Channel bus master capable of burst mode DMA. Due to the 24-bit addressing capabilities of the adapter, it was limited to using only the first 16MB of system address memory.

IBM LANStreamer Family of Adapter Cards: The LANStreamer technology employs a completely different design to previous IBM LAN adapters. The LANStreamer utilizes a revolutionary chip set that is capable of processing token-ring frames without using memory as a frame buffer. It does it *on-the-fly* as the frames are passing through the adapter. Therefore, the latency of assembling frames from an on-card buffer is eliminated.

This low latency chip set is the key to the small-frame performance characteristics of the LANStreamer adapter. The throughput for the LANStreamer Token-Ring MC 32 Adapter/A is quite high relative to its predecessors, especially for small frames. This is extremely important in client/server environments where research has shown that the vast majority of frames on the network are less than 128 bytes.

Another advantage of this technology is that since adapter memory buffers are no longer required, the adapter is less expensive to produce.

A consequence of the high LANStreamer throughput is that the LAN adapter is not usually the bottleneck in the system. Also, a side effect of using LANStreamer technology could be the higher CPU utilization. This sometimes happens because the LANStreamer adapter can pass significantly more data to the server than earlier adapters. This corresponds to more frames per second that must be processed by the server network operating system. Higher throughput is the desired effect but what this also means is that the bottleneck

sometimes moves quickly to the CPU when servers are upgraded to incorporate LANStreamer technology.

Of course, other components can emerge as the bottleneck as throughput increases. The wire (network bandwidth) itself can become a bottleneck if throughput requirements overwhelm the ability of the network technology being used. For example, if an application requires 3 MBps of throughput, then a token-ring at 16 Mbps will not perform the task. In this case a different network technology must be employed.

For more discussion of hardware performance tuning, please see 5.1, "Hardware Tuning" on page 167.

The LANStreamer technology is used in the IBM Auto LANStreamer Adapters for PCI and MCA as well as the EtherStreamer and Dual EtherStreamer MC 32 LAN adapters.

Note

The EtherStreamer LAN adapter supports full duplex mode, which allows the adapter to transmit as well as receive at the same time. This provides an effective throughput of 20 Mbps (10 Mbps on the receive channel and 10 Mbps on the transmit channel). To implement this feature, an external switching unit is required.

1.7.3 PeerMaster Technology

The PeerMaster technology takes LAN adapters one step forward by incorporating an on-board Intel i960 processor. This processing power is used to implement per port switching on the adapter without the need for an external switch. With this capability, frames can be switched between ports on the adapter, bypassing the file server CPU totally.

If more than one card is installed, packets can be switched both intra- and inter-card. The adapters utilize the Micro Channel to switch inter-card and can transfer data at the very high speed of 640 Mbps.

The IBM Quad PeerMaster Adapter is a four-port Ethernet adapter which utilizes this technology. It is a 32-bit Micro Channel bus master adapter capable of utilizing the 80 MBps data streaming mode across the bus either to/from system memory or peer to peer with another PeerMaster adapter.

It ships with 1MB of memory. Each port on an adapter serves a separate Ethernet segment. Up to six of these adapters can reside on a single server and up to 24 segments can be defined in a single server.

This adapter can also be used to create virtual networks (VNETs). Using VNETs, the NOS sees multiple adapter ports as a single network, eliminating the need to implement the traditional router function either internal or external to the file server.

The Ethernet Quad PeerMaster Adapter is particularly appropriate when there is a need for:

- Switching/Bridging traffic among multiple Ethernet segments
- Attaching more than eight Ethernet 10Base-T segments to the server

- Attaching more than four Ethernet 10Base-2 segments to the server
- Providing switching between 10Base-T and 10Base-2 segments
- Conserving server slots

An add-on to NetFinity provides an advanced Ethernet subsystem management tool. Parameters such as packets/second or total throughput can be monitored for each port, for traffic within an adapter, or for traffic between adapters.

By using NetFinity, you can graphically view the data, monitor for predefined thresholds, and optionally generate SNMP alerts.

1.8 Security Features

This section discusses some technologies used in IBM PC Servers to comply with the United States Department of Defense (DoD) security requirements. Security features in the IBM PC Server line vary by model and all models do not have all the security features described here. Check the User's Handbook that is shipped with the system, to see what features your system contains.

DoD requirements have been very influential in defining security standards used on computer system (both hardware and software) implementations around the world. The source for these requirements is the *Department of Defense, Trusted Computer System Evaluation Criteria, DoD 5200.28 STD*, dated 12/85. The essence of the requirements is contained in the Assurance section, Requirement 6: a "trusted mechanism must be continuously protected against tampering and/or unauthorized changes...". The National Computer Security Center (NCSC) evaluates computer system security products with respect to the criteria defined by the U.S. Department of Defense.

There are seven computer system security product classifications in the DoD requirements: A1, B3, B2, B1, C2, C1, and D. The requirements for these classifications fall into four basic groups: security policy, accountability, assurance, and documentation. Several criteria, which vary by security classification, are specified in each of these groups. Currently, A1 is the highest classification, followed by B3, B2, and so on. The C2 classification satisfies most of the security requirements for personal computing environments.

LogicLock: On the IBM MCA PC Servers, IBM implements a collection of security features referred to as the LogicLock security system. LogicLock is designed to be hardware compliant with the C2 security classification. It goes far beyond basic PC security systems in its design to protect data against unauthorized access.

LogicLock security features include:

- Tamper-evident switches
- Optional secure I/O cables
- Privileged-access password
- Optional secure removable media
- Selectable drive startup
- Unattended start mode

1.8.1 Tamper-Evident Cover

Systems equipped with a tamper-evident cover have a key lock for their covers and internal I/O devices. In the locked position, it mechanically prevents the covers from being removed. The key has been changed to a type that can be duplicated only by the manufacturer.

If the covers are forced open, an electro-mechanical switch and perimeter sensor detect the intrusion. If the computer was on during the break-in attempt, depending on options specified during system setup, it will either defer action until the next IPL, lock up, or pass a non-maskable interrupt (NMI) to the software.

The next time the computer is started, the power-on self-test (POST) routine displays a message informing the user of the intrusion, and requires that the automatic configuration program be run before the computer can be used. This is done to flag any configuration changes that may have occurred due to the intrusion (for example, removal of a disk drive). In addition, the system cannot be used without the privileged-access password if it has been set. There is a provision for maintenance that allows the system to be used without the covers in place. However, to use this feature, the key must have been used to remove the covers.

Other systems may have lockable covers. However, it is not that difficult to pry the system unit cover off, disable or unplug the key mechanism, and get inside the system. The tamper-evident mechanism is an important feature which flags the intrusion and prevents the operation of the system after a forced entry has occurred. This detection feature is very valuable for detecting the person most likely to break into the secured workstation, the user. Once the machine has been disabled, the system owner or administrator must be contacted to reset the system.

1.8.2 Secure I/O Cables

This rear-panel security option is an enclosure that is secured to the back of the computer by the cover lock. Its function is to prevent the cables from being removed and other cables from being attached. This effectively secures the serial, parallel, and SCSI cables, as well as other ports and cables provided by adapters. This is because it prevents someone from attaching a device through these connectors and gaining access to the data in the system.

The cable cover also has a tamper-evident feature.

1.8.3 Passwords

IBM PC Servers are equipped with several layers of password protection. The most basic is the power-on password. The power-on password must be entered correctly each time the system is turned on. After three incorrect attempts, the system must be turned off and back on in order to try again.

The keyboard password is another level of password protection and is used to lock the keyboard without turning the computer off. It also prevents rebooting the system by pressing the Ctrl+Alt+Del keys.

IBM PC Servers also provide an unattended server mode (or network server mode). This mode allows other computers to access a fixed disk drive on a server even though the keyboard is locked. This is useful, for example, when

there is a power failure; the machine is able to recover with the keyboard lock still in place.

1.8.3.1 Privileged-Access Password

Because the power-on and keyboard passwords can be defeated by deactivating the battery inside the system, another level of password protection is provided. This security feature is called the privileged-access password. It provides a much higher level of security. The privileged-access password restricts access to system programs, prevents the IPL source and sequence from being changed, and effectively deters unauthorized modifications to the hardware. Also, if a forced entry is detected by the tamper-evident cover switch, the privileged-access password (if it has been set) must be used in order to make the system operational again.

The privileged-access password is stored in a special type of read only memory called flash EEPROM. EEPROM is an acronym for electrically erasable programmable read only memory.

Systems are shipped with the privileged-access password disabled. To set this password, a jumper on the system board must be moved in order to put the system in the change state. Once this password is set, it cannot be overridden or removed by an unauthorized person.

Attention - Forgotten Password

If the administrator misplaces or forgets the privileged-access password, the system board will have to be replaced. There is no way to reset a forgotten privileged-access password.

1.8.4 Secure Removable Media

An optional 2.88MB diskette drive with security features is available for all IBM PC Server systems. The diskette drive is a 3.5-inch, one-inch high drive with media sense capability for the standard diskette capacities of 720KB, 1.44 MB, and 2.88MB. It can read and write data up to a formatted capacity of 2.88MB, while maintaining read and write capability with 720KB and 1.44MB diskette drives.

A control signal has been added to the diskette interface that supports LOCK, UNLOCK, and EJECT commands issued by the operating system. If the privileged-access password is not set, the diskette is unlocked during POST. If the password is set, the boot process does not unlock the diskette drive unless it is the designated IPL source. In this case, the LOCK and UNLOCK state is controlled by an operating system utility. For SCSI devices, there is a proposed standard UNLOCK command. In this case, the operating system will control the LOCK command if the privileged-access password is set. Access to the unlocking function with specific user authorization can be controlled by secured system software.

In the event of power loss, the system retains its state (secured or unsecured) independent of the state of the battery. A diskette can be inserted in the drive, but it cannot be removed if the power is off. When the drive is turned on and locked, the media cannot be inserted or removed.

1.8.5 Selectable Drive Startup

Selectable drive startup allows the system owner or administrator to select the IPL source and sequence. This allows the system owner to control the IPL source, but prevents the user from modifying the source and sequence. For example, the diskette drive can be excluded as an IPL source. This feature helps to ensure that the system owner's specified operating system is loaded.

The IPL sequence is stored in the system EEPROM, and can only be changed using the privileged-access password. Storage of the IPL sequence in the EEPROM protects it from being deactivated by removing the battery. The setup routine ensures that at least one IPL source is specified if the privileged-access password is used.

1.8.6 Unattended Start Mode

The unattended start mode automatically restarts the server after a power failure and resumes normal operation, without operator intervention.

It locks the keyboard when the system is powered on, but it allows the operating system and startup files to be loaded. The keyboard remains locked until the power-on password is entered.

This mode is useful for unattended operations because it allows authorized network user access to information on the server but prohibits unauthorized access via the system keyboard.

When the system is in the unattended mode, the password prompt will not appear unless an attempt to start the system from a diskette or other removable media is issued. If you start the system from a removable media, the password prompt will appear and you must enter the correct power-on password to continue.

1.9 Systems Management

Systems management is an important element of a successful LAN. The IBM PC Server brand ships with a very powerful systems and network management tool called NetFinity. In this section, we look at the capabilities of NetFinity; first, we need to take a look at some of the underlying technology which NetFinity has incorporated. NetFinity incorporates DMI which is an emerging standard for managing desktop machines and SNMP which is an established network management protocol. We take a look at each of these in the following sections.

1.9.1 DMI

The Desktop Management Interface (DMI) is a standard developed by an industry consortium that simplifies management of hardware and software products attached to, or installed in, a computer system. The computer system can be a stand-alone desktop system, a node on a network, or a network server. DMI is designed to work across desktop operating systems, environments, hardware platforms, and architectures.

DMI provides a way to obtain, in a standardized format, information about the hardware and software products installed in the system. Once this data is obtained, management applications written to the DMI specs can use this data to

manage those products. As DMI technology evolves, installation and management in desktops and servers will become easier.

It should be noted that the DMI specs say nothing about the transport protocol that is used between the manageable products and the management applications. Both of these elements of a DMI compliant system can be implemented using any native transport protocol available in the system.

The DMI architecture includes:

- Communicating service layer
- Management information format (MIF)
- Management interface (MI)
- Component interface (CI)

1.9.1.1 Communicating Service Layer

The service layer is the desktop resident program that is responsible for all DMI activities. Service layer communication is a permanent background task or process that is always ready for an asynchronous request.

The service layer is an information broker, handling commands from management applications, retrieving the requested information from the MIF database or passing the request on to manageable products as needed via the CI. The service layer also handles indications from manageable products and passes that information on to the management applications.

Management applications: These are remote or local programs for changing, interrogating, controlling, tracking and listing the elements of a desktop system and its components.

A management application can be a local diagnostic or installation program, a simple browser that walks through the MIF database on the local system or any other agent which redirects information from the DMI over a network.

Manageable products: These include hardware, software or peripherals that occupy or are attached to a desktop computer or network server, such as hard disks, word processors, CD-ROMs, printers, motherboards, operating systems, spreadsheets, graphics cards, sound cards, modems, etc.

Each manageable product provides information to the MIF database by means of a file which contains the pertinent management information for that product. Manageable products, once installed, communicate with the service layer through the component interface. They receive management commands from the service layer and return information about their status to the service layer.

1.9.1.2 Management Information Format (MIF)

A management information format (MIF) is a simple ASCII text file describing a product's manageable attributes, grouped in ways that make sense. The MIF has a defined grammar and syntax. Each product has its own MIF file.

When a manageable product is initially installed into the system, the information in its MIF file is added to the MIF database and is available to the service layer and thus to management applications.

The simplest MIF file contains only the component ID group, but MIFs can become as complex as needed for any given product.

1.9.1.3 Management Interface (MI)

The management interface (MI) shields management applications from the different mechanism used to obtain management information for products within a desktop system.

The MI allows a management application to query for a list of manageable products, access specific components and get and set individual attributes. Additionally, the MI allows a management application to tell the service layer to send back information about indications from manageable products.

The MI commands provide three types of operations to control manageable products:

- Get
- Set
- List

Get allows a management application to get the current value of individual attributes or group of attributes.

Set allows writeable attributes to be changed.

List allows management applications to read the MIF descriptions of manageable products, without having to retrieve the attribute values for that product. Thus, a management application can query a system and retrieve useful information about the contents of the system, with no previous knowledge of that system.

1.9.1.4 Component Interface (CI)

The component interface (CI) handles communication between manageable products and the service layer. The CI communicates with manageable products for get and set operations. It also receives indications from manageable products and passes those to the MI. Active instrumentation allows components to provide accurate, real-time information whenever the value is requested. A single component attribute can have a single value, or it can be obtained from a table using index keys.

1.9.2 SNMP

Simple Network Management Protocol (SNMP) is a network management protocol defined within the TCP/IP transport protocol standard. It is a rather generic protocol by which management information for a wide variety of network elements may be inspected or altered by logically remote users. It is a transaction-oriented protocol based on an interaction between managers and agents. The SNMP manager communicates with its agents. Agents gather management data and store it, while managers solicit this data and process it.

The SNMP architectural model has been a collection of network management stations and network elements such as gateways, routers and hosts. These elements act as servers and contain management agents which perform the network management functions requested by the network elements. The network management stations act as clients; they run the management applications which monitor and control network elements.

SNMP provides a means of communicating between the network management stations and the agents in the network resources. This information can be status information, counters, identifiers, etc.

The SNMP manager continuously polls the agents for error and statistical data. The performance of the network will be dependent upon the setting of the polling interval.

1.9.2.1 Management Information Base (MIB)

The Management Information Base is a collection of information about physical and logical characteristics of network objects. The individual pieces of information that comprise a MIB are called MIB objects and they reside on the agent system. These MIB objects can be accessed and changed by the agent at the manager's request.

The MIB is usually made up of two components:

- MIB II

This is a standard definition which defines the data layout (length of fields, what the field is to contain, etc.) for the management data for the resource. An example would be the resource name and address.

- MIB Extension

This incorporates unique information about a resource. It is defined by the manufacturer of the resource that is being managed. These are usually unique and proprietary in nature.

1.9.2.2 SNMP Agent

The SNMP agent is responsible for managed resources and keeps data about the resources in a MIB. The SNMP agent has two responsibilities:

1. To place error and statistical data into the MIB fields
2. To react to changes in certain fields made by the manager

1.9.2.3 SNMP Manager

An SNMP manager has the ability to issue the SNMP commands and be the end point for traps being sent by the agent. Commands are sent to the agent using the MIB as a communication vehicle.

1.9.2.4 Traps

In a network managed with SNMP, network events are called traps. A trap is generally a network condition detected by an SNMP agent that requires immediate attention by the system administrator. It is a message sent from an agent to a manager without a specific request for the manager.

SNMP defines six generic types of traps and allows definitions of enterprise-specific traps. This trap structure provides the following information:

- The particular agent object that was affected
- Event description(including trap number)
- Time stamp
- Optional enterprise-specific trap identification
- List of variables describing the trap

In summary, the following events describe the interactions that take place in an SNMP-managed network:

- Agents maintain vital information about their respective devices and networks. This information is stored in a MIB.
- The SNMP manager polls each agent for MIB information and stores and displays this information at the SNMP manager station. In this manner, the system administrator can manage the entire network from one management station.
- Agents also have the ability to send unsolicited data to the SNMP manager. This is called a trap.

1.9.3 NetFinity

IBM NetFinity is a family of distributed applications designed to enhance the system monitoring and management capabilities of a network. NetFinity has a flexible, modular design that allows for a variety of system-specific configurations. NetFinity is able to manage IBM and non-IBM desktops and servers and supports most client operating systems. It has management capabilities on Windows and OS/2. It is designed to work with the existing protocols on the network and includes support for NetBIOS, IPX, TCP/IP, and even ASYNC/Serial modem LAN protocols.

NetFinity is delivered as two components:

- NetFinity Services
- NetFinity Manager

1.9.3.1 NetFinity Services

NetFinity Services is the client portion of the system. This is a foundation that provides the underlying services for several levels of administration, including remote system and user management facilities. Figure 20 shows the main folder which opens for the NetFinity Services.

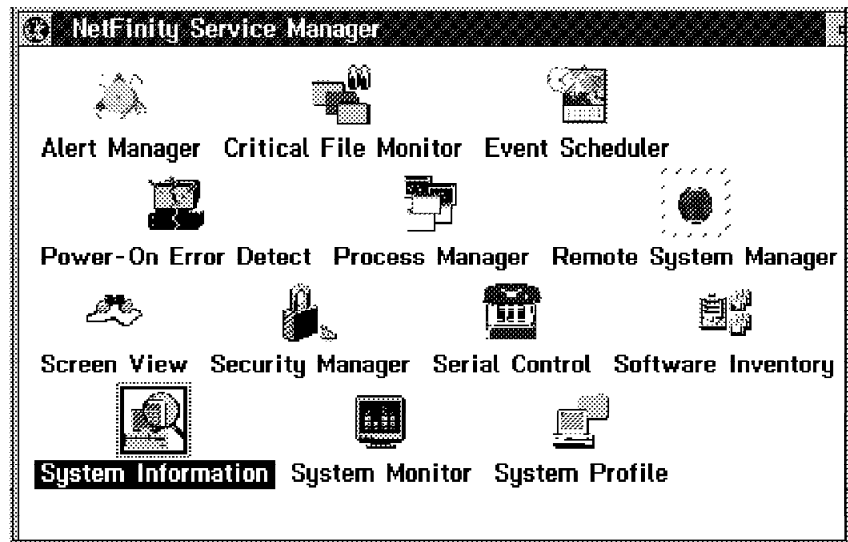


Figure 20. NetFinity Services Folder

NetFinity Services provides the following functions:

- **System Information**
Provides details regarding specific hardware and software configurations and user/application settings

- **System Profile**
Allows the systems administrator to define additional information for each system, such as location
- **System Monitor**
Provides system performance monitoring utilities, such as CPU, DASD, and Memory
- **Critical File Monitor**
Can generate alerts when critical files are changed or deleted
- **System Partition Access**
Allows an administrator to access the system partition on remote PS/2s
- **Predictive Failure Analysis**
Monitors PFA-enabled drives for errors
- **ServerGuard Service**
Can monitor server environmentals such as temperature and voltage and can remotely boot the system (requires the ServerGuard Adapter)
- **RAID Manager**
Allows the administrator to view the array configuration
- **ECC Configuration**
Allows the administrator to set thresholds for ECC scrubbing, counting, and NMI
- **Security Manager**
Controls which NetFinity Services each manager can access for a given system
- **Alert Manager**
Fully customizable alert generation, logging, and forwarding, and also has the ability to filter and forward SNMP alerts
- **Serial Control Service**
Supports an ASYNC connection between two NetFinity systems, either managers or clients

Note

The NetFinity Services installation is an intelligent process. It only installs the services for the configuration of your machine. Hence, you will only see icons for the services which are available to you. For example, on non-array machines, the RAID utility icon will not be present. It also allows you to update and add new services at a later date without reinstalling the base product.

NetFinity Services supports both IBM and non-IBM systems. It supports PCI, Micro-Channel, and EISA bus-based systems. It supports most client operating systems including DOS/Windows, Windows for Workgroups, OS/2 2.X, OS/2 Warp, OS/2 Warp Connect, and OS/2 SMP.

It also supports Novell NetWare. This means that there is a version of NetFinity Services which installs as a NetWare NLM on the file server and allows the server to be managed by a NetFinity Manager station.

NetFinity Services can also be installed on a Windows NT server and used to manage this platform as well.

NetFinity Services can be configured in three client modes of operation:

- **Stand-alone client**

Stand-alone mode allows an individual user, who is not connected to a network, to effectively manage or monitor their own system including hardware, resources and performance.

- **Passive client**

With the passive client installed on a LAN workstation, a NetFinity Manager is able to fully manage and monitor the resources and configuration setting of the workstation. However, with the passive mode installed, that same client is not able to perform its own management task locally. This mode is most effective for LAN administrators who do not want individual users to have management capability on an individual basis.

- **Active client**

The active client allows the NetFinity Manager to manage and monitor the resources and configuration setting of the workstation. In comparison to the passive client mode, the active client mode allows local users to perform their own subset of local system management tasks.

1.9.3.2 NetFinity Manager

The NetFinity Manager is the set of applications that is installed on the managing platform. It automates the collection of data from managed clients and archives it into a database, which maintains specific, unique workstation data and configuration settings. NetFinity also supports database exports into Lotus Notes or DB2/2.

In addition to logging the information in a database, an administrator may dynamically monitor performance on client workstations. An administrator may also identify resource parameters to monitor and maintain.

NetFinity Manager has the ability to discover LAN-attached NetFinity client workstations automatically. For example, if a new NetFinity client appears on the LAN, it will be sensed by the manager services and, from that point on, will be automatically included as a managed device within the profile.

A profile is a set of managed devices grouped by a set of unique attributes such as system processor types, operating systems, installed transport protocols, and administrator defined keywords. The keywords can be descriptors of systems, users or profiles. These NetFinity profiles can be dynamically declared, reset and maintained on an *as needed* basis by the administrator.

NetFinity Manager includes the following functions:

- **Remote Systems Manager**

This allows managers to access remote NetFinity-managed machines on a LAN, WAN, or a serial link. The manager can access the NetFinity services as if the manager was at that machine.

- **File Transfer**
Can send/receive files to the remote system.
- **Remote Session**
Can open a remote console to the managed device.
- **Screen View**
Can take a snapshot of any screen on the remote device.
- **DMI Browser**
Enables you to view information about DMI compliant hardware and software.
- **Process Manager**
Enables you to start/stop/view processes running on the managed device.
- **Software Inventory**
Can scan remote device for installed software using a software *dictionary*.
- **POST Error Detect**
Can detect and log errors at Power on System Test (POST) time on managed devices.
- **Event Scheduler**
Used to automate the execution of a service on one or multiple systems in a profile.

1.9.3.3 NetFinity Architecture

Each NetFinity service is comprised of two separate executables. One is a unique graphical user interface for the applicable operating system. The second is a native operating system executable, which is known as the base executable. The base executable is the code that performs the client management and monitoring tasks for each unique workstation. Communication between the GUI and the base executable is handled by the NetFinity IPC (inter-process communication) mechanism.

Using this IPC within the LAN, NetFinity was designed to provide a peer-to-peer platform architecture, which does not require a management server or a dedicated management console. From this design, a manager may take control of the NetFinity client system to perform all NetFinity administrative and problem reconciliation tasks as if they were the local user's tasks. Additionally, IBM has been able to isolate NetFinity from any network, protocol or operating system layer dependencies. In essence, IBM uses the existing transport layers within the installed network to allow NetFinity to communicate between NetFinity Manager and NetFinity Services. Since IPC resides on top of the Media Access Control (MAC) layer, it simply communicates between the installed NetFinity modules and services, utilizing the transport mechanism within the workstation.

If the transport layer between the two NetFinity workstations is dissimilar, then NetFinity utilizes a mapper (within a Manager), which receives data packets from one transport and, using NetFinity manager, is able to re-wrap the packets for transport into the foreign network.

When two NetFinity systems are connected in a networked environment, they communicate via the IPC into the mapper, and then subsequently into a NetFinity

Manager services and system module. This feature provide an extensive capability to merge dissimilar LANs into a single view of NetFinity managed assets.

1.9.3.4 DMI Support

NetFinity is the first product available to customers that includes DMI support. NetFinity implementation of DMI support provides instrumentation from its System Information Tool to the DMI service layer for both OS/2 and Windows clients. To accomplish this, IBM has delivered a DMI component agent that allows a NetFinity Manager to access a client desktop MIF database to deliver system specific information back into NetFinity DMI Browser. Today, NetFinity not only supports local DMI browsing capabilities, but also DMI alerting and a remote browser service.

1.9.3.5 Interoperability with Other Management Tools

NetFinity supports coexistence with almost any other LAN or enterprise management product, whether from IBM or other vendors. To provide for this integration, NetFinity Alert Manager was developed to allow its alerts to be captured and forwarded into any SNMP compliant management application. SNMP alerts are recognizable by the vast majority of different management tools in the market today. With this NetFinity feature, administrators can integrate sophisticated systems management functions with their existing SNMP-based systems management applications, such as IBM SystemView and HP OpenView. In direct support of heterogeneous LAN management environments, NetFinity is also launchable from within NetView for OS/2 and Novell NMS (NetWare Management Services).

1.9.4 SystemView

SystemView is an integrated solution for the management of information processing resources across heterogeneous environments. The objective of SystemView is to address customer requirements to increase systems and network availability, and to improve the productivity of personnel involved in system management.

The environments in which SystemView applies ranges from small stand-alone LANs to large multiprocessor systems. Depending on the environment, you might see OS/2, AIX/6000, or VM/ESA acting as a managing system with a consistent implementations on each platform. The benefit is the flexibility to deploy management functionality where it best suits the business needs. This also reduces management traffic, since the management is not implemented on a single platform.

1.9.4.1 SystemView Structure

An integral part of SystemView is a structure that enables the building of consistent systems management applications. The structure allows system management applications based on defined open architectures.

The SystemView structure consists of:

1. End-Use dimension

describes the facilities and guidelines for providing a consistent user interface to the management system. The end-use dimension provides a task-oriented, consistent look and feel to the user, through the powerful,

graphical drag and drop capability of OS/2 or AIX/6000. The primary benefit of the end-use dimension is the end-user productivity.

Some examples of products that have implemented SystemView conforming interfaces are:

- LAN Network Manager
- NetView for OS/2
- NetView for AIX
- DataHub Family of Database Management Products
- NetView Performance Monitor
- NetView Distribution Manager/2
- Service Level Reporter

2. Application dimension

Application dimension is a consistent environment for developing systems management applications. The primary benefits of the application dimension are automation of system management, integrated functions, and open implementation.

The application dimension consists of all system management tasks, grouped into six management disciplines:

- **Business:** Provides inventory management, financial administration, business planning, security management and management services for all enterprise-wide computer related facilities
- **Change:** Schedules, distributes, applies, and tracks changes to the enterprise information system
- **Configuration:** Manages the physical and logical properties of resources and their relationships such as connections and dependencies
- **Operations:** Manages the use of systems and resources to support the enterprise information processing workloads
- **Performance:** Collects performance data, tunes the information systems to meet service level goals and does capacity planning
- **Problem:** Detects, analyzes, corrects and tracks incidents and problems in system operations

IBM provides programming interfaces to achieve a cohesive systems management architecture. An example of this is implemented in the NetView for OS/2 product.

3. Data dimension

The data dimension provides systems management integration through the common modeling and sharing of management data. Data sharing among management application is a key systems management requirement. Enterprises want to enter management data into their systems only once. However, enterprises want the information in the information base to be accessible in an efficient form for applications needing it.

The SystemView data dimension provides a structure in which systems management applications and functions utilize a standardized set of data definitions and database facilities.

The following are the primary characteristics of the data dimension:

- **Common object definitions:** Products and applications share the data definitions in the SystemView data model. This allows the products and applications to utilize the data rather than replicate it.

- **Open and extendable data model:** This specifies the data definitions that represent the information processing data of an enterprise. The SystemView data dimension includes descriptions of the characteristics of resources and the relationships among them.
- **Heterogeneous access:** This structure provides for access of systems management data across heterogeneous platforms through the interfaces.
- **Support of existing applications:** Existing systems management applications are supported by the structure. Modifications to existing applications are required in order to participate in broad data sharing.

4. Managed resource dimension

Managed resource dimension allows these resources to benefit from SystemView applications to a greater degree. It allows management applications to be shared across multiple resources of similar types, by ensuring consistent definition and allowing classification of the resources. In this way, common attributes, behaviors, operations and notifications can be layered in a hierarchical classification, applying to the highest appropriate point in the classification hierarchy. This ensures the consistency and open approach required to deal with the large number and complexity of specific resources that have to be managed.

1.9.4.2 SystemView Management Protocols

The SystemView design has the objective to handle multiple management protocols. It allows multiple management protocol for resources such as SNA, SNMP and others, and CMIP for the management of agents.

The result is the ability to support currently available protocols, thus allowing the appropriate protocols to be selected. The benefit is the protection of the investment in current applications, and providing for growth in new technologies where appropriate.

1.10 Fault Tolerance

New hardware technologies have increased the reliability of computers used as network servers. RAID technology, hot swappable disk drives, error correcting memory and redundant power supplies are all effective at helping to reduce server down time caused by hardware failure.

Advances in software have also increased server availability. Disk mirroring, UPS monitoring software, and tape backup systems have further reduced the potential for server down time.

However, even with these advances, there are still many single points of failure which can bring a server down. For example, a failure on the system planar board will often result in a server crash, and there is no way to anticipate it. Also software products running on the server present an ever increasing chance for server failure as well.

For mission critical applications, there needs to be a way of protecting against such single points of failure.

Novell offers one such solution called NetWare System Fault Tolerance level III or SFT III.

1.10.1 NetWare SFT III

NetWare SFT III is a special version of the NetWare 3.x or 4.x NOS which adds a high degree of fault tolerance. It is composed of two servers, a primary and a secondary, which are mirrored together.

To clients on the network, only the primary server appears to be active. The secondary server remains in the background; however, it maintains the same essential memory image and the same disk contents as the primary server. If the primary server fails or halts, the secondary server automatically becomes the new primary server. This process is instantaneous and transparent to any client using the server.

NetWare SFT III has the following hardware requirements:

1. Two Servers (identical in make and model with at least an i386 processor)
2. 12 MB of RAM minimum in each server
3. Identical disk subsystems on each server
4. Identical video subsystems on each server
5. Two Mirrored Server Link Adapters (can be either fiber, coaxial or shielded twisted pair (STP) attached)

NetWare SFT III has the following software requirements:

1. NetWare SFT III (5, 10, 20, 50, 100, 250 user versions)
2. Identical DOS on each server (V3.1 or higher)

The two servers are not required to be on the same LAN segments or even required to have the same types of LAN adapters. The only requirement is that the servers be on the same internetwork and that all clients can get LAN packets to/from both of them.

1.10.1.1 SFT III Server Roles

A single SFT III server can be primary unmirrored, primary mirrored, or secondary mirrored at any given moment. When two servers are mirrored, they keep each other informed and aware of each other's role. When the secondary server finds that the primary server is no longer responding to its inquiries, it must determine whether the primary server is inoperable. If the primary server is inoperable, the secondary server takes over as the primary server, running without a mirrored partner.

The designation of primary or secondary server can change at any time. When the two servers are powered up and synchronized, the server that is activated first becomes the primary server. Once the roles of the servers have been defined, if the primary server fails, the secondary takes over as the primary and begins to service network clients.

1.10.1.2 Communication links

Each server is attached to two types of communication links:

- IPX internet link
- Mirrored Server Link (MSL)

IPX internet link: The IPX internet link is used for communicating with clients and for sending *server state* packets between the primary and secondary servers. The server state packets are also used to monitor the status of the internetwork.

Mirrored Server Link (MSL): The MSL is a bidirectional point-to-point connection that is used by the two servers to synchronize with each other. Information such as client requests and acknowledgments are passed back and forth on the MSL.

After a failure has occurred, the MSL is used to synchronize the memory and disk of the failed server. As it is being brought back up, the active server transfers the entire contents of its MS Engine memory image to the formally inactive server. After the contents of memory are transferred (a matter of seconds) normal activity resumes between the two servers.

If any disk activity occurs while the one server is down, the primary server sends all the disk updates to the formerly inactive server in a background mode, without impacting normal activity on the network. After the disks are fully re-mirrored, the system once again becomes protected and resilient to failures.

1.10.1.3 IOEngines and Mirrored Server (MS) Engines

Both the primary and secondary servers implement the operating system in two pieces: an IOEngine, which deals with the server hardware and is not mirrored, and a Mirrored Server (MS) Engine, which relies on the IOEngine for input.

The contents of the MSEngines are identical (mirrored), and both MSEngines have the same internal IPX number. The MSEngines mirror applications, data, and non-hardware related NetWare Loadable Modules (NLMs), such as Btrieve.

Any module which communicates directly with hardware, or makes explicit assumptions about the hardware, must reside in the IOEngine. Examples of utilities and NLMs that can be loaded into the IOEngine include LAN drivers, disk drivers, the MSL driver, print servers, and tape backup applications.

The IOEngine and the MSEngine have the following characteristics:

1. The two engines address the same memory space; however, each segment in memory is defined as belonging to either the IOEngine or the MSEngine. Except for rare instances, memory sharing is prohibited.
2. NLMs loaded in the MSEngine are mirrored automatically whenever the SFT III server is mirrored.
3. NLMs loaded in the IOEngine are never mirrored.
4. The primary server's IOEngine controls the entire logical server. It converts all network requests (packets), disk I/O completion interrupts, etc. into SFT III events, which are then submitted to both servers' MSEngines.

When a client needs to access a resource on the server, a request packet is sent over the network to the primary IOEngine. Clients always send their packets to the primary IOEngine because it advertises that it is the best route to the MSEngine. The primary IOEngine receives the request packet and sends a copy over the MSL to the secondary IOEngine. Both IOEngines send the request to their part of the MSEngine residing on separate machines. When each part of the MSEngine receives the event, it triggers processing (identical in both machines), resulting in identical replies to each IOEngine. Although both parts of the MSEngine reply to their IOEngines, only the primary IOEngine responds to the clients.

The MSEngines in both servers receive the same events and generate the same output. However, the secondary server's IOEngine discards all reply packets.

Consequently, clients only receive reply packets from the primary server's IOEngine; this is the same IOEngine to which they sent the original request packet. The clients view the mirrored server as any other NetWare server. Clients send a single request packet and receive a single reply packet from the same address. The duplication of requests to the secondary IOEngine and synchronization of events to both server engines happens transparently to network clients.

When the primary server fails and the secondary server takes over, network clients view the switch over as a simple routing change. The new primary (formerly the secondary) server's IOEngine begins advertising that it knows the route to the primary's MEngine internal network. The primary server sends a special packet to network clients, informing them of the new route. The primary server now sends response packets to clients rather than discarding them as it did when it was the secondary server.

This works in exactly the same way it would with regular NetWare if a route fails. The establishment of the new route is transparent to the client workstations and in the case of SFT III, to the MEngine as well.

When SFT III switches from the primary to the secondary server, clients may detect a slight pause, but server operations will continue because of SFT III's failure handling capabilities.

If a hardware failure causes the primary server to shut down, the secondary becomes the primary and immediately notifies clients on the network that it is now the best route to get to the MEngine. The client shell recognizes this notification and the route change takes place immediately.

IPX packets can be lost during the switch-over process, but LAN communication protocols are already set up to handle lost packets. When an acknowledgment packet is not received by the client, it simply does a retry to send the packet again. This happens quickly enough that the connection to the clients is maintained.

The following scenarios describe in more detail how SFT III failure handling works:

Scenario 1. Hardware fails in the primary server: Because the MSL times out from inactivity and no response is heard over the LAN, the secondary server infers that the primary server is down.

The secondary server displays a message on the console screen notifying the system administrator that the primary server is down and that the secondary server is taking over the primary server's role.

The secondary server sends a message to each client informing them that it is now the primary server.

Client packets are rerouted to the new primary (formerly the secondary) server.

SFT III keeps track of any disk changes following the server's failure.

The operator resolves the problem and restarts the server that has failed. The two servers synchronize memory images and begin running mirrored.

The primary server sends the disk changes over the mirrored server link to update the repaired server and to mirror the contents of the disk. Disk mirroring occurs in the background during idle cycles.

Scenario 2. Hardware fails in the secondary server: The primary server notifies the system administrator that the secondary server is down and that it is still the primary server.

SFT III keeps track of any disk changes following the secondary server's failure.

The system administrator repairs and brings up the secondary server.

The primary server sends the memory image and disk changes across the mirrored server link to update the secondary server and to re-synchronize the two servers.

Scenario 3. The Mirrored server link fails: The secondary server detects that information is not coming across the MSL.

However, server state packets are still coming across the IPX internet link, indicating that the primary server is still active.

The secondary server shuts down because the primary is still servicing the clients. Server mirroring is no longer possible.

SFT III keeps track of any disk changes following the secondary server's shutdown.

The system administrator repairs the mirrored server link and brings up the secondary server.

The primary server sends the memory image and disk changes across the mirrored server link to update the secondary server and to re-synchronize the two servers.

Scenario 4. A LAN adapter fails in the primary server: The operating system detects a LAN adapter failure in the primary server.

The operating system determines the condition of the secondary server's LAN adapters. If the secondary server's LAN adapters are more functional and the server is still mirrored, the secondary server takes over servicing clients on the network.

The system administrator is notified of the failure on the server's console screen and an entry is made in the log file.

SFT III keeps track of any disk changes following the failed server's shutdown.

The failed server will restart automatically and become the secondary server.

The active server sends the memory image and disk changes across the mirrored server link to update the failed server and to re-synchronize the two servers.

1.11 Uninterruptible Power Supply (UPS)

Digital computers require a clean source of direct current (DC). It is the computer's power supply which takes an alternating current (AC) from the input line and transforms it into clean DC voltages. However, problems on the input AC signal can often lead to DC voltages that are less than satisfactory for the digital circuits to operate properly.

There are five main types of AC line problems that can cause trouble for a computer system:

1. **Brownouts**

Brownouts are extended periods of low voltages often caused by unusually high demands for power such as that caused heavy machinery, air conditioning units, laser printers, coffee machines, and other high current electrical devices.

2. **Surges**

Surges are extended periods of high voltages and can be caused, for example, by some of the previously mentioned high current devices being turned off.

3. **Spikes**

Spikes are short duration, high voltages often due to a lightning strike, static, or faulty connections on the power line.

4. **Noise**

Noise is generally electromagnetic interference (EMI) or radio frequency interference (RFI) induced on the power line and can be caused by poor grounding.

5. **Blackouts**

Blackouts occur when the AC voltage levels are too low for the power supply of the computer to transform them into DC voltages for the digital circuits. At this point, the computer ceases to function. There are any number of causes for these to occur: a power failure, a short circuit, downed power lines, and human error, to name but a few.

