


SG24-4525-00

**IBM PC Server
Disk Subsystem Configuration and Sizing**

October 1995



IBM

International Technical Support Organization
This soft copy for use by IBM employees only.
Boca Raton Center



SG24-4525-00

**IBM PC Server
Disk Subsystem Configuration and Sizing**

October 1995

This soft copy for use by IBM employees only.

Take Note!

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

First Edition (October 1995)

This edition applies to IBM's PC Server hardware and software products currently announced at the date of publication.

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. JLPC Building 014 Internal Zip 5220
1000 NW 51st Street
Boca Raton, Florida 33431-1328

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.

This soft copy for use by IBM employees only.

Abstract

This document provides detailed information on how to select the right server system for your business needs and how to optimally configure the disk subsystem in IBM PC Server products. The document provides example configurations as well as reference material on the IBM PC Server disk subsystem.

This document is intended for IBM customers, dealers, system engineers and consultants who need to know the capabilities of the IBM PC Server subsystem.

(225 pages)

This soft copy for use by IBM employees only.

This soft copy for use by IBM employees only.

Contents

Abstract	iii
Figures	ix
Tables	xi
Special Notices	xiii
Preface	xvii
How This Document Is Organized	xvii
International Technical Support Organization Publications	xviii
ITSO Redbooks on the World Wide Web (WWW)	xviii
Acknowledgments	xix
Chapter 1. IBM PC Server Overview	1
1.1 IBM PC Server Positioning	2
1.1.1 Server Types	3
1.1.2 IBM Server Positioning and Selection Guidelines	8
1.2 The IBM PC Server 300 (EISA/PCI)	10
1.3 The IBM PC Server 310 (EISA/PCI or MCA/PCI)	13
1.4 The IBM PC Server 320 (EISA/PCI)	16
1.5 The IBM PC Server 320 (Micro Channel/PCI)	21
1.6 The IBM PC Server 500	23
1.7 The IBM PC Server 520 (EISA/PCI or MCA/PCI)	28
1.8 The IBM PC Server 720	34
1.9 IBM PC Server Summary	40
Chapter 2. IBM PC Server Disk Subsystem Options	45
2.1 Disk Subsystem Controllers Overview	45
2.1.1 IBM PC Server SCSI-2 Fast/Wide Adapter/A	45
2.1.2 IBM PC Server SCSI-2 Fast/Wide PCI Adapter	48
2.1.3 IBM Streaming RAID Adapter/A	50
2.1.4 IBM PC Server SCSI-2 Fast/Wide PCI-Bus RAID Adapter	52
2.1.5 RAID Adapters Features	53
2.1.6 SCSI-2 Fast Adapters	56
2.1.7 Summary of SCSI Adapters	56
2.2 Hard Disk Drives	59
2.2.1 The Ultrastar Family	59
2.2.2 The Deskstar Family	62
2.3 External Disk Subsystems	64

This soft copy for use by IBM employees only.

2.3.1	3516 Hot Swap Expansion Enclosure	64
2.3.2	3511 External Enclosure for SCSI Devices	67
2.3.3	3510 SCSI Storage Enclosure	69
2.3.4	SCSI Storage Enclosure Summary	71
2.4	Other IBM Options for the IBM PC Server	71
2.4.1	CD-ROM Drives	71
2.4.2	Tape Drives	75
Chapter 3. Disk Subsystem - Setup and Configuration		83
3.1	Internal SCSI Cabling	83
3.1.1	PC Server 300 Internal Cabling	83
3.1.2	PC Server 310 MCA or ISA Models Internal Cabling	86
3.1.3	PC Server 320 Internal Cabling	86
3.1.4	PC Server 500, 520 and 720 Internal Cabling	94
3.1.5	Internal Bays/Cabling Upgrades	108
3.1.6	Upgrade Trays and Hot Swap Bay Options	109
3.2	SCSI External Cabling	111
3.2.1	SCSI Adapter Connectors	111
3.2.2	Card to Option Cables	113
3.2.3	Option to Option Cables and Terminators	114
3.2.4	PC Server 300 and 310 External Cabling	119
3.2.5	PC Server 320, 320 Array, 520, 520 Array and 720 Non-Array External Cabling	120
Chapter 4. Disk Subsystem Sizing and Performance		127
4.1	Disk Subsystem Sizing	127
4.1.1	Introduction	127
4.1.2	Configuration Considerations	127
4.1.3	Summary	149
4.2	Disk Subsystem Performance	156
4.2.1	RAID Performance	156
4.2.2	Hardware versus Software RAID	159
4.2.3	Number of Drives	161
4.2.4	SCSI Channel Bandwidth	163
4.2.5	Stripe Size	165
4.2.6	RAM Disk Cache	168
4.2.7	Hardware Disk Cache	168
4.2.8	RAID-5 Degraded	171
Chapter 5. IBM PC Server High Availability Examples		173
5.1	NetWare SFT III by Novell	174
5.1.1	SFT III Components	175
5.1.2	Server Failure Handling	178

This soft copy for use by IBM employees only.

5.2 StandbyServer 32 by Vinca	179
5.2.1 Benefits	180
5.2.2 System Configuration	181
5.2.3 Installation Overview	182
5.2.4 Hard Disk Configuration Requirements	185
5.3 Dual Servers for High Availability	192
5.3.1 Objectives	192
5.3.2 Limitations	194
5.3.3 System Prerequisites	194
5.3.4 Design Specification	195
5.3.5 Setup Procedure	195
5.3.6 SCSI Switch	196
5.3.7 Configuration Examples	199
Appendix A. Part Number Equivalences for Europe	207
Appendix B. Description - Test Environment	211
B.1 Test Environment Used for the IBM PC Server 320	211
B.2 Test Environment Used for the IBM PC Server 500	212
Glossary	215
List of Abbreviations	217
Index	219

This soft copy for use by IBM employees only.

This soft copy for use by IBM employees only.

Figures

1.	Server Types and Disk Subsystem Load	7
2.	IBM PC Server Positioning and Selection	10
3.	IBM PC Server 300	11
4.	IBM PC Server 500 - Mechanical Design	25
5.	IBM PC Server 500 and 720	35
6.	SCSI-2 Fast/Wide Adapter/A	47
7.	SCSI-2 Fast/Wide PCI Adapter	49
8.	SCSI-2 Fast/Wide RAID Adapter Micro Channel	51
9.	SCSI-2 Fast/Wide RAID PCI-Bus Adapter	53
10.	RAID Advanced Functions Configuration Screen	54
11.	Ultrastar Hard Disk Drive Mechanical Design	60
12.	Deskstar Hard Disk Drive Mechanical Design	62
13.	3516 Hot Swap SCSI Enclosure	65
14.	3510 SCSI Enclosure	69
15.	Internal SCSI 2X CD-ROM Drive	72
16.	4/10GB 4mm DAT Tape Drive	79
17.	PC Server 300 Internal SCSI Cables	85
18.	PC Server 320 PCI/EISA Non Hot Swap Internal SCSI	89
19.	PC Server 320 PCI/MCA Non Hot Swap Internal SCSI	90
20.	IBM PC Server 320 Array Mechanical Design	92
21.	PC Server 320 Array or 320 with Hot Swap (SCSI-2 F/W PCI Adapter or SCSI-2 F/W PCI RAID Adapter)	93
22.	PC Server 500, 520, and 720 Internal Bay Structure	96
23.	PC Server 720 and 720 Array Mechanical Design	97
24.	IBM PC Server 500 - Mechanical Design	98
25.	PC Server 500/720 Array Models — SCSI-2 F/W RAID Adapter/A Connections	99
26.	PC Server 500/720 Array Models — Two SCSI-2 F/W RAID Adapters/A Connections to Three Backplanes.	100
27.	IBM PC Server 520 EISA/PCI - Mechanical Design	102
28.	PC Server 500 Non-Array with SCSI-2 F/W Adapter Connections	104
29.	IBM PC Server 520 MCA/PCI - Mechanical Design	107
30.	PC Server 720 Non-Array with SCSI-2 F/W PCI Adapter Connections	108
31.	Internal Cabling Upgrade Options Summary	109
32.	Centronics 50-pin Connector	111
33.	High Density PCI Controller SCSI Connectors	112
34.	IBM SCSI I/O External Interface Connectors on Micro Channel Adapters	113
35.	External Card to Option Cables Summary	114
36.	Option to Option Cables and Terminator Summary	115

This soft copy for use by IBM employees only.

37.	PC Server 300 External Cabling Example (PCI SCSI-2 Fast)	120
38.	PCI SCSI-2 Fast/Wide External SCSI (8-bit External Device)	121
39.	PCI SCSI-2 Fast/Wide External SCSI (16-bit External device)	122
40.	External 16 and 8-Bit SCSI Configuration	125
41.	External 8-Bit SCSI Configuration	126
42.	Disk Subsystem Cost for 4.5GB Total Capacity	132
43.	Disk Subsystem Cost (1.12GB, 2.25GB and 4.5GB Drives)	133
44.	IBM PC Server 500 with a Single 4.5GB Drive versus Four 1.12GB Drives	134
45.	IBM PC Server 500 with RAID-0 versus NetWare Multidisk Volume	138
46.	RAID-0, Disk Striping	141
47.	RAID-1, Disk Mirroring	143
48.	RAID-5, Striping Plus Parity Blocks	145
49.	Software versus Hardware Mirroring	147
50.	Disk Subsystem Cost (US Dollars/MB) for RAID-0, 1 and 5	149
51.	Disk Selection Guide	152
52.	Cost Software RAID-1 versus Hardware RAID-5	153
53.	Performance Software RAID-1 versus Hardware RAID-5	154
54.	IBM PC Server 500, RAID-0, 1 and 5 Throughput	158
55.	Software versus Hardware Mirroring	160
56.	IBM PC Server 320, Throughput versus Number of Drives	162
57.	Disk Arrays on Separate SCSI Channels	164
58.	Disk Array Spanning Two SCSI Channels	165
59.	IBM PC Server 320, RAID Controller Stripe Unit Size RAID-0	167
60.	IBM PC Server 320, Write-Back versus Write-Through	170
61.	IBM PC Server 320, RAID-5 with One Defunctional Drive	172
62.	SFT III Concept	175
63.	SFT III IO- and MSEngines	177
64.	StandbyServer 32 for LAN Server Concept	184
65.	StandbyServer 32 for LAN Server Minimum Configuration	186
66.	Drive Assignments in Both the Primary Server and in the Standby Server	189
67.	Switch Over to the Standby Machine	190
68.	Partition Assignment in Both Primary Server and Standby Machine	191
69.	APPLIED Concepts (4x2 Electronic Crosspoint Switch)	197
70.	APPLIED Concepts (4x2 Electronic Crosspoint Switch)	198
71.	APPLIED Concepts (Backing up Multiple Servers)	199
72.	High Availability Server (Minimum Configuration)	200
73.	High Availability Server (External RAID Tower)	201
74.	High Availability Server (External Mirrored RAID-1 Towers)	202
75.	High Availability Server (External Duplexed RAID-1 Towers)	203
76.	High Availability Server (NetFinity Implementation)	204
77.	Sample Cabling for two CPU's	205

This soft copy for use by IBM employees only.