A UPS is an external device which connects the AC input line to the computer's power supply. It contains several components which can alleviate most AC line problems. These are:

- Surge suppressors which protect against any large spikes on the input line
- Voltage regulators which ensure that the output voltage lies within an acceptable range for the computer input
- Noise filters which take out any EMI/RFI noise on the input line
- Batteries which can provide an instantaneous power source in the case of a power failure and also help to filter the input line

The blackout is often considered to be the most common type of failure. However, when monitoring the power line, users are often surprised to find that it is brownouts which occur far more frequently. It is also the brownouts that can cause the most damage since they are usually unobserved and unexpected. The UPS is critical here because it filters the input line providing a clean, stable input to the computer's power supply.

The primary service, however, that the UPS provides in the case of AC line problems is extra time. While a UPS can enable the server to continue operating even if there is a power loss, the primary benefit of a UPS is that the server software has time to ensure that all caches are written to disk, and to perform a tidy shutdown of the system.

Some UPSs also offer an automated shutdown and reboot facility for network operating systems. This is often provided via a serial link to the server and is commonly known as UPS monitoring.

1.11.1 APC PowerChute

American Power Conversion introduced PowerChute in 1988. PowerChute is software which interacts with the UPS to provide an orderly shutdown of a server in the event of an extended AC power failure. PowerChute offers user notification of impending shutdown, power event logging, auto-restart upon power return, and UPS battery conservation features.

The current version is PowerChute Plus V4.2. The PowerChute Plus software consist of two main components. The first is the UPS monitoring module that runs as a background process on the server. It communicates with the UPS and the user interface module, logs data and events, notifies user of impending shutdowns, and when necessary, shuts down the operating system.

The second component is the user interface module, which may also be known as the workstation module. The user interface can run either locally on the server or over a network on a workstation. It gathers real-time data such as UPS output, UPS temperature, output frequency, ambient temperature, humidity and UPS status.

When PowerChute Plus is used with a Smart UPS or Matrix UPS, the PowerChute monitoring features are augmented by sophisticated diagnostic and management features. These include:

- Scheduled server shutdowns
- Interactive/scheduled battery logging
- Detailed power quality logging
- Real-time graphical displays showing:
 1. Battery voltage
 2. Battery capacity
 3. UPS load
 4. Utility line voltage
 5. Run time remaining

1.11.1.1 Flex Events

Flex Events is a feature of PowerChute Plus. It logs UPS related events which have occurred and allows for actions to be taken based on these events.

Events can range in severity from informational (not severe) to critical (severe). For instance, there is an event called *UPS Output Overload*. This event is considered a critical event and will be generated when the rated load capacity of the UPS has been exceeded. It is critical because if the situation is not remedied by unplugging excess equipment from the UPS, the UPS will not support the load if the power fails.

Flex Events is programmable such that when an event occurs, you can configure PowerChute to take certain actions. Depending on the event you can:

- Log that event
- Send early warning pop-up messages to specified administrator
- Broadcast messages to users on the network
- Shut down the host computer
- Run a command file (an external executable file)
- Send E-mail to notify users

Chapter 2. IBM PC Server Family Overview

The IBM PC Server family contains three product lines which offer different features and capabilities:

- The PC Server 300 series

This series is targeted at small enterprises or workgroup LANs. These machines offer leading technology and are very price competitive. They are more limited in terms of upgrade and expansion capabilities than the other two lines in the family.

- The PC Server 500 series

This series is targeted for medium to large enterprises who need more power and more expansion capabilities. With 22 storage bays available (18 of which are hot swappable), these machines are very suitable for the enterprise workhorse server.

- The PC Server 700 series

The PC Server 720 is the only model in this line at the current time. It is a *super server* targeted for customers who need the maximum computing power in a PC server environment. With its multiprocessing capability, it is very suitable for the application server environment. It offers state of the art technology and also has wide expansion capabilities.

Figure 21 shows the products in relation to one another.




Figure 21. IBM PC Server Family of Products

2.1 IBM PC Server Model Specifications

The following tables show the specifications for each model in the current line. They are included for a reference of the standard features of each line.

2.1.1 IBM PC Server 300

System Model	8640 ONO	8640 ONJ	8640 OP0	8640 OPT
Processor	486DX2/66	486DX2/66	Pentium 60	Pentium 60
Bus Architecture	PCI/EISA	PCI/EISA	PCI/EISA	PCI/EISA
Disk Controller	PCI IDE ISA IDE	PCI IDE ISA IDE	PCI SCSI-2 Fast ISA IDE	PCI SCSI-2 Fast ISA IDE
STD Hard File Size	None	728MB IDE	None	1GB SCSI-2
Memory Std/Max (MB)	8/128	8/128	16/192	16/192
L2 Cache Std/Max (KB)	256/256	256/256	256/512	256/512
Graphics	VGA	VGA	VGA	VGA

2.1.2 IBM PC Server 310

<i>Table 9. IBM PC Servers 310 Models</i>		
System Model	8639 OXT	8639 MXT
Processor	Pentium 75	Pentium 75
Bus Architecture	PCI/ISA	PCI/MCA
Disk Controller	PCI SCSI-2 Fast	PCI SCSI-2 Fast
STD Hard File Size	1.08GB	1.08GB
Memory Std/Max (MB)	16/192	16/192
L2 Cache (KB)	256	256
Graphics	SVGA	SVGA

2.1.3 IBM PC Server 320 EISA

Table 10. IBM PC Servers 320 EISA Models

System Model	8640 ONO	8640 ONJ	8640 OP0	8640 OPT
Processor	Pentium 90	Pentium 90	Pentium 90	Pentium 90
SMP	1-2 P90	1-2 P90	1-2 P90	1-2 P90
Bus Architecture	PCI/EISA	PCI/EISA	PCI/EISA	PCI/EISA
Disk Controller	ISA IDE PCI SCSI-2 F/W	ISA IDE PCI SCSI-2 F/W	ISA IDE PCI SCSI-2 F/W Raid	ISA IDE PCI SCSI-2 F/W Raid
STD Hard File Size	None	1.12GB	None	2 * 1.12GB
Memory Std/Max (MB)	16/256	16/256	16/256	16/256
L2 Cache (KB)	256/512	256/512	256/512	256/512
Graphics	SVGA	SVGA	SVGA	SVGA

2.1.4 IBM PC Server 320 MCA

<i>Table 11. IBM PC Servers 320 MCA Models</i>					
System Model	8640 MX0	8640 MXT	8640 MYO	8640 MYT	8640 MYR
Processor	Pentium 75	Pentium 75	Pentium 90	Pentium 90	Pentium 90
SMP	1-2 P75	1-2 P75	1-2 P75	1-2 P75	1-2 P75
Bus Architecture	PCI/MCA	PCI/MCA	PCI/MCA	PCI/MCA	PCI/MCA
Disk Controller	PCI SCSI-2 F/W	PCI SCSI-2 F/W	PCI SCSI-2 F/W	PCI SCSI-2 F/W	PCI SCSI-2 F/W RAID
STD Hard File Size	None	1.08GB	None	1.12GB	2 * 1.12GB
Memory Std/Max (MB)	16/256	16/256	16/256	16/256	16/256
L2 Cache Std/Max (KB)	256/512	256/512	256/512	256/512	256/512
Graphics	SVGA	SVGA	SVGA	SVGA	SVGA

2.1.5 IBM PC Server 500

Table 12. IBM PC Server 500 Models

System Model	8641 0Y0	8641 1Y0	8641 0YT	8641 0YV	8641 0YR	8641 0YS
Processor	Pentium 90	Pentium 90	Pentium 90	Pentium 90	Pentium 90	Pentium 90
Bus Architecture	MCA	MCA	MCA	MCA	MCA	MCA
Disk Controller	SCSI-2 F/W	SCSI-2 F/W RAID	SCSI-2 F/W	SCSI-2 F/W	SCSI-2 F/W RAID	SCSI-2 F/W RAID
Hard File Size	None	None	1.12GB	2.25GB	1.12GB x 3	2.25GB x 3
Memory Std/Max (MB)	32/256	32/256	32/256	32/256	32/256	32/256
L2 Cache Std/Max (KB)	256/256	256/256	256/256	256/256	256/256	256/256
Graphics	SVGA	SVGA	SVGA	SVGA	SVGA	SVGA

2.1.6 IBM PC Server 520 EISA

<i>Table 13. IBM PC Servers 520 EISA Models</i>				
System Model	8641 EZ0	8641 EZV	8641 EZS	8641 EZE
Processor	Pentium 100	Pentium 100	Pentium 100	Pentium 100
SMP	1-2 P 100	1-2 P 100	1-2 P 100	1-2 P 100
Bus Architecture	PCI/EISA	PCI/EISA	PCI/EISA	PCI/EISA
Disk Controller	PCI SCSI-2 F/W	PCI SCSI-2 F/W	PCI SCSI-2 F/W RAID	PCI SCSI-2 F/W RAID
Hard File Size	None	2.25GB	2 * 2.25GB	4 * 2.25GB
Memory Std/Max (MB)	32/256	32/256	32/256	32/256
L2 Cache (KB)	512	512	512	512
Graphics	SVGA	SVGA	SVGA	SVGA

2.1.7 IBM PC Server 520 MCA

<i>Table 14. IBM PC Servers 520 MCA Models</i>					
System Model	8641 MZ0	8641 MZV	8641 MZS	8641 MZE	8641 MZL
Processor	Pentium 100	Pentium 100	Pentium 100	Pentium 100	Pentium 100
SMP	1-2 P100	1-2 P100	1-2 P100	1-2 P100	1-2 P100
Bus Architecture	PCI/MCA	PCI/MCA	PCI/MCA	PCI/MCA	PCI/MCA
Disk Controller	PCI SCSI-2 F/W	PCI SCSI-2 F/W	PCI SCSI-2 F/W RAID	PCI SCSI-2 F/W RAID	PCI SCSI-2 F/W RAID
Hard File Size	None	2.25GB	2 * 2.25GB	4 * 2.25GB	6 * 2.25GB
Std Memory (MB)	32/256	32/256	32/256	32/256	32/256
L2 Cache (KB)	512	512	512	512	512
Graphics	SVGA	SVGA	SVGA	SVGA	SVGA

2.1.8 IBM PC Server 720

<i>Table 15. IBM PC Servers 720 Models</i>				
System Model	8642 OZO	8642 1Z0	8642 2ZS	8642 4ZS
Processor	Pentium 100	Pentium 100	Pentium 100	Pentium 100
# of CPUs in base model	1	1	2	4
SMP	1 - 6 Pentium processors	1 - 6 Pentium processors	1 - 6 Pentium processors	1 - 6 Pentium processors
Bus Architecture	PCI/MCA/CBus	PCI/MCA/CBus	PCI/MCA/CBus	PCI/MCA/CBus
MCA Speed	80 MBps	80 MBps	80 MBps	80 MBps
CBus Speed	400 MBps	400 MBps	400 MBps	400 MBps
Disk Controller	PCI SCSI-2 F/W	PCI SCSI-2 F/W Raid	PCI SCSI-2 F/W Raid	PCI SCSI-2 F/W Raid
Hard File Size	None	None	2 * 2.25GB	4 * 2.25GB
Memory Std/Max	64MB/1GB	64MB/1GB	64MB/1GB	64MB/1GB
L2 Cache	512 KB each processor	512 KB each processor	512 KB each processor	512 KB each processor
Graphics	SVGA	SVGA	SVGA	SVGA

Chapter 3. Hardware Configuration

The different technologies used to implement the PC Server family require different methods for configuration. Unfortunately, there is no one common configuration program which can be run on a machine to completely configure it. In most cases, multiple programs will need to be run in order to complete this process.

This chapter gives instructions on using the various configuration programs and when to use each one. There are some model dependencies, however. If you see differences between what you see on your machine and what is documented here, consult your handbook which comes with the system.

The configuration programs and a brief explanation of each are listed below:

- Setup program

This program is used to configure system options on PCI/EISA/ISA machines. The system options include such things as diskette and hard disk options, video subsystem, and system memory.

- EISA configuration utility

This utility is used to configure I/O adapters on PCI/EISA machines.

- SCSI select utility

This utility allows you to configure the SCSI subsystem on PCI/EISA/ISA machines.

- System programs

These programs allow you to configure system options, I/O adapters, and the SCSI subsystem on Micro Channel machines.

- RAID utility

This utility allows you to configure the RAID subsystem on machines equipped with this feature.

The following flowchart shows the steps necessary to configure the server hardware:




Figure 22. Hardware Configuration Steps

3.1 The Setup Program

The setup program is used to configure system options on ISA and EISA machines. The system options include such things as diskette and hard disk options, video subsystem, and system memory. These parameters are controlled by system BIOS and, hence, need to be modified before the operating system boots.

To access the setup program:

1. Turn on the server and watch the screen.

The BIOS level appears.

2. When the message

Press <F2> to enter SETUP

appears, press F2.

3. Follow the instructions on the screen to view or change the configuration. Please see the following sections for detailed instructions on this process.

After completion of these operations:

1. Select **Exit menu** from the menu bar.

Don't forget to select **Saves changes and exit**.

The setup program consists of three panels which are selectable from the menu bar:

- Main
- Advanced
- Security

3.1.1 Main Menu

Figure 23 shows the main panel of the setup program.

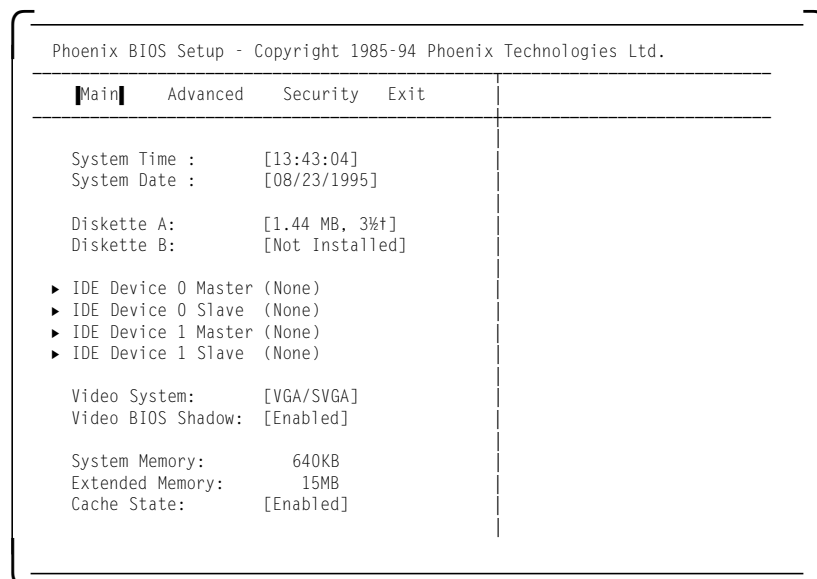


Figure 23. PC Server 320 Setup Program - Main Menu

The Main panel contains fields which allow the user to:

- Modify date and time
- Configure the diskette drives
- Configure the IDE disks
- Configure the video
- Enable/Disable level 2 system memory cache

Notes:

1. **Video BIOS Shadow**

This option allows the user to shadow the video BIOS into RAM for faster execution. The pre-installed SVGA Adapter supports this feature.

2. **IDE Devices**

If no IDE DASD devices are installed, you must set all the IDE devices to none.

Attention!

If a PCI SCSI card is installed, an PCI IRQ 5, 11 or 15 must be defined for this adapter and a DASD must be installed.

3.1.2 Advanced Menu

The Advanced option allows the user to:

- Change Boot Options
- Configure Integrated Peripherals

To reach the Advanced menu:

1. Press ESC to quit the main menu.
2. Use the arrows keys to select the **Advanced** option.

A screen like Figure 24 will appear.




Figure 24. PC Server 320 Setup Program - Advanced Menu

3.1.2.1 Advanced Menu - Boot Options

By pressing the Enter key, a screen like that shown in Figure 25 will appear.




Figure 25. PC Server 320 Setup Program - Boot Options Menu

Boot Sequence: Boot sequence allows the user to change the order the system uses to search for a boot device.

Other values can be:

- C: then A:, if the user wants to boot from the hard disk first
- C: only, if the user does not allow a boot from a diskette

Swap Floppies: This choice allows the floppy disk drives to be redirected. Normal is the default. When Swapped is selected, drive A becomes drive B and drive B becomes drive A.

Note

The option Swapped does not modify the boot sequence option. So if boot sequence is set to A: then C: and Swap floppies to Swapped, the user will get the following error message at IPL:

```
0632 Diskette Drive Error
```

Floppy Check: When enabled, the floppy drives are queried at boot time to verify the correct drive type. The machine boots faster when disabled.

Setup prompt: When enabled, the following message appears:

```
Press <F2> to enter SETUP
```

If disabled, the prompt message is not displayed but the function is still available.

POST errors: When enabled, if an error occurs during power on self-tests (POST) the system pauses and displays the following:

```
Press <F1> to continue or <F2> to enter SETUP
```

If disabled, the system ignores the error and attempts to boot.

3.1.2.2 Advanced Menu - Peripherals

To reach this menu:

- Press ESC to quit the Boot options.
- Use the arrows keys to select the integrated peripherals option.
- Press Enter.

A screen like the one in Figure 26 will appear:




Figure 26. PC Server 320 Setup Program - Integrated Peripherals Menu

Serial Port A: The port can be set to either COM1 or COM3 and uses the IRQ4 if enabled.

Serial Port B: The port can be set to either COM2 or COM4 and uses the IRQ3 if enabled.

Parallel Port: It can be set to LPT1 or LPT2 and uses the IRQ7.

Diskette Controller: It uses the IRQ 6 if enabled.

Integrated IDE adapter: Select:

- Primary, if IDE devices are present
- Secondary, if another IDE controller has been added to the system
- Disabled, otherwise

Note

You must set the W25 jumper on the planar board accordingly.

Large Disk DOS Compatibility: Select:

- DOS, if you have DOS
- Other, if you have another operating system including UNIX or Novell NetWare

DOS is the default value.

Memory Gap: Some ISA network adapters need to be mapped in system memory address space, normally at the upper end. Since the ISA bus is limited to 24-bit addressing (0-16 MB), systems with more than 16MB of memory installed will not accommodate these adapters.

When enabled, this selection will remap system memory such that a *gap* is created between the addresses of 15 and 16MB. This gap is then used to map the I/O space on the adapter into this area.

3.1.3 Security

The Security option allows the user to:

- Set the supervisor password.
- Set the power on password.
- Modify the diskette access rights.
- Set the fixed disk boot sector.

To reach the security panel from the Advanced - Peripherals panel:

- Press ESC to quit the Peripherals menu.
- Press ESC again to quit the Advanced menu.
- Use the arrows keys to select the **Security** menu.
- Press Enter.

A screen like the one in Figure 27 will appear:




Figure 27. PC Server 320 Setup Program - Security Menu

Set the Supervisor/User passwords: Two levels of passwords are available with IBM PC Server PCI/EISA systems:

- Supervisor password which enables all privileges
- User password which has all privileges except:
 - Supervisor password option
 - Diskette access option if supervisor selected for this option (see Diskette access rights below)
 - Fixed disk boot sector option

If either password is set, at boot time you will see:

Enter password

If you enter the wrong password, the following message appears on the screen, and you are prompted again:

Incorrect password

After 3 incorrect attempts, the following message appears and you must turn off the server and start again:

System disabled

Notes:

1. Before you set a supervisor password, you must first set your selectable drive-startup sequence.
2. Only the supervisor can change the supervisor password.
3. The set user password can be selected only if a supervisor password has been created.
4. The option for fixed disk boot sector is available only with supervisor privilege level.
5. If the supervisor password is removed, the user password is also removed.

Attention!

If a supervisor password is forgotten, it cannot be overridden or removed. If you forget your supervisor password, you must have the system board replaced to regain access to your server.

Password on boot: This option allows you to enable/disable the password on boot.

It must be enabled to have the user password operational.

Diskette access rights: This option may be set to:

- Supervisor, if access is only allowed to the supervisor
- User, if access is allowed to both supervisor and user

The supervisor password must be set for this feature to be enabled.

Fixed disk boot sector: This option may be set to:

- Normal
- Write protect

As a form of virus protection, if this option is enabled, the boot sector of the hard drive is write-protected.

Note

BIOS of PCI/EISA servers is located in a Flash ROM on the motherboard. If necessary, it can be updated with a bootable diskette which has the new BIOS (.BIN) file. This file will be named:

- M4PE_Txx.BIN for DX2-66 models
- M5PE_Txx.BIN for Pentium models

Where xx is the BIOS level as it appears when booting.

For more information on how to obtain BIOS updates, please reference Appendix B, "Hardware Compatibility, Device Driver, and Software Patch Information" on page 199.

Attention! Make sure you have the right file for your system. The process will allow you to install the wrong file. If you do, the server will not reboot successfully; instead, at power on, the screen will be blank and the system will beep twice. The only fix is to replace the motherboard.

3.2 EISA Configuration Utility

This utility is used when you add or remove an ISA or EISA adapter. We will use an example to illustrate the process. In our example we will add an Auto T/R 16/4 ISA adapter in slot 3 of a PC Server 320. The steps to complete the process are:

1. Boot with the EISA configuration utility diskette.
2. Answer Y to the question:

Do you want to configure your system now [Y,N]?

A screen like the one in Figure 28 will appear:

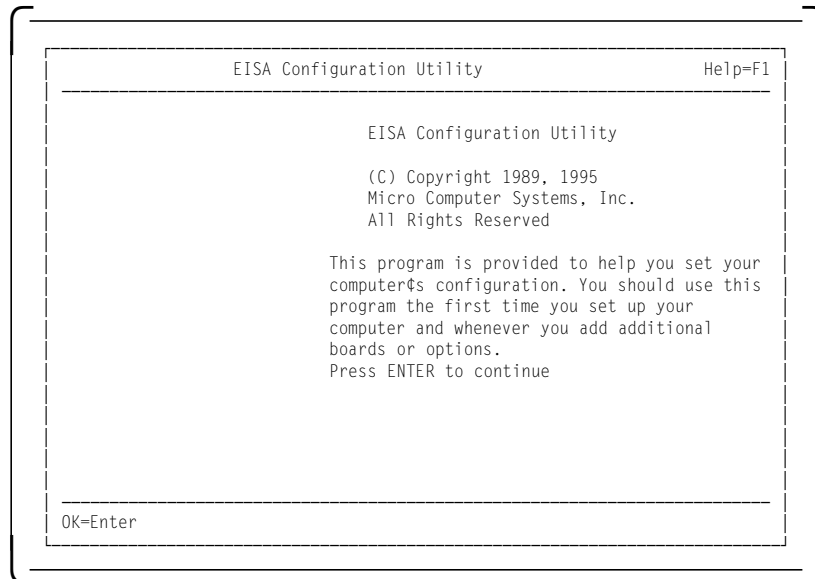


Figure 28. EISA Configuration Utility - Main Panel

3. Press Enter. A screen like the one in Figure 29 will appear.




Figure 29. EISA Configuration Utility - Steps

If you are not familiar with the ISA and EISA cards, you can read the information in step 1; otherwise you can skip to step 2.

A screen like the one in Figure 30 on page 75 will appear.




Figure 30. EISA Configuration Utility - Step 2

Note

EISA adapters ship with a diskette which contains a configuration file (a .CFG file) which the EISA configuration utility needs so that it knows what parameters are available for the adapter. This .CFG file should be copied to the EISA Configuration diskette. If the file has been copied to the diskette, the EISA adapter is added automatically.

If you have not copied the .CFG file, you will be prompted to insert the adapter's configuration diskette into the diskette drive during this process and the .CFG file will be copied to your diskette.

Our token-ring adapter is recognized but not in the correct slot. This is because it is not possible for EISA systems to determine what slot ISA adapters are in. So we must tell the system what slot it is in by *moving* the adapter to the correct slot.

4. To move the ISA adapter to the correct slot:
 - a. With the arrow key, select the desired adapter.
 - b. Press F7. A Move Confirmation panel appears:

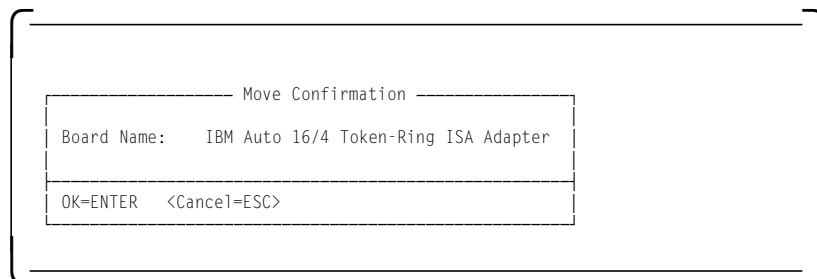


Figure 31. EISA Configuration Utility - Move Confirmation Panel

- c. Select **OK**.
 - d. With the arrow key, select the destination slot and press Enter.
 - e. Press F10 to return to the EISA Configuration menu.

5. View or Edit Details

After adding EISA or ISA adapters, you will often need to view and/or edit the settings for the adapter. To view or edit an adapter's details:

- a. From the Main menu, select step 3 (View or Edit Details) with the arrow key.
- b. Press Enter to view configuration details. You will see a screen similar to that shown in Figure 32




Figure 32. EISA Configuration Utility - Step 3

Use the Edit Resources option to change interrupt request levels, I/O addresses and other parameters whose settings may need to be changed to avoid conflicts with other devices.

Note

Sometimes changing a setting during this step requires you to change a switch or jumper setting on the system board or on an adapter.

When finished, press F10 to exit and return to the EISA Configuration menu.

6. Examine switches or print report

You can use this option to display the correct switch and jumper settings for the installed devices that have switches and jumpers. You can also choose to print a system configuration report. To do this:

- a. Use the arrow key to select step 4 and press Enter.
- b. Select the board(s) marked with an arrow and press Enter.
- c. The necessary switch/jumpers settings are displayed in a screen similar to the one shown in Figure 33 on page 77.




Figure 33. EISA Configuration Utility - Step 4

- d. Press F7 if you want to print configuration settings. You can print:
- Settings for selected board or option
 - Settings for selected board or option to a file
 - All configuration settings
 - All configuration settings to a file

Appendix A, "EISA Configuration File" on page 189 contains a sample configuration report which includes all configuration settings.

- e. When finished, press F10 to return to the Configuration menu.
7. Select step 5 and press Enter to save your configuration.

3.3 SCSI Select Utility Program

This utility is used on PCI/EISA models of the IBM PC Server line and allows the user to:

- View and modify parameters for the SCSI controller
- View and modify parameters of SCSI devices
- Perform low-level formatting of attached SCSI hard disks

To access the SCSI Select Utility Program:

- Turn on the server and watch the screen
- When the message Press <Ctrl><A> appears, press Ctrl and A simultaneously.

A screen like the one in Figure 34 will appear.




Figure 34. IBM PC Server SCSISelect Utility Program - Main Menu

- Press Enter to enter the Configure/View Host Adapter. option

A screen like the one in Figure 35 will appear.




Figure 35. IBM PC Server SCSI Select Utility Program - Host Adapter Settings

The fields on this panel are described as follows:

SCSI Parity Checking: Select this option to enable or disable SCSI Parity checking on the host adapter. If enabled, the host adapter will check parity when reading from the SCSI bus to verify the correct transmission of data from your SCSI devices. SCSI Parity checking should be disabled if any attached SCSI device does not support SCSI parity. Most currently available SCSI devices do support SCSI parity.

Host Adapter SCSI termination: All SCSI interfaces use daisy-chain cabling. The cable starts at the adapter and goes to the first device, and then out of that device to the next device and so on until it reaches the last device in the chain. The last device has an incoming cable and a *terminator*. The terminators are used to absorb potential signal reflections on the SCSI bus which would cause interference. The last device on the bus must always be terminated.

The SCSI-2 Fast/Wide PCI adapter that came with the PCI/EISA server has three connectors which can be the starting points for a daisy-chained cable: one 8-bit, 50-pin (SCSI-I) internal connector, one 16-bit, 68-pin (SCSI-II Wide) internal cable connector, plus another 16-bit, 68-pin external connector. The adapter has built-in terminators on these connectors.

The setting for the Host Adapter SCSI termination needs to be configured depending on which connectors are used. This option is comprised of two entries, a low and a high. You can think of them as software jumpers. Each entry, low and high, can take on either an on or off value, thereby giving four possible different combinations of the two entries. The chart below shows the proper values of these entries depending upon what connectors have been used.

Note

Only two of the three connectors can be used, either the two internal or one internal and one external.

Table 16. Host Adapter SCSI Termination Parameter

16-bit (68pin) internal connector	8-bit (50pin) internal connector	16-bit (68pin) external connector	Low value	High value
Yes			On	On
		Yes	On	On
	Yes		On	On
Yes		Yes	Off	Off
Yes	Yes		Off	On
	Yes	Yes	Off	On

After configuring the host adapter, you need to configure the SCSI devices. To do this:

- Use the arrow keys to select SCSI Device Configuration
- Press Enter

A screen like the one in Figure 36 on page 80 will appear:

SCSI Device Configuration								
SCSI Device ID	#0	#1	#2	#3	#4	#5	#6	#7
Initiate Sync Negotiation	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Max Sync Transfer Rate	10.0	10.0	10.0	10.0	10.0	10.0	10.0	10.0
Enable disconnection	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Initiate Wide negotiation	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Options listed below have NO EFFECT if the BIOS is disabled								
Send start init command	No	No	No	No	No	No	No	No
Include in BIOS Scan	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
SCSI Device ID	#8	#9	#10	#11	#12	#13	#14	#15
Initiate Sync Negotiation	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Max Sync Transfer Rate	10.0	10.0	10.0	10.0	10.0	10.0	10.0	10.0
Enable disconnection	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Initiate Wide negotiation	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Options listed below have NO EFFECT if the BIOS is disabled								
Send start init command	No	No	No	No	No	No	No	No
Include in BIOS Scan	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Figure 36. PC Server 320 SCSI Select Utility Program - SCSI Device Configuration

To modify settings on this screen:

- Use the arrow keys to select the parameter to modify.
- Press Enter to edit the value.
- Use the arrow keys to select the new value or press Esc to quit.
- Press Enter to validate the new value.

The fields in this screen are described below.

Initiate Sync Negotiation: The host adapter always responds to synchronous negotiation if the SCSI device initiates it. However, when this field is set to Yes, the host adapter will initiate synchronous negotiation with the SCSI device.

Some older SCSI-1 devices do not support synchronous negotiation. Set Initiate Sync Negotiation for these devices to avoid malfunction.

Maximum Sync Transfer Rate: The default value is 10.0 MBps for SCSI-II Fast devices. If you are using SCSI-II Fast/Wide devices, the effective maximum transfer rate is 20.0 MBps.

Older SCSI-1 devices do not support fast data transfer rates. If the transfer rate is set too high, this may cause your server to operate erratically or even hang. Select 5.0 Mbps for any SCSI-I devices.

Enable Disconnection: This option determines whether the host adapter allows a SCSI device to disconnect from the SCSI bus (also known as the Disconnect/Reconnect function).

You should leave the option set to Yes if two or more SCSI devices are connected to optimize bus performance. If only one SCSI device is connected, set Enable Disconnection to No to achieve better performance.

Send Start Unit Command: Enabling this option reduces the load on your server's power supply by allowing the SCSI devices to power-up one at a time when you boot the server. Otherwise, the devices all power-up at the same time.

The SCSI-2 Fast and Wide adapter issues the start unit command to each drive one at a time. The SCSI-2 Fast/Wide Streaming RAID adapter issues the start unit command to two drives at a time.

Note

In order to take advantage of this option, verify that the auto-start jumpers have been removed on hard drives. Otherwise, the drives will spin up twice: once at Power on Reset (POR) time and again when the adapter sends the start unit command.

Include in BIOS SCAN This option determines whether the host adapter BIOS supports devices attached to the SCSI bus without the need for device driver software. When set to Yes, the host adapter BIOS controls the SCSI device. When set to No, the host adapter BIOS does not search the SCSI ID.

Notes:

1. Send Start Unit Command and Include in BIOS Scan have no effect if BIOS is disabled in the Advanced Configuration Options panel (see Figure 37).
2. Disabling the host adapter BIOS frees up 8-10 KB memory address space and can shorten boot-up time. But you should disable this option only if the peripherals on the SCSI bus are all controlled by device drivers and do not need the BIOS (for example, a CD-ROM).

After completing the device configuration, there are a few more parameters which need to be configured. To do this:

- Press ESC to quit the SCSI Device Configuration menu.
- Use the arrow keys to select the Advanced Configuration Options menu
- Press Enter.

A screen like the one in Figure 37 will appear.




Figure 37. PC Server 320 SCSISelect Utility Program - Advanced Configuration

To modify the settings on this screen:

- Use the arrow keys to select the parameter to modify.

- Press Enter to edit the parameter.
- Use the arrow keys to select the new value or press Esc to quit.
- Press Enter to validate the new value.

When finished:

- Press Esc to quit the SCSI Advanced Configuration options menu.
- Press Esc to quit the Configuration menu.
- Use the arrow keys to select the SCSI Disk utility.
- Press Enter.

A screen like the one in Figure 38 will appear.

```

----- Select SCSI ID Disk and Press <Enter> -----
SCSI ID #0 : IBM      DPES-31080
SCSI ID #1 : No device
SCSI ID #2 : No device
SCSI ID #3 : IBM      CDRM 00203
SCSI ID #4 : No device
SCSI ID #5 : No device
SCSI ID #6 : No device
SCSI ID #7 : AHA-2940/AHA-2940W
SCSI ID #8 : No device
SCSI ID #9 : No device
SCSI ID #10 : No device
SCSI ID #11 : No device
SCSI ID #12 : No device
SCSI ID #13 : No device
SCSI ID #14 : No device
SCSI ID #15 : No device

```

Figure 38. PC Server 320 SCSISelect Utility Program - DASD Information

This screen shows the devices that are attached to the adapter and their SCSI IDs. It will also allow you to perform a low-level format of the disk or to scan it for media defects if desired. To do this:

- Use the arrow keys to select the DASD to format.
- Follow the directions on the screen.

When finished:

- Press Esc to quit the SCSI disk utility.
- Select Yes to confirm.

You have now completed the SCSI subsystem configuration.

Don't forget to save changes before you exit.

3.4 System Programs

If you have a PCI/MCA machine, you will run the system programs. The system programs are a set of utility programs you can use to configure the SCSI subsystem, system options, and I/O adapters. Also, you can use them to set passwords, change the date and time, and test the server. In effect, they are the equivalent of the SETUP, EISA CONFIG, and SCSI SELECT for an ISA/EISA machine.

These programs are obtainable in several ways:

- Shipped with the server on two diskettes called the reference diskette and the diagnostic diskette
- Created from images for these diskettes on the ServerGuide CD-ROM shipped with the system.
- On the system partition of the machine

Non-array systems are shipped with the system programs already installed in a protected area of the hard disk called the system partition. The system partition is protected against operating-system read, write, and format operations to guard against accidental erasure or modification. Disk-array systems do not have a system partition.

You can start the system programs in one of two ways:

1. Boot using the system partition
2. Boot using reference diskette

The system partition should be used if available. The reference diskette is normally used to:

- Configure and test disk-array models (since there is no system partition)
- Test non-array models if you can not start the system programs from the system partition
- Reconstruct the programs on the system partition of a non-array model when you replace the hard disk drive or if the programs are damaged
- To install the DOS keyboard-password program and other stand-alone utility programs

3.4.1 Starting From the System Partition

To start the system programs from the system partition:

1. Turn off the server.
2. Remove all media (diskettes, CDs, or tapes) from all drives.
3. Turn on the server. The IBM logo appears on the screen.
4. When the F1 prompt appears, press F1. A second IBM logon screen appears, followed by the system programs Main Menu. The Main Menu is shown in Figure 39 on page 84.

To select an option:

1. Use the up arrow key or down arrow key to highlight a choice.
2. Press Enter.

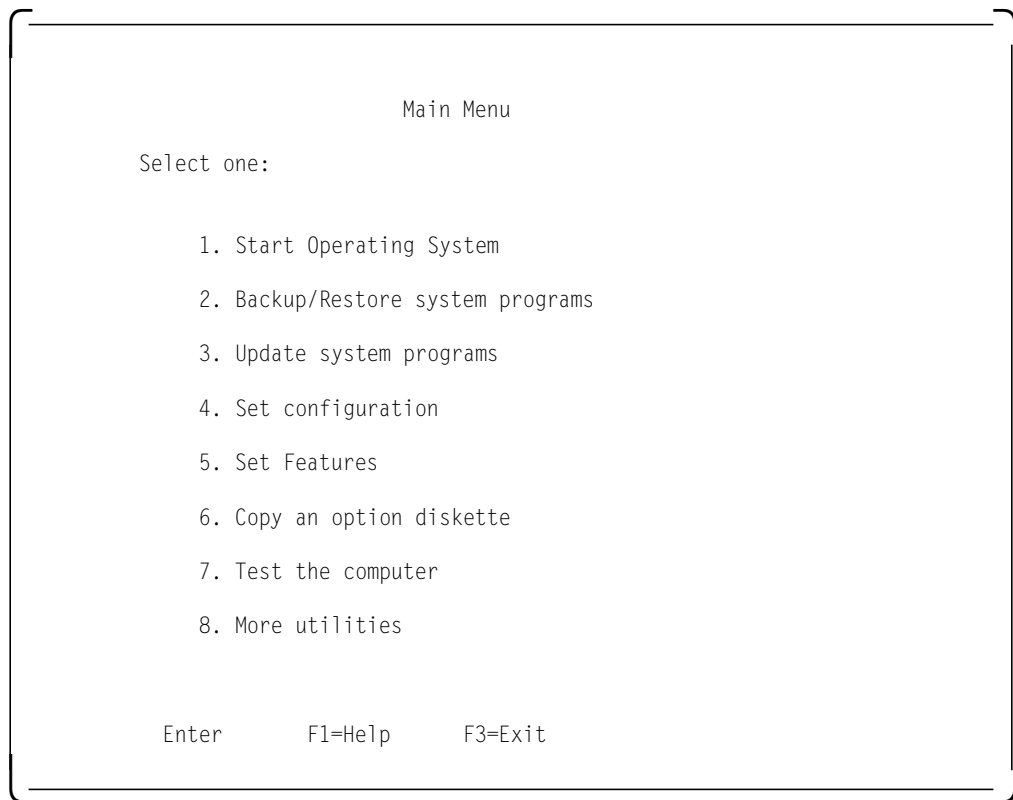


Figure 39. System Programs - Main Menu

3.4.2 Starting From the Reference Diskette

To start the system programs from the reference diskette:

1. Turn off the server.
2. Insert the reference diskette into your diskette drive.
3. Turn on the system.

After a few moments, the system programs Main Menu appears. It will look similar to the one in Figure 39.

To select an option:

1. Use the up arrow key or down arrow key to highlight a choice.
2. Press Enter.

3.4.3 Main Menu Options

The following are the options available on the Main Menu. Included with each option is a brief description of its purpose.

1. Start operating system

Exits from the system programs and loads the operating system.

2. Backup/restore system program

Makes a backup copy of the system programs from the hard disk to diskette or restores the system programs from the diskette to hard disk.

3. Update system programs

Periodically, updated versions of the reference diskette and diagnostic diskette are made available. This option copies a new version of the system programs to the system partition. This option does not apply to disk-array models.

Note

This utility will only install system programs that are a later version than the ones already installed on the system partition.

4. Set configuration

This option contains programs used to view, change, back up, or restore the configuration information. It also contains the Automatic Configuration program.

The configuration information consists of:

- Installed system options
- Memory size
- Adapter locations and assignments
- SCSI subsystem parameters

5. Set features

This option allows you to set system parameters such as date and time, type of console, startup sequence, fast startup mode, and passwords.