Tables

1.	Technical Characteristics of the PC Server 300	12
2.	Technical Characteristics of the PC Server 310	14
3.	Technical Characteristics of the PC Server 320 EISA/PCI	17
4.	Technical Characteristics of the PC Server 320 EISA/PCI Entry Models	19
5.	Technical Characteristics of the PC Server 320 MCA	21
6.	Technical Characteristics of the PC Server 500	26
7.	Technical Characteristics of the PC Server 520 PCI/EISA	30
8.	Technical Characteristics of the PC Server 520 PCI/MCA	32
9.	Technical Characteristics of the PC Server 720	37
10.	IBM PC Server Models Overview	40
11.	Cable Lengths	47
12.	IBM SCSI-2 Fast Adapter Summary	56
13.	IBM SCSI Fast/Wide Adapter Summary	57
14.	Technical Characteristics of UltraStar Disk	61
15.	Technical Characteristics of Deskstar Disk	63
16.	Key Benefits of 3516 Hot Swap Expansion Enclosure	66
17.	Key Benefits of 3511 External Enclosure for SCSI Devices	67
18.	Key Benefits of 3510 SCSI Storage Enclosure	70
19.	IBM SCSI External Enclosure Models Summary	71
20.	CD-ROM Feature Summary	74
21.	CD-ROM 2X and 4X Comparison Summary	74
22.	IBM Tape Options - Summary	82
23.	Options Installation Considerations for PC Server 300	84
24.	Options Installation Considerations for PC Server 320 Non Array	87
25.	Options Installation Considerations for PC Server 320 Array	91
26.	Options Installation Considerations for PC Server 500, 520 and 720	95
27.	Trays for Disk Drives and Bay Options for Hot Swap Upgrade	110
28.	IBM External Micro Channel SCSI Cabling Summary - (IBM 50 or 68 pin external connector)	116
29.	IBM External PCI SCSI Cabling Summary - (50 or 68 pin high density external connector)	117
30.	US, EMEA Part Number Equivalences	207

This soft copy for use by IBM employees only.

This soft copy for use by IBM employees only.

Special Notices

This document is intended to help IBM customers, dealers, systems engineers and consultants to understand the disk subsystem used in IBM PC Server products. The information in this publication is not intended as the specification of any programming interface that is provided by OS/2, OS/2 LAN Server, NetWare or any other product mentioned in this publication. See the PUBLICATIONS section of the IBM Programming Announcement for these products for more information about what publications are considered to be program documentation.

The information in this publication is not intended as the specification of any programming interfaces that are provided by OS/2, OS/2 LAN Server, NetWare or any other products mentioned in this publication. See the PUBLICATIONS section of the IBM Programming Announcement for these products for more information about what publications are considered to be product documentation.

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, 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

This soft copy for use by IBM employees only.

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.

The following document contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples contain the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

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	AIXwindows
AS/400	DB2/2
Distributed Application Environment	DualStor
IBM	LANStreamer
Micro Channel	Operating System/2
OS/2	Personal Computer AT
Personal Computer XT	Personal System/2
PS/2	S/370

The following terms are trademarks of other companies:

Windows is a trademark of Microsoft Corporation.

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.

This soft copy for use by IBM employees only.

C-bus is a trademark of Corollary, Inc.

Adaptec is a trademark of Adaptec Incorporated

CorelSCSI is a trademark of Corel Systems Corporation

Digital Data Storage is a trademark of Sony Corporation

i960 is a trade mark of Intel Corporation

Lotus is a trademark of Lotus Development Corporation

Macintosh is a trademark of Apple Computer, Incorporated

NetWare is a trademark of Novell, Incorporated

Novell is a trademark of Novell, Incorporated

PowerChute is a trademark of American Power Conversion

SFT is a trademark of Novell, Incorporated

Seagate is a trademark of Seagate Technology, Incorporated

Windows NT is a trademark of Microsoft Corporation

486 is a trademark of Intel Corporation

80386 is a trademark of Intel Corporation

ARCserve is a trade mark of Cheyenne Software Incorporated

Corel is a trademark of Corel Systems Corporation

Intel is a trademark of Intel Corporation

Kodak is a trademark of Eastman Kodak Company

Lotus Notes is a trademark of Lotus Development Corporation

Microsoft is a trademark of Microsoft Corporation

NetWare is a trademark of Novell, Incorporated

Pentium is a trademark of Intel Corporation

SCO is a trademark of The Santa Cruz Operation, Incorporated

Sytron is a trademark of Sytron, Incorporated

UnixWare is a trademark of Novell, Incorporated

X/Open is a trademark of X/Open Company Limited

486 is a trademark of Intel Corporation

Other trademarks are trademarks of their respective companies.

This soft copy for use by IBM employees only.

This soft copy for use by IBM employees only.

Preface

This document is intended to assist IBM customers, dealers, systems engineers and consultants in selecting the most suitable IBM PC server and especially how to configure this server's most critical components, the disk subsystem.

How This Document Is Organized

The document is organized as follows:

- Chapter 1, "IBM PC Server Overview"
This chapter explains the different available server types (file server, multimedia server, print server, communication server and database/application server). It also provides a positioning and an overview of the IBM PC Server product family as well as guidance in selecting the right server systems for your particular business needs.
- Chapter 2, "IBM PC Server Disk Subsystem Options"
This chapter provides an overview of available IBM PC Server disk subsystem option.
- Chapter 3, "Disk Subsystem - Setup and Configuration"
This chapter provides disk subsystem setup and configuration examples.
- Chapter 4, "Disk Subsystem Sizing and Performance"
This chapter is divided into two parts. The first part includes considerations on how to define the most appropriate disk subsystem configuration. The second part contains disk subsystem performance considerations, that will assist you in optimizing the disk subsystem performance.
- Chapter 5, "IBM PC Server High Availability Examples"
This chapter explains some sample solutions that will provide high availability in a LAN environment where access to a server is critical.
- Appendix A, "Part Number Equivalences for Europe"
This appendix provides a conversion table from US disk subsystem options into EMEA part numbers.
- Appendix B, "Description - Test Environment"

This soft copy for use by IBM employees only.

This appendix provides a description of the test environment used for Chapter 4, "Disk Subsystem Sizing and Performance" on page 127.

International Technical Support Organization Publications

- *IBM Personal Computer Disk Subsystem Considerations*, SG24-2510-00
- *Advanced PS/2 Servers Planning and Selection Guide*, GG24-3927-01

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 1-800-284-4721. Visa and MasterCard 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>

This soft copy for use by IBM employees only.

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:

Michael Koerner
International Technical Support Organization, Boca Raton Center

The authors of this document are:

Michael Koerner
International Technical Support Organization, Boca Raton Center

Nelson Pesce
IBM Uruguay

Michael Ringel
IBM Germany

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

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

Erwin van de Bos
IBM Netherlands

Sam Cheung
IBM Raleigh

Phil Horwitz
IBM Raleigh

David Laubscher
IBM Raleigh

John Michie
IBM Canada

This soft copy for use by IBM employees only.

Frank Morassuti
IBM Canada

Howard Sykes
IBM Raleigh

This soft copy for use by IBM employees only.

Chapter 1. IBM PC Server Overview

Over the last years, PC based server systems became more and more powerful, and a considerable increasing number of customers implemented business critical solutions using PC server and workstation based local area networks. In these networks a faulty server system may seriously impact an enterprise's operation and cause tremendous costs to recover from this failure situation and to regain normal operation. Thus, PC based server systems must provide a similar level of reliability, availability and serviceability as already provided by mainframes (IBM /390) and other systems (for example, the IBM AS/400, very commonly used in office environments and the IBM RS/6000 system family).

As functions like symmetric multiprocessing (SMP), ECC memory, powerful and reliable disk subsystems, and network and system management concepts become commonplace for PC servers, the task to select an appropriate server as well as to configure this server according to a customer's environment and requirements, becomes more and more complex and difficult.

To assist you in selecting the most suitable server and especially how to configure this server's most critical component, the disk subsystem, this book will provide you with the following:

- IBM PC Server family product overview and positioning
- IBM PC Server internal and external disk subsystems
- Disk subsystem cabling and configuration examples
- Disk subsystem sizing and performance information
- IBM PC Server high availability considerations
 - SFT III
 - Vinca
 - Standby server and SCSI-switch

1.1 IBM PC Server Positioning

To assist you in selecting the right server for your business, the following section will provide guidelines and hints to select the most suited server to meet your business needs. The members of the IBM PC Server family have been designed to cover a wide variety of requirements.

For server selection, often the first criteria used is the number of users the server should be capable to support. Although this is a very important factor during the server selection process, there are several other factors to consider. For example, in an application development or image processing environment even a small work group may have the need for a much more powerful server than an entire department or even an enterprise using a server as file and program repository only.

The following is a list of criteria you should take into consideration when selecting a server:

- Number of users attached to a server
- Purpose (file, print, database or communication server)
- Transaction frequency and size
- Amount of data to be stored on the server
- Application programs to be used
- Network operating system
- Reliability/fault tolerance
- System management and support
- Cost

Another important factor in the process of server selection is the fact that servers can be used for different purposes and dependent on the server type, different components are critical for maximum performance. Knowing about the potential bottlenecks will help you to select the right server and even combine several server types in a single multi-functional server, which in turn helps you to implement a cost effective solution.

This soft copy for use by IBM employees only.

1.1.1 Server Types

The following paragraphs provide a short overview and description about the various server types and their purpose. This will assist you in understanding and identifying potential bottlenecks as well as which server types may be combined in a single multi-function server.

Server types:

- *File Server*

This is currently the most frequently used server type. A file server holds data and program files on its hard disks and lets network attached workstations access these files as they would reside in each individual workstation to load programs and data files. In addition, both the programs, and especially the data can be shared among users to provide fast and efficient information flow and exchange.

The advantages are obvious:

- Shareable information for effective data access and interchange
- Data integrity:
Only a single copy may be accessed and used by all users, not one different version of the same file on each individual computer.
- Centralized data and program maintenance. An application program needs to be installed only once and can be shared by as many users as the license agreement permits.
- Data security and controlled access:
Access rights define the programs and data that can be used by an individual user to ensure data security and access control (license protection).
- Cost reduction:
Large capacity hard disks are required for the server system only. The individual workstations do not need any removable media such as diskette drives and hard disks at all (medialess workstations). The operating system can also be loaded over the network, or relatively small hard disks can be used only to hold the operating system and some local programs and data. This can lead to considerable cost reductions if a server will have to support large number of I/Os.

- *Multimedia Server*

Multimedia servers are basically file servers. However, due to the very large file sizes of, for example, digitized images, audio or video data,

This soft copy for use by IBM employees only.

they require powerful disk and LAN subsystems combined with a fast and flexible I/O bus to achieve a sufficient data transfer rate to the attached multimedia workstations.

- *Print Server*

The print server provides central printing capability to all attached workstations without the need for a local workstation printer. Print output created in a workstation will be redirected to the print server, spooled to a print queue and then printed on a server attached printer. No printer at all or only a low cost printer is required at the workstation itself. High quality or high capacity printers are only required for the print server. As different printers may be attached to the print server, you gain flexibility in directing your print output to the most appropriate printer for a particular print job.

Frequently, both the file and print server function are combined in a single server system. Both, the printer being connected to a central location and the combination of both server types, will further reduce your cost.

- *Communications Server*

The communications server provides connectivity and communication beyond the domain of a LAN. The most frequent implementations are:

- Remote LAN access:

The communication server usually maintains a modem pool, allowing multiple remote users to dial-in and access LAN resources as they would be directly connected to the LAN. This enables location-independent access to data and provides a new dimension of flexibility in accessing centrally kept and managed resources from either desktop or mobile workstations.

- Online services access:

Using a modem pool the same as remote LAN access, the LAN users can share the modems connected at the server to dial-out and access to popular online services, Bulletin Board Systems (BBS) or Internet.

- Mainframe access:

Here the communication server acts as a gateway to a mainframe oriented network (for example SNA networks) and permits users to access and use also host data and applications. The host access is usually accomplished by using terminal emulation programs (for example, 3270 terminal emulation in the workstations). The

This soft copy for use by IBM employees only.

communication server then emulates a display control unit and provides the communication path to a host based system.