6. Copy an options diskette

Micro-Channel machines use configuration files called Adapter Descriptor Files (.ADF files) in order to know what parameters and values are available for the adapter. This option copies configuration and diagnostic files from an option diskette to the system partition or to the backup copy of the system programs diskettes. The server needs these files to make the new options operational.

Attention!

This utility will prompt you for both the reference diskette and the diagnostic diskette so that the proper programs can be copied from the adapter option diskette to these diskettes. Make sure that you have copies of both diskettes before you select this utility. These diskettes can be obtained from Diskette Factory on the ServerGuide CD.

7. Test the computer

Run diagnostics on the system hardware. These tests show if the hardware is working properly. If a hardware problem is detected, an error message appears explaining the cause of the problem and the action to take.

8. More utilities

This option is a set of utilities that displays information which is helpful when service is required. Revision levels and the system error log are some of the utilities available in this option.

3.4.4 Backup/Restore System Programs Menu

When you select this option from the Main Menu, a screen like the one in Figure 40 will appear.

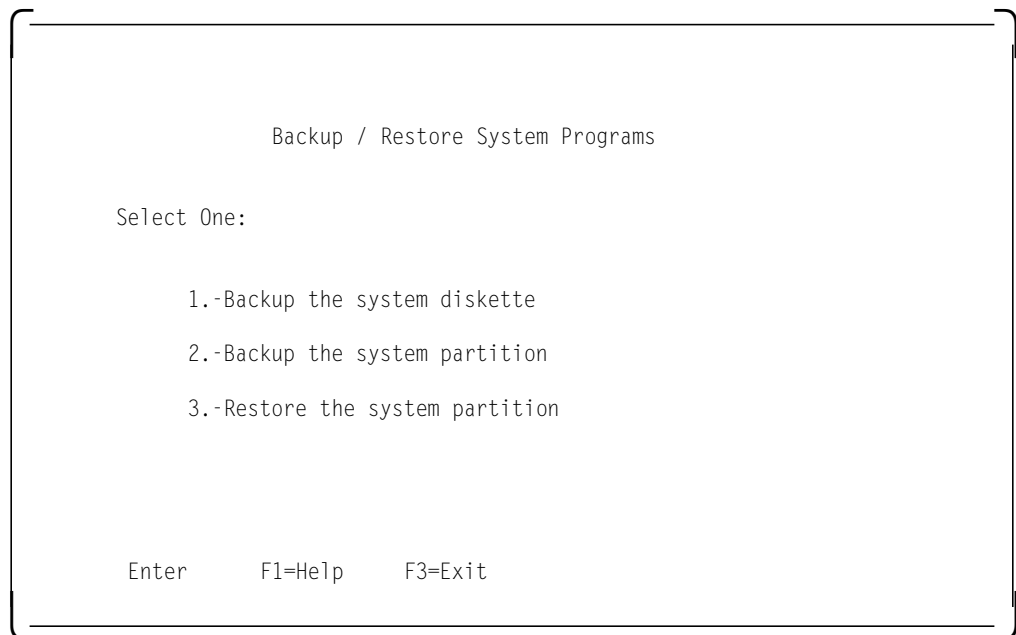


Figure 40. System Programs - Backup/Restore System Programs Menu

The following options are available:

1. **Backup the system diskettes**

Makes a backup copy of the Reference and Diagnostic diskettes.

2. **Backup the system partition**

Makes a backup of the system partition from the hard disk drive to diskettes. You need two diskettes to perform this procedure.

3. **Restore system partition**

Restores the system partition from the backup diskettes. Use this utility program to rebuild the system partition in case of accidental loss or damage.

Note

You can only use this option when the system programs are running from a diskette.

3.4.5 Set Configuration Menu

The Set Configuration menu allows you to work with the system configuration. Select this option to view, change, backup or restore the configuration. The Set Configuration menu is shown in Figure 41 on page 87.




Figure 41. System Programs - Set Configuration Menu

The following options are available:

1. View configuration

This displays the present configuration information which includes the size of the memory installed in the server, the assignment of the built-in features, and the location of installed adapters and their parameter settings.

Figure 42 on page 88 is an example of a configuration screen.

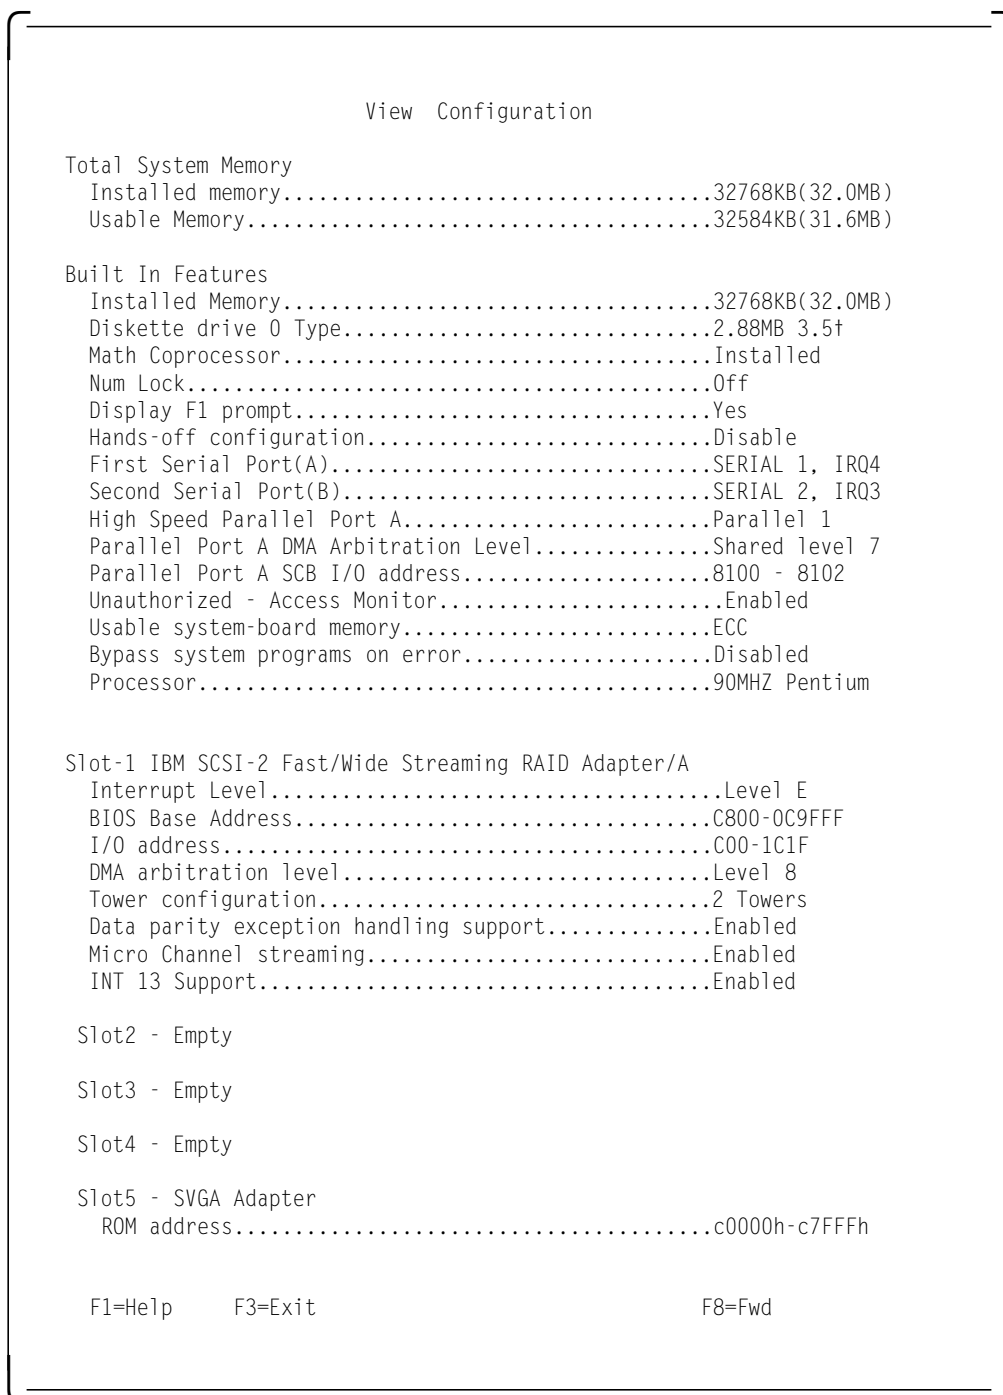


Figure 42. System Programs - View Configuration Screen

Note
 You can not make changes from the View Configuration screen.

2. Change configuration

Change configuration allows you to change configuration information used by the server. Use the Home, End, Page Up, Page Down, Up and Down arrow keys to move from field to field, and F5 and F6 to scroll through the choices. Use F10 to save the changes to non-volatile CMOS RAM.

If a setting conflicts with another, an asterisk(*) appears next to the conflicts and the word conflicts appears at the right corner of the windows. You must remove all conflicts before the configuration can be saved.

Note

Some settings are set by the system and can not be overridden. The user can change only those settings which have not been set by the system. These are the settings enclosed in brackets([..]).

3. Backup configuration

This makes a copy of the configuration information stored in CMOS RAM to the hard disk system partition or to the reference diskette. Back up the configuration of your server every time an adapter is added, removed or a configuration assignment is changed.

4. Restore configuration

This retrieves the configuration information that the user saved previously with Backup configuration program and restores this information to the CMOS RAM.

5. Run automatic configuration

This verifies and updates the configuration information. Use this program after an adapter is added or removed, to configure the server. The automatic configuration program sets system-board features first, and then adapters one at a time, starting with the lowest numbered slot.

After you use this choice, select Backup configuration to make a copy of the new configuration information.

Notes

Running automatic configuration resets parameters to their default values.

This option can be used whenever the default values for adapters will be adequate. If the default values need to be changed, then you must use the Change configuration option discussed previously. Often, the automatic configuration can be used as a starting point to ensure no settings conflict. Then, the change configuration can be used to fine tune any desired settings.

If you had a configuration where specific values were required and you experience problems after running automatic configuration, you need to change the parameters back to the required values.

6. Display memory map

This displays how memory in the address range 0C0000H through 0DFFFFH is being used by any adapters installed in the server. It identifies each device using this space by name and shows the addresses being used by that device.

This choice is helpful to determine what parts of the memory in the server are not being used (for example to find space available for an Expanded Memory Specification (EMS) page frame). Figure 43 on page 90 shows a Memory Map screen.

```

Memory Map
0C000H
  SVGA Adapter
  Micro Channel slot.....5
  Size(in KB).....32

0C8000H
  IBM SCSI-2 Fast/Wide Streaming RAID Adapter
  Micro Channel Slot.....1
  Size(in KB).....8

0CA000H
  Memory not being used
  Size(in KB).....88

F1=Help  F3=Exit

```

Figure 43. Set Configuration - Memory Map

7. Set and view SCSI device configuration

This displays the current SCSI device settings, and allows you to make changes for some devices. If a setting causes a conflict with another, an asterisk(*) appears next to the conflicts. Figure 44 on page 91 is an example of a SCSI configuration screen.

Note

For disk-array models you can *view* settings for the RAID Adapter. However, you must use the IBM RAID controller diskette to *change* the RAID adapter settings. The IBM RAID controller diskette can be obtained from Diskette Factory on the ServerGuide CD-ROM.

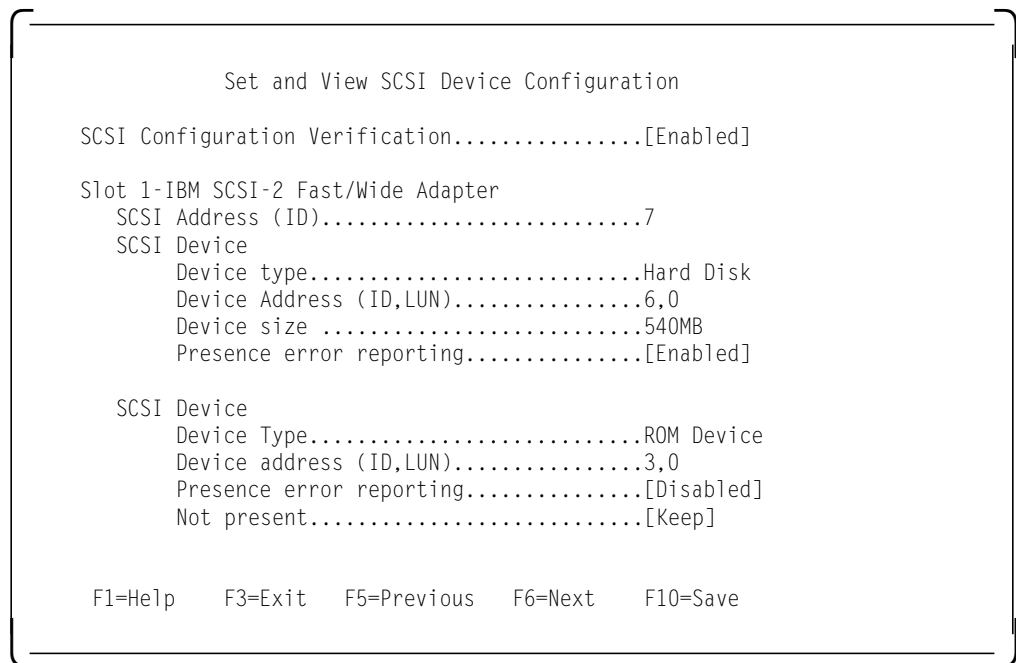


Figure 44. Set Configuration - SCSI Device Configuration

Presence error reporting works as follows:

- When presence error reporting is set to enabled on a device, POST reports a configuration error if the device is disconnected or turned off.
- When presence error is set to disabled on a device, POST does not check to see if the device is attached or turned on.

The Not present feature instructs the system on what actions to take for an individual SCSI device which was previously configured but was not found during this discovery process. Keep will cause the configuration program to leave the device in the configuration. Remove will cause the device to be permanently removed from the configuration. To remove a SCSI device:

1. Locate the Not present line for the SCSI device you want to remove.
2. Highlight the Keep setting; then press F5 or F6 to change it to Remove.
3. Press F10 to save the changed setting.

3.4.6 Set Features Menu

When the user selects this option from the Main Menu, a screen like the one in Figure 45 on page 92 will appear.




Figure 45. Set Features Menu

1. Set date and time

This allows you to set the current date and time. Once set, CMOS RAM keeps the date and time even when the system is powered off.

2. Set password and unattended start mode

This helps prevent the use of the server by unauthorized persons. Two types of passwords are available from this option:

- Power-on password
- Administrator password

When you select this option from the Set Features menu, a screen like the one in Figure 46 on page 93 will appear.




Figure 46. Set Passwords and Security Features

a. Set power-on password

This prevents unauthorized use of the server. After you set the power-on password, a password prompt will appear on the display each time the server is powered on.

To use the server, type the correct password at the password prompt and then press Enter. If the password was not entered correctly, you will see a message that says the password was incorrect and be prompted to try again. If you do not enter the correct password after three tries, you must turn the system off and then on again before further attempts are possible.

Note

If you forget the power-on password, you can remove it by moving jumper JMP1 on the system board. This jumper is located above the memory connectors.

b. Change power-on password

This option displays the instructions for changing the current power-on password which are:

1) Turn off the server and wait a few seconds; then turn it on.

2) When

Enter password

appears, type the current password followed by a space, the new password followed by a space, and the new password again.

3) Press Enter.

c. Remove power-on password

This option shows the instructions for removing the current power-on password which are:

- 1) Turn off the server and wait a few seconds; then turn it on.
- 2) When
Enter password

appears, type the current password followed by a space.
- 3) Press Enter.

Attention!

If the unattended start mode is active, you must remove it before you can change or remove the power-on password because the password prompt does not appear when the unattended start mode is set.

d. Set unattended start mode

The unattended start mode locks the keyboard and mouse when the system is powered on, but it allows the operating system and startup files to be loaded. The keyboard and mouse remain locked until you enter the power-on password.

This mode is useful for unattended operations and for network servers because it allows authorized network users access to information on the server but prohibits unauthorized access via the system keyboard and mouse.

When the server is in the unattended mode, the password prompt will not appear unless an attempt to start the system from a diskette or other removable media is issued. If you start the server from a removable media, the password prompt will appear and you must enter the correct power-on password to continue.

Note

A power-on password must be set on the server before you can set the unattended start mode because it locks the keyboard and mouse when the system is powered on.

e. Set privileged-access mode

This options allows you to set a privileged-access password (also called an administrator password). If set, this password is needed to start the system programs.

Before the privileged-access password can be set, the privileged-access switch at location JMP2 on the system board must be placed in the change state (position 0).

Attention!

Do not forget the privileged-access password. If you forget this password, the system board must be replaced before you can regain access to your system.

f. Change privileged-access password

This option allows you to change the administrator password. You must know the current privileged-access password in order to complete the procedure.

g. Remove privileged password

This option removes the administrator password. You must know the current privileged-access password in order to complete the procedure.

h. Set unauthorized-access monitor

If your system is equipped with a tamper-evident cover, you can select this option to enable or disable the unauthorized-access monitor feature. This feature works with the privileged-access password. If both features are enabled, the server will alert you if someone attempts to tamper with the locked covers. If tampering is detected, the actions taken are dependent upon the operating system. Some possible actions include halting the system and/or displaying and logging an error message.

If this feature is set and it then gets subsequently activated by cover tampering, on the next boot, the system will display a message and force the operator to enter the administrator password. If this password is not entered correctly, the system will not come up.

3. Set keyboard speed

This changes the speed at which characters respond when you type. There are two options:

• **Normal keyboard**

The normal keyboard speed displays the character on the screen at the rate of 10 characters per second.

• **Fast keyboard**

The fast keyboard speed displays the character on the screen at the rate of 30 characters per second.

4. Set console

This changes the server console. The server console is the hardware that you use to communicate with your server. There are three options for this setting:

a. Display and keyboard

Select this option if you want to use a display and keyboard as the server console. Errors will occur when the server is turned on if the display or the keyboard is not attached.

b. Display only

Select this option if you want to use a display without a keyboard as the server console.

c. No console

Select this option if you want to use the server without a system console. If you select this option, you can remove the monitor, the keyboard, and any attached ASCII terminal and the system will still come up without the errors normally associated with a missing keyboard or missing monitor.

5. Set startup sequence

This option lets you select the sequence of the drives on which the system searches for a valid boot record when you turn it on.

The default startup sequence is the first diskette drive followed by the first hard disk. You can alter this sequence with this utility. Figure 47 shows an example of the startup sequence when one hard disk is installed.

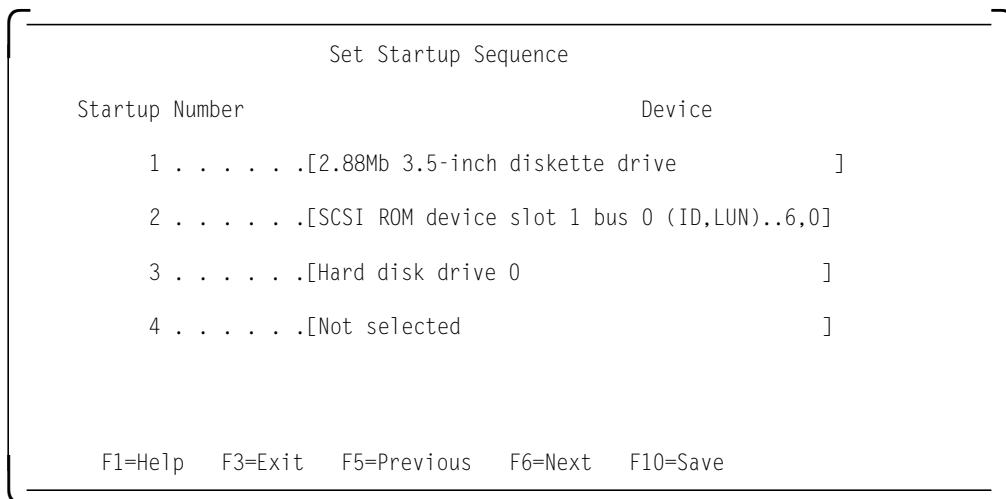


Figure 47. Set Startup Sequence Screen

To change the startup sequence, move the cursor to the device box and use the F5 and F6 keys to toggle the options. The devices that you can choose will be determined by the number of devices that you have installed. A maximum of four devices may be selected in the startup sequence.

6. Set fast startup mode

If enabled, the server runs a less extensive POST routine at cold boot time, thereby reducing the amount of startup time required by the system prior to loading the operating system.

Note

You will save some startup time in this mode. However, in the unlikely event of a memory failure, you might not be aware of it until you begin using your operating system.

7. Set power on features

This panel lets you set the options for when and how the server can be started. It can be turned on automatically at a set time, or be turned on from a remote location through a modem.

When you select this option from Set Features menu, a screen like the one in Figure 48 on page 97 will appear.

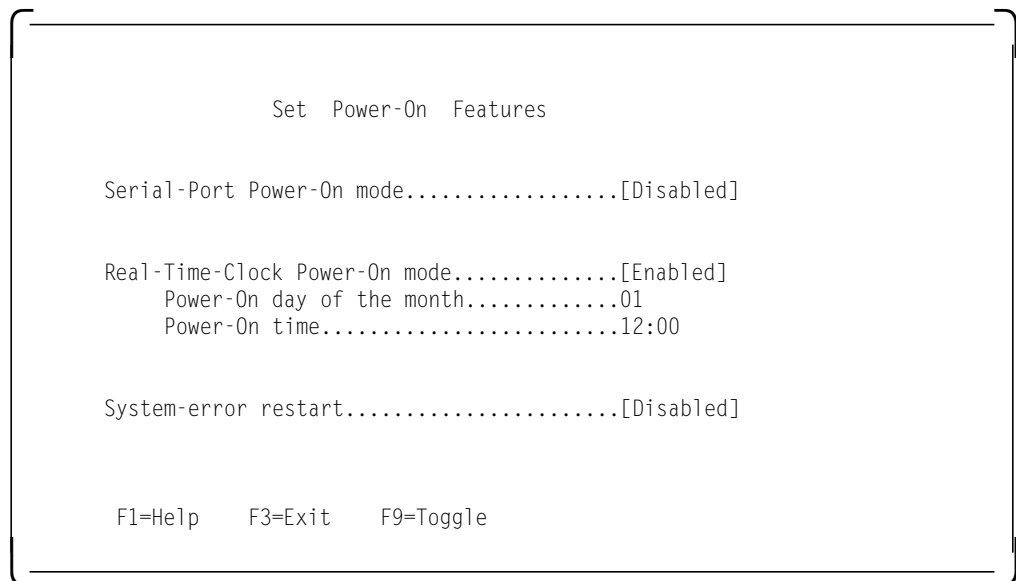


Figure 48. Set Power-On Features Screen

Serial-Port Power-On Mode: This option allows you to start the server using a modem. When this option is enabled, the server starts when the modem detects an incoming call on the attached phone line.

Real-Time Clock Power-On Mode: When this option is enabled, the server will turn on when the present time matches the day and time values you have entered.

If you want the server to start only on a certain day of the month, type the date in the space provided. If you want the server to start every day, you must leave this field without a date. Figure 48 shows an example where the server will be powered on the fifth day of every month at 12:00 PM.

System Error Restart: When the system-error restart feature is enabled, the server can automatically restart in situations that would normally result in the server stopping or locking up.

3.4.7 Test the Computer

This option is the primary method of testing the system. It can be used to test both the internal components of the system as well as the keyboard and some external devices. The amount of time required to test all components depends on the number of components. The more optional adapters and devices you have attached to the server, the longer the testing takes.

If you can not determine whether a problem is caused by the hardware or by software, you can run the diagnostic programs to confirm that the hardware is working properly. If a hardware problem is detected, an error message appears explaining the cause of the problem and the action to take.

Note

If you start the system programs from the reference diskette and select this option, a message prompts you to insert the diagnostic diskette.

3.4.8 More Utilities Menu

When you select this option from the Main Menu, a screen like the one in Figure 49 will appear.




Figure 49. More Utilities Menu

This panel has the following options:

1. Display revision levels

This displays information that you might need if you required service. If you have a printer attached, you can print the information presented on the screen by pressing the Print Screen key. Figure 50 is an example of a screen displayed with this option.




Figure 50. Display Revision Level Screen

2. Display system error log

This displays and resets the entries in the error log. As errors occur, they are recorded in this log. Error log information is helpful when service is required. Figure 51 on page 99 is an example of a System Error Log screen.

```
System Error Log           page 1 of 1

Type of error: Processor board replaced
Previous system identifier: 00001B4B
Previous manufacturing identifier: 30343236
Model: F8
Submodel: 47
BIOS revision level: 03
Number of occurrences: 1
Date of error: 01/30/95  Time of error:19:55

F1=More information      F3=Exit      F4=Clear Entry
```

Figure 51. System Error Log Screen

3. Stand-alone utility information

This describes how to use the additional utilities available on the diagnostic diskette. Some of these programs can be used only with specific operating systems.

4. Set and view system identification

This lets you view the system Vital Product Data (VPD). VPD includes the machine and model number, serial number, and part identification codes for the system board and processor board.

5. Set character font

This option allows to change the default font in the server. It displays examples of the available fonts and allows you to choose one. This will become the default font at the next cold boot time.

3.4.9 Advanced Diagnostic Program

The advanced diagnostic program is intended to be used by trained service technicians or very experienced users who can correctly interpret the advanced diagnostic screen messages. It is an alternative method of testing the server.

An advantage of running the Advanced Diagnostic program is that if you suspect a specific area of the server is causing a problem, you can choose to test only that area. This will take less time than testing all of the components in the server.

To start the Advanced Diagnostic program:

1. Start the system programs.
2. When the system programs main menu appears, press and hold Ctrl and then press A. A message appears telling you to insert the diagnostic diskette.
3. Insert the diagnostic diskette.

After a few seconds, a screen like the one in Figure 52 appears.

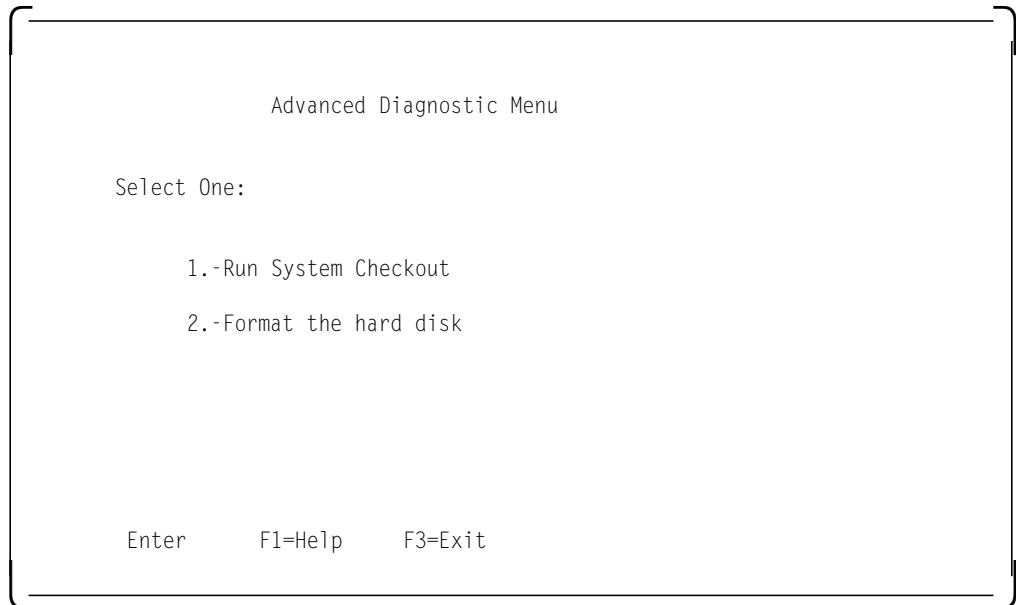


Figure 52. Advanced Diagnostic Menu

The following choices are available:

Run System Checkout: If you select this option the following menu appears:

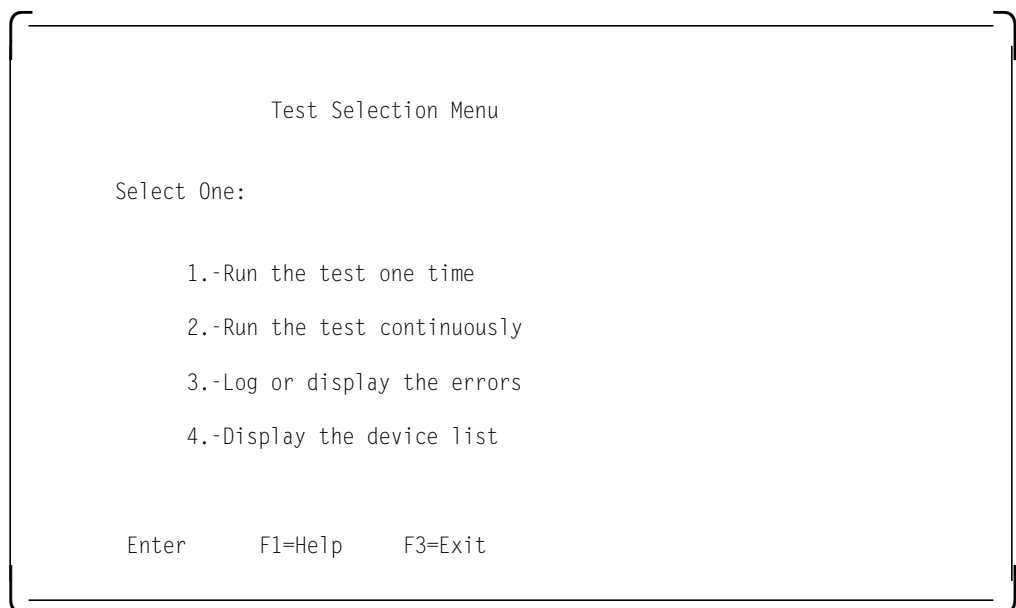


Figure 53. Test Selection Menu

1. **Run the test one time**

Select this option to test the devices one time. If you use it, you will be asked to respond to questions about areas of the server hardware. Press and hold Ctrl and then press C to stop the test.

2. Run the test continuously

Select this option to continuously test devices. Tests requiring a response are bypassed. Press and hold Ctrl and then press C to stop the test.

3. Log or display the errors

Select this option to start an error log, stop a log in process, or display an existing log. This can be useful for locating intermittent errors.

4. Display the device list

Select this option to display the list of installed devices. You can remove devices from this list.

3.4.9.1 Low-Level Format Program

On non-disk array models, you have the option of formatting your hard disk during the advanced diagnostic testing. The advanced diagnostic format program is a low-level format program. It will completely reformat the disk which includes removing the system partition.

Depending on the hard disk capacity, the low-level format program could take up to two hours.

Use the low-level format program in the following cases:

- When you are installing software that requires a low-level format
- When you get reoccurring messages from the diagnostic test directing you to run the low-level format program on the hard disk
- As a last resort before replacing a failing hard disk drive
- If you want to erase the system partition

Note

For non-disk array models, you must run the low-level format program from the diskette.

For disk-array models, you must use the programs provided on the RAID utility diskette to format the hard disk drives.

3.5 RAID Controller Utility

On RAID models, this utility is the only one that can discover and configure the attached SCSI devices.

This utility comes on a bootable diskette which is shipped with every RAID Server. Also, it can be rebuilt if necessary with the Diskette Factory option of ServerGuide.

The IBM RAID utility allows you to configure disk arrays as well as view the status of the arrays already defined. It provides a graphical display of the defined arrays and provides diagnostic information to assist you if a hard disk within the array malfunctions.

The common configuration tasks that will be described here are as follows:

1. Verify the configuration
2. Low-level format of the disks
3. Defining a hot-spare drive
4. Creating a disk array
5. Defining logical drives
6. Look at the drives information
7. Backup/restore the configuration

After booting with the RAID Controller program diskette, if more than one RAID adapter is installed, it will prompt you to select which adapter you want to work with (see Figure 54). Otherwise, a screen like the one in Figure 55 will appear.




Figure 54. RAID Configuration Program - Adapter Selection

Use the arrow keys to select the adapter you want to configure and press Enter.

A screen like the one in Figure 55 will appear.




Figure 55. RAID Configuration Program - Main Menu

3.5.1 Drive Information

During this first step, we will verify that all the installed units are recognized and available for future use.

Select **Drive Information** on the main menu.

The following screen appears:

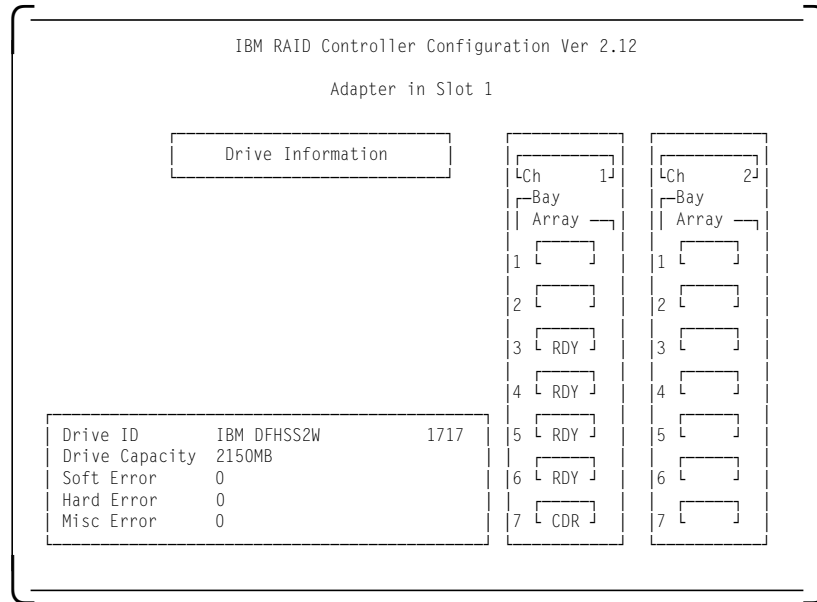


Figure 56. RAID Configuration Program - Drive Information

The columns on the right show any devices attached to each channel on the adapter. The box on the left shows information about a specific device. The unit whose information is displayed is highlighted on the right.

Use the up and down arrow keys to display the information for other devices.

Note

The RAID Utility shows device type and/or status for each device which it recognizes. The devices may take on several states which you may encounter when running the Utility.

Here is a list of indicators used along with a description of each:

- **CDR** = A CD-ROM device.
- **DDD** = A DASD device which is defunct.
- **FMT** = A DASD which is being formatted.
- **HSP** = A DASD which has been defined as hot-spare Drive.
- **OFL** = A DASD which is being used to rebuild a defunct drive. This is an intermediate state between DDD and ONL.
- **ONL** = A DASD which has been assigned to an Array.
- **RDY** = A DASD which is ready to use but not assigned to an array or defined as hot-spare drive.
- **TAP** = A Tape Drive.
- **UFM** = A DASD which is unformatted. The drive requires a low-level format before it can be used.

3.5.2 Formatting the Disks

If your DASD devices appear as UNF or if you have experienced some problems with them, you can low-level format the disks.

If your disks appear as RDY, and you don't want to reformat, then you can skip over this section.

To format a device:

1. From the main menu select **Advanced Functions**.

The following menu appears:

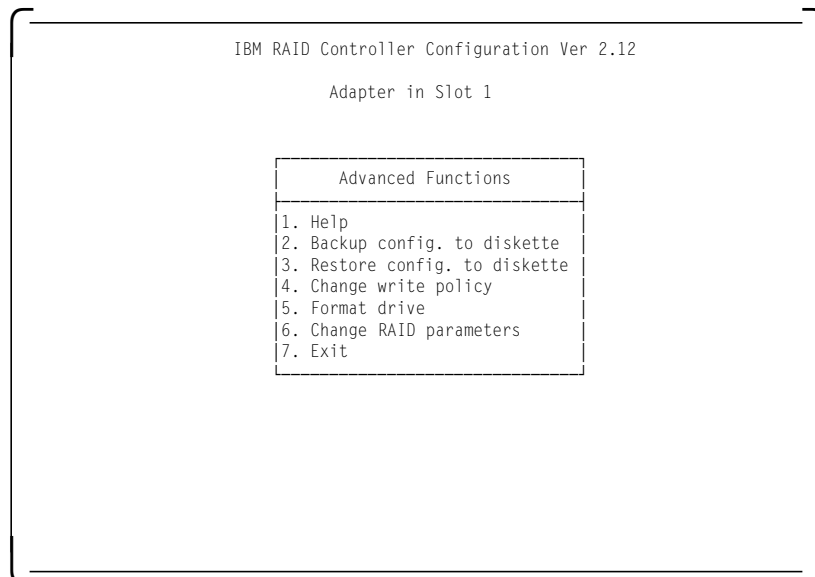


Figure 57. RAID Configuration Program - Advanced Functions Menu

2. Select **Format drive**.
3. With the spacebar, select the devices to format. They appear as FMT.

4. Press Enter to format the selected drives.

Notes

All the desired devices may be selected and formatted together to reduce the time necessary for formatting.

The operation can be time consuming. It takes approximately 20 minutes to format six 2GB disks attached to the same adapter.

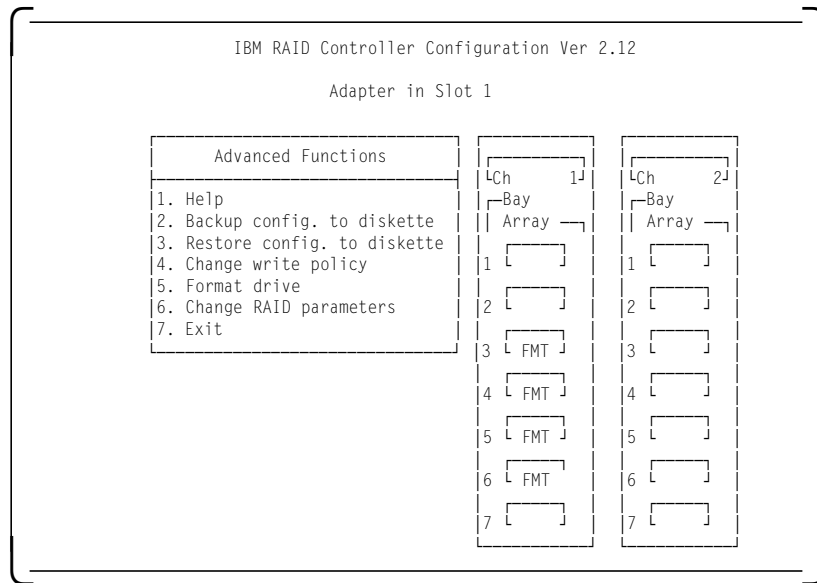


Figure 58. RAID Configuration Program - DASD Formatting

3.5.2.1 Change RAID Parameters

The panel for changing RAID parameters looks like Figure 59.

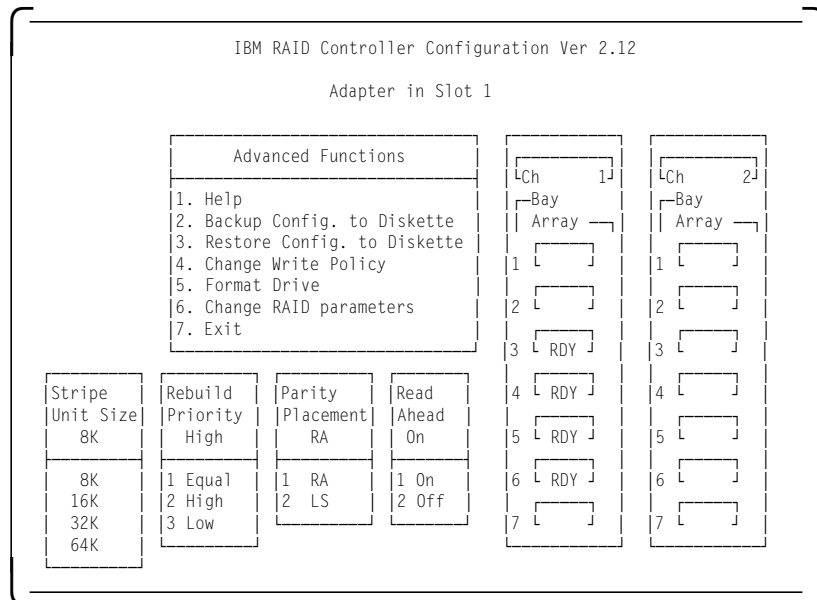


Figure 59. RAID Configuration Program - Change RAID parameters

The following fields are configurable:

Stripe Unit size: This parameter should be chosen such that the stripe unit size is close to the size of the system I/O requests. The default is set to 8KB.

Note: Once data is stored in the logical drive, the stripe unit cannot be changed without destroying data in the logical drives.

Rebuild Priority: This parameter sets the priority of the execution order of the rebuild I/O requests with respect to the system I/O requests. The rebuild priority can be changed without affecting data in the logical drive.

Parity placement: Parity placement defines how parity is placed in the disk array with respect to the data. The following figure shows the differences between Left Symmetric and Right Asymmetric parity placement.

AAA, BBB, CCC and so on are the data stripe units; PP1-3 are the corresponding parity.

Right Asymmetric (RA)				Left Symmetric (LS)			
Disk	Disk	Disk	Disk	Disk	Disk	Disk	Disk
1	2	3	4	1	2	3	4
PP0	AAA	BBB	CCC	AAA	BBB	CCC	PP0
DDD	PP1	EEE	FFF	EEE	FFF	PP1	DDD
GGG	HHH	PP2	III	III	PP2	GGG	HHH
JJJ	KKK	LLL	PP3	PP3	JJJ	KKK	LLL

The default value is Right Asymmetric. In some situations, you may want to try Left Symmetric to improve performance.

Note: Once a parity placement scheme is chosen and data is stored in the logical drives, it cannot be changed without destroying data in the logical drives.

Read-Ahead: Normally the IBM SCSI-2 Fast/Wide Streaming RAID adapter transfers data from a disk to its local cache in steps of stripe unit size. This provides excellent overall performance when workloads tend to be sequential. However, if the workload is random and system I/O requests are smaller than stripe unit size, reading ahead to the end of the stripe unit will result in a wasted SCSI bus bandwidth and wasted disk utilization.

When Read-Ahead is set to OFF, the size of data transfer from the disk to the local cache is equal to the system I/O request size, and no read-ahead to the end of the stripe unit size is performed.

Note: The Read-Ahead setting can be modified without destroying data in a logical unit.

3.5.3 Defining a Hot-Spare Disk

After selecting **Create/delete array** from the main menu of the IBM RAID Configuration program, the following screen appears:

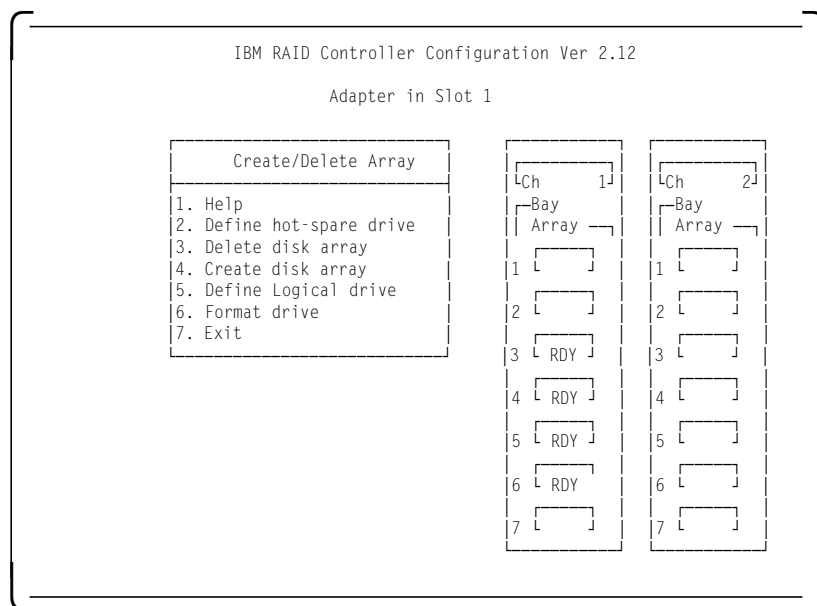


Figure 60. RAID Configuration Program - Create/Delete Array Menu

1. Select **Define hot-spare drive**. The cursor becomes active in the Bay/Array selection list.
2. Use the up and down arrow keys to select the RDY (ready) drive to be used as a hot-spare and press Enter. The RDY status changes to HSP (hot-spare). Figure 61 shows an example where a hot-spare disk has been defined.
3. Press Esc to validate.

Notes:

1. The hot-spare drive is not associated with any array. If more than one array is defined, you need only 1 hot-spare drive for all the arrays.
2. If the disks have different sizes, the hot-spare drive must be the largest.




Figure 61. RAID Configuration Program - Hot-Spare Disk Definition

3.5.4 Creating a Disk Array

1. Select **Create/Delete Array** from the Main Menu.
2. Select **Create Disk Array** from the Create/Delete Array menu and press Enter.
3. Using the Enter key, select each drive you want included in the array you are now defining. The devices will be highlighted as they are selected and the status changes from RDY to ONL (online). Also the array ID, (A for the first one, then B and so on) appears behind the drives that make up the array.
4. Press Esc to validate.

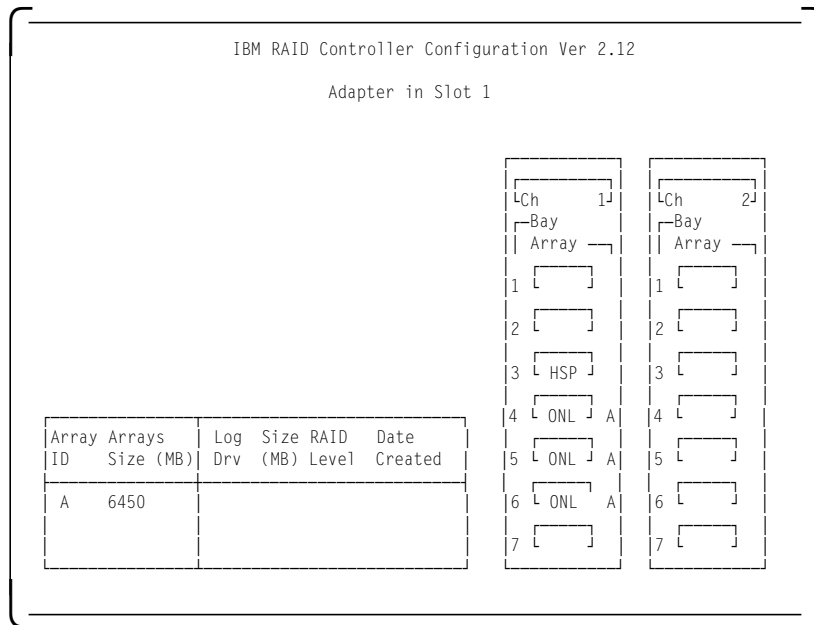


Figure 62. RAID Configuration Program - Disk Array Creation

5. Press any key to return.

Now that an array has been defined, it is necessary to create the logical drives on this array. Logical drive definitions for created arrays must be done before exiting the Create/Delete array menu.

Notes:

1. Up to four arrays can be defined per controller.
2. An array can be defined across the two channels of one controller.
3. You cannot define an array across two controllers.
4. For maximum performance and configuration flexibility, all installed drives should be configured into a single-disk array. Where data separation is desired, use separate logical drives on the same disk array. Each of these logical drives will appear as a separate physical disk to software programs like FDISK, and may be further divided into logical partitions. With this approach, if additional physical drives are added at a later date, they can be easily added as a second array without reconfiguring the original drives.

3.5.5 Defining Logical Drives

1. Select **Define logical drive** from the Create/Delete Array menu. The following screen appears:

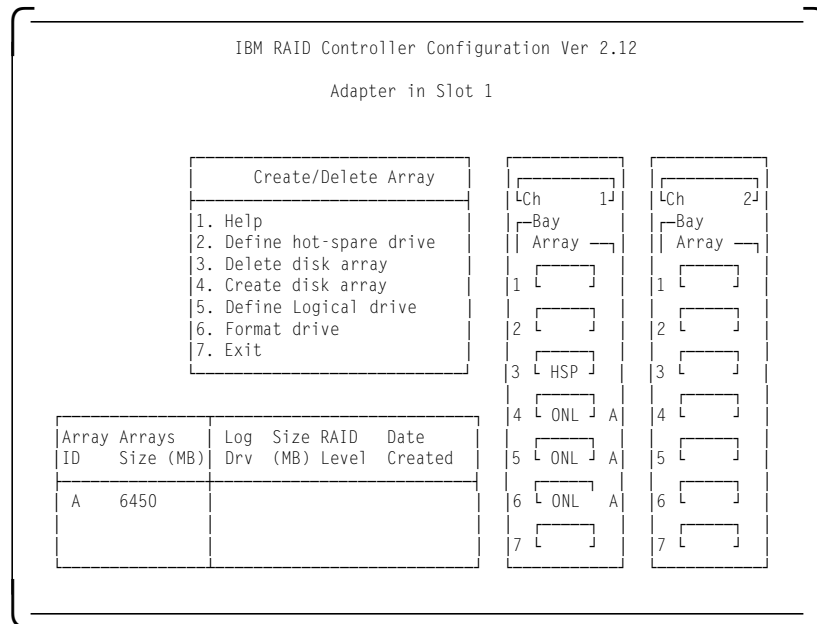


Figure 63. RAID Configuration Program - Logical Drive Definition

2. Highlight the array to be defined and press Enter.

The define logical drive pop-up window appears.

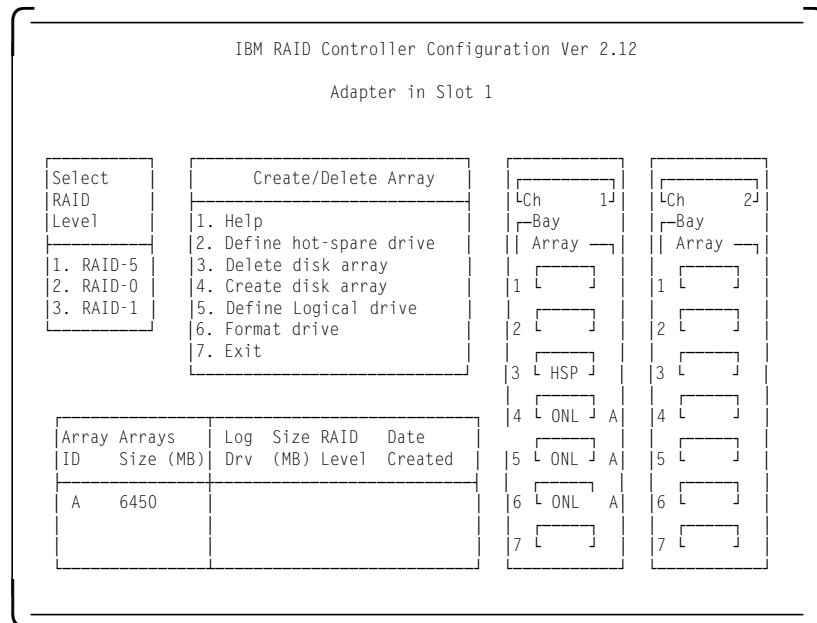


Figure 64. RAID Configuration Program - Logical Drive Definition - Array Selection

Notes:

- a. If only one hard disk has been assigned to the array, RAID level 0 will be automatically assigned and the pop-up menu does not appear.
 - b. If only two hard disks are assigned to the array, RAID-5 is not selectable, since at least three hard disk drives in an array are necessary to assign RAID level 5 to one of the logical drives.
 - c. Define any RAID 5 logical drives first. This is because the first logical drive seen will be the first RAID 5 drive created and will be assigned drive letter c:. In this way, the operating system will be installed on a RAID 5 drive.
 - d. The three RAID levels (0, 1 and 5) can be mixed in the same array allowing the user to chose the optimum RAID level for each data type.
 - e. It is, however, not recommended to mix RAID 0 with the others levels if a hot-spare drive is defined, since an automatic rebuild would not be possible.
3. Assign the RAID level wanted and press Enter.

The logical drive size pop-up menu shows the space available in this array for logical drives.

RAID 5 has been chosen in this example.




Figure 65. RAID Configuration Program - RAID level Selection

- 4. Enter the size needed, within the limitation (for example 400).
- 5. Press Enter to confirm.
- 6. As shown in Figure 66 on page 111, information about the new logical drive appears in the logical drive list.




Figure 66. RAID Configuration Program - Size Definition

7. Repeat the procedures for the other logical drives and arrays.

Note

It is not necessary to define multiple logical drive of same RAID level. You can partition these logical drives the same way you partition a real disk when you install the operating system.

8. Select **Exit** to leave this menu.
9. Select **Yes** to save the changes.

Figure 67 shows an example completed array definition which includes two logical drives: one 400MB RAID-5 array and one 2925MB RAID-1 array.

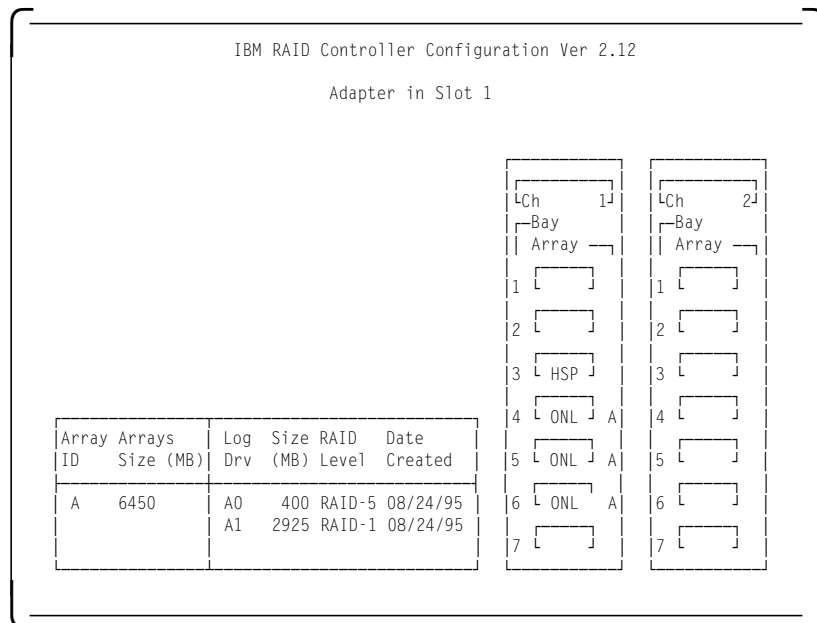


Figure 67. RAID Configuration Program - Result

3.5.6 Setting the Write Policy

Now that the logical drives have been created, you may wish to change the write policy for these drives.

If you want to change the write policy:

1. Select **Advanced functions** from the main menu.
2. Select **Change write policy**.
3. Use the up and down arrow keys to select the logical drive.
4. Press Enter to toggle the write policy.

The different options are:

WB for Write-Back: In the write-back mode, the completion status is sent before the data is actually written to the drives. While this mode will yield higher performance, it is critical to protect the DASD power to ensure the write operations actually complete, since the processor thinks the operation has been completed.

WT for Write-Through: In write-through mode, the data is written directly to disk and no acknowledgement is sent until the operation is complete.

3.5.7 Initializing the Array

If you are using RAID Level 1 or 5, you must select **Initialize logical drives** for proper operation. This sets the drive to a predetermined state. Any data existing on the drive is overwritten with zeros, and the corresponding parity is initialized to the correct value.

1. Select **Initialize/synchronize array** from the main menu.
2. Select **Initialize logical drive**.
3. Select the logical drive you want to initialize by pressing the spacebar.
4. Press Enter to start the initialization.
5. Select **Yes** to confirm.

3.5.8 Backup/Restoring the Configuration

Attention!

The array configuration is stored on EEPROM on the adapter. It is very important to back up this configuration on a diskette. If the adapter ever needs to be replaced, the configuration can be restored to the new adapter without the need to reconfigure the array. Otherwise, the array would need to be redefined and your data would be lost.

This is important not only at installation time but whenever the array configuration changes. For example, if a hot-spare drive is defined and subsequently, a drive fails and the hot-spare is enabled, the array configuration is updated. If you are unable to restore the configuration *exactly* as it was before, it would have the same effect as re-initializing the arrays: all data would be lost.

We *strongly* recommend that the configuration be backed up.

To back up the configuration:

1. Select **Advanced functions** from the Main Menu.

2. Select **Backup config. to diskette**.
3. Press Enter.
4. Enter the filename you want for backup (config is the default).
5. Insert a diskette to hold the backup file.
6. Press Enter.

To restore the configuration:

1. Select **Advanced functions** from the Main Menu.
2. Select **Restore config. from diskette**.
3. Press Enter.
4. Enter the filename you used for backup.
5. Insert the diskette used for backup.
6. Press Enter.

A screen like the one in Figure 68 will appear. This shows you the date when the backup was performed.




Figure 68. RAID Configuration Program - Advanced Functions Menu

7. Select **Yes** to confirm.

Chapter 4. Novell NetWare Installation

In this chapter, we examine the different steps required to install Novell NetWare.

Novell NetWare can be installed in the following three ways:

- With the ServerGuide CD-ROM
- With the original product diskettes
- With the original product CD-ROM

We will focus on the ServerGuide process.

4.1 ServerGuide Overview

ServerGuide is a CD-ROM library shipped with every IBM PC Server. It contains software products, device drivers, utility programs, and diskette images. The software products are encrypted and require activation keys which are used to *unlock* them. This provides a convenient method of distributing the software products and at the same time allows for an integrated, time-saving approach to product installation.

Each ServerGuide package contains a unique license number, which must be specified on your software order with the software product(s) part number. The ServerGuide license number and the software product(s) part number are used by IBM to generate activation keys, which are given to you. The activation keys can then be used to unlock the software product(s) for installation.

The following is the list of software products for which you will need an activation key in order to install:

- OS/2 Warp Full Pack
- OS/2 for SMP
- OS/2 LAN Server Advanced 4.0
- OS/2 LAN Server Entry 4.0
- System Performance Monitor (SPM/2)
- TCP/IP V2.0
- Novell NetWare V3.12
- Novell NetWare V4.1

4.2 Starting ServerGuide

You can start ServerGuide either from a bootable CD-ROM drive or from a diskette. To start ServerGuide from the CD-ROM drive:

1. Insert the CD entitled "ServerGuide - MAIN CD".
2. Turn on the system.

To start ServerGuide from the diskette drive:

1. Insert the CD entitled "ServerGuide - MAIN CD".
2. Insert the license diskette in the diskette drive.
3. Turn on the system.

After the POST tests run, ServerGuide boots, and the ServerGuide language selection screen appears (see Figure 69 on page 116).

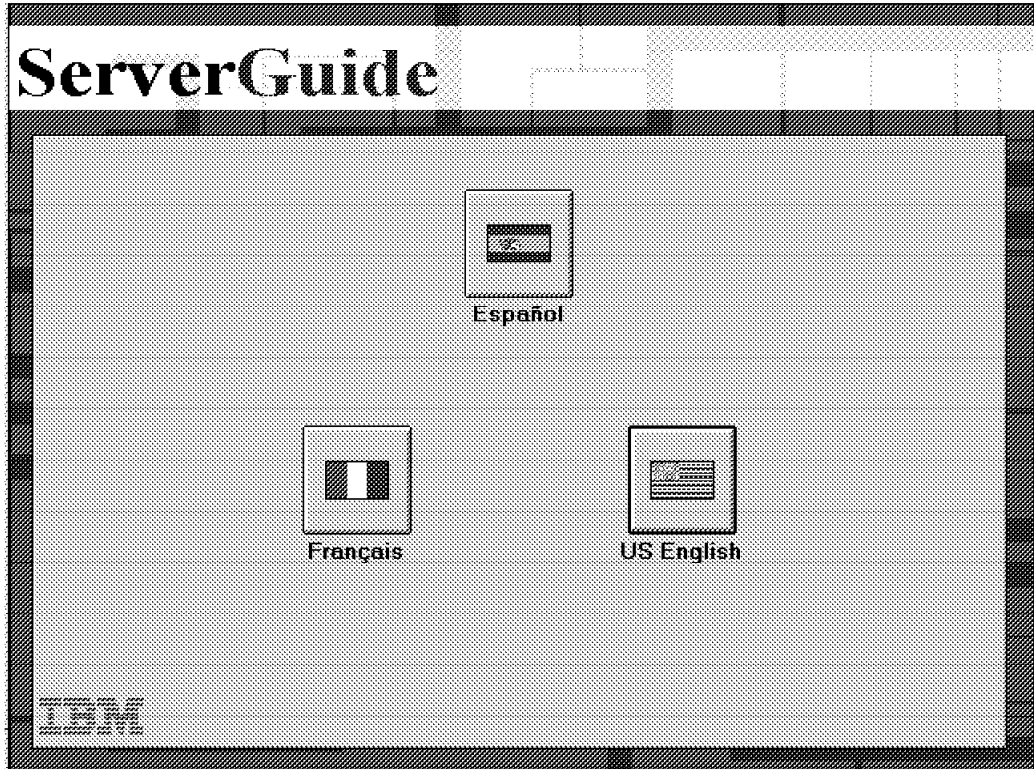


Figure 69. ServerGuide Language

Select the language that you want to use during the software installation process. There are three options:

- Espanol
- Francais
- US English

For this installation process, we selected **US English**.

The Main Menu appears See Figure 70 on page 117.




Figure 70. ServerGuide Main Menu

This menu contains seven options:

Start Here: This option contains information to help you become familiar with ServerGuide features and the contents of the CDs.

Software Installation: Select this option to install any of the software programs included with ServerGuide.

NetFinity Overview: This allows you to learn more about NetFinity and its capabilities.

Utility Programs: ServerGuide provides a variety of useful utility programs:

- Diskette Factory, used to create diskette images from CD
- Restore the System Partition, used to restore the system programs in the system partition (a copy of system programs needed for your model is stored on ServerGuide)
- Custom Settings, used to automatically set certain functions to enabled or disabled at startup
- Security Features, used to set your power-on password

Performance Tuning: Some models (IBM PC Server models 95 and 500 for example) allow you to select how you intend to use the server:

- Application/Database Server
- File Server
- Multimedia Server
- General Use (mixed environment)

The performance tuning programs can tune your server automatically based on this input.

Online Documentation: The online documentation includes:

- Product documentation
- Optional installation documentation
- NetFinity documentation
- NOS Integration Guides which provide detailed information on:
 - IBM PC Servers Technologies
 - Hardware Configuration
 - NOS Installation
 - Fault Tolerance
 - Network Management
 - Hardware Tuning
 - NOS Tuning

Product Information and Demonstrations: ServerGuide provides demonstrations to help you learn about some of the products available from IBM such as:

- NetFinity
- OS/2 Communication Manager
- OS/2 LAN Server

4.3 Installing NetWare 4.1 with ServerGuide

To install NetWare using ServerGuide, from the main menu:

1. Select **Software Installation**.

The software installation screen appears. It includes two options: OS/2 and NetWare.

2. Select **Novell NetWare**.

The select a keyboard screen appears. This feature includes several kinds of keyboards that you can use in your server.

3. Select **United States**.

The select a country screen appears.

4. Select **United States**.




Figure 71. Installing NetWare

The installing NetWare screen appears. This screen allows you to select the NetWare version to install. You also can specify whether you want to install NetFinity Services for NetWare on the server.

Note

For NetWare 4.1, ServerGuide provides the ability for you to transfer your existing license during the software installation process. If you transfer your NetWare license, you do not need to purchase a software activation key.

5. Select **Install NetWare 4.1**.
6. Select the number of connections. This corresponds to the activation key or license that you purchased previously.




Figure 72. Configuring NetWare

A screen like the one in Figure 72 will appear. This screen confirms your choice of country and keyboard and asks you to confirm the LAN adapter and Disk Controller that ServerGuide detected in the server.

7. Accept the default or modify the values if you want to use a different configuration.

If you selected to install NetFinity Services, the Configuring IBM NetFinity screen appears (see Figure 73 on page 121).

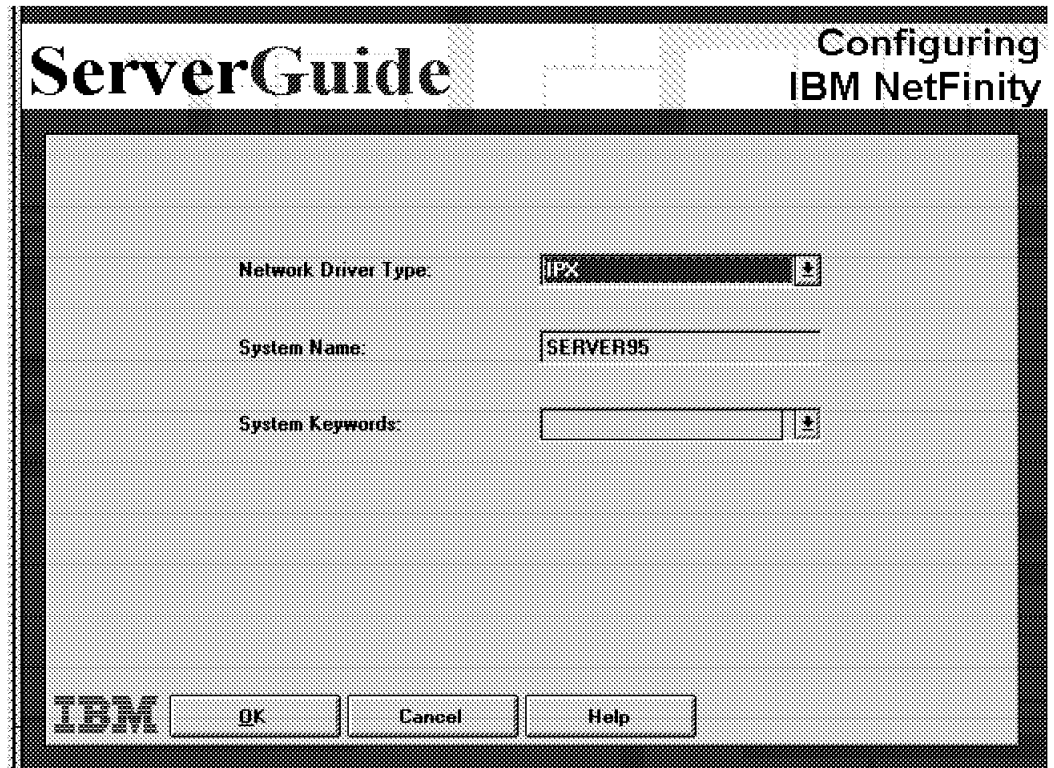


Figure 73. Configuring IBM NetFinity

8. Specify the following information:
 - Network Driver Type (Protocol)
 - System Name
 - System Keywords
9. Select **OK**.

Note

You can also use the NF config command from the server console to configure NetFinity after the initial installation.

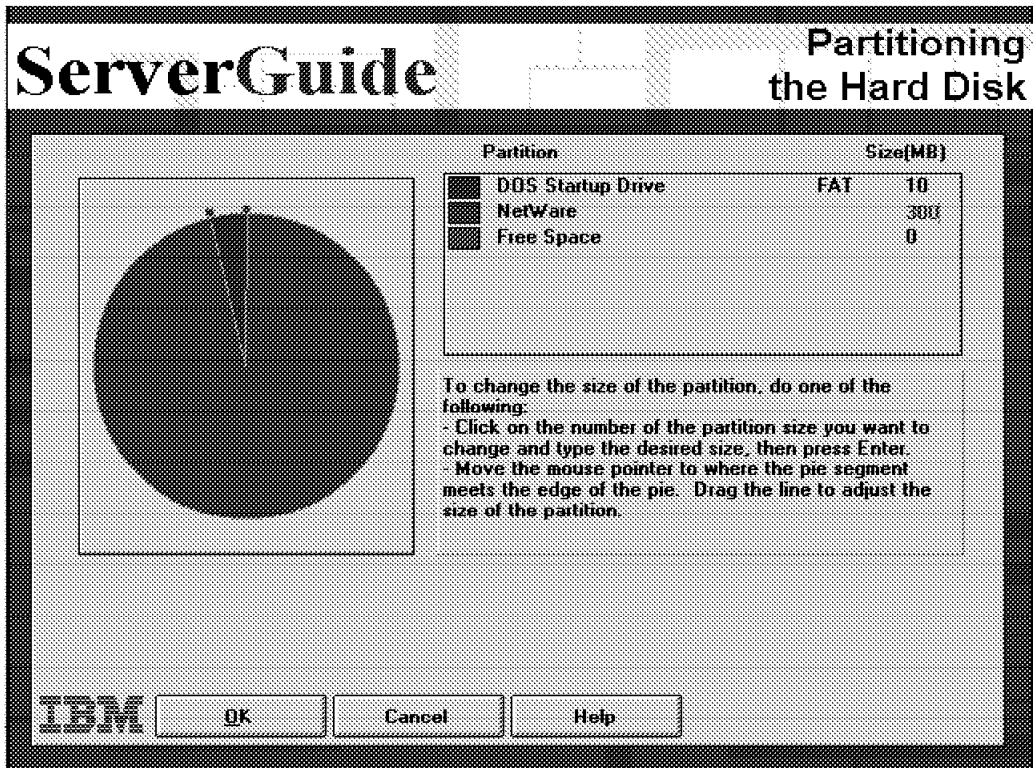


Figure 74. Partitioning the Hard Disk

A screen like the one in Figure 74 appears. ServerGuide defaults to create the following partitions:

- DOS Startup Drive
- NetWare
- Free Space

10. Select **OK** if you want to accept the values or change them.




Figure 75. Reviewing Configuration

At this point, you will see a screen like the one in Figure 75. You can review or change the selected options you have just entered.

11. Select **OK** after you have verified any parameters you wish to recheck.

At this point, the server is rebooted and the hard disk is partitioned by ServerGuide.




Figure 76. Unlocking Programs

12. When the server comes back up, you will see a screen like the one in Figure 76. Enter the activation keys for the NetWare version selected as shown in Figure 71 on page 119.
13. Select **Unlock**.
14. You are prompted to insert the NetWare CD and remove the license diskette. ServerGuide proceeds with the installation.

Attention

While ServerGuide is installing software, be sure not to touch any keys on the keyboard unless prompted by the install program. You will see the NetWare installation panels come up and change quickly as the ServerGuide installation progresses. However, do not touch any keys unless *ServerGuide* prompts you to do so.

15. If you selected to install NetFinity Services for NetWare, you are prompted to change the NetWare CD to the NetFinity CD.
16. Press Enter to exit from ServerGuide when all files have been transferred from the CD to hard disk.
17. Type DOWN at the colon prompt and then press Enter.
18. Type EXIT at the colon prompt and then press Enter to exit to DOS.

Note

At this point, the NetWare operating system has been installed. However, NetWare Directory Services (NDS) has not been installed. The following steps take you through this process.

19. Type SERVER from the NWSERVER directory to restart NetWare.
20. Install NetWare Directory Services.
 - a. Type LOAD INSTALL and press Enter to start the installation utility.
 - b. Select **Directory Options(Install NetWare Directory Services)** from the main menu and press Enter.
 - c. Select **Install Directory Services onto this server** and then press Enter.

A screen like the one in Figure 77 will appear.



Figure 77. Installing NetWare Directory Services (NDS)

- d. Select **Yes, this is the first NetWare 4 server** if this is the only NetWare 4.X server on your network. Otherwise, if you have other NetWare 4.X servers, check your network connections to make sure that you can communicate with them over the LAN connection where you are installing this server.
- e. A screen like the one in Figure 78 will appear. Enter a name for this directory tree and press Enter.



Figure 78. Assigning a Name to a Directory Tree

- f. Select a time zone and press Enter.
- g. Verify the time configuration and press F10 to save and continue.

A screen like the one in Figure 79 on page 126 will appear. This is where the context for the server will be defined.

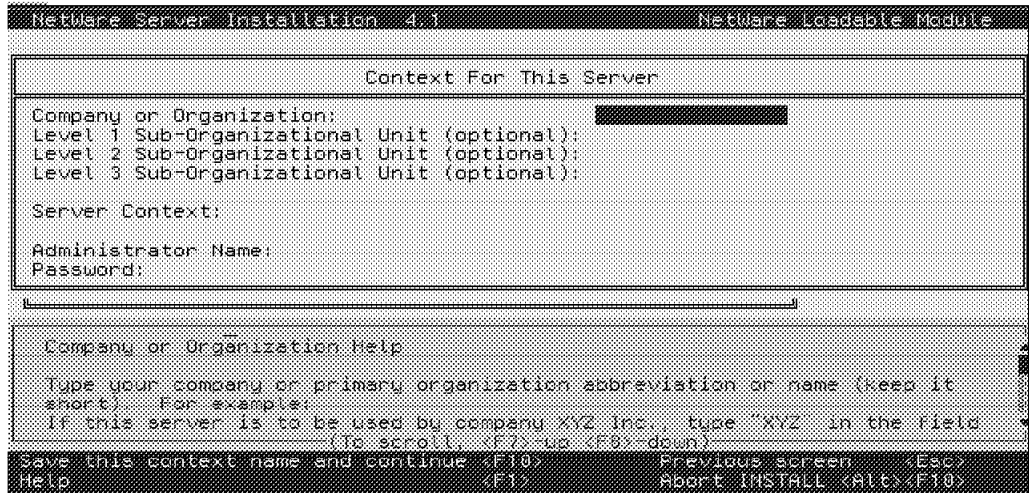


Figure 79. Assigning a Context for the Server

- h. Specify the server context and type an administrator password.
- i. Retype the administrator password.
- j. Select **Yes** when the system asks you to save the directory information.

A screen like the one in Figure 80 will appear. All the information that you have previously entered for NetWare Directory Services is shown in this screen.

Attention

Be sure to remember this information. You will need the context of the server to log on as well as to perform many administrator and end-user tasks. Please also remember that the context is *case-sensitive*.

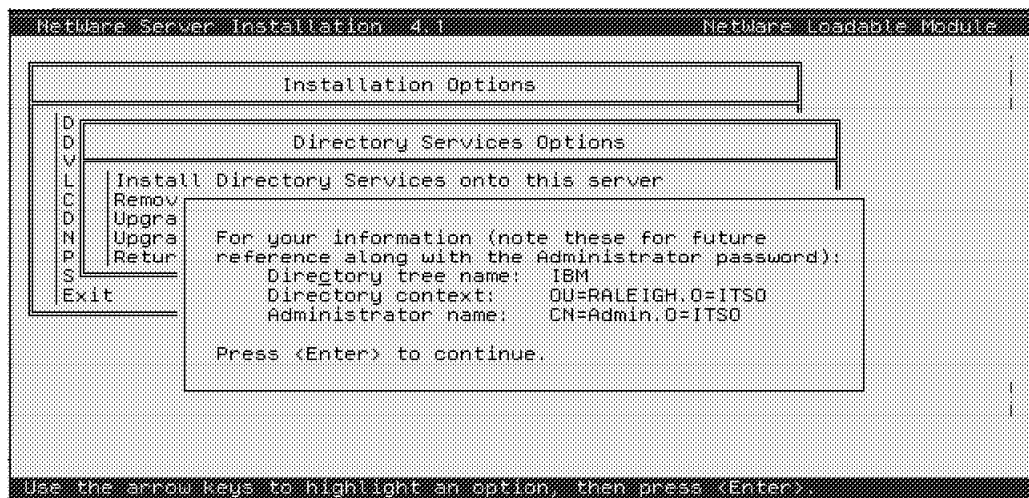


Figure 80. Server Context Information.

- k. Check the directory information and then press Enter.
- l. Press Esc twice and select **Yes** to exit the installation utility.

21. Type `DOWN` at the colon prompt.
22. Type `EXIT` at the colon prompt to return to DOS and finish the installation.
23. Reboot the server and verify that the server comes up properly.

The installation is complete. You are now ready to log on as the administrator and add users and applications.

4.4 Installing NetWare 3.12 with Diskettes

In this section, we detail the procedures necessary to install NetWare V3.12 via the original product diskettes.

4.4.1 Hardware Requirements

The following hardware is required to install NetWare 3.12:

- One file server with an Intel 386 processor or higher
- A minimum of 6 MB of RAM
- A minimum of 10 MB in the DOS partition
- At least one network adapter card for the server

4.4.2 Software Requirements

The following software is required to install NetWare 3.12:

- NetWare 3.12 diskettes
- DOS 3.3 or later
- Drivers for disk controller and LAN adapter cards

4.4.3 Information Requested at Time of Installation

When you are installing NetWare 3.12, you will have to provide the following information:

- Server name
- IPX internal network number
- Locale information including country code, code page, and keyboard mapping
- Filename format
- Disk type
- I/O port address and interrupt setting of hard disk controller
- LAN drive type
- I/O port address and interrupt value of network board

4.4.4 Installation Files

The NetWare 3.12 installation uses the following files:

1. **INSTALL.BAT**

INSTALL.BAT is a utility that is used to automate the installation of the server boot files.

2. **SERVER.EXE**

This is the primary network operating system file that identifies the server as a NetWare 3.12 server. During installation it is copied to a directory on the DOS partition.

3. **Disk Drivers**

These are specialized NLMs (NetWare Loadable Modules) that work with the disk controller and act as the interface to the disk subsystem to provide the server's disk operating system. Disk drivers are loaded from the directory where NetWare was installed after SERVER.EXE is executed. Disk driver files have a .DSK extension.

4. **INSTALL.NLM**

INSTALL.NLM is the main utility used during the server installation process. It is used to create and manage NetWare partitions, hard disks and volumes, create and manage NetWare configuration files, and upload NetWare diskettes.

5. **LAN Drivers**

These are specialized NLMs that provide the control interface to the network adapter card. The LAN drivers available with NetWare 3.12 are installed in the SYS:SYSTEM directory.

6. **STARTUP.NCF**

STARTUP.NCF is a server configuration boot file; it contains the commands to load the disk driver for the server. It is stored on the disk from which you boot. It mounts volume SYS, and then turns control over to the AUTOEXEC.NCF file.

7. **AUTOEXEC.NCF**

AUTOEXEC.NCF is a server configuration boot file; it provides the server with commands to complete the boot process after SERVER.EXE and STARTUP.NCF are executed. It runs from the SYS:SYSTEM directory.

4.4.5 **Installation Procedure**

1. Install DOS on a small DOS partition.

A minimum of 10 MB of disk space is required because server boot files are copied to this partition during installation.

2. Reboot the server.
3. Insert diskette_1 into the server's diskette drive.
4. Start the NetWare install utility.

Type `A:INSTALL` and then press Enter.

5. A screen like the one in Figure 81 on page 129 will appear.




Figure 81. NetWare Installation

6. Select **Install new NetWare 3.12** and then press Enter.
7. Select **Retain Current Disk Partition** and then press Enter.
8. Enter a server name and the press Enter.
9. Enter an IPX internal network number or accept the default and then press Enter.

Note

An IPX internal network number is a number used to identify the server from any other server on the network. The IPX internal network number must be different than the network number used for the LAN adapter card.

10. Specify to which directory the boot files will be installed or accept the directory default. The default is SERVER.312. Press Enter.
11. The system will prompt you for the SYSTEM_1 diskette; insert it into the diskette drive and press Enter.
12. The system will prompt you for the SYSTEM_2 diskette; insert it into the diskette drive and press Enter.
13. The system will prompt you for the UNICODE diskette; insert it into the diskette drive and press Enter.
14. Select country code, code page, and keyboard mappings and then press Enter.
15. Select a filename format and then press Enter.
16. Set up any special startup set commands (optional).