– Facsimile (FAX) server:

A FAX server adds FAX capability to its LAN attached workstations. It acts as a concentrator who receives and stores FAX messages, for later retrieval and processing by LAN attached workstations. It also stores FAX messages created by its LAN attached workstations, and sends them independently from the workstation status at a predefined date and time (for example, at night, when there is little network traffic and lower rates are charged). In either case, a workstation does not need to be operational at the time a FAX message is sent or received.

The communication server function is usually Central Processing Unit (CPU) bound and does not use the disk subsystem very much, except for loading the communication server program when starting such a system or in the case of FAX spooling to disk for later transmission. Afterwards, the code is memory resident and no further significant disk access is required. However, fast communication networks like T1 lines¹ or Integrated Services Digital Network (ISDN) may require a fast communication adapter and an I/O bus delivering sufficient bandwidth for data transfer.

• *Database and Application Server*

This type of server is used when *in addition to the network operating system* application program code is executed by the servers processor(s). Typically, the application program is split between the server and the workstation or client. The workstation (or client) runs a Graphical User Interface (GUI) for end user presentation services and the server itself handles the back end tasks that are CPU, disk I/O, and memory intensive.

Such applications are also called Client/Server or distributed applications. As this is a new and emerging way to optimize process execution in network environments, it also creates new requirements for server systems.

Typical applications in this area are:

- Groupware applications such as Lotus Notes
- Large database applications as like IBM DB2/2

¹ 1.544 Mbps transmission rate lines that contains 24 circuits

This soft copy for use by IBM employees only.

Servers used in such a distributed application environment need to deliver maximum performance and reliability by all their major components and functions.

The load created by sophisticated operating and network operating systems, combined with the load created by large and complex applications, exceed more and more the capabilities of a single processor chip. They also require large and fast memory.

Large databases require high capacity, very fast, and fault tolerant disk subsystems for effective data storage and retrieval.

The increasing number of attached workstations also require a fast and flexible LAN subsystem in order to take full benefit of the server's functionality and power.

Finally, the I/O bus used to connect all the components such as CPU, memory, disk, and LAN subsystems needs to be fast, reliable and flexible to avoid bottlenecks, even if the individual subsystems are very fast.

Servers designated as application servers in a distributed application (client/server) environment must be carefully selected. Combining an application server's function with other server types, for example a communication server, has to be carefully evaluated to avoid potential bottlenecks negatively impacting a server's performance.

Summary Server Types:

As you have learned from the previous paragraphs, there are different server types, each for a particular purpose. However, to reduce your investment in a server system, you may combine several server functions that don't contend for the same subsystem into one system, and may still offer good overall performance on both server functions.

Figure 1 on page 7 provides assistance in identifying which server types may be combined in a single system. It shows the server types and the typical load to each of its subsystems. When, for example, selecting a file server, the disk subsystem is identified as the most critical component, both from a performance and capacity point of view. On the other hand, there is almost no load at all to the printer port. Therefore, you may decide to combine both the file and the print server function into a single system. In addition, if your processor subsystem provides sufficient computing power, and your I/O bus bandwidth as well as expandability in terms of number of expansion slots, you may even combine a third server function in your system, that is a communication server.

This soft copy for use by IBM employees only.

Please be aware, when using your server in a real distributed application environment, you carefully need to evaluate the load created to this server. If the application running on this server is very much CPU bound, then you may add the function of a file server. If on the other hand, the function is more disk I/O oriented, you need to make sure that there will be no bottleneck when two functions, that of a database server and that of a file server, compete for the same resources.

To ease this decision, IBM offers server systems, providing sufficient computing power for the most demanding applications as well as multiple I/O bus systems (Micro Channel, Enhanced Industry Standard Architecture (EISA) and Peripheral Component Interconnect (PCI)). Therefore, you may even split your disk subsystem, so that one set of drives and adapter handles the database and another disk subsystem handles the file I/O requests, thus eliminating competition for the same resources which could lead under heavy load to a performance degradation of the affected subsystems.

Subsystem Performance Requirements					
Server type	Disk Subsystem	I/O Bus	LAN I/O	Processor	Printer Port
File server	very high	high	high	low	none
Communi- cations server	very low	med./high	med./high	high	none
Print server	medium	medium	low to high	very low	med./high
Application server <small>(Distributed Appls.)</small>	very high	medium to high	medium to high	very high	none

Figure 1. Server Types and Disk Subsystem Load

This soft copy for use by IBM employees only.

1.1.2 IBM Server Positioning and Selection Guidelines

The following paragraph will help you in selecting a server product from the IBM PC Server family that will meet your business needs.

The key questions during the process of server selection are:

- *What will be the purpose of this server?*

This will lead you to the type of server you need, and whether you may need to combine several server types into a single system.

- *Where do I want to use this server, at a central site or remotely, for example in a branch office?*

This also implies the question about support and service in case of a problem. If there is no onsite support available, you need to think about remote maintenance facilities, as well as investing in server reliability features such as an Uninterrupted Power Supply (UPS), Error Correction Code (ECC) memory, and fault tolerant disk subsystems. Even when your initial investment may seem to be a bit high, please remember the loss on productivity and time in case of an (avoidable) server crash. The costs of a single crash may, by far, exceed your initial investment to prevent such events.

- *How many users are to be supported by this server?*

Here you decide, what this server will support, a small work group, a department or small enterprise or is it to be used for a large enterprise.

This helps you to select the basic server model. It also assists you in sizing your disk and LAN subsystem.

- *The disk subsystem*, because usually an increasing number of users also require an increasing amount of disk space to hold the user data.
- *The LAN subsystem*, because too many users connected to one LAN adapter (LAN segment) could result in performance degradation caused by network congestion. You may install multiple LAN adapters, divide your LAN in segments and therefore avoid performance degradation due to overloaded network segments.

- *What will be the average transaction frequency and size of them?*

The transaction frequency and size will provide you with disk subsystem performance requirements, for example the number of drives, disk controller, and RAID level to be used. This will help you to select the appropriate adapter, the size and number of hard disks. Chapter 4, "Disk Subsystem Sizing and Performance" on page 127 of this book will

This soft copy for use by IBM employees only.

provide detailed considerations to assist you in selection the best solution.

- *Are the data held by the server or the application mission critical, and how much server downtime can I effort?*

Here you have to decide what functions and components you need to include in your server configuration to meet your business requirements and to avoid severe impacts due to server downtime. Functions and components to provide maximum availability are, for example:

- UPS support to protect you against power failures
- ECC memory to avoid server crashes due to memory errors
- Fault tolerant disk subsystems to protect your programs and data
- Entire backup server solutions in case an entire server system fails

Functions, products and implementations, to achieve full server fault tolerance, will be described in Chapter 5, "IBM PC Server High Availability Examples" on page 173 of this book.

IBM Server Positioning:

Figure 2 on page 10 will assist you in selecting the most appropriate server dependent on the server type, such as file, print, communication and application server (shown on the X-axis). In addition, dependent on the environment the server is intended to support, from work group servers up to large enterprise servers (shown on the Y-axis). The overlap on the Y-axis is caused due to the fact that there is rarely a clear separation between the different environments. Under normal circumstances, the various environments overlap to a certain extent, where you may choose different server types equally serving your environment. However, if your budget does allow, it is recommended to choose the larger server system, because the demand for computing power and disk storage capacity tends to increase as business grows over time.

This soft copy for use by IBM employees only.




Figure 2. IBM PC Server Positioning and Selection

1.2 The IBM PC Server 300 (EISA/PCI)

The IBM PC Server 300 is designed and priced as an entry file server, to efficiently handle 5-50 or more clients, using primarily industry standard system and subsystem technology. While this server is positioned as an entry server, it contains many of the performance features, expansion capabilities, system and power management software of more expensive LAN servers. The IBM PC Server 300 (EISA/PCI) participates in the non-interconnected, stand alone file/print server market. This segment primarily consists of two major environments, small business and departmental LANs of large and medium enterprises.

This soft copy for use by IBM employees only.




Figure 3. IBM PC Server 300

Because the IBM PC Server 300 has two bus architectures, the Peripheral Component Interconnect (PCI) and Extended Industry Standard Architecture (EISA) buses, it offers an additional opportunity. Large and medium sized customers who have standardized on EISA servers for departmental LANs can now migrate to the new 132Mbps PCI architecture with an IBM server.

This soft copy for use by IBM employees only.

As an open architecture system, the PC Server 300 has been tested with numerous non-IBM adapters and devices to assure industry standard compatibility.

The PC Server 300 features full-tower packaging (shown in Figure 3 on page 11) the same case as the PC Server 320, refer to figure 1.4, "The IBM PC Server 320 (EISA/PCI)" on page 16. For further information about its advantages.

Table 1 shows the features of the different PC Server 300 models.

<i>Table 1 (Page 1 of 2). Technical Characteristics of the PC Server 300</i>				
Model Type Number	US, LA	8640-0P0	8640-0PT	N/A
	EMEA	8640-1P0	8640-1PT	8640-CPT
Processor	Type	Pentium	Pentium	Pentium
	Speed	60MHz	60MHz	60MHz
	Upgrade	ZIF Upgrade	ZIF Upgrade	ZIF Upgrade
	L2 Cache	256KB	256KB	256KB
System Bus	Type	EISA/PCI	EISA/PCI	EISA/PCI
	Speed (Mbps) EISA/PCI	33/132	33/132	33/132
Memory	Standard	16MB	16MB	16MB
	Maximum	192MB	192MB	192MB
	Parity	Industry Standard, 72-pin, 70ns	Industry Standard, 72-pin, 70ns	Industry Standard, 72-pin, 70ns
	EOS (ECC on SIMM)	Optional	Optional	Optional
Hard Disk Controller	Type	Optional	SCSI-2 Fast	SCSI-2 Fast
	I/O Bus	PCI/EISA	PCI	PCI
	SCSI-speed	N/A	10MBps	10MBps
	Channels	N/A	1 1	1 1
	Devices per Channel	N/A	7	7

This soft copy for use by IBM employees only.

<i>Table 1 (Page 2 of 2). Technical Characteristics of the PC Server 300</i>				
Model Type Number	US, LA	8640-0P0	8640-0PT	N/A
	EMEA	8640-1P0	8640-1PT	8640-CPT
Hard Drives	Standard	0	1GB	1GB
	Access Time	N/A	8.6ms	8.6ms
	Maximum	13.5GB	13.5GB	13.5GB
CD-ROM (SCSI)	Standard	No	No	Yes
	Speed	N/A	N/A	2X
Free Expansion Bays	3 ½"	2	1	1
	5 ¼" ½ H	6	6	5
	Hot Swap (1/3H)	Optional	Optional	Optional
Free Expansion Slots	PCI	2	1	1
	ISA/EISA	5	5	5
	PCI or EISA 2	1	1	1
Note:				
1 One shared channel. Seven 8-bit devices can be connected to the internal or external connector.				
2 This slot can be used either by a PCI or EISA adapter.				

1.3 The IBM PC Server 310 (EISA/PCI or MCA/PCI)

As an economical entry level server, and to protect your investment as you migrate to the PCI bus, look no further than the IBM PC Server 310.

In your choice of PCI-equipped Micro Channel and ISA models, the PC Server 310 mini-tower gives you the features you need to support up to twelve users at a very competitive price. For example, the PC Server 310 model 8639-MXT comes with a SCSI-2 Fast controller on the PCI bus as standard, along with three hard drives for a maximum capacity of more than 12.8GB of high-performance storage. Combine this capacity with the 75MHz Pentium processor and Ethernet adapter that are also standard and the PC Server 310 delivers all the power you need to support a small work group LAN.

Standard features include:

- Pentium 75MHz microprocessor

This soft copy for use by IBM employees only.

- 256KB L2 adaptive write-back cache
- PCI SCSI-2 Fast controller card
- Quad speed CD-ROM drive
- 1.44 diskette drive
- Ethernet LAN Card
- Local bus SVGA video with 1MB of video RAM
- ServerGuide (a CD library to ease installation of network operating systems/applications)
- NetFinity (comprehensive system management software)
- Mini tower packaging with five device storage bays and five card slots supported by a 200 watts switchable world-wide power supply

To make setup, installation and management as easy as possible, the PC Server 310 also comes with the ServerGuide CD-ROM library and a quad-speed CD-ROM drive. ServerGuide automatically installs and configures your selected network operating system, along with Netfinity, the network management software that works with IBM hard disks to alert you to drive problems before they occur as well as monitor and manage traffic, and more.

For economy, and simplicity of network setup and management, choose the IBM PC Server 310.

Table 2 provides an overview of the features of the PC Server 310 MCA and EISA models.

Model Type Number	US, LA	8639-0XT	8639-MXT
	EMEA	8639-0XT	8639-MXT
Processor	Type	Pentium	Pentium
	Speed	75MHz	75MHz
	Upgrade	ZIF Upgrade	ZIF Upgrade
	L2 Cache	256KB	256KB
System Bus	Type	ISA/PCI	MCA/PCI
	Speed (Mbps) ISA/PCI or MCA/PCI	8/132	20/132

This soft copy for use by IBM employees only.

<i>Table 2 (Page 2 of 2). Technical Characteristics of the PC Server 310</i>			
Model Type Number	US, LA	8639-0XT	8639-MXT
	EMEA	8639-0XT	8639-MXT
Memory	Standard	16MB	16MB
	Maximum	192MB	192MB
	Parity	Industry Standard, 72-pin, 70ns	Industry Standard, 72-pin, 70ns
	EOS (ECC on SIMM)	Optional	Optional
Hard Disk Controller	Type	SCSI-2 Fast	SCSI-2 Fast
	I/O Bus	Onboard PCI	Onboard PCI
	SCSI-speed	10MBps	10MBps
	Channel	1 1	1 1
	Devices per Channel	7	7
Hard Drives	Standard	1.08GB	1.08GB
	Access Time	10.5ms	10.5ms
	Maximum	12.8GB 2	12.8GB 2
CD-ROM (SCSI)	Standard	Yes	Yes
	Speed	4 X	4 X
Free Expansion Bays	3 1/2"	1	1
	5 1/4" 1/2H	1	1
	Hot Swap (1/3H)	No	No
Free Expansion Slots	PCI	1	1
	ISA	2	0
	MCA	0	2 3
Note:			
1 Seven devices can be connected to the internal or external connector			
2 Using two 5.31GB drives and one 2.25GB drive.			
3 These slots support 20MBps data transfer rate.			

1.4 The IBM PC Server 320 (EISA/PCI)

The IBM PC Server 320 is designed primarily as an entry enterprise LAN server that efficiently connects a large number of clients for file and database server applications. These models are the better way to increase the LAN server performance, function and fault tolerance at a very attractive price. Three models provide outstanding configuration flexibility to help customers tailor solutions for their LAN server needs.

Features:

- High performance obtained with single 90MHz Pentium and 256KB L2 write-back cache, or 90MHz Pentium symmetric multiprocessor (SMP) upgrade and 512KB write-back L2 cache option
- Advanced architecture featuring 132MBps PCI bus and either a SCSI-2 Fast/Wide PCI Adapter, or SCSI-2 Fast/Wide RAID PCI Adapter
- Large capacity (1.12GB, 2.25GB, 4.51GB), SCSI-2 Fast/Wide ultra-performance (7200 rpm) drives
- Two high speed NS16550 compatible serial ports and one super fast (2MBps) parallel port which supports devices using Enhanced Capabilities/Enhanced Parallel Port (ECP/EPP) protocols adhering to the IEEE 1284 standard
- Double speed CD-ROM drive
- ServerGuide, a CD library to ease installation of network operating systems/applications
- NetFinity, comprehensive system management software

The 90MHz Pentium processor subsystem with SMP upgrade ability, coupled to a PCI advanced bus and state of the art PCI SCSI-2 Fast/Wide or SCSI-2 Fast/Wide RAID controller, create a high performance LAN server platform to handle heavy file server applications or moderate database applications.

PC Server 320 features full-tower packaging with the following customer benefits:

- A generous amount of adapter slots and device bays to satisfy the storage and functional requirements of an entry enterprise LAN server. Nine device storage bays for a non hot swap system, or ten storage bays for a hot swap system are provided.
- Eight adapter card slots (three PCI and five EISA) are supported.

This soft copy for use by IBM employees only.

- Six hot swap hard drive bays available optionally or standard to minimize business disruption in the event of a hard drive failure.
- A cooling design, called FloThru. This cooling system helps keep internal electronic components of the PC Server 320 running cool to improve reliability. FloThru’s air flow design utilizes dual fans, one of which is temperature sensing.
- A 300 watts switchable world wide power supply.

The PC Server 320 is recommended to customers with large to medium sized businesses who have previously purchased EISA based systems. These new servers enable this customer set to migrate to the faster, industry standard, advanced PCI bus while being able to run their existing applications that may relay on EISA/ISA adapters.

The secondary market segment for PC Server 320 is small businesses that require a high degree of processor power for database applications. This segment will find the PC Server 320 price point within their budget and the function matched to their needs.

Table 3 provides an overview of the features of the PC Server 320 EISA models.

<i>Table 3 (Page 1 of 3). Technical Characteristics of the PC Server 320 EISA/PCI</i>					
Model Type Number	US, LA	8640-0Y0	8640-0YT	8640-1Y0	8640-1YT
	EMEA	8640-CY0	8640-CYT	N/A	8640-CYR
Processor	Type	Pentium	Pentium	Pentium	Pentium
	Speed	90MHz	90MHz	90MHz	90MHz
	L2 Cache (std./max.)	256/512KB	256/512KB	256/512KB	256/512KB
	Std. #	1	1	1	1
	Max. #	2	2	2	2
	Upgrade	Chip	Chip	Chip	Chip
System Bus	Type	EISA/PCI	EISA/PCI	EISA/PCI	EISA/PCI
	Speed (MBps) EISA/PCI	33/132	33/132	33/132	33/132

This soft copy for use by IBM employees only.

<i>Table 3 (Page 2 of 3). Technical Characteristics of the PC Server 320 EISA/PCI</i>					
Model Type Number	US, LA	8640-0Y0	8640-0YT	8640-1Y0	8640-1YT
	EMEA	8640-CY0	8640-CYT	N/A	8640-CYR
Memory	Standard	16MB	16MB	16MB	16MB
	Maximum	192MB	192MB	192MB	192MB
	Parity	Industry Standard, 72-pin, 70ns			
	EOS (ECC on SIMM)	Optional	Optional	Optional	Optional
Hard Disk Controller	Type	SCSI-2 F/W	SCSI-2 F/W	SCSI-2 F/W RAID	SCSI-2 F/W RAID
	I/O Bus	PCI	PCI	PCI	PCI
	SCSI-speed	20MBps	20MBps	20MBps	20MBps
	Channels	1 1	1 1	2 2	2 2
	Devices per Channel	15	15	7/15 3	7/15 3
Hard Drives	Standard	0	1.12GB	0MB	2 x 1.12GB
	Access Time	N/A	8.0ms	N/A	8.0ms
	Maximum	13.5GB	13.5GB	13.5MB	13.5GB
CD-ROM (SCSI)	Standard	Yes	Yes	Yes	Yes
	Speed	2X	2X	2X	2X
Free Expansion Bays	3 1/2"	2	1	2	2
	5 1/4" 1/2 H	5	5	0	0
	Hot Swap (1/3H)	Optional	Optional	6	4

This soft copy for use by IBM employees only.

<i>Table 3 (Page 3 of 3). Technical Characteristics of the PC Server 320 EISA/PCI</i>					
Model Type Number	US, LA	8640-0Y0	8640-0YT	8640-1Y0	8640-1YT
	EMEA	8640-CY0	8640-CYT	N/A	8640-CYR
Free Expansion Slots	PCI	1	1	1	1
	ISA/EISA	4	4	4	4
	PCI or EISA 4	1	1	1	1
<p>Note:</p> <p>1 One shared SCSI channel:</p> <ul style="list-style-type: none"> Both internal 8-bit and internal 16-bit connectors can be used or One internal connector (8 or 16-bit) and one external 16-bit connector can be used <p>2 Two independent SCSI channels:</p> <ul style="list-style-type: none"> One shared 16-bit or external SCSI channel One dedicated SCSI-channel with the internal 8-bit connector <p>3 Number of drives per SCSI-channel supported by the PCI RAID controller:</p> <ul style="list-style-type: none"> Up to 6 internal hot swap drives Up to 7 external hot swap drives when using the external SCSI Enclosure IBM 3516 Up to 15 external SCSI-devices when using non hot swap enclosures <p>4 Shared slot, can be used for a PCI or EISA adapter</p>					

Table 4 provides technical details for the PC Server 320 EISA/PCI entry models.

<i>Table 4 (Page 1 of 2). Technical Characteristics of the PC Server 320 EISA/PCI Entry Models</i>			
Model Type Number	US, LA	8640-0X0	8640-0XT
	EMEA	8640-0X0	8640-0XT
Processor	Type	Pentium	Pentium
	Speed	90MHz	90MHz
	L2 Cache (std./max.)	256/512KB	256/512KB
	Std. #	1	1
	Max. #	2	2
	Upgrade	Chip	Chip
System Bus	Type	EISA/PCI	EISA/PCI
	Speed (Mbps) EISA/PCI	33/132	33/132

This soft copy for use by IBM employees only.

<i>Table 4 (Page 2 of 2). Technical Characteristics of the PC Server 320 EISA/PCI Entry Models</i>			
Model Type Number	US, LA	8640-0X0	8640-0XT
	EMEA	8640-0X0	8640-0XT
Memory	Standard	16MB	16MB
	Maximum	192MB	192MB
	Parity	Industry Standard, 72-pin, 70ns	Industry Standard, 72-pin, 70ns
	EOS (ECC on SIMM)	Optional	Optional
Hard Disk Controller	Type	SCSI-2 F/W	SCSI-2 F/W
	I/O Bus	PCI	PCI
	SCSI-speed	20MBps	20MBps
	Channels	1 1	1 1
	Devices per Channel	15	15
Hard Drives	Standard	0	1.08GB
	Access Time	N/A	10.5ms
	Maximum	13.5GB	13.5GB
CD-ROM (SCSI)	Standard	Yes	Yes
	Speed	4 X	4 X
Free Expansion Bays	3 ½"	2	1
	5 ¼" ½H	5	5
	Hot Swap (1/3H)	Optional	Optional
Free Expansion Slots	PCI	1	1
	ISA/EISA	4	4
	PCI or EISA 2	1	1
<p>Note:</p> <p>1 One shared SCSI channel:</p> <ul style="list-style-type: none"> • Both internal 8-bit and internal 16-bit connectors can be used or • Either internal connector (8 or 16-bit) and external connector 16-bit can be used connector <p>2 Shared slot, can be used for a PCI or EISA card connection.</p>			

This soft copy for use by IBM employees only.

1.5 The IBM PC Server 320 (Micro Channel/PCI)

This system complements the IBM PC Server 320 EISA/PCI, shows IBM's commitment to protect customers investments, and provides a highly scalable and flexible server platform. This server features both, Peripheral Component Interconnect (PCI) and Micro Channel architecture for fast reliable LAN performance, and easy configuration of adapters. The PC Server 320 models combine advanced bus architectures with powerful 75MHz and 90MHz Pentium processors that can be upgraded to symmetrical multiprocessing (SMP) by simply plugging in a second processor.