Use this to enter any specific set commands which might be required for the LAN or DISK drivers upon server startup. Most adapters will work with the defaults to get you up and running. See the documentation that came with your LAN adapter if you are not sure.

17. Assign AUTOEXEC.BAT to load SERVER.EXE.

When the system asks whether the DOS AUTOEXEC.BAT should load the SERVER.EXE, select **Yes** and press Enter.

At this point, SERVER.EXE is automatically invoked and the system console screen will appear.

18. Load the disk driver.

Enter the appropriate disk driver at the colon prompt. For example, for an IBM PC server 320 with a RAID controller, you must use the IBMRAID.DSK driver.

Type `LOAD C:SERVER.312IBMRAID.`

19. Type `LOAD INSTALL.NLM` and press Enter. You will see a screen like the one in Figure 82.




Figure 82. NetWare V3.12 Installation - Main Menu

20. Select **Disk Options** from the Installation Options menu and then press Enter.

21. Select **Partition Tables** from the available disk options menu. You will see a screen like the one in Figure 83.

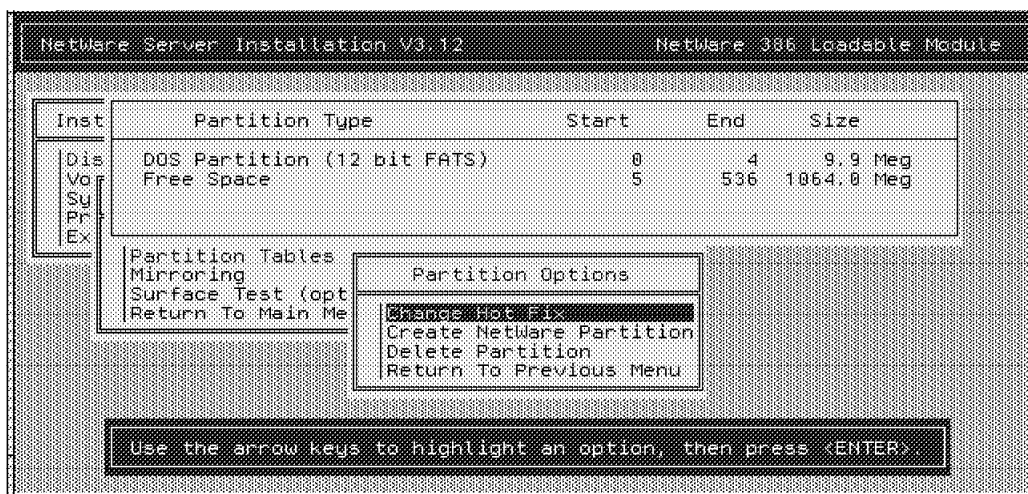


Figure 83. Create Partition

22. Select **Create NetWare Partition** and press Enter. You will see a screen like the one in Figure 84 on page 131.
23. Press Esc. This will accept the default values, namely the partition size, which is the remaining free space on the drive.




Figure 84. Partition Information

24. Select **Yes** to create the partition and then press Esc twice to return to the Installation Options menu.
25. Select **Volume Options** from the Installation Options menu.
26. Press the Insert key to create a volume. You will see a screen like the one in Figure 85.




Figure 85. Creating a New Volume

The volume with the name of SYS: will be created by default using all the available free space in the partition. If you want to create more volumes, you need to change the Initial Segment Size of the SYS: volume.

27. Press Esc to create the SYS: volume.
28. Select **Yes** and press Enter.
29. Highlight SYS: and press Enter.

30. Place the cursor on the status field and press Enter. A screen like the one in Figure 86 on page 132 will appear.




Figure 86. Volume Status

31. Select **Mount Volume** and then press Enter.
32. Press Esc to return to the Installation Options menu.
33. Select **System Options**.
34. Select **Copy System and Public Files** and press Enter. A screen like the one in Figure 87 will appear.




Figure 87. Copy System and Public Files

The system will prompt you to insert some diskettes. The diskettes contain files which make all the NetWare commands and utilities available. When the system is finished copying all the files, you will return to the Installation Options menu.

35. Press Alt + Esc to leave the installation module in memory and return to the system console prompt.
36. Load the LAN driver at the prompt.

For example if you are using an IBM 16/4 ISA Adapter you would type LOAD TOKEN, and then press Enter.

37. Press Enter to accept the default value for the I/O port address, or change the address if your LAN adapter needs a specific value.
38. When the system prompts you for the interrupt setting, press Enter to accept the default value, or change the value to the correct interrupt setting for your LAN adapter.
39. You now must bind the IPX protocol to your LAN adapter. In this example, TOKEN is the driver for the LAN adapter, so you would type the following:

```
BIND IPX TO TOKEN
```

If you have a different adapter, you must substitute the appropriate LAN driver for your adapter.

40. Enter the network number.

The network number is an eight digit hexadecimal number which is the network address for the LAN segment that is attached to the network board.

41. Press Alt + Esc to return to the Installation Options menu.
42. Select **System Options** from the Installation Options menu.
43. Select **Create STARTUP.NCF File**.




Figure 88. Path for STARTUP.NCF File

44. Press Enter to accept the default path (see Figure 88).

A screen like the one in Figure 89 on page 134 will appear. As the installation has progressed, NetWare has kept track of the options you have chosen for your disk driver. It automatically places this command in the STARTUP.NCF file for you as well as the default value for the minimum number of packet receive buffers.




Figure 89. STARTUP.NCF File

45. Press Esc and select **YES** to save the STARTUP.NCF file.
46. Press Esc to return to the System Options menu.
47. Select **Create AUTOEXEC.NCF File**.

A screen like the one in Figure 90 will appear. All the console commands you have previously entered have been placed in the AUTOEXEC.NCF file.

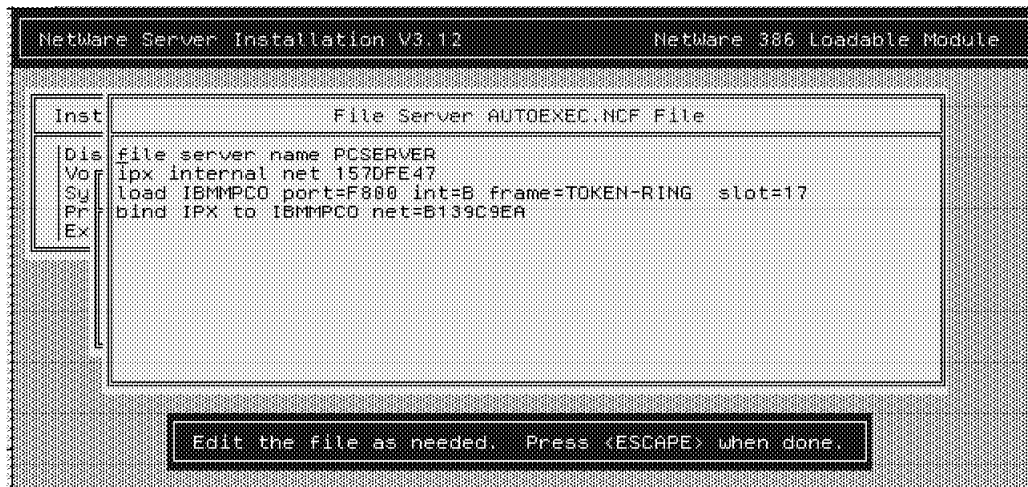


Figure 90. AUTOEXEC.NCF File

48. Press Esc and select **Yes** to save the AUTOEXEC.NCF file.
49. Press Esc twice to exit the installation utility.
50. Type DOWN at the colon prompt.
51. Type EXIT at the colon prompt to return to DOS and finish the installation.
52. Reboot the server and verify that the server comes up properly.

All of the commands included in the STARTUP.NCF and AUTOEXEC.NCF files should execute correctly and return you to a colon prompt indicating that the server is up and running.

4.5 Installing NetWare 4.1 with the Original CD-ROM

In this section, we detail the procedures necessary to install NetWare 4.1 using the original product CD-ROM.

4.5.1 Hardware Requirements

The following hardware is required to install NetWare 4.1:

- One file server with an Intel 386 processor or higher
- A minimum of 8 MB of RAM
- A minimum of 20 MB in a DOS disk partition
- At least one network adapter card
- A CD-ROM drive

4.5.2 Software Requirements

The following software is required to install NetWare 4.1:

- NetWare 4.1 CD
- DOS 3.3 or later
- Drivers for the adapter cards

4.5.3 Installation Procedure

1. Install DOS in the primary partition.

A minimum of 20 MB of hard disk is required; this is because during the installation, boot files are copied to this partition.

2. Install the drivers for CD-ROM access. See Appendix C, "Configuring DOS CD-ROM Support" on page 203 for more information on configuring the system for CD-ROM support.

During this process the CONFIG.SYS and AUTOEXEC.BAT files are modified and a drive letter is assigned to the CD-ROM drive.

3. Reboot the server.
4. Insert the NetWare 4.1 CD into the CD-ROM drive.
5. Change to the CD-ROM drive.
6. Type `Install` and press Enter.

The installation menu appears as shown in Figure 91 on page 136.

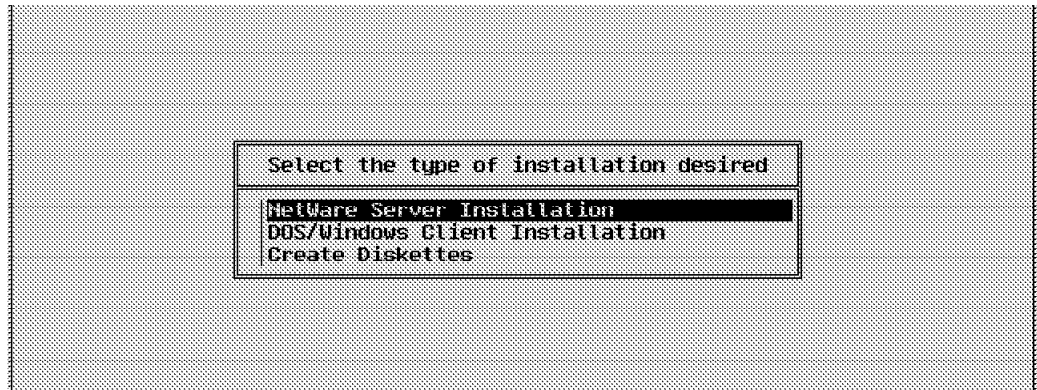


Figure 91. Installation Menu

7. Select **NetWare Server Installation** and press Enter.
8. Select **NetWare 4.1** and press Enter.
9. Select the kind of installation that you want perform. There are two options, as follows:
 - **Simple Installation**

If you select this option, the installation program sets the default values. You must specify only the server name. The installation then proceeds to copy the boot files to the DOS partition.
 - **Custom Installation**

If you select this option, the installation program will ask you for values to complete the installation.
10. Select **Custom installation of NetWare 4.1** and press Enter.
11. Specify the server name and press Enter.
12. Enter an internal IPX network number or accept the default and then press Enter.

Note

The IPX internal network number is a number used to identify the server from any other server on the network. It must be different than the network number used in the bind statement which assigns a network number to the LAN adapter.

13. Specify in which directory the server boot files will be installed or accept the default. The default is NWSERVER. Press Enter.

The server boot files are transferred from the CD to the DOS partition.
14. Select a locale configuration for the server. This includes a country code, code page, and a keyboard mapping. Press Enter.
15. Select the filename format and then press Enter.
16. Set up special set commands.

This is optional and it is used only if you need specific parameters for your LAN adapter(s). Most will operate with the defaults.
17. Select **Yes** to assign AUTOEXEC.BAT to load SERVER.EXE.

18. Specify the path where the AUTOEXEC.BAT file is located or press Enter to accept the default. The default is C:\AUTOEXEC.BAT.

A screen like the one in Figure 92 will appear.



Figure 92. Disk Driver Options

19. Select a disk driver.

If the driver that you need is not in the list, insert the disk driver diskette appropriate for your adapter and press the Insert key.

A screen like the one in Figure 93 will appear.




Figure 93. Network Driver Options

20. Select a network driver.

If the driver that you need is not in the list, insert the LAN driver diskette appropriate for your adapter and press the Insert key.

21. Verify the drivers you have selected and select **Continue Installation**.

22. Assign a network number for the LAN adapter card(s) in your server. This number is used in the bind IPX statement. Press Enter to accept a random network number.

Attention

This number must be different from the IPX internal network number.

23. Select **Manually** when the system asks you to choose the method for setting up the NetWare partition(s).

24. Select **Create, delete and modify disk partitions** and press Enter.

A screen like the one in Figure 94 will appear.



Figure 94. Create Partition

25. Select **Create NetWare disk partition**.

A screen like the one in Figure 95 will appear.

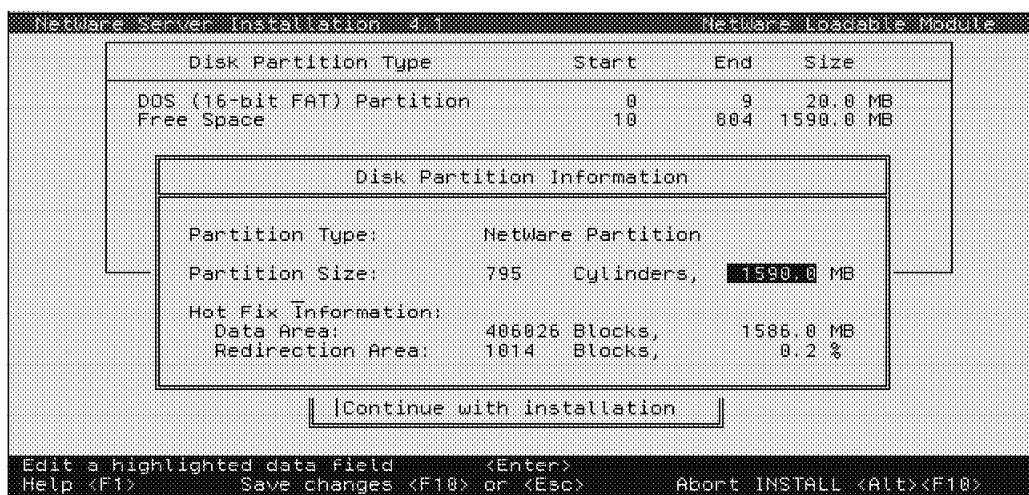


Figure 95. Disk Partition Information

26. Press F10 and select **Yes** to create the NetWare partition.

27. Select **Return to Previous Menu** and press Enter.

28. Select **Continue with Installation** from the Disk Partition and Mirroring menu and press Enter.

29. Press Enter to continue when the information for available disk space appears.

A screen like the one in Figure 96 will appear.



Figure 96. New Volume Information

30. Press F10 to accept the defaults for volume SYS:.
31. Select **YES** when the system asks whether you want to save the volume changes.
32. The system asks for a license diskette. Insert the license diskette in drive A: and press Enter.

A screen like the one in Figure 97 will appear.

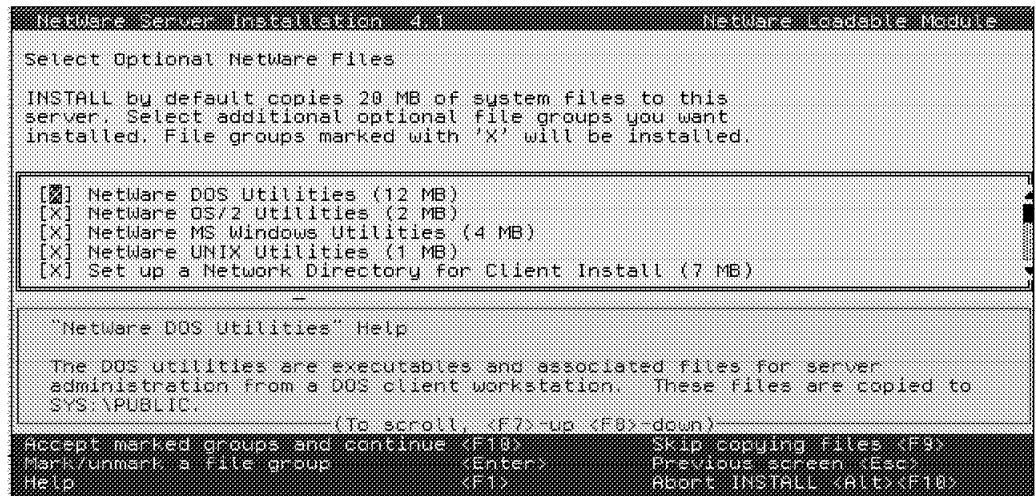


Figure 97. Optional NetWare Files

33. Select the optional NetWare files that you want to install.
34. Press F10 to accept the marked group.

The selected files are copied immediately from the CD.

When this completes, a screen like the one in Figure 98 on page 140 will appear.

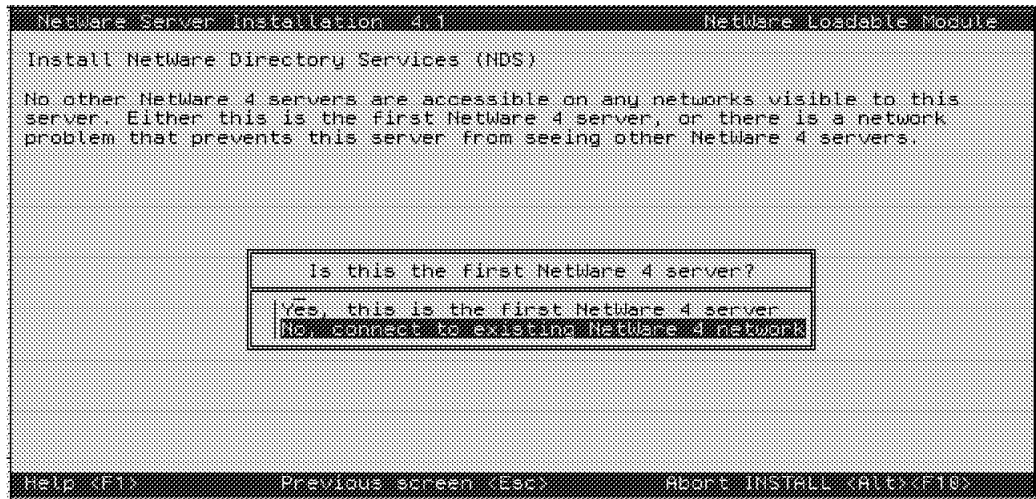


Figure 98. Install NetWare Directory Services(NDS)

35. Select **Yes, this is the first NetWare 4 server** if this is the only NetWare 4.X server on your network. Otherwise, if you have other NetWare 4.X servers, check your network connections to make sure that you can communicate with them over the LAN connection where you are installing this server.

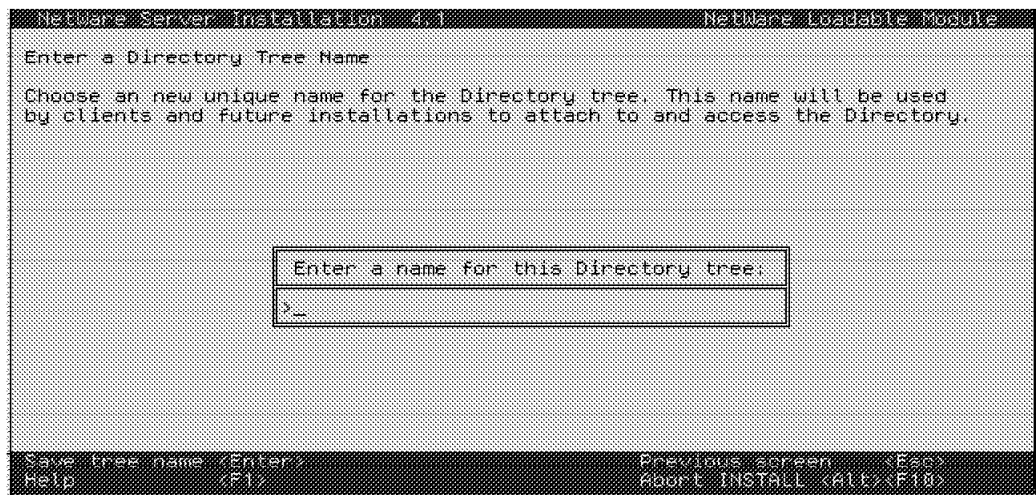


Figure 99. Assigning a Name to the Directory Tree

36. Enter a name for this directory tree and press Enter.
37. Select a time zone and press Enter.
38. Verify the time configuration and press F10 to save.

A screen like the one in Figure 100 on page 141 will appear. This is where the context for the server is defined.

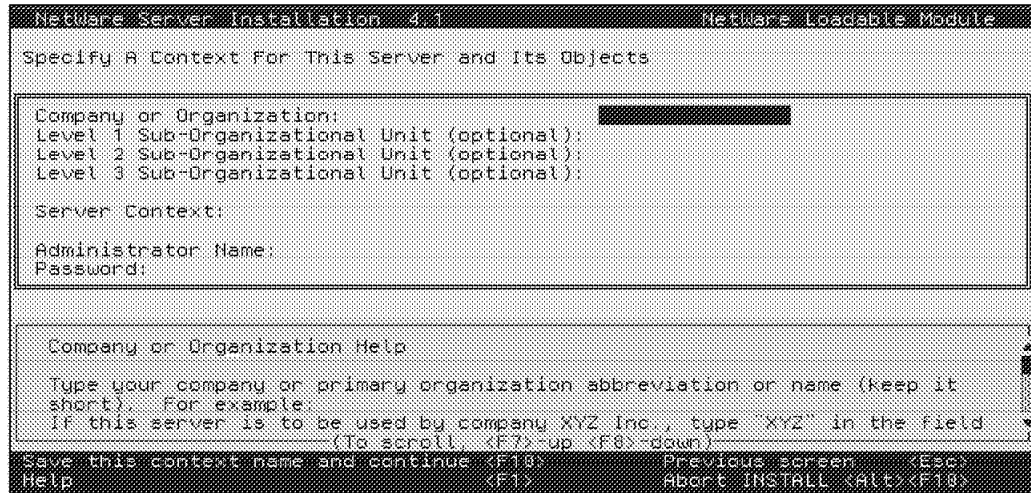


Figure 100. Context for the Server

39. Specify the server context and type an administrator password.
40. Retype the administrator password.
41. Select **Yes** when the system asks whether you want to save the directory information.

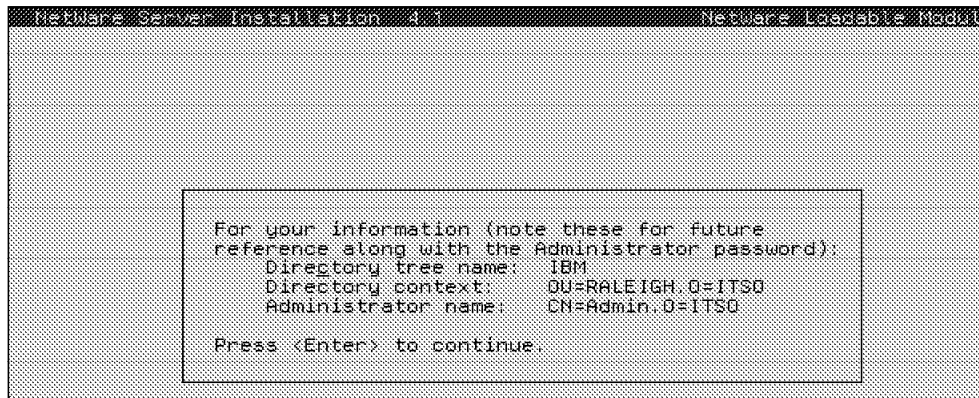


Figure 101. Server Context Information

Attention

Be sure to remember this information. You will need the context of the server to log on as well as to perform many administrator and end-user tasks. Please also remember that the context is *case-sensitive*.

42. Press Enter.

A screen like the one in Figure 102 on page 142 will appear. As the installation progressed, NetWare kept track of the options you chose for your disk driver. It automatically places this command in the STARTUP.NCF file for you.

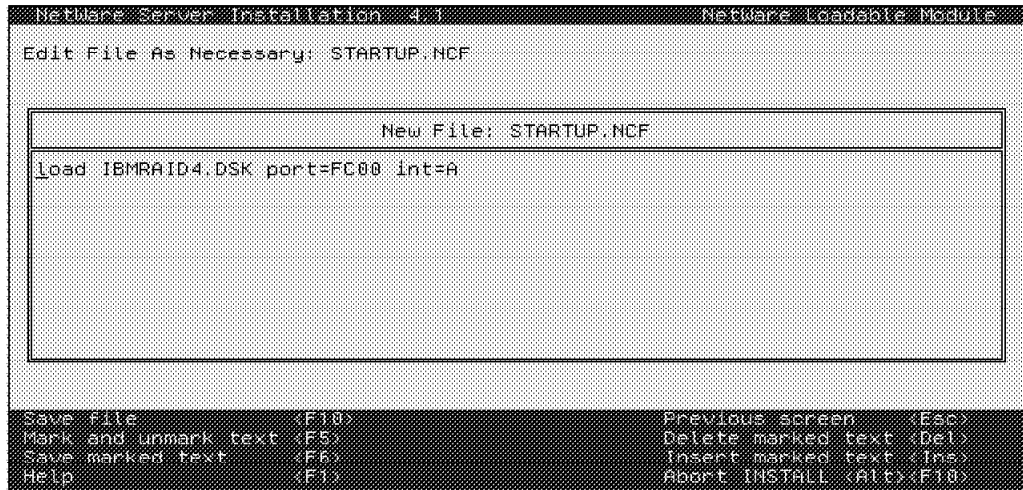


Figure 102. Editing STARTUP.NCF File

43. Press F10 and select **Yes** to save the STARTUP.NCF file.

A screen like the one in Figure 103 will appear. Many of the configuration parameters you have previously entered have been placed in the AUTOEXEC.NCF file for you.

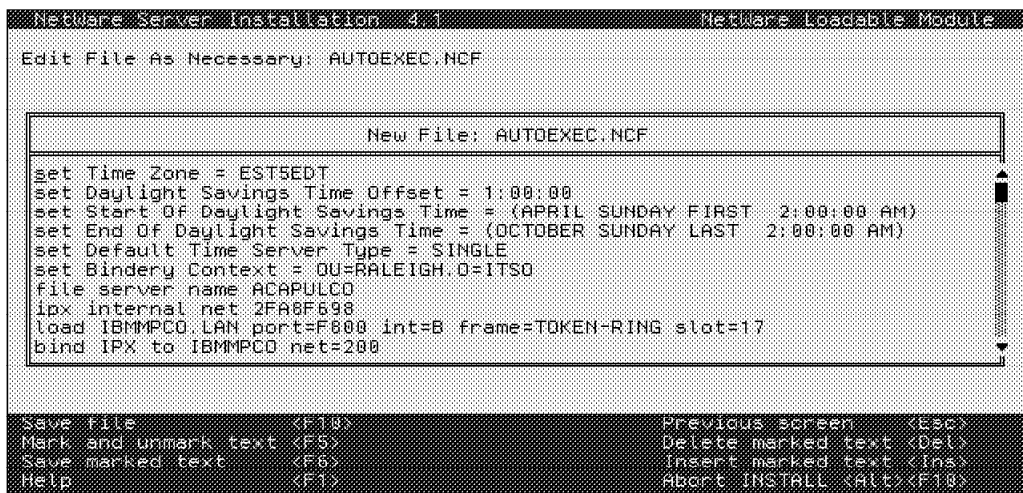


Figure 103. Editing AUTOEXEC.NCF File

44. Press F10 and select **Yes** to save the AUTOEXEC.BAT file.

The files are transferred from the CD to hard disk. You will see a screen similar to Figure 104 on page 143 which gives you progress of the installation.

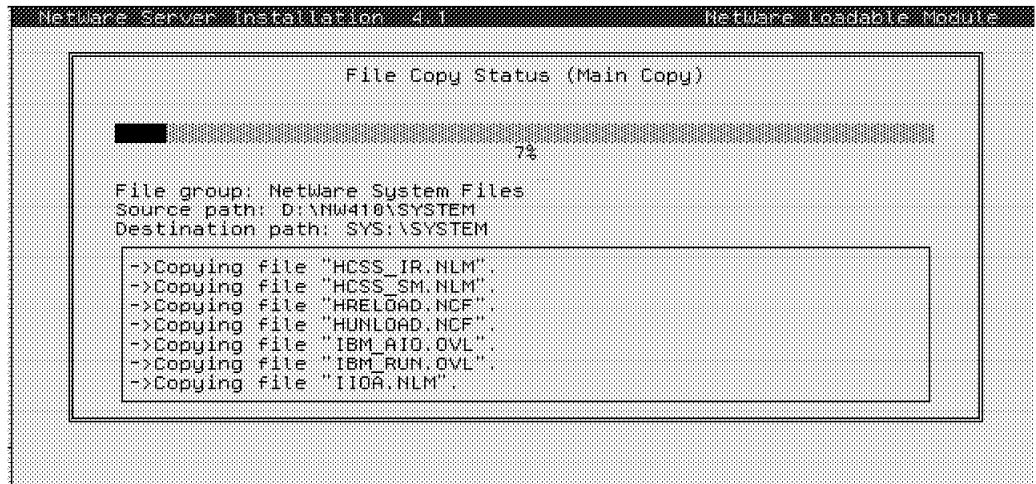


Figure 104. File Copy Status

45. When all files have been transferred, a screen like the one in Figure 105 will appear.




Figure 105. Other Installation Options

46. Select **Continue Installation** and press Enter.
47. Press Enter to exit to the system console screen and finish the installation.
48. Type DOWN at the colon prompt.
49. Type EXIT at the colon prompt to return to DOS.
50. Reboot the server and verify that it comes up properly.

The installation is complete. You are now ready to log on as the administrator and add users and applications.

4.6 NetFinity Services for NetWare

NetFinity Services for NetWare allows you to monitor and manage NetWare servers remotely. With the NetFinity programs running on the server, the administrator can anticipate and correct many problems before they become serious enough to affect the server. NetFinity Services for NetWare will also enhance other hardware systems management tasks, including hardware inventory and resource allocation.

NetFinity Services for NetWare does not include a user interface for use at the server console. Because NetWare is designed to be most effective when used in a client/server environment, monitoring and maintenance of the NetWare server must be performed remotely from a NetFinity Manager system.

4.6.1 System Requirements

The minimum system requirements for the server in which the NetFinity Services for NetWare will be installed are:

- Novell NetWare 3.11 or later
- 4 MB hard disk space
- A LAN adapter card
- IPX communications protocol
- CLIB.NLM dated 5/19/94 or later

4.6.2 Installing NetFinity Services for NetWare

NetFinity Services for NetWare can be installed in the following three ways:

- Using ServerGuide CD (see 4.3, "Installing NetWare 4.1 with ServerGuide" on page 118)
- Using diskettes on the NetWare Server
- Using diskettes from a workstation

You can make the diskettes from the ServerGuide CD using the diskette factory feature.

4.6.2.1 Installing with Diskettes on the NetWare Server

To install NetFinity Services for NetWare on your NetWare server:

1. Turn on the server.
2. Start the NetWare operating system.
3. Insert NetFinity Services for NetWare Diskette #1 into the diskette drive.
4. Start the NetFinity Installation program.
Type `LOAD A: NETFINST` and then press Enter.
5. Specify a new target path for the program files or accept the default path.

The default path is `SYS:NETFIN`. Press Enter.

Note

If the new path includes a directory that does not exist, the installation program will ask you whether you want this directory to be created. Type `Y` and press Enter to confirm this and continue the installation.

The files are transferred from the diskette to hard disk.

6. The installation program will prompt you to insert NetFinity Services for NetWare, Diskette #2.
7. Configure the Network Driver.

After the program files have been copied onto your NetWare server, the following screen appears. Please note that in this example, the name of the NetWare server is PCSERVER.

```
NetFinity Network Driver Configuration

1: System Name : PCSERVER
2: Network Driver : Novell IPX
3: Driver Enabled : TRUE
4: System Keywords :
  <not used>
  <not used>
  <not used>
  <not used>
  <not used>
  <not used>
  <not used>
  <not used>
  <not used>
  <not used>
5: Network Time-out: 15
6: Options
7: Save
8: Reset
9: Exit

Select an Entry (1-9) ->
```

Figure 106. NetFinity Network Driver Configuration

8. Enter a System Name.
 - a. Type 1 and then press Enter.

The following message appears on the display:
Type a name for your server->
 - b. Type a name for your server, and then press Enter.
9. Select a Network Driver.
 - a. Type 2 and then press Enter.

A list of network communications drivers supported by your server appears, followed by this message:
Select a driver number->
 - b. Type the number of a network driver and then press Enter.
10. Enable the Network Driver.
 - a. Type 3 and then press Enter.

The following message appears:
Set the driver status(1=False, 2=True)->
 - b. Type 2 and then press Enter to enable the driver.

11. Set System Keywords.

System keywords can help the NetFinity Manager system organize the systems on the network more effectively. Each system can have up to eight keywords.

- a. Type 4 and then press Enter.

A numbered list of keywords and the following message appears:

Choose keyword number to edit->

- b. Type the number of the system keyword you want to use and press Enter.

The following message appears:

Type new keyword (1-16 characters) or type DELETE to remove->

- c. Type the new system keyword and then press Enter.

12. Set the Network Time-out value (optional).

The network time-out is the number of seconds that NetFinity will attempt to communicate with a remote system that is not responding. If NetFinity does not establish contact with the remote system within this time, it cancels the communication attempt.

To change the Network Time-out setting:

- a. Type 5 and then press Enter.

The following text appears:

Type time-out value (5-500 in 5 second intervals)->

- b. Type the new Network Time-out value and then press Enter. If you do not change this value the system will use the default setting. The Network Time-out default setting is 15 seconds.

13. Enable Options.

This allows you to enable/disable NetFinity features such as the service execution alerts option. If you enable the service execution alerts option, NetFinity will generate an alert whenever one of your NetFinity services is started by a remote user that is accessing the server. The alert includes the name of the service that was run and information about the user that started the service.

To enable this NetFinity option:

- a. Type 6 and then press Enter.

The following message appears:

Service Execution Alerts: False(1=False, 2=True)->

- b. Type the number of the option you want to enable and press Enter.

14. Save the configuration.

Type 7 and then press Enter.

15. Exit the Network Driver configuration program.

Type 9 and then press Enter.

16. Add statements to the AUTOEXEC.NCF file.

When you have finished configuring the network communication driver, the following message appears:

The following lines need to be added to AUTOEXEC.NCF:

```
SEARCH ADD SYS:NETFIN\  
LOAD NETFBASE
```

If you choose to update AUTOEXEC.NCF, AUTOEXEC.NCF will be backed up to AUTOEXEC.RPS. Otherwise, additions will be saved to the file AUTOEXEC.NEW in the path SYS:NETFIN\
Update AUTOEXEC.NCF(Enter Y or N)?

SYS:NETFIN is the default target path or the new path you entered during the installation procedure.

17. Type **Y** if you want NetFinity Services for NetWare to start automatically when the server is started. Press **Enter**.
18. Press any key to close the screen.
The NetFinity Services for NetWare installation is now complete.
19. Type **DOWN** at the colon prompt.
20. Type **EXIT** at the colon prompt to return to DOS.
21. Reboot the server to start NetFinity Services for NetWare.

4.6.2.2 Installing with Diskettes from a Workstation

To install NetFinity Services for NetWare from a workstation, the workstation must have NetFinity Manager for OS/2, NetFinity Services for OS/2, or NetFinity Services for Windows installed.

To install NetFinity Services for NetWare from a workstation:

1. Log on to the NetWare server from the workstation.
2. Map a drive letter to a directory on the server where NetFinity Services for NetWare will be installed.
3. Insert NetFinity Services for NetWare Diskette #1 into diskette drive.
4. Start the NetFinity installation program.

If the workstation is running OS/2:

- a. Start an OS/2 full screen or OS/2 windows session.
- b. Type **A:INSTALL** and then press **Enter**.

If the workstation is running windows:

- a. Select **File** from the program manager menu bar.
- b. Select **Run...** from the File pull-down menu.
- c. Type **A:NETFINS** in the command line field and then select **OK**.

A window like the one in Figure 107 will appear.




Figure 107. NetFinity Installation

5. Enter the source drive and directory name where the NetFinity program files are located or accept A: as the default.
6. Enter the drive and directory name where the NetFinity program files will be copied. The default is C:\NETFIN. This should be the drive letter that you mapped to earlier.
7. Select **OK**.

If the path includes a directory that does not exist, the installation program will tell you that this directory will be created. Select **OK** to continue the installation.

A window like the one in Figure 108 will appear.




Figure 108. NetFinity Services for NetWare

8. Select **Install**.

The NetFinity installation program copies the program files to the NetWare server. A window like the one in Figure 109 will appear. This displays the name of the file currently being copied as well as the percentage of the installation that is completed.




Figure 109. NetFinity Installation - Copying Files

The installation program will prompt you to insert NetFinity Services for NetWare Diskette #2.

9. Insert NetFinity Services for NetWare Diskette #2 and select **OK**.

After the program files have been copied onto your NetWare server, a window like the one in Figure 110 will appear.




Figure 110. Network Driver Configuration

10. Enter a name for the server in the System Name field. This name identifies the server on the network.
11. Select **Novell IPX** from the Network Drivers field.

Attention

If the workstation supports additional communications protocols, the names of these protocols will appear in the Network Drivers field. Do not enable any protocols other than IPX because IPX is the only supported protocol by NetFinity Services for NetWare.

12. Select the **Driver Enabled** check box to activate the driver on startup.
13. Enter the System Keywords.

System keywords are not necessary for NetFinity services functions, but they can be used to organize the network by allowing the NetFinity Manager to discover machines using the keywords as filters.

14. Set the Network Time-Out value.

The network time-out is the number of seconds that NetFinity will attempt to communicate with a remote system that is not responding. If NetFinity does not establish contact with the remote system within this time, it cancels the communication attempt.

If you do not change this value, the system will use the default setting which is 15 seconds.

15. Select NetFinity Options (optional).

This allows you to enable/disable NetFinity features, such as the service execution alerts option. If you enable the service execution alerts option,

NetFinity will generate an alert whenever one of your NetFinity services is started by a remote user that is accessing the server. The alert includes the name of the service that was run and information about the user that started the service.

To activate this option:

- a. Select **Options...**
- b. Select **Service Execution Alerts**.
- c. Select **OK**.

16. Select **Save** to save the configuration.

A screen like the one in Figure 111 will appear.




Figure 111. Configuration Update

17. Select **OK** to return to the Network Driver Configuration window.

18. Select **Exit**.

A screen like the one in Figure 112 will appear.



Figure 112. NetFinity Installation Complete

19. Select **OK** to complete the NetFinity Services for NetWare installation.

4.6.2.3 Starting NetFinity

You can configure the server so that NetFinity Services for NetWare starts automatically whenever the server is started, or you can start the program manually.

- If you want NetFinity Services for NetWare to start automatically, you must add the following statements to the AUTOEXEC.NCF file:

```
SEARCH ADD SYS:NETFIN  
LOAD NETFBASE
```

SYS:NETFIN is the default target path or the path for the program files that you specified during installation.

After you save these changes, the NetFinity base program will start automatically when the server is started.

- If you do not want the NetFinity base programs to start automatically, you can start them manually. To start NetFinity manually:
 1. Turn on the server.
 2. When the command prompt appears type `SEARCH ADD SYS:NETFIN\` and then press Enter.
 3. Type `LOAD NETFBASE`, and then press Enter.

`SYS:NETFIN` is the default target path or the path for the program files that you specified during installation.