The IBM PC Server 320 MCA has the same characteristics as described in 1.4, "The IBM PC Server 320 (EISA/PCI)" on page 16, the same tower case design, the same type of memory, and the same hard drives. The differences between both systems are:

- Six 40MBps Micro Channel and two PCI slots providing maximum I/O throughput
- All eight slots available for I/O adapters due to on board PCI SVGA and PCI SCSI-2 F/W adapters
- Micro Channel bus parity checking to data protection and integrity
- Entry model equipped with Pentium 75MHz processors upgradable to 90 MHz Pentium processors
- C2 security option for protection against unauthorized access

Table 5 provides technical details for the PC Server 320 MCA models.

<i>Table 5 (Page 1 of 3). Technical Characteristics of the PC Server 320 MCA</i>						
Model Type Number	US, LA	8640-MX0	8640-MXT	8640-MY0	8640-MYT	8640-MYR
	EMEA	8640-MX0	8640-MXT	8640-MY0	8640-MYT	8640-MYR
Processor	Type	Pentium	Pentium	Pentium	Pentium	Pentium
	Speed	90MHz	90MHz	90MHz	90MHz	90MHz
	L2 Cache (std./max.)	256/512KB	256/512KB	256/512KB	256/512KB	256/512KB
	Std. #	1	1	1	1	1
	Max. #	2	2	2	2	2
	Upgrade	Chip	Chip	Chip	Chip	Chip

This soft copy for use by IBM employees only.

<i>Table 5 (Page 2 of 3). Technical Characteristics of the PC Server 320 MCA</i>						
Model Type Number	US, LA	8640-MX0	8640-MXT	8640-MY0	8640-MYT	8640-MYR
	EMEA	8640-MX0	8640-MXT	8640-MY0	8640-MYT	8640-MYR
System Bus	Type	MCA/PCI	MCA/PCI	MCA/PCI	MCA/PCI	MCA/PCI
	Speed (Mbps) MCA/PCI	40/132	40/132	40/132	40/132	40/132
Memory	Standard	16MB	16MB	16MB	16MB	16MB
	Maximum	256MB	256MB	256MB	256MB	256MB
	Parity	Industry Standard, 72-pin, 70ns				
	EOS (ECC on SIMM)	Optional	Optional	Optional	Optional	Optional
Hard Disk Controller	Type	SCSI-2 F/W (onboard)	SCSI-2 F/W (onboard)	SCSI-2 F/W (onboard)	SCSI-2 F/W (onboard)	SCSI-2 F/W RAID
	I/O Bus	PCI	PCI	PCI	PCI	PCI
	SCSI-speed	20MBps	20MBps	20MBps	20MBps	20MBps
	Channels	1 1	1 1	1 1	1 1	2 2
	Devices per Channel	15	15	15	15	7/15 3
Hard Drives	Standard	0	1.08GB	0MB	1.12GB	2 x 1.12GB
	Access Time	N/A	10.5ms	N/A	6.9ms	6.9ms
	Maximum	13.5GB	13.5GB	13.5GB	13.5GB	13.5GB
CD-ROM (SCSI)	Standard	Yes	Yes	Yes	Yes	Yes
	Speed	2X	2X	2X	2X	2X
Free Expansion Bays	3 1/2"	1	0	1	1	1
	5 1/4" 1/2 H	5	5	0	0	0
	Hot Swap (1/3H)	Optional	Optional	6	5	4

This soft copy for use by IBM employees only.

<i>Table 5 (Page 3 of 3). Technical Characteristics of the PC Server 320 MCA</i>						
Model Type Number	US, LA	8640-MX0	8640-MXT	8640-MY0	8640-MYT	8640-MYR
	EMEA	8640-MX0	8640-MXT	8640-MY0	8640-MYT	8640-MYR
Free Expansion Slots	PCI	2	2	2	2	2
	MCA	6	6	6	6	6
	PCI or MCA 4	0	0	0	0	0

Note:

1 One shared SCSI channel:

- Both; internal 8-bit and 16-bit connectors can be used
- or
- Either; one internal connector (8-bit or 16-bit) and one external 16-bit connector

Up to 15 devices can be connected.

2 Two independent SCSI channels:

- One shared channel with internal 16-bit and external 16-bit connectors
- One dedicated channel with internal 16-bit connector

3 Number of drives per SCSI-channel supported by the PCI RAID controller:

- Up to 6 internal hot swap drives
- Up to 7 external hot swap drives when using the external SCSI Enclosure IBM 3516
- Up to 15 external SCSI-devices when using non hot swap enclosures

4 A slot that is shared and can be used for a PCI or MCA card connection.

1.6 The IBM PC Server 500

The PC Server 500 has a unique design that combines proven Micro Channel technology with the latest Intel Pentium 90MHz processor, improved hard disk, SCSI-2 or RAID subsystems, and ECC memory. The PC Server 500 is a follower of the very successful Server 95 family² of systems. Designed for more demanding environments, the PC Server 500 holds the tradition of upgradeable modular design, advanced features and expandability.

IBM PC Server 500 is designed and priced to be a professional server, to efficiently handle 10-300+ clients, using a combination of industry technologies and IBM's focused leadership at key subsystem technology.

² The PC Server 95 family is not covered by this book as it is no longer offered by IBM.

This soft copy for use by IBM employees only.

While the IBM PC Server 500 is positioned as mid-level server, it contains many of the performance, reliability, systems management features and expansion capabilities of more expensive mission critical LAN servers.

Technical highlights:

- Processor complex design using Pentium 90MHz processor, IBM's SynchroStream technology (to obtain a zero wait state Pentium), standard 256KB write-back L2 cache, 256MB ECC memory, and Vital Product Data (VPD) support.
- Standard 16MB Error Checking and Correcting (ECC) memory expandable to 256MB ECC memory.
- SCSI-2 Fast/Wide controllers matched with performance leadership disk drives (1-inch 1.12GB and 2.25GB hot swap ready, high performance drives).
- Optional Busmaster Token-Ring or Ethernet data streaming adapter.
- High Speed ExpressPrint parallel port that supports multiple high-speed laser printers to work multiplexed.
- Reliability features like ECC memory, Micro Channel bus parity, hot swappable drive bays and Predictive Failure Analysis (PFA) with NetFinity support are included standard with the system.

Figure 4 on page 25 shows the mechanical design of the Server 500, and the positioning of bays, processor card and power supply.

This soft copy for use by IBM employees only.




Figure 4. IBM PC Server 500 - Mechanical Design

With the PC Server 500 comes a mechanical design that offers outstanding configuration flexibility with a dramatic tower packaging. PC Server 500 advantages include:

- Extensive expansion capability and power to meet advanced server needs, with room for future growth, plenty bays and many slots.
- A total of twenty two media bays, eighteen of them equipped for high-performance hot swap drives.
- Eight adapter card slots using the parity checked 32-bit Micro Channel Bus Technology.
- Powerful 434 watts auto-sensing, worldwide power supply handles full configuration power requirements. Optional 220 watts modular power supply supports peak levels of internal disk drive capacity (required when installing drives in the PC Server 500's two lower drive banks).

This soft copy for use by IBM employees only.

- A new FloThru cooling design. Three fans positioned to provide full length front to back system cooling.
- Integrated CD-ROM drive with ServerGuide, an automated network operating system installation and configuration tool.

With these improved features the IBM Server 500 is the ideal server for more complex and advanced interconnected LANs found in medium and large businesses running critical and non-critical business applications. The IBM Server 500 configurability capability for RAID 1 or 5 with ECC memory offers a very cost effective solution for high availability and data protection, necessary for the business environment.

The Table 6 shows the technical characteristics for the PC Server 500 models, including the array and non-array models.

Table 6 (Page 1 of 3). Technical Characteristics of the PC Server 500

Model Type Number	US, LA	8641-0Y0	8641-0YT	8641-0YV	8641-1Y0	8641-0YR	8641-0YS
	EMEA	Build to Order (Based 8641-S0C) 					
Processor	Type	Pentium	Pentium	Pentium	Pentium	Pentium	Pentium
	Speed	90MHz	90MHz	90MHz	90MHz	90MHz	90MHz
	L2 Cache	256KB	256KB	256KB	256KB	256KB	256KB
	Upgrade	Complex	Complex	Complex	Complex	Complex	Complex
System Bus	Type	Micro Channel	Micro Channel	Micro Channel	Micro Channel	Micro Channel	Micro Channel
	Speed (MBps)	40	40	40	40	40	40
Memory	Standard	32MB	32MB	32MB	32MB	32MB	32MB
	Maximum	256MB	256MB	256MB	256MB	256MB	256MB
	Type	ECC, 70ns	ECC, 70ns	ECC, 70ns	ECC, 70ns	ECC, 70ns	ECC, 70ns

This soft copy for use by IBM employees only.

<i>Table 6 (Page 2 of 3). Technical Characteristics of the PC Server 500</i>							
Model Type Number	US, LA	8641-0Y0	8641-0YT	8641-0YV	8641-1Y0	8641-0YR	8641-0YS
	EMEA	Build to Order (Based 8641-S0C) 1					
Hard Disk Controller	Type	SCSI-2 F/W	SCSI-2 F/W	SCSI-2 F/W	SCSI-2 F/W RAID	SCSI-2 F/W RAID	SCSI-2 F/W RAID
	I/O Bus	Micro Channel	Micro Channel	Micro Channel	Micro Channel	Micro Channel	Micro Channel
	SCSI-speed	20MBps	20MBps	20MBps	20MBps	20MBps	20MBps
	Channel	1 2	1 2	1 2	2 3	2 3	2 3
	Devices per Channel	15	15	15	7/15 4	7/15 4	7/15 4
Hard Drives	Standard	0	1.12GB	2.25MB	0	3 x 1.12GB	3 x 2.25GB
	Access Time	N/A	6.9ms	7.5ms	N/A	6.9ms	7.5ms
	Maximum	40.5GB	40.5GB	40.5GB	40.5GB	40.5GB	40.5GB
CD-ROM (SCSI)	Standard	Yes	Yes	Yes	Yes	Yes	Yes
	Speed	2X	2X	2X	2X	2X	2X
Free Expansion Bays	3.5"	0	0	0	0	0	0
	5 ¼" ½H	2	2	2	2	2	2
	Hot Swap (1/3H) 5	6-18	5-17	5-17	6-18	3-15	3-15

This soft copy for use by IBM employees only.

<i>Table 6 (Page 3 of 3). Technical Characteristics of the PC Server 500</i>							
Model Type Number	US, LA	8641-0Y0	8641-0YT	8641-0YV	8641-1Y0	8641-0YR	8641-0YS
	EMEA	Build to Order (Based 8641-S0C) 1					
Free Expansion Slots	MCA	6	6	6	8	6	6
	Size	32-bit	32-bit	32-bit	32-bit	32-bit	32-bit
<p>Note:</p> <p>1 When choosing the 8641-SOC (BTO), memory, disks, disk and LAN adapter must be ordered separately and will be factory installed.</p> <p>2 One shared SCSI channel:</p> <ul style="list-style-type: none"> • Both internal 8-bit and internal 16-bit connectors can be used or • Either internal (8 or 16-bit) connector and external 16-bit connector can be used <p>3 Two independent SCSI channels:</p> <ul style="list-style-type: none"> • One shared between Internal 16-bit and External 16-bit connectors • One dedicated channel with internal 16-bit connector <p>4 Number of drives per SCSI-channel supported by the Streaming RAID Adapter /A:</p> <ul style="list-style-type: none"> • Up to 6 internal hot swap drives • Up to 7 external hot swap drives when using the external SCSI Enclosure IBM 3516 • Up to 15 external SCSI-devices when using non hot swap enclosures <p>5 The first value indicates number of free bays with the standard configuration and the second value is the maximum number of free bays achievable by using additional backplanes.</p>							

1.7 The IBM PC Server 520 (EISA/PCI or MCA/PCI)

Step up to your organization's LAN server requirements for increased data storage and high speed processing with the new, easy to install PC Server 520. Today's enterprise LAN environments demand the performance of single and dual 100MHz Pentium processors, high performance bus architectures and large disk storage capacities to handle the needs of advanced applications. The industry standard PC Server 520 systems can do this, and more.

Data hungry symmetrical multiprocessing (SMP) applications not only require larger amounts of memory, but also multiple high-speed hard disks to maximize performance. The PC Server 520 maxi-tower package with twenty-two storage bays is capable of storing over 40GB of data on eighteen, hot swap, high performance hard drives. This is the cost effective SMP

This soft copy for use by IBM employees only.

server solution your team needs to be competitive in today's fast paced business environment.

Technical highlights:

- Powerful LAN solutions with single 100MHz Pentium with 512KB L2 write-back cache and 100MHz Pentium SMP upgrade option.
- Advanced architecture featuring 132MBps PCI bus and either a SCSI-2 Fast/Wide PCI Adapter or SCSI-2 Fast/Wide PCI RAID Adapter standard.
- Large capacity (1.12GB, 2.25GB, 4.51GB, 5.31GB), SCSI-2 Fast/Wide ultra-performance (up to 7200 rpm) hard drives available standard or as options. Hard drives that can alert you when something is about to go wrong so they normally can be replaced before failure occurs.
- 4X CD-ROM drive, mouse, ServerGuide and NetFinity allow easy installation of network applications and provide comprehensive system management tools.
- PC Server 500 case design, featuring FloThru air flow design and up to eighteen hot swap drive bays to allow replacement of hard drives without powering down the server.
- Easy installation of adapters with plug and play support.

The reliability of a server is critical to your business. The PC Server 520 systems achieve rugged dependability by meeting high quality standards and design points, such as stringent IBM systems assurance testing of both IBM and non-IBM adapters and certification with many popular network operating systems.

Choice of PCI/EISA or PCI/MCA systems are offered with this family.

Table 7 on page 30 shows the different features for the PC Server 520 EISA models, and Table 8 on page 32 shows the same information for the MCA models.

This soft copy for use by IBM employees only.

<i>Table 7 (Page 1 of 2). Technical Characteristics of the PC Server 520 PCI/EISA</i>						
Model Type Number	US, LA	8641-EZ0	8641-EZV	8641-EZ1	8641-EZS	8641-EZE
	EMEA	8641-EZ0	8641-EZV	8641-EZ1	8641-EZS	8641-EZE
Processor	Type	Pentium	Pentium	Pentium	Pentium	Pentium
	Speed	100/66MHz	100/66MHz	100/66MHz	100/66MHz	100/66MHz
	Type of Multiproc.	2-way SMP				
	Std. #	1	1	1	1	1
	Max. #	2	2	2	2	2
	L2 Cache (std./max.)	256/512KB	256/512KB	256/512KB	256/512KB	256/512KB
	SMP upgrade	Chip	Chip	Chip	Chip	Chip
System Bus	Type	EISA/PCI	EISA/PCI	EISA/PCI	EISA/PCI	EISA/PCI
	Speed (Mbps) EISA/PCI	33/132	33/132	33/132	33/132	33/132
Memory	Standard	32MB	32MB	32MB	32MB	32MB
	Maximum	256MB	256GB	256MB	256MB	256MB
	Type	EOS, 70ns				
Hard Disk Controller	Type	SCSI-2 F/W	SCSI-2 F/W	SCSI-2 F/W RAID	SCSI-2 F/W RAID	SCSI-2 F/W RAID
	I/O bus	PCI	PCI	PCI	PCI	PCI
	SCSI-speed	20MBps	20MBps	20MBps	20MBps	20MBps
	Channels	1 1	1 1	2 2	2 2	2 2
	Devices per Channel	15	15	7/15 3	7/15 3	7/15 3
Hard Drives	Standard	0	2.25GB	0	2 x 2.25GB	4 x 2.25GB
	Access Time	N/A	7.5ms	N/A	7.5ms	7.5ms
	Maximum internal capacity	40.5GB	40.5GB	40.5GB	40.5GB	40.5GB

This soft copy for use by IBM employees only.

<i>Table 7 (Page 2 of 2). Technical Characteristics of the PC Server 520 PCI/EISA</i>						
Model Type Number	US, LA	8641-EZ0	8641-EZV	8641-EZ1	8641-EZS	8641-EZE
	EMEA	8641-EZ0	8641-EZV	8641-EZ1	8641-EZS	8641-EZE
CD-ROM (SCSI)	Standard	Yes	Yes	Yes	Yes	Yes
	Speed	4X	4X	4X	4X	4X
Free Expansion Bays	3 1/2"	0	0	0	0	0
	5 1/4" 1/2 H	2	2	2	2	2
	Hot Swap (1/3H) 4	6-18	5-17	6-18	4-16	2-14
Free Expansion Slots	PCI	1	1	1	1	1
	EISA	5	5	5	5	5
	PCI or EISA 5	1	1	1	1	1

Note:

1 One shared SCSI channel:

- Both internal (8-bit and 16-bit) connectors can be used
- or
- Either internal connector (8 or 16-bit) and external 16-bit connector can be used

2 Two independent SCSI channels:

- One shared between internal 16-bit and external 16-bit connectors
- One dedicated channel with internal 16-bit connector

3 Number of drives per SCSI-channel supported by the PCI RAID controller:

- Up to 6 internal hot swap drives
- Up to 7 external hot swap drives when using the external SCSI Enclosure IBM 3516
- Up to 15 external SCSI-devices when using non hot swap enclosures

4 The first value indicates number of free bays with the standard configuration and the second value is the maximum number of free bays achievable by using additional backplanes.

5 A slot that is shared and can be used for a PCI or EISA card connection.

This soft copy for use by IBM employees only.

<i>Table 8 (Page 1 of 2). Technical Characteristics of the PC Server 520 PCI/MCA</i>					
Model Type Number	US, LA	8641-MZ0	8641-MZV	8641-MZS	8641-MZE
	EMEA	8641-MZ0	8641-MZV	8641-MZS	8641-MZE
Processor	Type	Pentium	Pentium	Pentium	Pentium
	Speed	100/66MHz	100/66MHz	100/66MHz	100/66MHz
	Type of Multiproc.	2-way SMP	2-way SMP	2-way SMP	2-way SMP
	Std. #	1	1	1	1
	Max. #	2	2	2	2
	L2 Cache (std./max.)	256/512KB	256/512KB	256/512KB	256/512KB
	SMP upgrade	Chip	Chip	Chip	Chip
System Bus	Type	Micro Channel/PCI	Micro Channel/PCI	Micro Channel/PCI	Micro Channel/PCI
	Speed (MBps) Micro Channel/PCI	40/132	40/132	40/132	40/132
Memory	Standard	32MB	32MB	32MB	32MB
	Maximum	256MB	256GB	256MB	256MB
	Type	EOS, 70ns	EOS, 70ns	EOS, 70ns	EOS, 70ns
Hard Disk Controller	Type	SCSI-2 F/W (onboard)	SCSI-2 F/W (onboard)	SCSI-2 F/W (onboard) SCSI-2 F/W RAID	SCSI-2 F/W (onboard) SCSI-2 F/W RAID
	I/O bus	PCI	PCI	PCI	PCI
	SCSI-speed	20MBps	20MBps	20MBps	20MBps
	Channels	1 1	1 1	2 2	2 2
	Devices per Channel	15 3	15 3	15 3 , 7/15 4	15 3 , 7/15 4
Hard Drives	Standard	0	2.25GB	2 x 2.25GB	4 x 2.25GB
	Access Time	N/A	7.5ms	7.5ms	7.5ms
	Maximum internal capacity	40.5GB	40.5GB	40.5GB	40.5GB

This soft copy for use by IBM employees only.

<i>Table 8 (Page 2 of 2). Technical Characteristics of the PC Server 520 PCI/MCA</i>					
Model Type Number	US, LA	8641-MZ0	8641-MZV	8641-MZS	8641-MZE
	EMEA	8641-MZ0	8641-MZV	8641-MZS	8641-MZE
CD-ROM (SCSI)	Standard	Yes	Yes	Yes	Yes
	Speed	4X	4X	4X	4X
Free Expansion Bays	3 ½"	0	0	0	0
	5 ¼" ½ H	2	2	2	2
	Hot Swap (1/3H) 7	6-18	5-17	4-16	2-14
Free Expansion Slots	PCI	2 6	2	1	1
	MCA	6	6	6	6
	PCI or MCA 5	0	0	0	0

Note:

1 One shared SCSI channel is provided:

- Can use the internal 16-bit connector
- or
- Can use the external 16-bit connector

2 Two independent SCSI Channels are provided:

- One shared between internal and external 16-bit connectors
- One dedicated internal 16-bit connector

3 The onboard SCSI-2 F/W controller supports up to 15 devices

4 Number of drives per SCSI-channel supported by the PCI RAID controller:

- Up to 6 internal hot swap drives
- Up to 7 external hot swap drives when using the external SCSI Enclosure IBM 3516
- Up to 15 external SCSI-devices when using non hot swap enclosures

5 A slot that is shared and can be used for a PCI or EISA card connection.

6 The first value indicates number of free bays with the standard configuration and the second value is the maximum number of free bays achievable by using additional backplanes.

This soft copy for use by IBM employees only.

1.8 The IBM PC Server 720

The IBM PC Server 720 is the latest member of the very successful PC Server family of systems. Designed for more demanding environments, the PC Server 720 expands the tradition of upgradeable modular design, advanced features, and expandability by offering symmetrical multiprocessing with up to six processors.

This soft copy for use by IBM employees only.




Figure 5. IBM PC Server 500 and 720

Designed to satisfy a customer's most demanding requirements this SMP system offers performance, Micro Channel and PCI subsystems, up to six Pentium 100MHz processors, and high performance SCSI-2 F/W all housed in an attractive and highly functional, floor-standing tower. PC Server 720 is a unique design that combines proven technology with an innovative SMP

This soft copy for use by IBM employees only.

architecture, integrated system management functions, high-performance hard disks, and RAID subsystem. Included with the comprehensive PC Server 720 systems are a built-in CD-ROM drive, ServerGuide, NetFinity system management, LogicLock security, and ECC memory. The outstanding physical configuration capability of the PC Server 720 system offers the customer excellent future system growth to match growth in business requirements.

The PC Server 720 is designed and priced as an advanced professional server that can efficiently handle 10 to 500 plus clients, using a combination of industry technologies and IBM's leadership focused on key subsystem technology.

The PC Server 720 uses a full 64-bit Pentium 100/66MHz processor, running with zero wait states and an optimized write-back L2 cache as well as an L3 Victim cache designed to improve the cache hit rate and enhance performance even more. This system is ideally suited to processor-intensive environments such as database, transaction processing, and application serving. In addition, the PC Server 720 offers higher levels of price/performance with the inclusion of higher performing, larger capacity hard disk drives.

Technical highlights:

- Symmetrical multi processing. Up to six P54C 100MHz processors cards each with their own 512KB of write-back L2 cache and a L3 write-back Victim cache for enhanced performance.
- Mix and match processor speeds. Not constrained to all processor complexes having the same speed.
- Fully symmetric processor, I/O and interrupt architecture.
- Micro Channel and PCI card support, with bus transfers protected by parity checking for both address and data lines.
- Standard 64MB, 60ns ECC memory is expandable to 1GB.
- Four-way memory interleave.
- Full tower packaging with twenty-two device storage bays and seven adapter card slots (32-bit Micro Channel or PCI).
- SCSI-2 Fast/Wide controllers matched with performance leadership disk drives (1-inch 1.12GB and 2.25GB hot swap ready, high RPM drives).