4.7 The RAID Administration for NetWare Utility

The RAID Administration for NetWare Utility is a set of NetWare loadable modules which aids in managing RAID devices. It provides a graphical screen of the arrays and allows you to obtain information about arrays which have been previously defined. It also monitors the arrays and sends you alerts if a hard disk is experiencing errors or is damaged. This utility also allows you to replace a defunct disk.

The RAID Administration for NetWare Utility can manage up to four IBM RAID adapters.

4.7.1 Installing the Utility

The Raid Administration for NetWare Utility can be installed in the following three ways:

- Using ServerGuide
- Using the RAID adapter option diskette from a workstation
- Copying NLMs to the server

If you use ServerGuide to install the NetWare operating system, the RAID Administration for NetWare Utility is loaded automatically during the installation process.

You can make the RAID adapter option diskettes from the ServerGuide CD using the diskette factory feature.

4.7.1.1 Installing from a Workstation

To install RAID Administration for NetWare from a workstation:

1. Log on to the NetWare server from a workstation.
2. Map a drive letter on the server to `SYS:SYSTEM`.
3. Insert the RAID Adapter Options Diskette into the diskette drive.
4. Change to the diskette drive prompt.
5. Type `INSNWADM d:` and press Enter.

Where `d` is the drive mapped to `SYS:SYSTEM`. The NetWare Loadable Modules will be copied to this drive.

`INSNWADM` will copy the following files to this mapped drive:

- `RAIDASPI.NLM`
- `RAIDADM.NLM`

6. Type `LOAD RAIDADM` at the NetWare server console prompt and press Enter to start the RAID Administration for NetWare Utility.

Note

Loading RAIDADM.NLM will automatically load RAIDASPI.NLM.

A screen like the one in Figure 113 will appear.

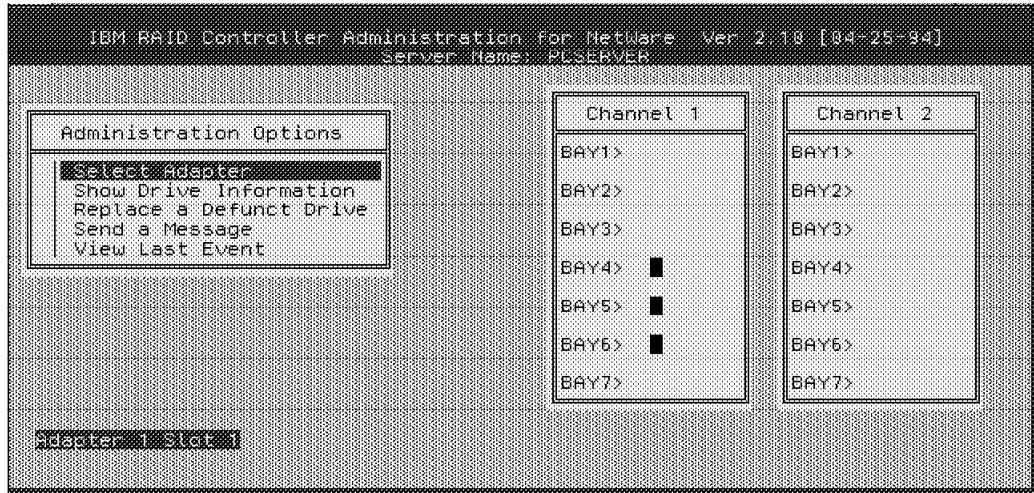


Figure 113. Raid Administration for NetWare - Main Menu

4.7.1.2 Copying NLMs to the NetWare Server

To install RAID Administration for NetWare:

1. Turn on the server.
2. Insert the RAID Adapter Options diskette into the diskette drive.
3. Change to the diskette drive prompt.
4. Type `CD NW311` to change to the NW311 directory.
5. Copy RAIDASPI.NLM and RAIDADM.NLM to the NWSERVER directory. NWSERVER is the directory where boot files for the NetWare operating system were installed.
6. Start the NetWare operating system.
7. Type `LOAD C:NWSERVERRAIDASPI` at the NetWare server console and press Enter.
8. Type `LOAD C:NWSERVERRAIDADM` at the NetWare server console to start the RAID administration for NetWare Utility.

A screen like the one in Figure 113 will appear.

4.8 Hard Disk Failure Simulation

We strongly suggest that you practice the rebuild process before it actually becomes necessary to use it, so that you become familiar with the process. The procedures below allow you to simulate a failure and take you step-by-step through the process of the rebuild. It is assumed that an array has already been

configured. It is also assumed that the RAID administration in the server and the NetFinity Manager in a workstation are running.

4.8.1 Simulating with a Hot Spare Drive

1. Start the RAID Administration for NetWare Utility.

A screen like the one in Figure 114 will appear.

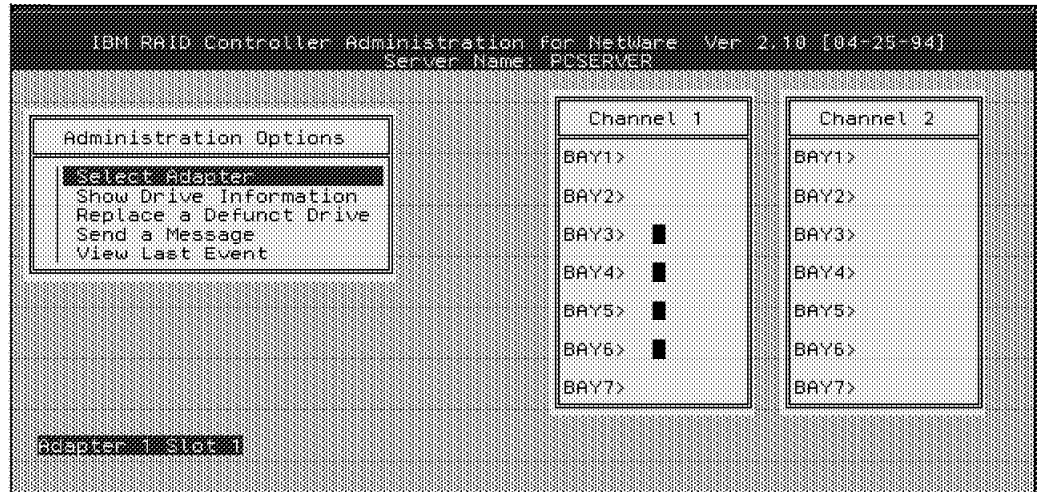


Figure 114. RAID Administration Utility - Main Menu

Notes

If any of the disks appears to be a slightly different shade of green than the rest, this indicates that some errors have been encountered by the RAID Controller with this device but the retries have been successful.

To get information about these errors, you can use the NetFinity Services Manager. To do this:

1. Start the RAID Manager from the NetFinity Services Manager folder.
2. Select the drive.
3. Click on the drive with mouse button 2.
4. Select Device statistics.

2. Select **Show Drive Information** from the Administration Options menu to verify the array configuration.

A screen like the one in Figure 115 on page 154 will appear.

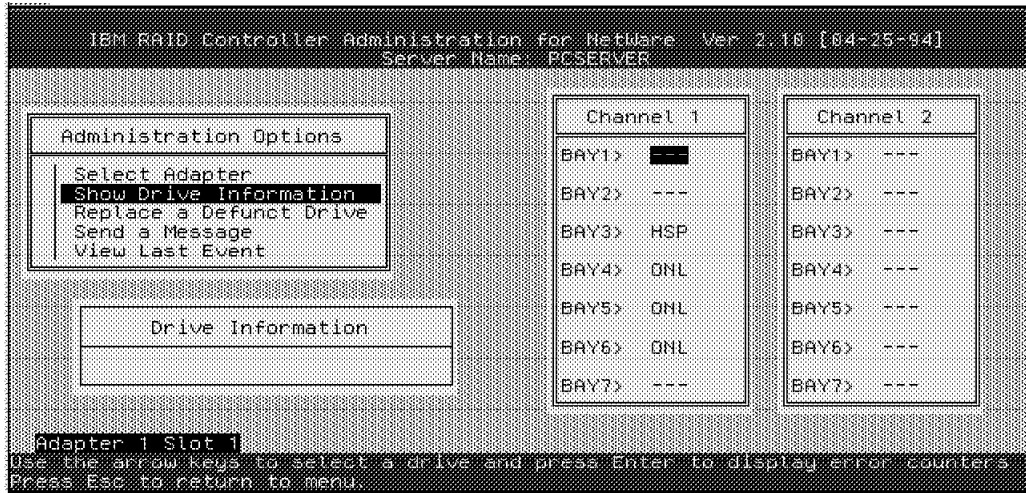


Figure 115. Verifying Array Configuration

The array configuration can also be verified from a workstation with NetFinity Manager installed using RAID manager (see Figure 116).

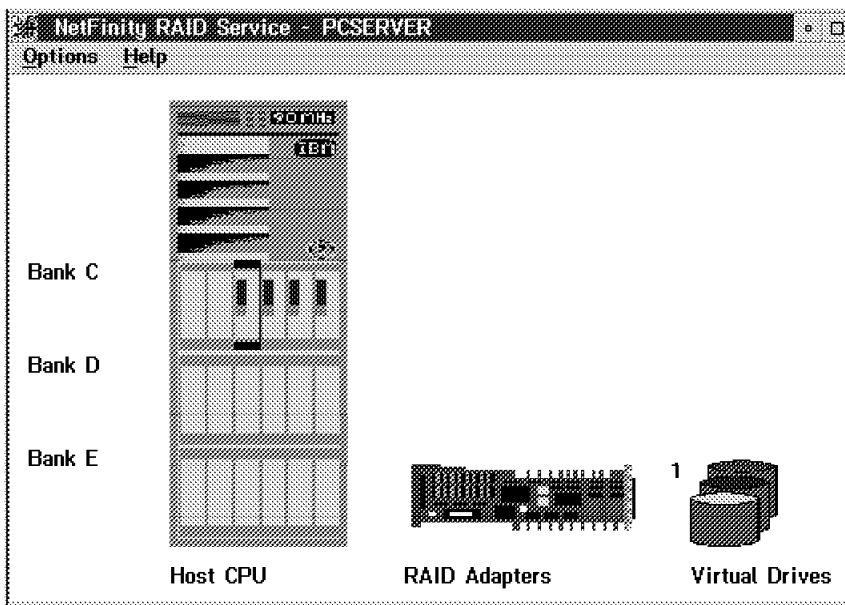


Figure 116. Using RAID Manager to View Array Configuration

The hot spare drive is enclosed in a box.

3. Simulate a disk failure by unplugging it with the machine powered on.

In this example, the disk in bay 5 was unplugged. When you do this, you will see an asterisk next to bay 5 and a flashing asterisk in the View Last Event option. This is depicted in Figure 117 on page 155.




Figure 117. Detecting the Disk Failure

If NetFinity has been configured properly, it will generate an alert. Figure 118 shows the NetFinity Alert Manager Log where such an alert has been generated.




Figure 118. Disk Failure - NetFinity Alert

In this example, two alerts were received. The first alert refers to the logical drive which was defined using the failed disk. This has been logged as a critical alert. The second alert relates to the physical drive which has failed. It shows up as a dead device on the alert.

This failure can also be checked from a workstation with NetFinity Manager installed using the RAID Service feature. This is shown in Figure 119 on page 156.

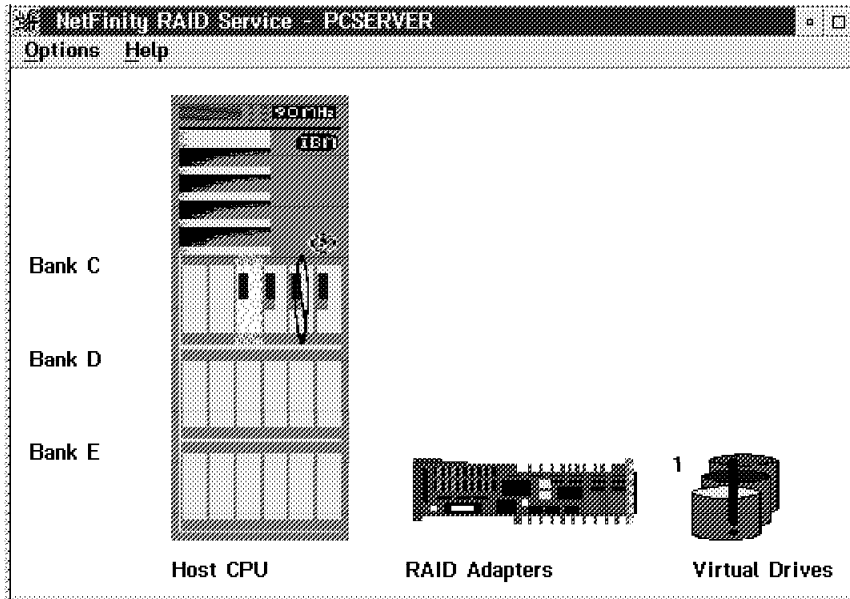


Figure 119. Disk Failure - NetFinity RAID Service

4. Select the **View Last Event** option from the Administration Options menu.
A screen like the one in Figure 120 will appear.




Figure 120. View Last Event Message

5. At the end of the rebuild process, the RAID Administration for NetWare Utility sends a message. This message is shown in Figure 121 on page 157.




Figure 121. RAID Administration - Recovery Message

NetFinity also sends some alerts. Figure 122 shows the NetFinity Alert Manager Log for these alerts.




Figure 122. NetFinity Recovery Alert

Again, two alerts are received. The first refers to the logical drive which is now online and the second to the physical disk which is also now online.

6. Select **Show Drive Information** to review the changes in the array configuration. A screen like the one in Figure 123 on page 158 will appear.




Figure 123. Changes in Array Configuration

The changes in the array are as follows:

- The original hot spare disk in bay 3 changed to status ONL (online).
- The faulty disk in bay 5 changed to status DDD (defunct).

We will now make the simulated failed disk the new hot spare disk.

7. Plug the unplugged disk back in.
8. Select **Replace a Defunct Drive** from the Administration Options menu.
9. Select **DDD** and press Enter.

A screen like the one in Figure 124 will appear.

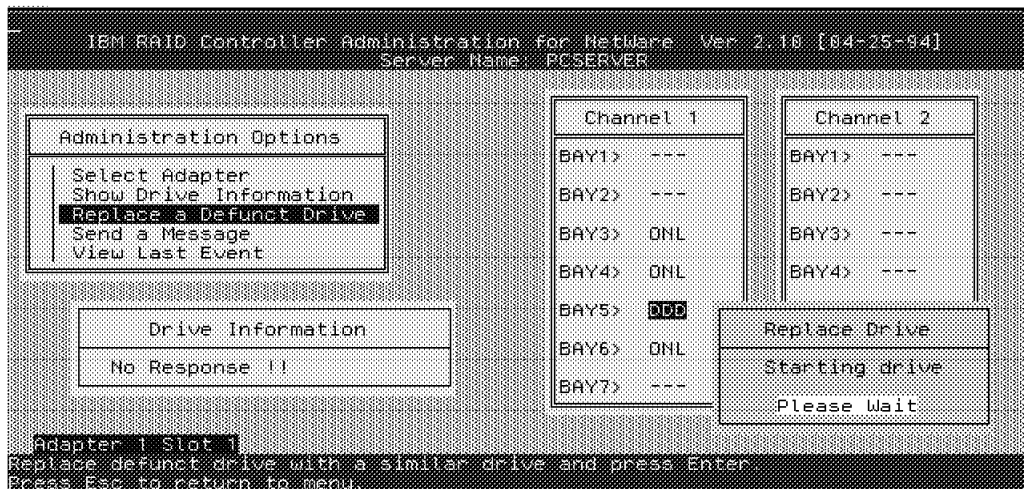


Figure 124. RAID Administration - Replace a Defunct Drive

10. Select **Show Drive Information** to verify whether the process was successful.

A screen like the one in Figure 125 on page 159 will appear.




Figure 125. RAID Administration - Verifying the Replacement of a Defunct Drive

The new array configuration is as follows:

- The disk status in bay 3 remains as ONL.
- The disk status in bay 5 changed to HSP.

11. NetFinity sends another alert. The NetFinity Alert Manager looks like Figure 126.




Figure 126. NetFinity New Hot Spare Drive Alert

This alert indicates that a new hot spare drive is available.

12. Verify the process using the NetFinity RAID service.

A screen like the one in Figure 127 on page 160 appears.




Figure 127. NetFinity RAID Service - New Hot Spare

You can see that a new disk is enclosed in a box. It is the new hot spare drive in the array configuration.

This simulation is complete.

4.8.2 Simulating without a Hot Spare Drive

To do the same test without a hot spare drive:

1. Start the RAID Administration for NetWare Utility.
2. Select **Show Drive Information** to review the array configuration.

A screen like the one in Figure 128 will appear.

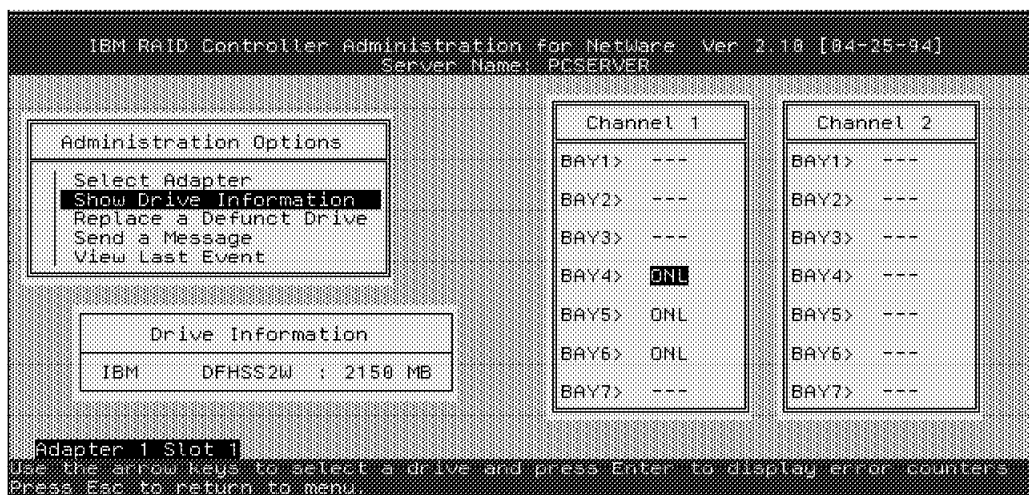


Figure 128. RAID Administration - Array Configuration

The array configuration can also be verified using the NetFinity RAID service (see Figure 129 on page 161).

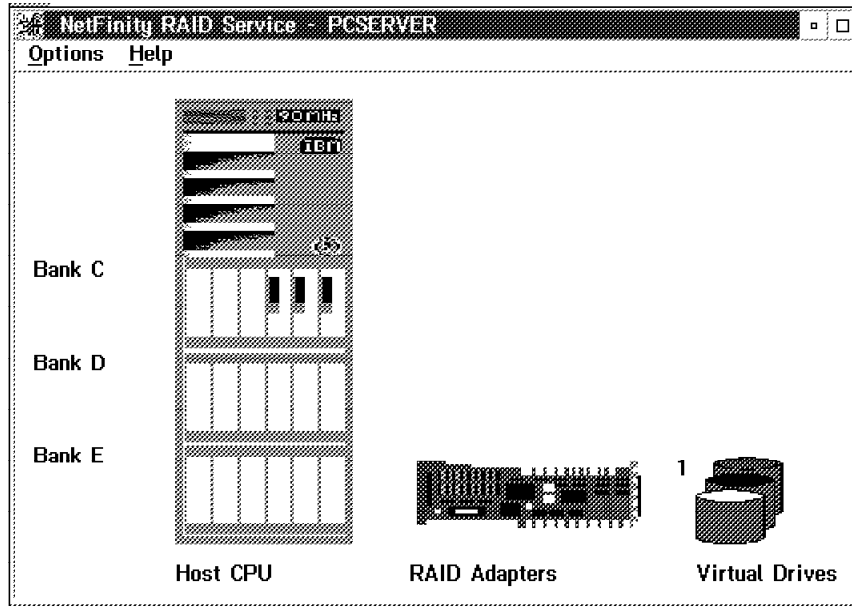


Figure 129. NetFinity RAID Service - Verifying Configuration

3. Simulate a disk failure by unplugging it with the machine powered on.

In this example the disk in bay 5 was unplugged. If you do this, you will see an asterisk next to bay 5 and a flashing asterisk in the View Last Event option. This is depicted in Figure 130.




Figure 130. Detecting the Disk Failure

4. NetFinity will send an alert.

The Alert Manager screen will look like Figure 131 on page 162.

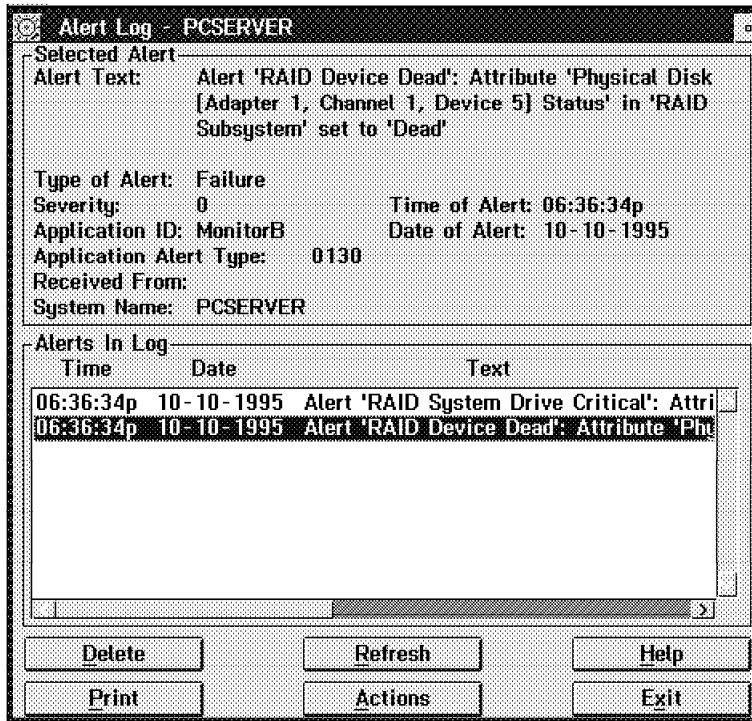


Figure 131. NetFinity Alert Log

In this example, two alerts were received. The first alert refers to the logical drive which was defined using the failed disk. This has been logged as a critical alert. The second alert relates to the physical drive which has failed. It shows up as a dead device on the alert.

5. The failure also can be checked using the NetFinity RAID service (see Figure 132).




Figure 132. NetFinity RAID Service - Disk Failure

6. Select **View Last Event** from the RAID Administration Utility to see the message.

A screen like the one in Figure 133 will appear.




Figure 133. Last Event Message

Since there is no hot spare drive defined in this example, the disk can not be automatically replaced.

7. Select **Show Drive Information** to review the disks' status in the array configuration.

A screen like the one in Figure 134 will appear.

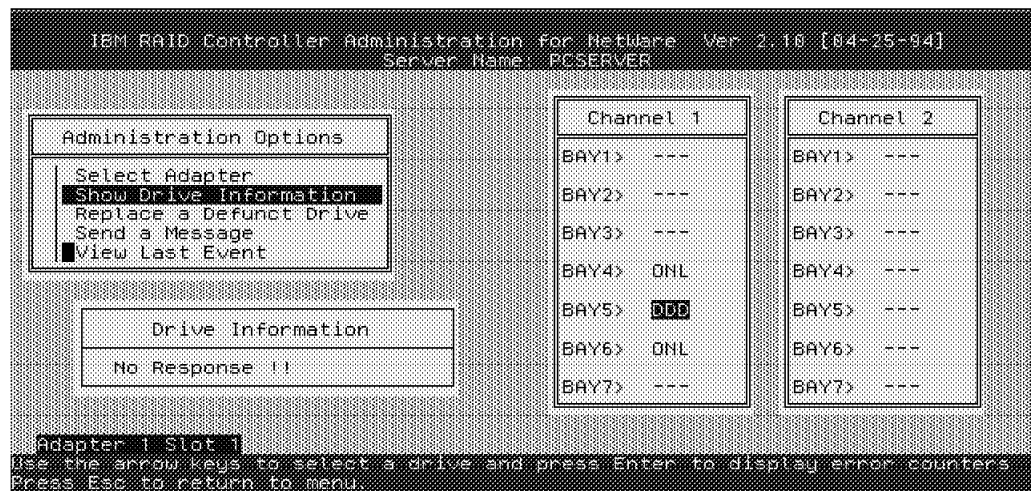


Figure 134. RAID Administration Utility - Reviewing Disk Status

The status of the faulty disk in bay 5 changed to DDD (defunct).

8. Plug a new disk in bay 5.
9. Select **Replace a Defunct Drive** from the Administration Options menu.
10. Select the **DDD** drive in bay 5 and press Enter.

A screen like the one in Figure 135 on page 164 will appear.

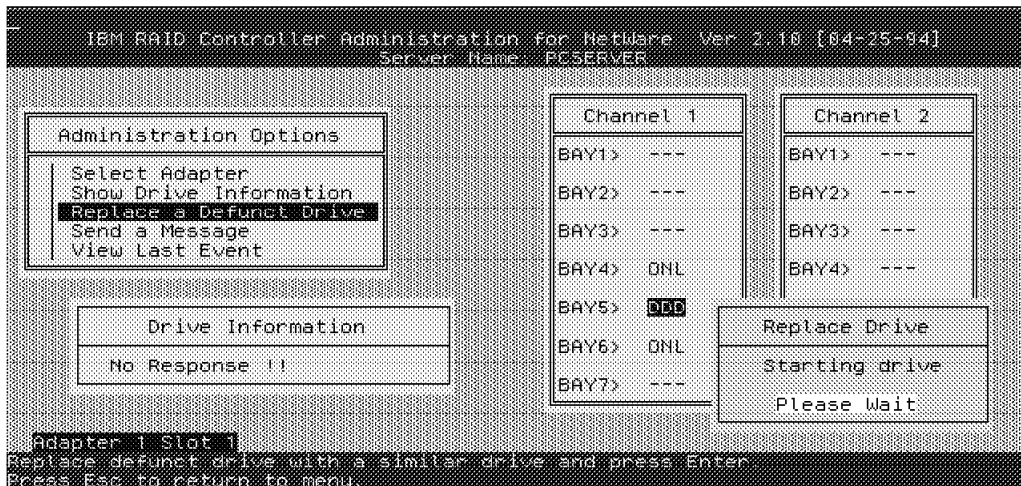


Figure 135. RAID Administration - Replace a Defunct Drive

After the drive is started, the RAID Administration for NetWare Utility attempts to rebuild the new disk. The progress of the rebuild is shown in Figure 136.

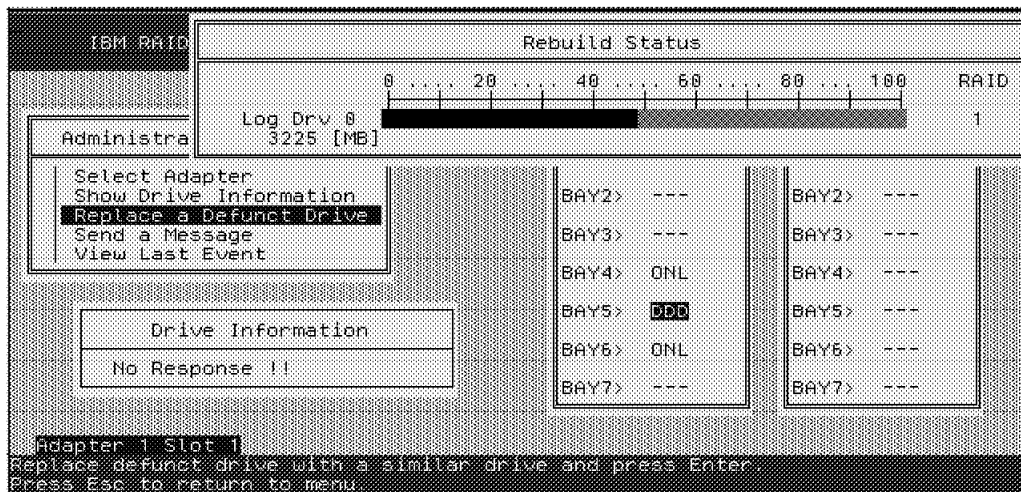


Figure 136. RAID Administration - Rebuild Progress

When the process is finished, a screen like the one in Figure 137 on page 165 will appear.

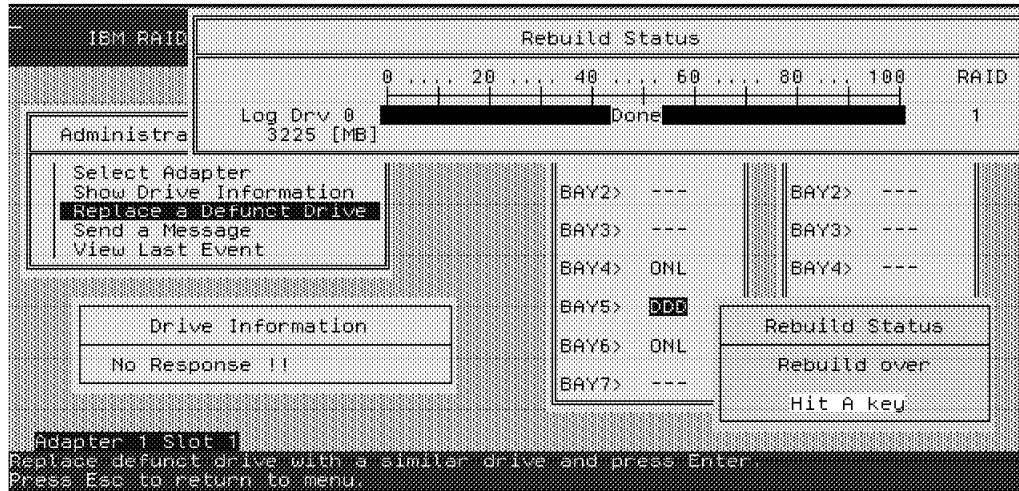


Figure 137. RAID Administration - Verifying the Rebuild Status

11. Press any key to complete the process.
12. NetFinity sends two alerts. The NetFinity Manager looks like Figure 138.




Figure 138. NetFinity Alert - New Disk Online

These alerts indicate that the new disk is online.

13. The simulation is complete.

Chapter 5. Performance Tuning

To get the optimal performance from your NetWare server, you need to consider the performance of both the hardware and the software. In this chapter, we examine the general performance characteristics of file server hardware. We also discuss several performance monitoring tools. We then discuss the NetWare operating system and its relevant performance parameters.

5.1 Hardware Tuning

A server is comprised of several key subsystems which must be balanced in order to optimize performance. Any one element which is undersized or over-utilized can create a bottleneck and affect the entire system performance.

The key subsystems are:

- Central Processing Unit (CPU)
- System memory
- I/O bus
- Disk subsystem
- LAN subsystem
- Controlling software

Figure 139 depicts this interaction among the subsystems in a server.




Figure 139. LAN Server Controlled Subsystems

Overall system performance is a function of each individual subsystem and can vary based on the application of the server. For example, file servers need fast

LAN adapters and fast disk subsystems. In contrast, database servers typically need fast processors (or multiple processors) and fast disk subsystems. Both file and database servers require large amounts of memory for operating system caching.

Tuning system performance involves determining where the current bottleneck is and then alleviating this bottleneck. For example, if the bottleneck is the processor, then a faster processor can be installed or alternatively, the CPU load can be reduced by using bus master LAN and disk adapters.

If the bottleneck is memory, then additional memory can be installed.

If the bottleneck is the disk subsystem, then either additional disks or disk adapters can be installed, or a specialized high performance disk subsystem such as an IBM Array Controller with IBM Fast SCSI-2 drives can be used. This allows an overlap of disk I/O requests.

If the bottleneck is the LAN adapter, then a faster LAN interface such as the IBM LANStreamer MC 32 adapter with 40 MBps data streaming support can be installed. Another optimization technique that can be employed is to utilize multiple LAN adapters in the server increasing throughput onto one or multiple segments.

The network operating system can create bottlenecks in its use of system resources. For example, inefficient use of buffers will waste the memory resources and can cause a memory bottleneck.

The performance of a subsystem is often inter-related to that of another subsystem. For example, memory constraints often cause excessive disk I/O as the operating system accesses virtual memory. High disk utilization can also occur if there is insufficient memory available for disk caching algorithms. Ineffective disk caching algorithms for a particular application workload can create memory constraints by wasting cache that is never accessed.

5.1.1 General Performance Characteristics

The disk subsystem is a major component of server performance and is frequently the location of the server bottleneck. Performance efficiency of this subsystem depends on whether the disk cache is sufficiently large and intelligent enough to contain the correct data required to satisfy most network requests from cache. In addition, the hard disks themselves must be fast enough to provide read data for the cache and empty the cache quickly when data must be written to disk. Thus, the disk subsystem must offer performance that can keep the cache filled with useful data. Failure to do so results in a lower disk cache hit rate, slower user access times and decreased performance.

The LAN adapter must move data between the network and system memory as fast as possible. For applications that issue large read and write requests, the width of the adapter data bus is an important factor. For example, applications such as multimedia, which require high bandwidth data can benefit greatly by using 32-bit adapters and data streaming technology.

On the other hand, when an application transfers small data frames, *latency* or the time it takes to move a frame through the LAN adapter becomes a critical factor. The IBM LANStreamer technology can be of great benefit in these applications due to its low latency design.

Systematic measurements of server performance employing adequate resolution typically produce a graph shaped like the one shown in Figure 140 on page 169. The characteristics that shape this graph are important for understanding potential bottlenecks in a file server.

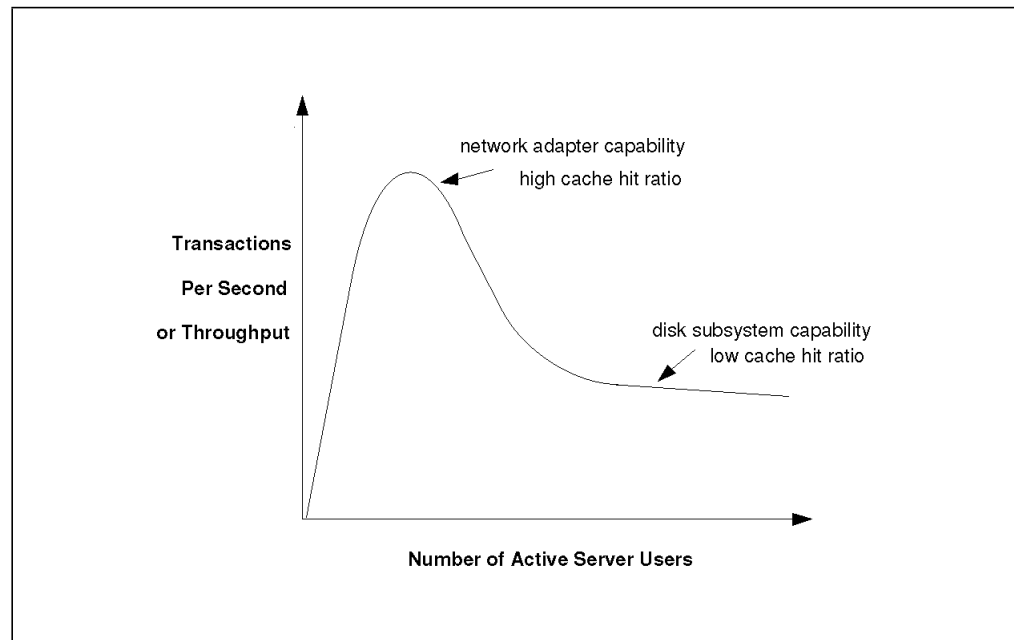


Figure 140. File Server Performance - General Characteristics

The horizontal axis shows the number of simultaneous active users. The vertical axis indicates server throughput, or number of transactions per second. This is the total number of transactions per second at the server, rather than the individual transaction rate for each workstation.

For the purpose of measuring throughput, a transaction can be an arbitrary unit of work. However, it is more meaningful to consider a transaction to be a typical user task. For example, a transaction can be defined as a networked client loading a spreadsheet, processing it at the workstation and then saving it back to the server. This process can be executed by a large number of workstations attached to a single file server to determine the maximum number of load-and-save transactions that can be accomplished per second.

Note

It is important to note that the number of transactions per second will be greatly dependent upon the transaction definition. Obviously, a one page document load-and-save transaction can complete much faster than a 100 page document load-and-save transaction. Thus, file server transaction rates are meaningless unless the specific type of transaction is clearly defined.

Figure 140 shows a graph of file server performance where throughput initially increases at a constant rate as users are added to the server. As the total number of users is increased, the network operating system is able to maintain a sufficiently high disk cache hit rate. Most user I/O requests are being serviced directly from the cache. The curve continues to climb until it reaches a peak,

which represents the maximum server transactions per second or throughput rate.

After the peak throughput is reached, the curve begins to slope downward. As the number of users is increased, the *caching engine* of the operating system begins to break down. The reduction in the disk cache hit rate is caused by the increasing amount of data that each additional user requests the server to access. The ratio of requested data to the size of the network operating system disk cache increases to a point where the network operating system disk caching is no longer effective. Furthermore, the disk cache is often storing data that was requested to be written to disk by application users. This write caching, when employed, causes write-data to occupy disk cache memory space until it is finally written to the disk surface. Thus, slow disk write operations can compound the performance degradation of servers that support write caching because the amount of useable disk cache memory is reduced by write-data waiting for service by the slower disk system.

The curve continues to break down until it reaches the transaction rate sustainable by the disk subsystem. In this state, few I/O requests can be serviced directly from disk cache and most require direct disk access. At this point, the server has reached saturation.

The initial slope of the curve is dependent upon how quickly transactions can be processed by the server, which, in turn, depends primarily on how quickly the LAN adapter is able to transfer data.

The peak of the curve is the maximum server transaction rate that the particular server configuration is capable of sustaining for a specific transaction type. The maximum transaction rate is primarily dependent upon performance of the network adapter and disk subsystem combination. When the graph flattens out, performance of the disk subsystem has the greatest influence on overall server performance.

We can gain a better understanding of how to tune server hardware by studying the curve dynamics as we modify various hardware components. Changing the LAN adapter or the disk subsystem can alter the height or width of the graph.

For example, in Figure 141 on page 171, adding a faster network adapter will increase the initial slope of the graph and also provide additional throughput. Therefore, the total number of transactions that can be processed per second will be improved. Furthermore, response time seen by any individual application user will be reduced. However, the reduction in response time is a limited benefit. As more users are added, the effects of the improved network adapter will be offset by increased disk I/O and increased write-data, which reduces the disk cache hit rate and overall server performance.




Figure 141. Differences in LAN Adapters

Upgrading the disk subsystem will tend to flatten out the top of the curve as it will provide a higher sustainable data transfer rate. In most cases, the disk subsystem becomes the bottleneck when a large number of users becomes active. Since most disk subsystems are significantly slower than a cache-hit operation, the throughput curve begins to decline. High performance disk subsystems, such as an IBM RAID Controller with Fast SCSI-2 drives, offer such a high level of performance that for many applications it allows the peak transaction rate to be sustained indefinitely. The general shape of the curve will be similar to the one shown in Figure 142 on page 172. Where the peak performance in transactions per second occurs, the line will continue horizontally rather than dipping as it would when the bottleneck is the disk subsystem. That is, for this workload the disk subsystem is no longer the system bottleneck and the peak transaction rate is sustained.




Figure 142. Differences in Disk Subsystems

Improving performance of the disk subsystem will usually prolong the maximum transactions per second rate shown in Figure 142. The effect of a faster disk can be to improve the disk cache hit rate. This is because write-data can be flushed from the cache quicker, so that the amount of memory available for caching of new data is increased.

5.2 Performance Analysis Tools

Using performance analysis tools is an effective method of understanding server performance.