This soft copy for use by IBM employees only.

- Array capability offered with high performance SCSI-2 Fast/Wide Streaming-RAID adapter with matched cabling and disk drives. RAID 0, 1, and 5 supported.

The mechanical design of this system is almost identical to the PC Server 500. The differences between both systems are:

- A 470 watts power supply instead of the standard 434 watts power supply for the PC Server 500.
- Seven 32-bit MCA/PCI slots.
- Onboard SVGA adapter.

The PC Server 720 system with Micro Channel and PCI architecture, upgradeable processor complex cards and NetFinity are designed to take advantage of today's technology, and are well positioned to take advantage of future technologies providing higher performance, availability, and increased fault-tolerant computing.

Table 9 describes the technical specifications to the PC Server 720 family.

Model Type Number	US, LA	8642-0Z0	8642-1Z0	8642-2ZS	8642-4ZS	N/A
	EMEA	8642-0Z0	8642-1Z0	8642-2ZS	8642-4ZS	8642-M0Z
Processor	Type	Pentium	Pentium	Pentium	Pentium	Pentium
	Speed	100/66MHz	100/66MHz	100/66MHz	100/66MHz	100/66MHz
	Type of Multiproc.	N-way SMP				
	L2 Cache per processor	512KB	512KB	512KB	512KB	512KB
	L3 Victim cache per processor	Yes	Yes	Yes	Yes	Yes
	Std. #	1	1	2	4	BTO
	Max. #	6	6	6	6	6
	Upgrade	Proc. complex				

This soft copy for use by IBM employees only.

Model Type Number	US, LA	8642-0Z0	8642-1Z0	8642-2ZS	8642-4ZS	N/A
	EMEA	8642-0Z0	8642-1Z0	8642-2ZS	8642-4ZS	8642-M0Z
System Bus	Type	Micro Channel, PCI				
	Speed (Mbps) MCA/PCI	80/132	80/132	80/132	80/132	80/132
Memory	Standard	64MB	64MB	64MB	64MB	BTO
	Maximum	1GB	1GB	1GB	1GB	1GB
	Type	ECC, 60ns				
Hard Disk Controller	Type	SCSI-2 F/W	SCSI-2 F/W RAID	SCSI-2 F/W RAID	SCSI-2 F/W RAID	BTO
	I/O bus	PCI	Micro Channel	Micro Channel	Micro Channel	BTO
	SCSI-speed	20MBps	20MBps	20MBps	20MBps	20MBps
	Channels	1 1	1 1	2 2	2 2	BTO
	Devices per Channel	15	7/15 3	7/15 3	7/15 3	BTO
Hard Drives	Standard	0	0	2 x 2.25GB	4 x 2.25GB	BTO
	Access Time	N/A	N/A	7.5ms	N/A	BTO
	Maximum internal capacity	40.5GB	40.5GB	40.5GB	40.5GB	40.5GB
CD-ROM (SCSI)	Standard	Yes	Yes	Yes	Yes	Yes
	Speed	2X	2X	2X	2X	2X
Free Expansion Bays	3 1/2"	0	0	0	0	0
	5 1/4" 1/2 H	2	2	2	2	BTO
	Hot Swap (1/3H) 4	6-18	6-18	4-16	2-14	BTO

This soft copy for use by IBM employees only.

<i>Table 9 (Page 3 of 3). Technical Characteristics of the PC Server 720</i>						
Model Type Number	US, LA	8642-0Z0	8642-1Z0	8642-2ZS	8642-4ZS	N/A
	EMEA	8642-0Z0	8642-1Z0	8642-2ZS	8642-4ZS	8642-M0Z
Free Expansion Slots	PCI	0	0	0	0	0
	MCA	0	0	0	0	0
	PCI or MCA 5	6	6	6	6	6
<p>Note:</p> <p>1 One shared SCSI channel is provided:</p> <ul style="list-style-type: none"> • Can use both internal 8-bit and 16-bit connectors or • Can use either internal connector (8 or 16-bit) and the external 16-bit connector <p>2 Two independent SCSI channels are provided:</p> <ul style="list-style-type: none"> • One shared between the internal and external 16-bit connector • One dedicated internal 16-bit connector <p>3 Number of drives per SCSI-channel supported by the RAID controller (Streaming RAID Adapter /A):</p> <ul style="list-style-type: none"> • Up to 6 internal hot swap drives • Up to 7 external hot swap drives when using the external SCSI Enclosure IBM 3516 • Up to 15 external SCSI-devices when using non hot swap enclosures <p>4 The first value indicates number of free bays with the standard configuration and the second value is the maximum number of free bays achievable by using additional backplanes.</p> <p>5 Each of the slots a the 720 can support a MCA orPCI card, all of them are shared.</p>						

1.9 IBM PC Server Summary

The following options are common for all IBM PC Server systems and thus are not listed in the table below:

- Keyboard
- Mouse
- ServerGuide (except PC Server 300, 8640-0P0)
- NetFinity (one Manager and one Client license)
- APC Powerchute Plus (except PC Server 300, 8640-0P0)
- Serial UPS cable

Table 10 provides a summary of the IBM PC Server family including part numbers and a brief description of the technical highlights.

<i>Table 10 (Page 1 of 4). IBM PC Server Models Overview</i>	
Part Number	Description
PC Server 300 PCI/EISA	
86400P0	Pentium 60, 256K-WB, 8MB Parity, 0MB IDE, 1 Enhanced IDE Port (2 drive support), No SCSI adapter or cable included (must be added as an option), VGA EISA Adapter, FloThru, 1.44MB, 2 PCI slots, 1 Shared PCI/EISA slot, 5 EISA slots, 300 watts power supply
86400PT	Pentium 60, 256K-WB, 8MB Par, 1GB SCSI-2 Fast, 1 Enhanced IDE Port (2 drive support), 1 SCSI-2 Fast PCI Adapter, 4-drop 8-bit SCSI-2 Fast cable, VGA EISA Adapter, Flothru, 1.44MB diskette drive, Double-speed CD-ROM, 2 PCI slots, 1 Shared PCI/EISA slot, 5 EISA slots, 300 watts power supply
PC Server 310	
86390XT	Pentium 75, with ZIF upgrade 256WB L2 cache, 16MB, 1.08GB SCSI-2 Fast, SCSI-2 Fast PCI Adapter, SVGA onboard, 1.44MB, Quad-speed CD-ROM, C2-enabled 1 PCI free slots, 2 ISA free slots, 200 watts power supply
8639MXT	Pentium 75, with ZIF upgrade 256WB L2 cache, 16MB, 1.08GB SCSI-2 Fast, SCSI-2 Fast PCI Adapter, SVGA onboard, 1.44MB, Quad-speed CD-ROM, C2-enabled, 1 PCI free slots, 2 MCA free slots, 200 watts power supply; support for popular network operating systems
PC Server 320 PCI/EISA	
86400X0	Pentium 75, Shared-Cache Dual-SMP capable, 256K-WB/WT, 16MB, Diskless, SCSI-2 Fast/Wide PCI Adapter, 4-drop SCSI-2 Fast cable, 1 Enhanced IDE port, SVGA EISA adapter, Flothru 2-fans, 1.44MB, Quad-speed CD-ROM, 2 PCI slots, 1 Shared PCI/EISA slot, 5 EISA slots, 300 watts power supply
86400XT	Pentium 90, Shared-Cache Dual-SMP capable, 256K-WB/WT, 16MB, 1.08GB SCSI-2 Fast, SCSI-2 Fast/Wide PCI Adapter, 7-drop SCSI-2 Fast/Wide cable, 1 Enhanced IDE port, SVGA EISA Adapter, Flothru 2-fans, 1.44MB, Quad-speed CD-ROM, 2 PCI slots, 1 Shared PCI/EISA slot, 5 EISA slots, 300 watts power supply

This soft copy for use by IBM employees only.

<i>Table 10 (Page 2 of 4). IBM PC Server Models Overview</i>	
Part Number	Description
86401Y0	Pentium 90, Shared-Cache Dual-SMP capable, 256K-WB/WT, 16MB, Diskless, SCSI-2 Fast/Wide RAID PCI Adapter, 2-drop SCSI-2 Fast/Wide cable (6 hot swap connectors, 1 non hot swap), 1 Enh. IDE port, SVGA EISA Adapter, Flothru 2-fans, 1.44MB, Double-speed CD-ROM, 2 PCI slots, 1 Shared PCI/EISA slot, 5 EISA slots, 300 watts power supply
86401YT	Pentium 90, Shared-Cache Dual-SMP capable, 256K-WB/WT, 16MB, 2x1.12GB SCSI-2 Fast/Wide, SCSI-2 Fast/Wide RAID PCI Adapter, 2-drop SCSI-2 Fast/Wide cable (6 hot swap connectors, 1 non hot swap), 1 Enh. IDE port, SVGA EISA Adapter, Flothru 2-fans, 1.44MB, Double-speed CD-ROM, 2 PCI slots, 1 Shared PCI/EISA slot, 5 EISA slots, 300 watts power supply
PC Server 320 PCI/MCA	
8640MX0	Pentium 75, Shared-Cache Dual-SMP capable, 256K-WB/WT, 16MB, Diskless, Integrated SCSI-2 Fast/Wide PCI Controller, 4-drop non hot swap SCSI-2 Fast cable, Integrated SVGA/1MB controller, Flothru 2-fans, 1.44MB, Double-speed CD-ROM, 2 PCI slots, 6 Microchannel slots, 300 watts power supply
8640MXT	Pentium 75, Shared-Cache Dual-SMP capable, 256K-WB/WT, 16MB, 1.08GB SCSI-2 Fast, Integrated SCSI-2 Fast/Wide PCI Controller, 4-drop non hot swap SCSI-2 Fast cable, Integrated SVGA/1MB controller, Flothru 2-fans, 1.44MB, Double-speed CD-ROM, 2 PCI slots, 6 Microchannel slots, 300 watts power supply
8640MY0	Pentium 90, Shared-Cache Dual-SMP capable, 256K-WB/WT, 16MB, Diskless, Integrated SCSI-2 Fast/Wide PCI Controller, 16-bit 2-drop SCSI-2 Fast/Wide cable connecting 6 hot swap bay enclosure and Double speed CD-ROM, Integrated SVGA/1MB controller, Flothru 2-fans, 1.44MB, 2 PCI slots, 6 Microchannel slots, 300 watts power supply
8640MYT	Pentium 90, Shared-Cache Dual-SMP capable, 256K-WB/WT, 16MB, 1.12GB SCSI-2 F/W hot swap, Integrated SCSI-2 Fast/Wide PCI Controller, 2-drop SCSI-2 Fast/Wide cable connecting 6 hot swap bay enclosure and Double speed CD-ROM, Integrated SVGA/1MB controller, Flothru 2-fans, 1.44MB, 2 PCI slots, 6 Microchannel slots, 300 watts power supply
8640MYR	Pentium 90, Shared-Cache Dual-SMP capable, 256K-WB/WT, 16MB, 2x1.12GB SCSI-2 F/W hot swap, Integrated SCSI-2 Fast/Wide PCI Controller, PCI SCSI-2 Fast/Wide RAID Adapter, 2-drop SCSI-2 Fast/Wide cable connecting 6 hot swap bay enclosure and Double speed CD-ROM to the RAID adapter, Integrated SVGA/1MB controller, Flothru 2-fans, 1.44MB, 2 PCI slots, 6 Micro channel slots, 300 watts power supply
PC Server 500 MCA	
86410Y0	Pentium 90/60Mhz, 256K-WB, 32MB ECC, 0GB, SCSI-2 Fast/Wide Adapter, 40MBps Data Streaming, SVGA, C2, 2.88MB, Processor Complex, SynchroStream, CD-ROM 2X, 434 watts power supply
86411Y0	Pentium 90/60Mhz, 256K-WB, 32MB ECC, 0GB hot swap, SCSI-2 Fast/Wide RAID Adapter (supports 2 channels that can be externally or internally), 40MBps Data Streaming, SVGA, C2, 2.88MB, Processor Complex, SynchroStream, CD-ROM 2X, 434 watts power supply
86410YT	Pentium 90/60Mhz, 256K-WB, 32MB ECC, 1.12GB SCSI-2 Fast/Wide hot swap, SCSI-2 Fast/Wide Adapter, 40MBps Data Streaming, SVGA, C2, 2.88MB, Processor Complex, SynchroStream, CD-ROM 2X, 434 watts power supply
86410YV	Pentium 90/60Mhz, 256K-WB, 32MB ECC, 2.25GB SCSI-2 Fast/Wide hot swap, SCSI-2 Fast/Wide Adapter, 40MBps Data Streaming, SVGA, C2, 2.88MB, Processor Complex, SynchroStream, CD-ROM 2X, 434 watts power supply

This soft copy for use by IBM employees only.