5.2.1 DatagLANce

IBM DatagLANce is one such tool that can give valuable insight into server performance. DatagLANce operates at the physical layer and monitors the actual frames being transmitted on the network.

You can use the DatagLANce analyzer to get an accurate picture of the current activity on your network or a historical record of network activity over a specified period of time. You can design your own screens and save them from the menu. If you wish, you can create 32 different bar charts of real-time statistics, launch an analysis session and call up frame summary views, protocol interpreted views, and hexadecimal views (all color coded, highlighted, and tracked simultaneously). You can rearrange the statistics, add to them, and save them under a new name. You can use the DatagLANce analyzer's alarms to let you know when certain statistical thresholds, such as network utilization, are reached. Since DatagLANce runs on top of OS/2, you can even use your machine for other applications while the analyzer is monitoring your network.

The following is a summary of DatagLANce features:

- 88 source and destination address pairs filtered in real time
- Real-time frame capture while monitoring
- Eight fast, super-powerful, programmable event detectors

- Flexible, user-definable interface
- Reliable, accurate information
- Continuous reports of top talkers, ring map (token-ring only), error conditions, statistics, and selected network data
- Broad data import and export support
- Extensive protocol decode coverage
- 10-millisecond time stamp (token-ring) or 32-millisecond time stamp (Ethernet)
- Optional 840-nanosecond, high-resolution time stamp
- Fully windowed, graphical, multitasking user interface

DatagLANce supports all or portions of the following protocol suites:

- FDDI Protocol Suite
- Token-Ring Protocol Suite
- Ethernet/802.3 Protocol Suite
- IBM Protocol Suite
- TCP/IP Protocol Suite
- SUN NFS Protocol Suite
- XNS Protocol Suite
- Novell NetWare Protocol Suite
- DECnet Protocol Suite
- AppleTalk Protocol Suite
- Banyan VINES Protocol Suite
- ISO Protocol Suite
- X.25 Protocol Suite

5.2.2 NetWare Monitoring Tools

NetWare provides several useful tools which allow you to monitor your existing network server, tune it and make full use of the available resources. The following are the three most important tools:

- MONITOR.NLM
- Set command
- SERVMAN.NLM

5.2.2.1 MONITOR.NLM

MONITOR.NLM is an effective administration tool, providing details about many aspects of the network file server. Figure 143 on page 174 shows a screen from the monitor utility.




Figure 143. MONITOR Utility

The monitor utility allows you to check many important aspects of file server performance. It allows you check resource allocations and utilization of these resources. For example, you can:

- List system modules loaded
- View memory statistics
- View tracked resources
- Check status of LAN drivers and statistics
- List active connections
- Clear a connection
- List open files
- List record locks

5.2.2.2 Set Command

Use the set command at the file server system console to verify and adjust parameters of the server. To verify server parameters using the set command:

1. Type SET at the server system console.
2. Press Enter.

At this point, a full-screen menu appears listing the general category of parameters to choose from. You can select from the following general categories of parameters:

- Communications
 - Memory
 - File caching
 - Directory caching
 - File system
 - Record locks
 - Transaction tracking
 - Miscellaneous
3. Select the category you wish to check.

4. You are given a choice of parameters to either view or set.

In the following sections, we look at many of these parameters and make recommendations for setting them.

Note

Many of the set parameters can be adjusted from the server console while the server is up and running. However, several parameters must be set in either the AUTOEXEC.NCF or STARTUP.NCF files. The information displayed using the set command will indicate which parameters must be set in this way.

5.2.2.3 SERVMAN.NLM

SERVMAN is a server system console utility, which can be used to view and configure system parameters. Figure 144 shows the SERVMAN utility.




Figure 144. SERVMAN Utility

5.3 Tuning NetWare

As was discussed in 5.1, "Hardware Tuning" on page 167, tuning a server can be a complicated task because it is not a single device, but is made up of several subsystems all working in unison. The subsystems include the following:

- Disk subsystem
- Network subsystem
- System memory
- System processor and I/O bus

Each of the above subsystems affects the performance of the other subsystems. For example, the lack of available system memory will have an adverse effect on the performance of the network communications and disk storage subsystems.

In this section, we look at NetWare performance parameters which relate to these subsystems and we will also give recommendations for setting them.

5.3.1 Disk Subsystem

The file server disk subsystem contains the files which are accessed by the LAN workstations. The efficiency at which access is performed is directly related to the file server performance. The following section describes some of the findings of the IBM PC Server performance laboratory on the major factors that affect hard disk subsystem performance.

5.3.1.1 Block Size and Cache Buffer Size (NetWare 3.12)

NetWare V3.12 allows you to set the cache buffer size in system memory and the block size on the physical hard disk. Matching these values can help balance the workload of the CPU and the DASD subsystem, thereby yielding an overall higher throughput. It also improves file read/write retrieval time for a user on a LAN workstation.

Block Size: The volume block size can be set when the volume is created using the INSTALL.NLM facility.

Default = 4KB

Range = 4KB, 8KB, 16KB, 32KB and 64KB (NetWare 3.12)

Cache Buffer Size: The cache buffer size can be set in the STARTUP.NCF file using the following format:

```
SET CACHE BUFFER SIZE=VALUE
```

Default value = 4KB

Range = 4KB - 16KB (NetWare 3.12)

Recommendation: Set the block size to 8192 (8KB). The cache buffer size should be set to 8192 bytes also. These two values are closely related and complement each other with respect to performance. Matching their size helps to match buffer-to-buffer data transfers between disk and disk controller and NetWare file cache in system memory.

In some cases, the adjusting of the block size parameter to match the cache buffer size might not be as important as other considerations for the setting of the block size. For example, a 64KB volume block size might be needed for installing a multimedia application. This would be more important to users' overall performance than would the matching of the block size to the cache buffer size.

In the absence of these kind of requirements, use the values in the following table to relate volume block size to cache buffer size.

Volume Block Size in KB	Cache Buffer Size in bytes
4	4096
8	8192
16	16384
32	16384
64	16384

5.3.1.2 Block Size (NetWare 4.1)

In all versions of NetWare prior to NetWare 3.X, the block allocation unit was fixed at 4KB. In NetWare 3.X, larger allocation units are possible, but they result in a lot of unused disk space at the end of files. For example, if the last part of a file took up 1KB of a 16KB block, the remaining 15KB was wasted.

Sub-allocation is implemented in NetWare 4.1 to overcome the problem of wasted disk space due to under-allocated disk blocks. Sub-allocation allows multiple file endings to share a disk block. The unit of allocation within a sub-allocated block is a single sector (512 bytes). This means that as many as 128 file ends can occupy one 64KB block. Using sub-allocation, the maximum loss of data space per file is 511 bytes. This would occur when a file has one more byte than can be allocated in a full 512-byte sector.

NetWare 4.1 defaults to a block size other than 4KB during the creation of volumes. It determines default block sizes based on volume size rather than on performance criteria. This is done in order to conserve server cache memory. The default sizes are shown in the following table.

Volume Size	Default
Less than 32MB	4KB
32 to 150MB	8KB
150 to 500MB	16KB
500 to 2000MB	32KB
2000MB and up	64KB

Recommendation: IBM recommends a 64KB block size for all NetWare 4.1 volumes. The larger 64KB allocation unit allows NetWare to use the disk channel more efficiently by reading and writing more data at once. This results in faster access to mass storage devices and improved response times for network users.

5.3.1.3 Turbo Fat Indexing

The file allocation table (FAT) is an index. The table entries point to the allocated disk blocks in which a file is located. When a file is opened for reading or writing, the FAT is read to find the actual location of the file on the disk. To improve performance, the entire FAT is cached in server memory, allowing the server to quickly access disk allocation blocks of the data volume.

When a file exceeds 64 blocks, NetWare creates a Turbo FAT index that groups all the FAT entries corresponding to the file. Subsequently, the file is accessed through a Turbo FAT. Upon completion, the Turbo FAT index is closed but not flushed from server memory. Instead, an aging process starts, and after a specified period, the Turbo FAT index is flushed from memory.

The time that the index is allowed to remain in main system memory can be varied using the set command in the AUTOEXEC.NCF file as follows:

```
SET TURBO FAT RE-USE WAIT TIME=VALUE
```

NetWare 3.12 and 4.1: Default = 5 minutes 29.6 seconds

Increasing the Turbo FAT re-use wait time helps ensure that the index structures are not flushed if the file is reused quickly. For example, a large database file might be accessed frequently throughout the day, with many file load and save operations. By increasing this parameter there is an increased likelihood of servicing the file request from a Turbo FAT index, thereby improving performance.

5.3.1.4 Directory Caching

Directory caching allows fast access to frequently used directories. When NetWare receives a read request from a workstation, it begins by searching the directory and FAT cache to determine the location of the file which will ultimately be either on disk or in cache. A directory cache buffer is a portion of the file server memory that holds entries from the directory table. A directory entry stays in a cache buffer as long as it's being accessed frequently.

The server will allocate a minimum number of directory cache buffers at boot time and then increase this value automatically as the demand grows for more buffers.

NetWare has parameters to set both the minimum and maximum number of buffers available for directory caching. These parameters can be increased to improve performance as directory service requests are more likely to be serviced from the cache. However, the following items should be considered:

- If the minimum value is set too low, this will cause an unnecessary delay in the allocation of additional directory cache buffers, thus slowing response time to directory searches.
- If the maximum value is set too high, this can consume too much memory and affect the file cache memory pool size. Therefore, the server must have enough memory to accommodate this increase.

There is also a parameter to affect the time delay before allocation takes place and this can be decreased to improve performance.

You can modify the values using set commands in the AUTOEXEC.BAT file as follows:

NetWare 3.12 and 4.1

```
SET MAXIMUM DIRECTORY CACHE BUFFERS=VALUE
```

Default = 500

Range = 20 to 4000

```
SET DIRECTORY CACHE ALLOCATION WAIT TIME=VALUE
```

Default = 2.2 seconds

Range = 0.5 to 120 seconds

5.3.1.5 File Caching

File caching has the most affect on overall file server performance. Because accessing data from cache is much faster than accessing data from disk, the more data than can be held in cache, the better the server performance. Hence, the number of file cache buffers available can greatly impact server performance.

Use the set command to change the default values as follows:

```
SET MINIMUM FILE CACHE BUFFERS = VALUE
```

Default = 20

```
SET MAXIMUM FILE CACHE BUFFERS = VALUE
```

Default = 1000

Recommendation: Increase the minimum value to 250 as a starting point and make sure your file server has enough memory to accommodate this value.

NetWare 4.1 provides a way to warn you if the number of cache buffers drops too low. The SET MINIMUM FILE CACHE REPORT THRESHOLD parameter allows you to set a threshold for the number of file cache buffers. The value of this parameter is added to the value of the MINIMUM FILE CACHE BUFFERS parameter and when the number of available cache buffers drops below this value, NetWare generates a warning. For example, if your MINIMUM FILE CACHE BUFFERS is 200 and you set the MINIMUM FILE CACHE REPORT THRESHOLD to 25, NetWare generates a warning when there are 225 buffers available for cache.

Original Cache Buffers (NetWare 3.12 and 4.1): The original cache buffers are the number of cache buffers allocated at the network server startup. You can view this value with the MONITOR.NLM utility.

By multiplying the original cache buffers by the cache buffer size, you can calculate the amount of system memory available to the cache upon server startup. For example, if the number of cache buffers is 3000 and the cache buffer size is 4096, then the cache occupies 12MB of system memory at the time of startup:

$$3000 \times 4096 = 12,288,000 \text{ bytes (12MB)}$$

Total Cache Buffer (NetWare 3.12 and 4.1): The total cache buffers are the number of cache buffers *currently available* to cache the disk I/O requests. They can be viewed in the MONITOR.NLM information screen.

As you load NLMs on the server, the total number of available cache buffers decreases. You can calculate the memory requirements of NLMs and set parameters by subtracting the total cache buffers from the original cache buffers.

Let us illustrate with an example. Let's say that the number of original cache buffers allocated at server startup was 3000. Let us also say that the value of the total cache buffers is 2500 after the file server is up and running and quiesced. By subtracting the total cache buffer value from the original cache buffer value, you can establish how many cache buffers have been allocated to the loading of NLMs and set parameters. Also, if you know that the cache buffer

size is equal to 4096, the amount of memory used in megabytes can be calculated as follows:

$$\begin{aligned} \text{system memory used} &= (\text{original cache buffers} - \text{total cache buffers}) \\ &\quad \times \text{cache buffer size} \\ (3000 - 2500) \times 4096 &= 2,048,000 \text{ bytes (2MB)} \end{aligned}$$

From this calculation, you can determine if the file server requires more memory, particularly when adding NLMs or changing the server configuration with set parameters.

Recommendation: Novell recommends that you maintain the total number of cache buffers at 70% of the total server memory. If you regularly monitor this parameter and notice that it is lower than 70%, then it is time to think about adding more memory.

5.3.1.6 Maximum Concurrent Disk Cache Writes

The maximum concurrent disk cache writes parameter can be changed to accommodate more write requests. This may be necessary if the file server is subject to a high number of disk writes (for example, when copying a large number of files to the file server). You can change this parameter in the AUTOEXEC.NCF file using the following command:

```
SET MAXIMUM CONCURRENT DISK CACHE WRITES=VALUE
```

NetWare 3.12: Default = 50

Range = 10 to 1000

NetWare 4.1: Default = 50

Range = 10 to 4000

Recommendation: By changing the default value to 100, the performance of the file server in write requests (copying files up to server) was increased by over 30% in this single type of operation. Furthermore, with the value set at 100 there were no other areas of the testing routine that suffered significantly from a degradation in performance and throughput did not suffer. It is worth mentioning, however, that not every conceivable combination of possible file server loads was tested.

In server environments with heavy read requests, this parameter should be adjusted down from this level. Although the performance of write operations will be degraded, it will leave more cache buffers available for read requests. In these environments, start with a value of 30 and work upward.

In summary, tune this value higher for write intensive file requests (start with 80), and lower for read intensive file server requests (start with 30).

To fine tune this parameter, use the MONITOR.NLM utility at the server console and observe the number of dirty cache buffers. If it is above 70% of the total cache buffers, then increase concurrent disk writes until the percentage of dirty cache buffers decreases.

5.3.1.7 Dirty Disk Cache Delay Time

The dirty cache delay time controls the time taken to hold the write requests in the cache when the cache is not full. If your system has unused cache memory, and your typical server workload contains a high number of sequential write operations, then you may improve performance by increasing this value. The dirty cache delay time value can be set in the AUTOEXEC.NCF file using the following format:

```
SET DIRTY DISK CACHE DELAY TIME=VALUE
```

NetWare 3.12 and 4.1: Default value = 3.3 seconds

Range = 0.1 to 10 seconds

5.3.1.8 Spanning Volumes Over Multiple Drives

When a volume spans more than one disk, NetWare automatically stripes data across the disks. When performing a disk read request, it will get the first piece of data from disk 1, the next one from disk 2 and so on, until it reads the amount equal to the volume block size. This yields similar performance gains to a disk array.

On write requests, when multiple disks are chained together, the disk controller can initiate a write to disk 1 and, without waiting for a completion code, service the write request to disk 2 before coming back to where it left off with the first write request to disk 1. This helps to optimize performance of the hard drive subsystem and can give a marked increase in performance.

While these techniques do provide considerable performance improvements, there is no inherent fault tolerance in this scheme. If one disk fails, then data will be lost, as the fault will destroy the entire NetWare volume.

As discussed previously in 1.6, "Disk Subsystem" on page 16, some models of the IBM PC Server line provide RAID support. This is the recommended way to gain this same kind of performance while ensuring an added level of fault tolerance.

5.3.1.9 Multiple Volumes

Where possible, divide the physical disk space into multiple volumes, even on disks as small as 400MB. Do not define one large SYS: volume, but instead try to allocate other volumes. For example, a volume called DATA could be used for user data files while a volume named APPS could be used for user application files. This will ease administration as well as make it easier to perform disk reorganizations when necessary. Also, try to plan your use of name space on volumes in advance as they will perform faster when installed in advance of loading volume data files.

5.3.1.10 Immediate Purge of Files

Setting Immediate Purge of Files to OFF allows the file system to maintain deleted files on the disk to be purged as necessary. This is very useful as these deleted files can be recovered using the salvage command. You can modify this parameter by using the set command in the AUTOEXEC.NCF file as follows:

```
SET IMMEDIATE PURGE OF FILES=VALUE
```

NetWare 3.12 and 4.1: Default=OFF

Attention

Set this parameter to ON only if you are absolutely sure your backup strategy enables you to restore a file which has been deleted.

5.3.1.11 Enable Disk Read After Write Verify

In NetWare, this read-after-write verify capability complements the hot fix function. This allows the system to prevent any file from being written to a bad disk block by attempting to read the data after it has been written to disk. NetWare's read-after-write capability and hot fix function are implemented in software.

This value is normally determined by the device driver and normally defaults to ON.

On all IBM disk controllers, this read-after-write operation is completed at the hardware level. Leaving this parameter ON will degrade disk subsystem performance.

You can modify this value using the set command in the AUTOEXEC.NCF file as follows:

```
SET ENABLE READ AFTER WRITE VERIFY=VALUE
```

NetWare 3.12 and 4.1: Default = ON

This value can be viewed and set by using the MONITOR.NLM utility.

5.3.2 Network Subsystem

Poor network design or network traffic congestion can have an adverse effect on the server and workstation performance. The following sections highlight some important factors to consider in order to optimize the performance of the network subsystem.

5.3.2.1 LAN Device Driver Efficiency

The device driver is responsible for communicating between the NetWare operating system and the LAN adapter hardware. It assembles and disassembles the IPX packets on the LAN as they come through the driver.

The NetWare device drivers produced by IBM and Novell are checked for performance before they are released. If they do not perform better (or at least as fast as the old driver), they are not released. Therefore, always use the latest device driver unless otherwise stated.

Some of the software LAN adapter device drivers are written with command line options which can be used to change the behavior of the device driver. For example, some third party manufacturers have a command line option to switch the early token release option on or off. This option is available on most 16 Mbps token-ring adapters.

5.3.2.2 Packet Size

Different network topologies will use different packet sizes. For example, a token-ring LAN adapter set to 16 Mbps can use packets of up to 4202 bytes. If the packet size is left at the network default value, which is 1514 bytes, this could unnecessarily slow down the LAN adapter performance. The reason for this is that it takes just as many instructions to process a small frame as it does a larger frame. Therefore, for a given amount of data to transfer, the larger frames will be transferred more quickly. You can adjust the packet size using set commands.

For **NetWare 3.12 and 4.1**, the syntax is:

```
SET MAXIMUM PHYSICAL RECEIVE PACKET SIZE = VALUE
```

Default = 1514

Range = 618 to 24682

Recommendation:

1. For token-ring 16 Mbps = 4202 packet size
2. For token-ring 4 Mbps = 2154 packet size
3. For Ethernet 10 Mbps = 1514 packet size

In the case of 4 Mbps token-ring adapters, a packet size of 4202 can be supported on some of the newer IBM LAN adapters.

Note

If you are in a multi-segment LAN or WAN environment, you must make sure that all inter-networking devices (bridges and routers for example) are capable of passing the packet size that you choose.

5.3.2.3 Packet Receive Buffers

Packet receive buffers hold communication packets that are waiting to be processed or transmitted. You can adjust the minimum and maximum number of buffers that NetWare generates by changing the following parameters:

1. MINIMUM PACKET RECEIVE BUFFERS (STARTUP.NCF)
2. MAXIMUM PACKET RECEIVE BUFFERS (STARTUP.NCF or AUTOEXEC.NCF)

The syntax is as follows:

```
SET MINIMUM PACKET RECEIVE BUFFERS=VALUE
```

```
SET MAXIMUM PACKET RECEIVE BUFFERS=VALUE
```

Minimum Packet Receive Buffers: NetWare 3.12

Default = 10

Range = 10 to 100

Recommendation: This value can vary depending on the number of LAN adapters and the number of users on the LAN. Allow at least 10 buffers for each LAN adapter card and one buffer per active LAN user. Active LAN users can be

monitored using MONITOR.NLM. Use this value rather than the number of defined LAN users.

For example:

1. For 40 LAN users with 2 LAN adapters in the file server:
 $40 \text{ (users)} + ((10 \times 2) \text{ LAN adapters}) = 60 \text{ buffers minimum; set the maximum to } 160.$
2. For 100 LAN users with 4 LAN adapters in the file server:
 $100 \text{ (users)} + ((10 \times 4) \text{ LAN adapters}) = 140 \text{ buffers minimum; set the maximum to } 240.$

Generally, if you have more memory available in the file server, setting the minimum packet receive buffers slightly higher can help during peaks of LAN activity.

Note

The IBM LANStreamer MC32 requires minimum packet receive buffers set to 200 or above. It may need to be higher depending on the number of LAN users.

Minimum Packet Receive Buffers: NetWare 4.1

Default = 10

Range = 10 to 1000

Recommendation

- For networks without wide area (WAN) segments, set the minimum packet receive buffers to 100.
- For networks that use WANs, set the minimum packet receive buffers to 500. This is because of the additional buffering sometimes needed for WAN links.

Maximum Packet Receive Buffers: NetWare 3.12

Default = 400

Range = 50 to 2000

Recommendation: For the maximum packet receive buffers, set this value at least 100 buffers above the minimum packet receive buffers value.

Some NetWare NLM applications, or LAN environments require you to increase the number of buffers beyond this point.

Maximum Packet Receive Buffers: NetWare 4.1

Default = 100

Range = 50 to 2000

Recommendation

- For networks without wide area (WAN) segments, set the maximum packet receive buffers to 300.
- For networks that use WANs, set the maximum packet receive buffers to 1000. This should be adequate to handle the additional buffering needed for WAN links.

5.3.2.4 New Packet Receive Buffer Wait Time

The new packet receive buffer wait time is the wait time taken before spawning a new packet receive buffer.

This parameter must be set in the AUTOEXEC.NCF file to change the value. The syntax is:

```
SET NEW PACKET RECEIVE BUFFER WAIT TIME=VALUE
```

For NetWare 3.12 and 4.1, the following is true:

- Default value = 0.1 second
- Range = 0.1-20 seconds

5.3.3 System Memory

The way in which NetWare manages system memory has a direct bearing on file server performance. Because of this, it is important to have an understanding of how the server allocates memory and the use of set parameters to control the system memory.

5.3.3.1 NetWare 3.12 Memory Pools

NetWare V3.12 is implemented as 32-bit code which means that it uses 32-bit addressing and the flat memory model. This gives it access to one flat segment of RAM of up to 4 GB in size. It loads the operating system, LAN and disk drivers and NLMs into one large flat memory space. When an application requests memory, the RAM is scanned for the first available block of memory large enough to accommodate the request.

NetWare V3.12 manages its RAM through the use of memory pools. These are designed to be allocated and optimized based on how they are used, long term or short term. Rather than pre-allocating a lot of memory at installation time, most resources are allocated on an as-needed basis. Any pre-assignment of memory is specified in the server's STARTUP.NCF and AUTOEXEC.NCF files.

The table below shows the different NetWare memory pools and how they are used.

Memory Pool Name	Data Stored
Cache Buffers Pool	Cache files, data and applications
Cache Movable Memory Pool (non-fragmented)	NetWare inter FAT tables
Cache Non-movable Memory Pool (fragmented)	CLIB and NLMs
Permanent Memory Pool	Directory cache buffers, LSLs and Turbo FATs.
Alloc Memory Pool	Menu screens, connection tables, print queues
Semi-Permanent Memory Pool	LAN and disk controller drivers

The size of the cache buffer pool is dynamic and determined by the amount of memory left over after the other pools have been allocated. As more buffers are allocated for NLMs and other system resources, this reduces the size of the cache buffers pool. Both the movable and non-movable pools return memory to the cache buffer pool when it is no longer needed.

Checking the Allocation of Memory to Each Memory Pool: You can use the MONITOR.NLM utility to view the allocation of system memory by pool. This helps you to understand how the server system memory is being used by the operating system to service the various requests. This is a dynamic process and will change as the network load on the server changes.

It is also important to monitor the pools as you are tuning other parameters. For example, when adjusting the number of packet receive buffers, the total cache buffers will be reduced as a percentage of the total system memory. Using the MONITOR.NLM utility can help you determine the effect of changing the number and size of the buffers. It may be necessary to add more memory to the server if the performance tuning results in the depletion of the total cache buffer memory available to efficiently service file requests.

Memory Fragmentation: As previously discussed, in NetWare 3.12, the cache buffer pool gives up memory to other memory pools such as the cache movable pool, the cache non-movable pool, and the permanent memory pool.

The cache movable and cache non-movable pools return system memory to the cache buffers pool, but anything allocated to the permanent memory pool does not return memory until you bring down the server.

The cache movable pool can be moved around so that contiguous blocks are always returned to the file cache buffer pool. This keeps them from becoming fragmented. This is not true of the cache non-movable memory pool which frees memory in non-contiguous chunks. This causes fragmentation and eventually degrades network performance.

If cache non-movable memory is allocated and de-allocated frequently due to the loading and unloading of NLMs, it is possible to end up with RAM which is so fragmented that there may not be enough free contiguous RAM to load additional NLMs. The file server will report that there is not enough memory to load the NLM and the server will have to be downed so the memory can reallocate and clear the fragmentation.

5.3.3.2 NetWare 4.1 System Memory Pools

NetWare 4.1 system memory architecture is quite different than V3.12. In NetWare 4.1, system memory is allocated to a single global memory pool of 4KB pages that NetWare can allocate to different processes. Each NetWare Loadable Module (NLM) has allocated memory pages that are returned to the global pool when they are not used. When an NLM initially requests memory allocation, it receives a specific set of memory pages called a node and uses that node for all subsequent calls. Because the NLM uses the same node repeatedly during the time that the NLM is loaded, allocation becomes more efficient and fragmentation is reduced.

NetWare 4.1 also uses more cache than earlier versions of NetWare. Using more cache enhances operating system performance but drives up the system memory requirements.

5.3.4 Memory Requirements

The following sections present some guidelines for calculating total server memory requirements for both the NetWare V3.12 and V4.1 environments.

5.3.4.1 Calculating Server Memory Requirements (V3.12)

To calculate the total system memory required, use the following guidelines:

1. Allocate a minimum of 4MB of RAM for loading the operating system, disk drivers, LAN drivers and the INSTALL.NLM.
2. Add 2MB for each of the following you plan to install:
 - Additional NetWare products(such as NetWare for Macintosh, or NFS)
 - A print server on this server (PSERVER.NLM)
 - Any of the following NLMs:
 - a. CBLIB.NLM
 - b. STREAMS.NLM
 - c. BTRIEVE.NLM
3. Allocate memory for FAT and directory caching. To calculate this requirement, multiply the amount of your server's hard disk space (in MB) by 0.008. For example, if you have a total of 1 GB of disk space, multiply 1 GB by 0.008. This results in a requirement for an additional 8MB of RAM.
4. Add at least 4MB for additional cache buffers to optimize performance. The more memory you can allot to cache buffers, the better your system's performance will be.

5.3.4.2 Calculating Server Memory Requirements (V4.1)

The memory requirements for NetWare 4.1 can be obtained using the same formula used for version 3.12 with the following changes:

- The operating system needs 6MB instead of 4MB.
- NetWare 4.1 can benefit even more from cache memory. Start with at least 6MB for cache buffers.

5.3.5 System Processor

In a pure file serving environment, the processor is usually not the bottleneck. However, in many NetWare implementations, there are other functions loaded on the server which can create a heavier load on the CPU. For example, a database engine loaded as a NetWare NLM can create a heavy load on the system CPU.

A high CPU utilization on a sustained basis is usually, but not always, an indication of a CPU bottleneck. By using the MONITOR.NLM utility, you can identify whether you have a CPU bottleneck or if you have any inefficient or CPU-intensive resources in the server.

For example, when monitoring the CPU utilization you may observe a high level of utilization. This does not necessarily mean that the system processor is the bottleneck. Other potential focus areas are the server hardware subsystems such as the disk controller and the LAN adapter.

One way to measure the load on the CPU from any one subsystem is to monitor the frequency of interrupt servicing. If, for example, a LAN adapter is constantly servicing packets and utilizing a large percentage of the CPU time for interrupt

processing, this is an indication that the LAN adapter could be a bottleneck in the file server.

In this case, the LAN adapter could be upgraded or another one added to balance the LAN load and use the system processor more efficiently. The use of bus mastering LAN adapters and disk controllers also helps to spread the load on the file server more evenly and thereby uses the system processor in a more efficient manner.

To use the MONITOR.NLM utility to examine subsystem load on the CPU, invoke it with the -P parameter as follows:

```
LOAD MONITOR -P
```

A detailed explanation of the monitor performance statistics is documented in the Monitor help screen.

Appendix A. EISA Configuration File

Board Information

IBM Dual Pentium PCI-EISA System Board

System

Manufacturer Micronics Computers, Inc.
ID MIC0054
Category SYS
Board slot type Embedded
Readable ID Yes

SYSTEM BOARD MEMORY

System Base Memory 640 KB Base Memory
Total System Memory. 16 MB Total Memory
Memory Gap between 15-16Megs. Memory Gap Enabled.

System Board I/O Resource Allocation

Serial Port A COM1 OR COM3 - Enabled
Serial Port B COM2 OR COM4 - Enabled
Parallel Port Parallel Port LPT1 - Enabled
Floppy Controller Floppy Controller - Enabled
ISA IDE Controller Secondary IDE IRQ 15 -
Enabled
Reserved System Resources. Reserved System Resources.

System - IBM Dual Pentium PCI-EISA System Board

Jumper Name: W1 - Level 1 Cache.

Default factory settings:
OFF

Change settings to:
OFF

```
+-----+
| o   o |
+-----+
  1   2
```

This jumper sets the level 1 (CPU's internal) cache to write-back or write-through mode. Install a jumper for write-through mode. Do not jumper for write-back mode. Default is the faster write-back mode.

=====

Jumper Name: W5 - Level 2 Cache.

Default factory settings: Change settings to:

```

                                     +---+
                                     | o | 1
OFF      OFF      |  |
                                     | o | 2
ON       ON       | | |
                                     | o | 3
                                     +---+
```

This jumper along with W10 selects the level 2 cache size. Jumper pins 1-2 for 256K cache size. Jumper pins 2-3 for 512K cache size. Default is pins 2-3; 512K cache.

=====

Jumper Name: W6 - CPU clock.

Default factory settings:
ON OFF

Change settings to:
ON OFF

```
+-----+
| o--o  o |
+-----+
  1   2  3
```

Jumper pins 1-2 for 90Mhz Pentiums or pins 2-3 for 100Mhz Pentiums. Default is pins 1-2; 90Mhz Pentiums.

=====

Jumper Name: W10 - Level 2 Cache.

Default factory settings:
OFF ON

Change settings to:
OFF ON

```
+-----+  
| o  o---o |  
+-----+  
1 2 3
```

This jumper along with W5 selects the level 2 cache size. Jumper pins 1-2 for 256K cache size. Jumper pins 2-3 for 512K cache size. Default is pins 2-3; 512K cache.

=====

Jumper Name: W20 - Monitor Type.

Default factory settings:
ON

Change settings to:
ON

```
+-----+  
| o---o |  
+-----+  
1 2
```

This jumper will set the monitor type. The default, installed, is for color monitors. Leave off the jumper for monochrome monitors.

=====

Jumper Name: W25 - ISA IDE IRQ.

Default factory settings:
OFF ON

Change settings to:
OFF ON

```
+-----+  
| o  o---o |  
+-----+  
1 2 3
```

Jumper pins 1-2 to set the ISA IDE controller as the PRIMARY IDE controller at IRQ 14. (The PCI IDE below must be disabled.) Jumpering pins 2-3 will set the ISA IDE as the SECONDARY IDE controller at IRQ 15. (A primary IDE must be present first.)

=====

Jumper Name: W28 - PCI IDE IRQ.

Default factory settings:

ON OFF

Change settings to:

ON OFF

```
+-----+
| o---o  o |
+-----+
  1   2   3
```

Jumper the pins to set the PCI IDE controller as the PRIMARY IDE controller at IRQ 14. Remove the jumper to disable the PCI IDE. The PCI IDE can not be set up as a SECONDARY IDE controller.

Board Information

IBM Auto 16/4 Token-Ring ISA Adapter

Slot 3

Manufacturer PNP
ID IBM0001
Category OTH
Board slot type Other
Readable ID No
Skirt No

IBM Auto 16/4 Token-Ring ISA Adapter Enabled

Board Information

PCI SCSI Controller

Embedded

Manufacturer PCI
ID 90047178
Category MSD
Board slot type PCI
Readable ID No
Skirt No

PCI Function 1 Enabled

Used Resources

Resource	Slot	Function
IRQ 0.....	System	Reserved System Resources.
IRQ 1.....	System	Reserved System Resources.
IRQ 3.....	System	Serial Port B
IRQ 4.....	System	Serial Port A
IRQ 6.....	System	Floppy Controller
IRQ 7.....	System	Parallel Port
IRQ 8.....	System	Reserved System Resources.
IRQ 2(9).....	Slot 3	IBM Auto 16/4 Token-Ring ISA Ada
IRQ 11.....	Embedded	PCI Function 1
IRQ 12.....	System	Reserved System Resources.
IRQ 13.....	System	Reserved System Resources.
IRQ 15.....	System	ISA IDE Controller
DMA 2.....	System	Floppy Controller
Port 2F2h.....	Slot 3	IBM Auto 16/4 Token-Ring ISA Ada
Port 0A20h - 0A23h.....	Slot 3	IBM Auto 16/4 Token-Ring ISA Ada
Port 8000h - 80FFh.....	Embedded	PCI Function 1
Memory		
Address	Amount	
# 0.....	640K....	System
0C0000h.....	32K....	Embedded
0C8000h.....	16K....	Slot 3
0CC000h.....	8K....	Slot 3
1M.....	14M....	System
16M.....	4K....	Embedded
		System Base Memory
		PCI Function 1
		IBM Auto 16/4 Token-Ring ISA Ada
		IBM Auto 16/4 Token-Ring ISA Ada
		Total System Memory.
		PCI Function 1

= Caching

Available Resources

IRQs	DMAs	ISA I/O Ports	Memory	Amount	Address
5	0	100h - 2F1h		64K	0A0000h
10	1	2F3h - 400h		64K	0B0000h
14	3			8K	0CE000h
	5			64K	0D0000h
	6			64K	0E0000h
	7			64K	0F0000h
				1M	15M

System Specifications

Slot Name	Slot Type	Board ID	Accept Skirted	Max Length	Bus- master	Slot Tag(s)
Slot 1	EISA	(Empty)	Yes	341mm	Yes	
Slot 2	EISA	(Empty)	Yes	341mm	Yes	
Slot 3	EISA	IBM0001	Yes	341mm	Yes	
Slot 4	EISA	(Empty)	Yes	341mm	Yes	
Slot 5	EISA	(Empty)	Yes	341mm	Yes	
Slot 6	EISA	(Empty)	Yes	341mm	Yes	

Nonvolatile memory 8K

Appendix B. Hardware Compatibility, Device Driver, and Software Patch Information

One of the challenges that LAN administrators face is ensuring that the adapters they want to use are supported and then finding the latest level of device drivers for these adapters. Another problem they face is trying to find information about known software bugs and their corresponding fixes. This appendix is included to help readers solve these problems.

B.1 Finding Compatibility Information on the World Wide Web

IBM does extensive testing on each new system before it is introduced to ensure that it is compatible with the wide variety of IBM and non-IBM adapters which exist in the marketplace. Recently, this information has been made public and is now maintained and updated on a regular basis on a World Wide Web (WWW) site. The Universal Resource Locator (URL) for this site is:

- <http://www.pc.ibm.com/compat/compat.html>

This site has a wealth of valuable information. It contains information by server (for example IBM PC Server 320, IBM PC Server 500, and so on). For each machine, it has reports on:

- LAN adapters
- WAN adapters
- Peripheral interface adapters
- Storage devices
- Other devices such as video chip sets

Figure 145 on page 200 shows a sample report on Ethernet LAN adapters tested on the PC Server 320.

Ethernet adapters

Mfg	Description	Model/PN	Class	Type	Bus
IBM	LAN Adapter for Ethernet	48G7169	LAN	Ethernet	ISA
IBM	LAN Adapter for Ethernet TP	60G0605	LAN	Ethernet	ISA
IBM	LAN Adapter for Ethernet CX	60G0615	LAN	Ethernet	ISA
3Com**	Etherlink** III 32-Bit EISA	3C579	LAN	Ethernet	EISA
3Com	Etherlink II	3C503	LAN	Ethernet	ISA
3Com	Etherlink II	3C503-16	LAN	Ethernet	ISA
3Com	Etherlink III	3C507-TP	LAN	Ethernet	ISA
3Com	Etherlink III	3C509-TP	LAN	Ethernet	ISA
Gateway**	Ethernet 16	Ethernet 16	LAN	Ethernet	ISA
Intel**	EtherExpress** 16C	.	LAN	Ethernet	ISA
Novell**	NE1000	NE1000	LAN	Ethernet	ISA
Novell	NE2000	NE2000	LAN	Ethernet	ISA
Novell	NE2000 Plus	NE2000 Plus	LAN	Ethernet	ISA
Novell	NE3210	NE3210	LAN	Ethernet	EISA
SMC**	Ethercard Elite** 32	LANEP32	LAN	Ethernet	EISA
SMC	Elite 16TP	LANET16	LAN	Ethernet	ISA

Figure 145. Sample Compatibility Report Showing Ethernet LAN Adapters

In addition, it contains reports by NOS and by machine and model number which shows the testing done to certify the hardware against the operating system. These reports show detailed information on hardware configurations, adapters, device drivers, and code levels used in the testing.

B.2 Finding Device Drivers on the World Wide Web

In addition to this web site, IBM has another site that contains a file which has LAN adapter device driver information. This file is a matrix of adapters and drivers and version levels. It is updated frequently as new drivers become available. The file is named DRVRLIST.DOC and is in the MS Word format. It can be obtained from the PCC FTP file library at the following URL:

- <ftp.pcco.ibm.com/pub/network>

or with a web browser at:

- <http://www.pc.ibm.com/files.html>

If it is an IBM adapter, the driver itself can be obtained from the same FTP or web sites. In addition, the IBM Networking Hardware Division maintains an FTP site at the following URL:

- <lansupport.raleigh.ibm.com/pub/products/lanprods/adapters>

For OEM adapters, many manufactures have web and FTP sites which work very similar to the process outlined above.

There is a one very useful web site which maintains an index to many OEM web sites which have device drivers for many popular boards. The site is called the LAN Drivers Page and the URL for this site is:

- <http://sunsite.unc.edu/~towfiq/lan-drivers.html>

B.3 Finding Software Patches on the World Wide Web

In addition to information on device drivers, the IBM PC FTP file library contains information on software patches in the OS/2 LAN Server Environment. It contains the latest fixes for OS/2 and LAN Server as well as system BIOS fixes for several IBM PC Server models.

Appendix C. Configuring DOS CD-ROM Support

This appendix is included to aid in configuring CD-ROM support in the DOS environment.

C.1 Installing CD-ROM Support for PCI Adapters.

To install CD-ROM support for an IBM PC Server with a PCI SCSI adapter:

1. Copy the following files from the ServerGuide license diskette to the hard disk (C:SERVER directory):

- IBMPRAID.SYS
- IBMASPI.SYS
- DAFCD.A.SYS
- MSCDEX.EXE

2. Add the following statements to the CONFIG.SYS file:

```
DEVICE=C:SERVERIBMPRAID.SYS  
DEVICE=C:SERVERIBMASPI.SYS  
DEVICE=C:SERVERDAFCDA.SYS
```

3. Add the following statement to the AUTOEXEC.BAT file:

```
C:SERVERMSCDEX /D:IBMCD001 /M:10 /L:E
```

C.2 Installing CD-ROM Support for Adaptec Adapters

To install CD-ROM support for an IBM PC Server with an Adaptec SCSI adapter:

1. Copy the following files from the ServerGuide license diskette to the hard disk (C:SERVER directory):

- ASPI8DOS.SYS
- ASPICD.SYS
- MSCDEX.EXE

2. Add the following statements to the CONFIG.SYS file:

```
DEVICE=C:SERVERASPI8DOS.SYS /D  
DEVICE=C:SERVERASPICD.SYS /D:ASPICD0
```

3. Add the following statement to the AUTOEXEC.BAT file:

```
C:SERVERMSCDEX /D:IBMCD001 /M:10 /L:
```

C.3 Installing CD-ROM Support for Micro-Channel Adapters

To install CD-ROM support for an IBM PC Server with a Micro-Channel SCSI adapter:

1. Copy the following files from the ServerGuide license diskette to the hard disk (C:SERVER directory):

- IBMRAID.SYS
- RDASPI.SYS
- IBMCDROM.SYS
- RDCDROM.SYS
- RAMDRIVE.SYS

- MSCDEX.EXE
2. Add the following statements to the CONFIG.SYS file:
DEVICE=C:SERVERIBMRAID.SYS
DEVICE=C:SEVERRDASPII.SYS
DEVICE=C:SERVERIBMCDROM.SYS /Q
DEVICE=C:SEVERRDCDROM.SYS
DEVICE=C:SEVERRAMDRIVE.SYS 2800 512 128 /E
 3. Add the following statement to the AUTOEXEC.BAT file:
C:SERVERMSCDEX /D:IBMCD001 /M:10 /L:

List of Abbreviations

ALU	Arithmetic-Logic Unit	ITSO	International Technical Support Organization
APA	all points addressable	MCA	Micro Channel Architecture
ASCII	American Standard Code for Information Interchange	MFM	Modified Frequency Modulation
ASPI	Advanced SCSI Programming Interface	MI	Management Interface (as in Desktop Management Interface)
BIOS	Basic Input/Output System	MIB	Management Information Base
CI	Component Interface (as in Desktop Management Interface)	MIF	Management Information Format
CMIP	Common Management Information Protocol	NetBIOS	Network Basic Input/Output System
CMOS	Complementary Metal Oxide Semiconductor	NLM	NetWare Loadable Module
CPU	Central Processing Unit	NVRAM	Non-Volatile Random Access Memory
DASD	Direct Access Storage Device	PCI	Peripheral Component Interconnect
DIMM	Dual In-Line Memory Module	PROFS	Professional Office System
DMA	Direct Memory Access	RAID	Redundant Array of Inexpensive Disks
DMI	Desktop Management Interface	RAM	Random Access Memory
DRAM	Dynamic Random Access Memory	RLL	Run Length Limited
ECC	Error Correcting Code	SCB	System Control Block
ECC-P	ECC w/ Parity	SCSI	Small Computer System Interface
EEPROM	Electrically Erasable Programmable Read-Only Memory	SIMM	Single In-Line Memory Module
EISA	Extended Industry Standard Architecture	SMP	Symmetric Multiprocessing
EMS	Expanded Memory Specification	SNA	Systems Network Architecture
EOS	ECC on SIMM	SNMP	Simple Network Management Protocol
ESDI	Enhanced Small Device Interface	SRAM	Static Random Access Memory
FTP	File Transfer Protocol	TCP/IP	Transmission Control Protocol/Internet Protocol
IBM	International Business Machines Corporation	TCQ	Tagged Command Queueing
IDE	Integrated Drive Electronics	UMB	Upper Memory Block
IPC	Inter-Process Communication	UPS	Uninterruptible Power supply
IPL	Initial Program Load	URL	Uniform Resource Locator
IPX	Internet Packet eXchange		
ISA	Industry Standard Architecture		

Index

Special Characters

.ADF files 85
.CFG file 75

Numerics

16/4 Token-Ring Bus Master Adapter 32
32-bit addressing 2, 168
80486DX2 2
80C186 19

A

abbreviations 205
acronyms 205
active client, NetFinity 43
adapter descriptor files, Micro-Channel 85
address bus 2
administering RAID arrays 151
advanced diagnostics, MCA machines 99
alert manager, NetFinity 42, 155
alerts, NetFinity 146
APC PowerChute 53
application dimension, SystemView 46
arithmetic logic unit (ALU) 1
asymmetric multiprocessing 3
asynchronous SCSI device 19
AUTOEXEC.NCF file 128, 134, 142
automatic configuration, Micro-Channel 85
availability 22
average access time 21

B

base executable, NetFinity 44
battery voltage 53
BIOS Scan 81
blackouts, AC power 52
block size 176
boot options, configuring 68
boot password 72
bottleneck 32, 168
bridging 33
brownouts 52
burst data transfer 19
bus architectures
 EISA 13
 ISA 13
 MCA 13
 PCI 15
bus master 31, 32, 168

C

cable cover 35
cache buffer size 176
cache memory
 cache hit 4, 171
 cache miss 4
 coherency of 6
 enabling/disabling 68
 hit rate 168
 L1 cache 4
 L2 cache 4
 L3 cache 5
 lazy-write cache 4
 MESI protocol 6
 tags 6
 write-back cache 4
 write-through cache 4
caching, on disk subsystems 22, 179
calculating memory requirements 187
CD-ROM, Installing NetWare 4.1 135
central processing unit (CPU) 1, 167
clock doubling 1, 2
clock rate 1
CMOS RAM 88
common command set 18
Component Interface (CI) 39
configuration, hardware 65
configuring NetFinity Services 120, 145
connections, NetWare 119
console 95
context, NetWare 126, 141
coping system and public files 132
CPU 187
CPU overlap 18
creating a NetWare partition 131
creating a NetWare volume 131
critical file monitor, NetFinity 42

D

DASD
 See Direct Access Storage Device (DASD)
data bus 2, 168
data dimension, SystemView 46
data spanning 24
data streaming 14, 16, 168
data stripe mirroring 25
data striping 24
DatagLANce 172
date and time, EISA machines 68
date and time, MCA machines 92
decoding protocols 172
dedicated cache 5

- defunct drives 158
- design considerations
 - availability 22
 - cost 23
 - performance and capacity 22
- Desktop Management Interface (DMI) 37
 - browser in NetFinity 44
 - communicating service 38
 - Component Interface (CI) 39
 - manageable products 38
 - management applications 38
 - Management Information Format (MIF) 38
 - Management Interface (MI) 39
 - MI commands
 - get 39
 - list 39
 - set 39
- destination address 172
- detecting disk failures 155
- diagnostics 97
- diagnostics, running 85
- Direct Access Storage Device (DASD) 16
 - average access time 21, 28
 - block size 176
 - cache buffer size 176
 - caching 22
 - failures 152
 - file caching 179
 - latency 22, 24, 168
 - maximum transfer rate 21
 - on-board cache size 21
 - read after write verify 182
 - rotational delay 21
 - seek time 22
- Direct Memory Access (DMA) 32
- directory caching 178
- directory tree, NDS 126, 141
- dirty disk cache delay time 181
- disconnect/reconnect 18, 80
- discovering workstations 146
- disk cache 168, 170
- disk driver 128
- disk duplexing 24
- disk failure 152
- disk mirroring 24
- disk partitions 122
- disk read after write verify 182
- disk subsystem 16, 168, 172
- diskette drive
 - ejecting diskettes 36
 - locking 36
 - unlocking 36
- diskette factory 117
- DMI
 - See Desktop Management Interface (DMI)
- double-bit errors 9
- DRAM
 - See dynamic random access memory (DRAM)

- driver efficiency 182
- driver for token-ring, installing 132
- dual path bus 7
- duplexing 24
- dynamic random access memory (DRAM) 3

E

- ECC
 - See Error Correcting Code (ECC)
- ECC Memory on SIMMs (EOS) 9, 11
- ECC on SIMMs memory 11
- ECC-P
 - See Error Correcting Code Parity (ECC-P)
- EISA
 - See Extended Industry Standard Architecture (EISA)
- ejecting diskettes 36
- Electrical Erasable Programmable Read Only Memory (EEPROM) 36
- encryption, of software products 115
- end-use dimension, SystemView 45
- Enhanced Small Device Interface (ESDI) 16
- EOS memory
 - See ECC Memory on SIMMs (EOS)
- Error Correcting Code (ECC) 9
 - ECC configuration, NetFinity 42
- Error Correcting Code Parity (ECC-P) 9
- error logs, NetFinity 155
- ESDI
 - See Enhanced Small Device Interface (ESDI)
- EtherStreamer LAN adapter 33
- event scheduler, NetFinity 44
- Extended Industry Standard Architecture (EISA) 12
 - configuration 65, 74

F

- fast SCSI interface 18
- fast startup mode, enabling 96
- fast/wide SCSI interface 18
- fat indexing 177
- file caching 179
- file purging 181
- file transfer, NetFinity 44
- flash EEPROM 36
- flash ROM, updating 73

H

- hard disk interfaces
 - bus level interface 16
 - device level interface 16
 - disk cache 168, 170
 - ESDI 16
 - SCSI 16
 - ST506 16
- hardware configuration 65

high availability 22
hot spare drive 154
HP OpenView, interoperability with NetFinity 45

I

IBM LANStreamer MC32 32
IBM Quad PeerMaster Adapter 33
IDE
 See Integrated Drive Electronics (IDE)
immediate file purging 181
Industry Standard Architecture (ISA) 12
INSTALL.BAT file 127
INSTALL.NLM 128
installing NetFinity Services 144
Installing NetWare 4.1 with CD-ROM 135
installing system and public files 132
installing the operating system 115
installing the RAID Administration utility 151
Integrated Drive Electronics (IDE) 16, 31
 configuring 68, 70
interleaved memory 7
internal network number, IPX 129
intrusion detection 35
IPL sequence 37, 69, 95
IPL source 36
IPX internal network number 129
ISA
 See Industry Standard Architecture (ISA)

K

keyboard speed 95

L

L1 cache 4
L2 cache 4
L3 cache 5
LAN adapter 31
LAN driver 128
LAN driver efficiency 182
LAN drivers, installing 132
LANStreamer adapter 32
large disk DOS compatibility 70
latency 22, 24, 168
latency, of LAN adapters 32
lazy-write cache 4
level 1 cache 4
license 119
license, for software products 115
local bus 15
locking diskette drives 36
logging failures, NetFinity 155
LogicLock 34
low-level format 82, 101

M

managed resource dimension, SystemView 47
Management Information Base (MIB) 40
Management Information Format (MIF) 38
Management Interface (MI) 39
maximum concurrent disk cache writes 180
maximum transfer rate 21
MCA
 See Micro-Channel Architecture (MCA)
memory gap 71
memory map 89
memory pools, NetWare 185
memory requirements, NetWare 187
MESI protocol 6
Micro-Channel Architecture (MCA) 12
 options diskette 85
 system programs 65, 82
Micro-Channel SCSI Adapter 19
microprocessor 1
mirroring 24
Modified Frequency Modulation (MFM) 16
MONITOR.NLM 173
mounting volumes 132
movable memory 186
multiple volumes 181
multiprocessing 3

N

NDS
 See NetWare Directory Services (NDS)
NetFinity
 Manager 41, 43
 Services 41
NetFinity Services 144
NetWare Directory Services (NDS) 126, 141
network adapter 170
network driver, NetFinity 145
network number, IPX 129
network subsystem 182
network time-out, NetFinity 146
new packet receive buffer wait time 185
NF config command 121
noise filters 52
noise, AC power 52
non-movable memory 186
Non-Volatile Random Access Memory (NVRAM) 10
NVRAM
 See Non-Volatile Random Access Memory (NVRAM)

O

online status, disk drive 165
operating system, installing 115
options diskette, Micro-Channel 85
original cache buffers 179

P

- P5 1
- P54C 1
- packet receive buffers 183
- packet size 183
- parallel port, configuration of 70
- parity 9
- parity placement, in RAID arrays 106
- passive client, NetFinity 43
- password jumper 36
- password on boot 72
- PCI
 - See Peripheral Component Interconnect (PCI)
- PeerMaster technology 33
- Pentium 2
- Peripheral Component Interconnect (PCI) 12, 15
- permanent memory 186
- Personal System/2 Micro Channel SCSI Adapter 19
- Personal System/2 Micro Channel SCSI Adapter with Cache 19
- POST error detect, NetFinity 44
- POST errors, prompting 69
- power quality 53
- PowerChute Plus 53
- predictive failure analysis, NetFinity 42
- presence error reporting 91
- privileged-access password 36
- process manager, NetFinity 44
- processor 1, 167, 187
- processor performance 168
- product demonstrations 118
- protocol analysis 172
- public files, installing 132
- purging files 181

Q

- Quad PeerMaster Adapter 33

R

- RAID
 - See Redundant Array of Inexpensive Disks (RAID)
- RAM
 - See Random Access Memory (RAM)
- RAM paging 31
- Random Access Memory (RAM)
 - access time 4, 12
- read after write verify 182
- read-ahead 106
- real-time clock power-on 97
- rebuild priority, of RAID array 106
- Redundant Array of Inexpensive Disks (RAID)
 - array initialization 112
 - backup and restoring configuration 112
 - creating disk arrays 102
 - data stripe mirroring 25
 - data striping 24

- Redundant Array of Inexpensive Disks (RAID)
 - (continued)

- defining hot-spare drives 102
- device status 102
- parity placement 106
- performance 171
- RAID Administration utility 151
- RAID Manager, NetFinity 42
- RAID utility 65, 101
- RAID-0 24, 30
- RAID-1 24, 30
- RAID-1 enhanced 25
- RAID-2 27, 30
- RAID-3 28, 30
- RAID-4 28, 30
- RAID-5 29, 30
- skewing 29
- stripe unit size 106
- write policy 112
- reference diskette 83
- registers 1
- reliability, of disk subsystems 22, 24
- remote session, NetFinity 44
- remote systems manager, NetFinity 43
- response time 170
- ring map 173
- rotational delay 21
- Run Length Limited (RLL) 16

S

- scatter/gather 18
- SCB
 - See Subsystem Control Block (SCB)
- scheduled server shutdowns 53
- scheduling events 44
- screen view, NetFinity 44
- SCSI
 - See Small Computer System Interface (SCSI)
- SCSI-2 Fast/Wide PCI Adapter 79
- second-level cache 4
- secure removable media 36
- security 71
 - administrator password, MCA 92
 - C2 level 34
 - cable cover 35
 - cabling 35
 - intrusion detection 35
 - LogicLock 34
 - password forgotten -warning 36
 - password jumper on mother board 36
 - power-on password, EISA 72
 - power-on password, MCA 92
 - preventing diskette removal 36
 - privileged-access password 36
 - privileged-access password, MCA 94
 - secure removable media 36
 - selectable drive startup 37
 - supervisor password, EISA 71

security (*continued*)

- tamper-evident switches 34
- unattended start mode 37
- unattended start mode, MCA 94
- user password, EISA 71

security manager, NetFinity 42

seek time 22

selectable drive startup 37

selecting a time zone 125

serial control service, NetFinity 42

serial port, configuration of 70

serial-port power-on 97

server components 167

SERVER.EXE file 127

ServerGuard 42

ServerGuide

- diskette factory 117
- overview 115
- performance tuning utility 117
- product demonstrations 118
- software encryption 115

SERVMAN.NLM 173

Set command 173

setting the console, MCA machines 95

setup program 65

shadow, enabling Video BIOS 68

shared cache 5

shared RAM 31

SIMM

- See Single In-line Memory Modules (SIMM)

Simple Network Management Protocol (SNMP) 39

- agent 40
 - Management Information Base (MIB) 40
 - manager 40
 - MIB extension 40
 - traps 40

simulating a disk failure 152

Single In-line Memory Modules (SIMM) 7

- matched sets 7

single-bit errors 9

skewing 29

Small Computer System Interface (SCSI)

- adapters 19
- asynchronous device 19
- command command set 18
- controller 17
- device configuration 80
- device configuration, MCA 90
- disconnect/reconnect 18, 80
- enabling parity checking 79
- fast device 18
- scatter/gather 18
- SCSI select utility 65, 78
- SCSI-2 Fast/Wide Adapter/A 19
- SCSI-2 Fast/Wide PCI Adapter 20
- SCSI-2 Fast/Wide RAID PCI Adapter 20
- SCSI-2 Fast/Wide Streaming-RAID Adapter 20
- SCSI-I 17

Small Computer System Interface (SCSI) (*continued*)

- SCSI-II 17, 19
 - setting adapter termination 79
 - start unit command 80
 - sync negotiation 80
 - synchronous device 19
 - Tagged Command Queuing 18
 - wide device 18

SMP

- See symmetric multiprocessing (SMP)

SNMP

- See Simple Network Management Protocol (SNMP)

software encryption 115

software inventory, NetFinity 44

software license 115

source address 172

spanning volumes over multiple drives 181

spikes, AC power 52

SRAM

- See static random access memory (SRAM)

ST506 16

stand-alone client, NetFinity 43

standard memory 9

start unit command, SCSI 80

startup sequence 95

STARTUP.NCF FILE 128, 133

STARTUP.NCF file. 142

static RAM

- See static random access memory (SRAM)

static random access memory (SRAM) 3

stripe unit size 106

Subsystem Control Block (SCB) 19

superscalar architecture 2

supervisor password 71

surge suppressors 52

surges, AC power 52

swapping floppy disk drives 69

switching 33

symmetric multiprocessing (SMP) 3

sync negotiation, SCSI devices 80

synchronous SCSI device 19

SynchroStream 8

SYS: volume 131

system and public files 132

system error log, MCA machines 99

system error restart 97

system information, NetFinity 41

system keywords, NetFinity 121, 146

system memory 185

system monitor, NetFinity 42

system name, NetFinity 145

system partition 83, 86

system partition access, NetFinity 42

system profile, NetFinity 42

system programs, Micro-Channel 65, 82

SystemView 45

T

- Tagged Command Queuing (TCQ) 18
- tamper-evident switches 34
- TCP/IP
 - See Transmission Control Protocol/Internet Protocol (TCP/IP)
- TCQ
 - See Tagged Command Queuing (TCQ)
- terminator, SCSI bus 79
- testing the computer 97
- throughput 169
- time zone, NetWare 125
- TOKEN.LAN driver 132
- tools, protocol analysis 172
- top ten talkers 173
- total cache buffers 179
- transaction rate 170
- transferring licenses 119
- Transmission Control Protocol/Internet Protocol (TCP/IP) 39
- traps, in SNMP 40
- Tuning NetWare 175
- turbo fat indexing 177

U

- unattended operations, security 37
- unattended start mode 37, 94
- Uninterruptible Power Supply (UPS) 52
 - monitoring 53
- unlocking diskette drives 36
- updating flash ROM 73
- Upper Memory Block (UMB) 31
- user password 71
- utilization, processor 187

V

- video configuration 68
- virtual network (VNET) 33
- voltage 53
- voltage regulators 52

W

- wait state 3
- wide SCSI interface 18
- write policy 112
- write-back 4, 112
- write-through 4, 112

Z

- zero wait state 3

**International Technical Support Organization
IBM PC Server and Novell NetWare
Integration Guide
December 1995**

Publication No. SG24-4576-00

Your feedback is very important to help us maintain the quality of ITSO Bulletins. **Please fill out this questionnaire and return it using one of the following methods:**

- Mail it to the address on the back (postage paid in U.S. only)
- Give it to an IBM marketing representative for mailing
- Fax it to: Your International Access Code + 1 914 432 8246
- Send a note to REDBOOK@VNET.IBM.COM

**Please rate on a scale of 1 to 5 the subjects below.
(1 = very good, 2 = good, 3 = average, 4 = poor, 5 = very poor)**

Overall Satisfaction	_____		
Organization of the book	_____	Grammar/punctuation/spelling	_____
Accuracy of the information	_____	Ease of reading and understanding	_____
Relevance of the information	_____	Ease of finding information	_____
Completeness of the information	_____	Level of technical detail	_____
Value of illustrations	_____	Print quality	_____

Please answer the following questions:

- a) If you are an employee of IBM or its subsidiaries:
- | | | |
|--|--------|-------|
| Do you provide billable services for 20% or more of your time? | Yes___ | No___ |
| Are you in a Services Organization? | Yes___ | No___ |
- b) Are you working in the USA? Yes___ No___
- c) Was the Bulletin published in time for your needs? Yes___ No___
- d) Did this Bulletin meet your needs? Yes___ No___

If no, please explain:

What other topics would you like to see in this Bulletin?

What other Technical Bulletins would you like to see published?

Comments/Suggestions: (THANK YOU FOR YOUR FEEDBACK!)

Name

Address

Company or Organization

Phone No.




Fold and Tape

Please do not staple

Fold and Tape



NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES



BUSINESS REPLY MAIL

FIRST CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

IBM International Technical Support Organization
Department HZ8, Building 678
P.O. BOX 12195
RESEARCH TRIANGLE PARK NC
USA 27709-2195



Fold and Tape

Please do not staple

Fold and Tape



Printed in U.S.A.

SG24-4576-00



Artwork Definitions

<u>id</u>	<u>File</u>	<u>Page</u>	<u>References</u>
ITSLOGO	4576SU	1	i

Figures

<u>id</u>	<u>File</u>	<u>Page</u>	<u>References</u>
ASYM	CH1A	5	1 5
SYMM	CH1A	6	2 6
INTERL	CH1A	7	3 7
DUALP	CH1A	8	4 7
ECCMEM	CH1A	10	5 9
ECCP	CH1A	11	6 11
STRM20	CH1A	14	7 13
STRM40	CH1A	14	8 14
STRM80	CH1A	15	9 15
DISK2	CH1B	17	10 17
STRIP4	CH1B	24	11
MIRROR	CH1B	25	12
DUPLEX	CH1B	25	13
STR1	CH1B	26	14
RAID6	CH1B	26	15 26
STRIP1	CH1B	27	16 27, 27
STRIP3	CH1B	28	17 28
STRIP5	CH1B	29	18
STRIP6	CH1B	29	19
NETFM3	CH1E	41	20 41
PCSRVRS	CH2A	55	21 55
CH301	CH3	66	22
CH302	CH3	67	23 67
CH303	CH3	68	24 68
CH304	CH3	69	25 69
CH305	CH3	70	26 70
CH306	CH3		

		71	27	71
CH313	CH3			
		74	28	74
CH314	CH3			
		74	29	74
CH315	CH3			
		75	30	74
CH316	CH3			
		75	31	
CH317	CH3			
		76	32	76
CH318	CH3			
		77	33	76
CH308	CH3			
		78	34	78
CH309	CH3			
		78	35	78
CH310	CH3			
		80	36	79
CH311	CH3			
		81	37	81, 81
CH312	CH3			
		82	38	82
RDMAIN	CH3			
		84	39	83, 84
BACKUP	CH3			
		86	40	86
SCMENU	CH3			
		87	41	86
VCMENU	CH3			
		88	42	87
MEMAP	CH3			
		90	43	89
VSCSI	CH3			
		91	44	90
SETFEAT	CH3			
		92	45	91
SETPAS	CH3			
		93	46	92
SSSEQ	CH3			
		96	47	96
POFEAT	CH3			
		97	48	96, 97
MUMENU	CH3			
		98	49	98
DRLEV	CH3			
		98	50	98
ERRLOG	CH3			
		99	51	99
AVDIAG	CH3			
		100	52	100
TEST	CH3			
		100	53	
CH319	CH3			
		102	54	102
CH320	CH3			
		102	55	102, 102
CH321	CH3			
		103	56	
CH322	CH3			

CH323	CH3	104	57	
CH324	CH3	105	58	
		105	59	105
CH325	CH3			
		107	60	
CH326	CH3			
		107	61	107
CH327	CH3			
		108	62	
CH328	CH3			
		109	63	
CH329	CH3			
		109	64	
CH330	CH3			
		110	65	
CH331	CH3			
		111	66	110
CH332	CH3			
		111	67	111
CH333	CH3			
		113	68	113
CH401	SGOVER			
		116	69	116
CH402	SGOVER			
		117	70	116
SVINET	4576CH4			
		119	71	124
SGCONF	4576CH4			
		120	72	120
SGCNETF	4576CH4			
		121	73	120
SVPARHD	4576CH4			
		122	74	122
SVREV	4576CH4			
		123	75	123
SVUNL	4576CH4			
		124	76	124
SGINST1	4576CH4			
		125	77	125
SGINST2	4576CH4			
		125	78	125
SGINST3	4576CH4			
		126	79	126
SGINST4	4576CH4			
		126	80	126
3INST	4576CH4			
		129	81	128
312INS	4576CH4			
		130	82	130
312CP	4576CH4			
		130	83	130
312PIF	4576CH4			
		131	84	131
312CV	4576CH4			
		131	85	131
312VS	4576CH4			
		132	86	132
312CPY	4576CH4			
		132	87	132
312SF	4576CH4			
		133	88	

312STF	4576CH4			133
		134	89	
312AEX	4576CH4			133
		134	90	
41NW01	4576CH4			134
		136	91	
NW41DD	4576CH4			135
		137	92	
NW41ND	4576CH4			137
		137	93	
NW41CP1	4576CH4			137
		138	94	
NW41V	4576CH4			138
		138	95	
NW41VOL	4576CH4			138
		139	96	
NW41OF	4576CH4			139
		139	97	
NW41FS	4576CH4			139
		140	98	
NW41TN	4576CH4			139
		140	99	
NW41DN	4576CH4			140
		141	100	
NW41SI	4576CH4			140
		141	101	
NW41ES	4576CH4			141
		142	102	
NW41AF	4576CH4			141
		142	103	
NW41MC	4576CH4			142
		143	104	
NW41IO	4576CH4			142
		143	105	
NETFND	NETF			143
		145	106	
WNETFI1	NETF			147
		147	107	
WNET	NETF			147
		148	108	
WNETFI2	NETF			148
		148	109	
WNETFI3	NETF			148
		149	110	
WNETFI5	NETF			149
		150	111	
WNETFI6	NETF			150
		150	112	
WRAID01	RAIDADM			150
		152	113	
WRAID02	RAIDADM			152, 152
		153	114	
WRAID03	RAIDADM			153
		154	115	
WRAID04	RAIDADM			153
		154	116	
WRAID05	RAIDADM			154
		155	117	
WRAID06	RAIDADM			154
		155	118	

				155
WRAID07	RAIDADM			
		156	119	155
WRAID08	RAIDADM			
		156	120	156
WRAID09	RAIDADM			
		157	121	156
WRAID10	RAIDADM			
		157	122	157
WRAID11	RAIDADM			
		158	123	157
WRAID12	RAIDADM			
		158	124	158
WRAID13	RAIDADM			
		159	125	158
WRAID14	RAIDADM			
		159	126	159
WRAID15	RAIDADM			
		160	127	159
WRAID16	RAIDADM			
		160	128	160
WRAID17	RAIDADM			
		161	129	160
WRAID18	RAIDADM			
		161	130	161
WRAID19	RAIDADM			
		162	131	161
WRAID20	RAIDADM			
		162	132	162
WRAID21	RAIDADM			
		163	133	163
WRAID22	RAIDADM			
		163	134	163
WRAID23	RAIDADM			
		164	135	163
WRAID24	RAIDADM			
		164	136	164
WRAID25	RAIDADM			
		165	137	164
WRAID26	RAIDADM			
		165	138	165
INTRO1	4576CH5			
		167	139	167
PERFMCE	4576CH5			
		169	140	169, 169
PRFMA	4576CH5			
		171	141	170
PERFME2	4576CH5			
		172	142	171, 172
MONITOR	4576CH5			
		174	143	173
SERVMAN	4576CH5			
		175	144	175
LANCOMP	AX3			
		200	145	199

Headings			
----------	--	--	--

<u>id</u>	<u>File</u>	<u>Page</u>	<u>References</u>
NOTICES	4576FM	xv	Special Notices ii
BIBL	4576PREF	xviii	Related Publications
CH1	CH1A	1	Chapter 1, IBM PC Server Technologies xvii
PROCESR	CH1A	1	1.1, Processors
TYPES	CH1A	2	1.1.3, Processor Types
MP	CH1A	3	1.2, Multiprocessing
ASYMUL	CH1A	3	Asymmetric Multiprocessing
SMP	CH1A	3	Symmetric Multiprocessing (SMP)
MEMORY	CH1A	3	1.3, Memory 2
SMP SH	CH1A	5	Shared Cache
SMPDED	CH1A	6	Dedicated Cache
INTERL	CH1A	6	1.3.2, Memory Interleaving
DUAL	CH1A	7	1.3.3, Dual Path Buses
ECC	CH1A	9	1.4, Memory Error Detection and Correction
BUS	CH1A	12	1.5, Bus Architectures
ISA	CH1A	13	1.5.1, ISA Bus
EISA	CH1A	13	1.5.2, EISA Bus
MCA	CH1A	13	1.5.3, Micro Channel Bus
STREAM	CH1A	13	1.5.3.1, Data Streaming
PCI	CH1A	15	1.5.4, PCI Bus
DISKSS	CH1B	16	1.6, Disk Subsystem 181
DISKIF	CH1B	16	1.6.1, Hard Disk Interfaces
SCSI	CH1B	17	1.6.2, SCSI Technology
SCSIADA	CH1B	19	1.6.3, SCSI Adapters
ADAPSUM	CH1B	21	1.6.3.1, Summary
DISKS	CH1B	21	1.6.4, Hard Disk Drives
RAID	CH1B	22	1.6.5, RAID Technology 21
LAN	CH1C	31	1.7, LAN Subsystem
SECURE	CH1D	34	1.8, Security Features
SECIO	CH1D	35	1.8.2, Secure I/O Cables
MANAGMT	CH1E	37	1.9, Systems Management
COMSL	CH1E	38	1.9.1.1, Communicating Service Layer
MIF	CH1E	38	1.9.1.2, Management Information Format (MIF)
MI	CH1E	39	1.9.1.3, Management Interface (MI)
CI	CH1E	39	1.9.1.4, Component Interface (CI)
MIB	CH1E	40	1.9.2.1, Management Information Base (MIB)
AGENT	CH1E	40	1.9.2.2, SNMP Agent
MANAGER	CH1E	40	1.9.2.3, SNMP Manager
TRAPS	CH1E		

SERVS	CH1E	40	1.9.2.4, Traps
MGR	CH1E	41	1.9.3.1, NetFinity Services
NARCH	CH1E	43	1.9.3.2, NetFinity Manager
DMI	CH1E	44	1.9.3.3, NetFinity Architecture
INTEROP	CH1E	45	1.9.3.4, DMI Support
STRUC	CH1E	45	1.9.3.5, Interoperability with Other Management Tools
PROTOD	CH1E	47	1.9.4.1, SystemView Structure
PCHUTE	CH1F	53	1.9.4.2, SystemView Management Protocols
FLEX	CH1F	53	1.11.1, APC PowerChute
CH2	CH2A	53	1.11.1.1, Flex Events
FAMSPEC	CH2A	55	Chapter 2, IBM PC Server Family Overview xvii
HWCONF	CH3	56	2.1, IBM PC Server Model Specifications
SYSPRO	CH3	65	Chapter 3, Hardware Configuration xvii
STPART	CH3	82	3.4, System Programs
REFDISK	CH3	83	3.4.1, Starting From the System Partition
MMENU	CH3	84	3.4.2, Starting From the Reference Diskette
BRSYS	CH3	84	3.4.3, Main Menu Options
CONFIG	CH3	86	3.4.4, Backup/Restore System Programs Menu
SFEAT	CH3	86	3.4.5, Set Configuration Menu
TEST	CH3	91	3.4.6, Set Features Menu
MOREUT	CH3	97	3.4.7, Test the Computer
ADGPR	CH3	98	3.4.8, More Utilities Menu
DASDDEF	CH3	99	3.4.9, Advanced Diagnostic Program
NOSINST	SGOVER	106	3.5.3, Defining a Hot-Spare Disk
NETSG	4576CH4	115	Chapter 4, Novell NetWare Installation xvii
STUNING	4576CH5	118	4.3, Installing NetWare 4.1 with ServerGuide 144
HTUNING	4576CH5	167	Chapter 5, Performance Tuning xvii
EISACFG	AX1	167	5.1, Hardware Tuning 33, 175
DDINFO	AX3	189	Appendix A, EISA Configuration File xvii, 77
WWWCOMP	AX3	199	Appendix B, Hardware Compatibility, Device Driver, and Software Patch Information xvii, 73
WWWDD	AX3	199	B.1, Finding Compatibility Information on the World Wide Web
WWWFIX	AX3	200	B.2, Finding Device Drivers on the World Wide Web
CDINFO	AXCDSUP	201	B.3, Finding Software Patches on the World Wide Web
		203	Appendix C, Configuring DOS CD-ROM Support xvii, 135

Index Entries

<u>id</u>	<u>File</u>	<u>Page</u>	<u>References</u>
CPU	CH1A	1	(1) central processing unit (CPU)
ALU	CH1A	1	(1) arithmetic logic unit (ALU)
SMP	CH1A	3	(1) symmetric multiprocessing (SMP) 3
DRAM	CH1A	3	(1) dynamic random access memory (DRAM) 3
SRAM	CH1A	3	(1) static random access memory (SRAM) 3, 3
RAM	CH1A	4	(1) Random Access Memory (RAM) 4, 4, 12
CACHE	CH1A	4	(1) cache memory 4, 4, 4, 4, 4, 4, 5, 6, 68, 168, 171
SIMM	CH1A	7	(1) Single In-line Memory Modules (SIMM) 7, 7
ECC	CH1A	9	(1) Error Correcting Code (ECC) 9, 42
ECCP	CH1A	9	(1) Error Correcting Code Parity (ECC-P) 9
EOS	CH1A	9	(1) ECC Memory on SIMMs (EOS) 9
NVRAM	CH1A	10	(1) Non-Volatile Random Access Memory (NVRAM) 10
ISA	CH1A	12	(1) Industry Standard Architecture (ISA) 12
EISA	CH1A	12	(1) Extended Industry Standard Architecture (EISA) 12, 65, 74
MCA	CH1A	12	(1) Micro-Channel Architecture (MCA) 12, 65, 82, 85
PCI	CH1A	12	(1) Peripheral Component Interconnect (PCI) 12
BUSES	CH1A	13	(1) bus architectures 13
I004	CH1B	16	(1) hard disk interfaces 16, 16, 16, 16, 16, 168, 170
DASD	CH1B	16	(1) Direct Access Storage Device (DASD) 16, 21, 21, 21, 21, 22, 22, 22, 24, 28, 152, 168, 176, 176, 179, 182
ESDI	CH1B	16	(1) Enhanced Small Device Interface (ESDI) 16
IDE	CH1B	16	(1) Integrated Drive Electronics (IDE) 16, 31, 68, 70
SCSI	CH1B	16	(1) Small Computer System Interface (SCSI) 16, 17, 17, 17, 18, 18, 18, 18, 18, 18, 19, 19, 19, 19, 20, 20, 20, 65, 78, 79, 79, 80, 80, 80, 80, 90
TCQ	CH1B	18	(1) Tagged Command Queuing (TCQ) 18
SCB	CH1B	19	(1) Subsystem Control Block (SCB) 19
DISKC	CH1B	22	(1) caching, on disk subsystems 179
RAID	CH1B	23	(1) Redundant Array of Inexpensive Disks (RAID) 23, 24, 24, 24, 25, 25, 27, 28, 28, 29, 29, 30, 30, 30, 30, 30, 42, 65, 101, 102, 102, 102, 106, 106, 112, 112, 112, 151, 171
SEC	CH1D	34	(1) security

			71, 71, 72, 92, 92, 94, 94
DMI	CH1E	37	(1) Desktop Management Interface (DMI) 37, 38, 38, 38, 38, 39, 39, 44
MIF	CH1E	38	(1) Management Information Format (MIF)
MI	CH1E	39	(1) Desktop Management Interface (DMI) (2) Management Interface (MI) 39
MIC	CH1E	39	(1) Desktop Management Interface (DMI) (2) MI commands 39, 39, 39
SNMP	CH1E	39	(1) Simple Network Management Protocol (SNMP) 39, 40, 40, 40, 40, 40
TCPIP	CH1E	39	(1) Transmission Control Protocol/Internet Protocol (TCP/IP) 39
MIB	CH1E	40	(1) Management Information Base (MIB)
NETF	CH1E	41	(1) NetFinity 41, 41
NETFMGR	CH1E	41	(1) NetFinity (2) Manager 43
SYSVIEW	CH1E	45	(1) SystemView
UPS	CH1F	52	(1) Uninterruptible Power Supply (UPS) 53
SG	SGOVER	115	(1) ServerGuide 115, 115, 117, 117, 118
TREE	4576CH4	126	(1) directory tree, NDS 141
CXT	4576CH4	126	(1) context, NetWare 141
NDS	4576CH4	126	(1) NetWare Directory Services (NDS) 126, 141

Tables

<u>id</u>	<u>File</u>	<u>Page</u>	<u>References</u>
ECCPERF	CH1A	12	1 12
MEMOP	CH1A	12	2 12
ARRYTYP	CH1B	23	6 23
MOD300	CH2A	56	8
MOD310	CH2A	57	9
MOD320E	CH2A	58	10
MOD320M	CH2A	59	11
MOD500	CH2A	60	12
MOD520E	CH2A	61	13
MOD520M	CH2A	62	14
MOD720	CH2A	63	15
TERMIN	CH3	79	16

Processing Options

Runtime values:

Document fileid	EG244576 SCRIPT
Document type	USERDOC
Document style	IBMXAGD
Profile	EDFPRF30
Service Level	0029
SCRIPT/VS Release	4.0.0
Date	95.12.22
Time	10:25:44
Device	3820A
Number of Passes	4
Index	YES
SYSVAR D	YES
SYSVAR G	INLINE
SYSVAR V	ITSCEVAL

Formatting values used:

Annotation	NO
Cross reference listing	YES
Cross reference head prefix only	NO
Dialog	LABEL
Duplex	YES
DVCF conditions file	(none)
DVCF value 1	(none)
DVCF value 2	(none)
DVCF value 3	(none)
DVCF value 4	(none)
DVCF value 5	(none)
DVCF value 6	(none)
DVCF value 7	(none)
DVCF value 8	(none)
DVCF value 9	(none)
Explode	NO
Figure list on new page	YES
Figure/table number separation	YES
Folio-by-chapter	NO
Head 0 body text	Part
Head 1 body text	Chapter
Head 1 appendix text	Appendix
Hyphenation	NO
Justification	NO
Language	ENGL
Layout	OFF
Leader dots	YES
Master index	(none)
Partial TOC (maximum level)	4
Partial TOC (new page after)	INLINE
Print example id's	NO
Print cross reference page numbers	YES
Process value	(none)
Punctuation move characters	,
Read cross-reference file	(none)
Running heading/footering rule	NONE
Show index entries	NO
Table of Contents (maximum level)	3
Table list on new page	YES
Title page (draft) alignment	RIGHT
Write cross-reference file	(none)

Imbed Trace

Page 0	SG244576
Page 0	4576SU
Page 0	4576VARS
Page 1	4576FM
Page i	4576EDNO
Page ii	4576ABST
Page xv	4576SPEC
Page xv	4576TMKS
Page xvi	4576PREF
Page xix	4576ACKS
Page xx	CH1A
Page 16	CH1B
Page 31	CH1C
Page 34	CH1D
Page 37	CH1E
Page 47	SFTCH1
Page 51	CH1F
Page 54	CH2A
Page 63	CH3
Page 113	4576CH4
Page 113	SGOVER
Page 143	NETF
Page 151	RAIDADM
Page 165	4576CH5
Page 188	AX1
Page 197	AX3
Page 201	AXCDSUP
Page 204	4576ABRV
Page 212	4576EVAL
Page 212	RCFADDR
Page 212	ITSCADDR FILE
Page 213	RCFADDR
Page 213	ITSCADDR FILE