<i>Table 10 (Page 3 of 4). IBM PC Server Models Overview</i>	
Part Number	Description
86410YR	Pentium 90/60Mhz, 256K-WB,32MB ECC, 3x1.12GB SCSI-2 Fast/Wide hot swap, SCSI-2 Fast/Wide RAID Adapter (supports 2 channels that can be externally or internally), 40MBps Data Streaming,SVGA,C2,2.88MB, Processor Complex, SynchroStream, CD-ROM 2X, 434 watts power supply
86410YS	Pentium 90/60Mhz, 256K-WB,32MB ECC, 3x2.25GB SCSI-2 Fast/Wide hot swap, SCSI-2 Fast/Wide RAID Adapter (supports 2 channels that can be externally or internally), 40MBps Data Streaming,SVGA,C2,2.88MB, Processor Complex, SynchroStream, CD-ROM 2X, ServerGuide, 434 watts power supply
PC Server 520 PCI/EISA	
8641EZ0	Pentium 100/66Mhz, Shared-Cache Dual-SMP capable, 256K-WB/WT, 32MB EOS, Diskless, Integrated SCSI-2 Fast/Wide PCI Controller, 16-bit 2-drop SCSI-2 Fast/Wide cable connecting 6 hot swap bay enclosure and Quad speed CD-ROM, Integrated SVGA/1MB controller, Flothru 2-fans, 1.44MB, 2 PCI slots, 5 EISA and 1 PCI/EISA shared slot, 434 watts power supply
8641EZV	Pentium 100/66Mhz, Shared-Cache Dual-SMP capable, 256K-WB/WT, 32MB EOS, 2.25GB SCSI-2 F/W hot swap, Integrated SCSI-2 Fast/Wide PCI Controller, 16-bit 2-drop SCSI-2 Fast/Wide cable connecting 6 hot swap bay enclosure and Quad speed CD-ROM, SVGA EISA Adapter, Flothru 3-fans, 1.44MB, 2 PCI slots, 5 EISA and 1 PCI/EISA shared slot, 434 watts power supply
8641EZ1	Pentium 100/66Mhz, Shared-Cache Dual-SMP capable, 256K-WB/WT, 32MB, Diskless, SCSI-2 Fast/Wide RAID PCI Adapter, 2-drop SCSI-2 Fast/Wide cable(6 hot swap connectors, 1 non-hot), 1 Enh. IDE port, SVGA EISA Adapter, Flothru 3-fans, 1.44MB, Quad-speed CD-ROM, 2 PCI slots, 1 Shared PCI/EISA slot, 5 EISA slots, 434 watts power supply
8641EZS	Pentium 100/66Mhz, Shared-Cache Dual-SMP capable, 256K-WB/WT, 32MB, 2 x 2.25GB SCSI-2 F/W hot swap, SCSI-2 Fast/Wide RAID PCI Adapter, 2-drop SCSI-2 Fast/Wide cable(6 hot swap connectors, 1 non hot swap), 1 Enh. IDE port, SVGA EISA Adapter, Flothru 3-fans, 1.44MB, Quad-speed CD-ROM, 2 PCI slots, 1 Shared PCI/EISA slot, 5 EISA slots, 434 watts power supply
8641EZE	Pentium 100/66Mhz, Shared-Cache Dual-SMP capable, 256K-WB/WT, 32MB, 4 x 2.25GB SCSI-2 F/W hot swap, SCSI-2 Fast/Wide RAID PCI Adapter, 2-drop SCSI-2 Fast/Wide cable (6 hot swap connectors, 1 non hot swap), 1 Enh. IDE port, SVGA EISA Adapter, Flothru 3-fans, 1.44MB, Quad-speed CD-ROM, 2 PCI slots, 1 Shared PCI/EISA slot, 5 EISA slots, 434 watts power supply
PC Server 520 PCI/MCA	
8641MZ0	Pentium 100/66Mhz, Shared-Cache Dual-SMP capable, 256K-WB/WT, 32MB, Open bay, Integrated SCSI-2 Fast/Wide PCI adapter, 4-drop non hot swap SCSI-2 Fast cable, Integrated SVGA/1MB controller, Flothru 3-fans, 1.44MB, Quad-speed CD-ROM, 2 PCI slots, 6 Micro channel slots, 434 watts power supply
8641MZV	Pentium 100/66Mhz, Shared-Cache Dual-SMP capable, 256K-WB/WT, 32MB, 2.25GB SCSI-2 F/W hot swap, Integrated SCSI-2 Fast/Wide PCI adapter, 4-drop non hot swap SCSI-2 Fast cable, Integrated SVGA/1MB controller, Flothru 3-fans, 1.44MB, Quad-speed CD-ROM, 2 PCI slots, 6 Micro channel slots, 434 watts power supply
8641MZS	Pentium 100/66Mhz, Shared-Cache Dual-SMP capable, 256K-WB/WT, 32MB, 2x2.25GB SCSI-2 F/W hot swap, Integrated SCSI-2 Fast/Wide PCI Controller, PCI SCSI-2 Fast/Wide RAID Adapter, 2-drop SCSI-2 Fast/Wide cable connecting 6 hot swap bay enclosure and Quad speed CD-ROM to the RAID adapter, Integrated SVGA/1MB controller, Flothru 3-fans, 1.44MB, 2 PCI slots, 6 Micro channel slots, 434 watts power supply

This soft copy for use by IBM employees only.

<i>Table 10 (Page 4 of 4). IBM PC Server Models Overview</i>	
Part Number	Description
8641MZE	Pentium 100/66Mhz, Shared-Cache Dual-SMP capable, 256K-WB/WT, 32MB, 4x2.25GB SCSI-2 F/W hot swap, Integrated SCSI-2 Fast/Wide PCI Controller, PCI SCSI-2 Fast/Wide RAID Adapter, 2-drop SCSI-2 Fast/Wide cable connecting 6 hot swap bay enclosure and Quad speed CD-ROM to the RAID adapter, Integrated SVGA/1MB controller, Flothru 3-fans, 1.44MB, 2 PCI slots, 6 Micro channel slots, 434 watts power supply
PC Server 720 PCI/MCA	
86420Z0	1 x (Pentium 100Mhz, 512KB-WB L2/L3 Cache) 6-way SMP capable, 400MBps C-Bus II Multi Processor Bus, 64MB ECC (1 Bank of 8 x 8MB 60ns ECC memory simm kits, 1 Bank empty), 7 I/O slots(PCI, MCA, or mix), Diskless, SCSI-2 Fast/Wide PCI Adapter (1 channel, 15 devices), 80MBps Data Streaming on MCA, Planar SVGA, C2 Security, 2.88MB, Flothru, 3 fans, SynchroStream, Double-speed CD-ROM, 470 watts power supply
86421Z0	1 x (Pentium 100Mhz, 512KB-WB L2/L3 Cache) 6-way SMP capable, 400MBps C-Bus II Multi Processor Bus, 64MB ECC(1 Bank of 8 x 8MB 60ns ECC memory simm kits, 1 Bank empty), 7 I/O slots(PCI, MCA, or mix), Diskless, SCSI-2 Fast/Wide RAID MCA Adapter (2 channels internally or 1 int/1 ext), 80MBps Data Streaming on MCA, Planar SVGA, C2 Security, 2.88MB, Flothru, 3 fans, SynchroStream, double-speed CD-ROM, 470 watts power supply
86422ZS	2 x (Pentium 100Mhz, 512KB-WB L2/L3 Cache) 6-way SMP capable, 400MBps C-Bus II Multi Processor Bus, 64MB ECC (1 Bank of 8 x 8MB 60ns ECC memory simm kits, 1 Bank empty), 7 I/O slots(PCI, MCA, or mix), 2 x 2.25GB Fast/Wide hot swap drives, SCSI-2 Fast/Wide RAID MCA Adapter (2 channels internally or 1 int/1 ext), 80MBps Data Streaming on MCA, Planar SVGA, C2 Security, 2.88MB, Flothru 3 fans, SynchroStream, double-speed CD-ROM, 470 watts power supply
86424ZS	4 x (Pentium 100Mhz, 512KB-WB L2/L3 Cache) 6-way SMP capable, 400MBps C-Bus II Multi Processor Bus, 64MB ECC (1 Bank of 8 x 8MB 60ns ECC memory simm kits, 1 Bank empty), 7 I/O slots(PCI, MCA, or mix), 4 x 2.25GB Fast/Wide hot swap drives, SCSI-2 Fast/Wide RAID MCA Adapter (2 channels internally or 1 int/1 ext), 80MBps Data Streaming on MCA, Planar SVGA, C2 Security, 2.88MB, Flothru 3 fans, SynchroStream, double-speed CD-ROM, 470 watts power supply

This soft copy for use by IBM employees only.

This soft copy for use by IBM employees only.

Chapter 2. IBM PC Server Disk Subsystem Options

The following sections provides an overview of important options available for IBM PC Server systems.

2.1 Disk Subsystem Controllers Overview

Optimal throughput for a disk subsystem is very much dependent on the kind of controller technology used in the server system. This section provides a description of disk subsystem controllers used for the IBM PC Server family.

2.1.1 IBM PC Server SCSI-2 Fast/Wide Adapter/A

The IBM SCSI-2 Fast/Wide Adapter/A is a high performance SCSI-2 Fast and Wide (16-bit SCSI bus) adapter capable of 20MBps SCSI bus data transfer rates which provides an interface to PC Server SCSI Options. The adapter supports a data transfer rate of 40MBps (data streaming) over the Micro Channel bus. Support for SCSI Tagged Command Queuing (TCQ) increases performance in DASD-intensive server environments.

Fast refers to a data transfer method. Data is moved to fixed disks at 10MBps, twice the speed of SCSI 1. *Wide* refers to the bus width which is increased from 8 to 16 bits, enabling transfers of up to 20MBps.

The IBM SCSI-2 Fast/Wide Adapter/A also supports standard 8-bit SCSI devices using either asynchronous, synchronous, or fast synchronous (10MBps) SCSI data transfer rates per ANSI Small Computer System Interface 2 (X3T9.2/375R Rev. 10K) for SCSI-2 features.

Key Features:

- 32-bit 40MBps Data Streaming Micro Channel Adapter
- Two independent SCSI-2 16-bit Fast/Wide Channels (with two internal and one external connectors)
- Supports SCSI Tagged Command Queuing Protocol (TCQ)
- Overlapped command processing up to 250 commands
- Compatible with SCSI-1 Adapters and Device Drivers
- Dedicated 80C186 local processor on board
- Fits short Micro Channel slots

This soft copy for use by IBM employees only.

- Attachment of up to fifteen physical SCSI devices

This combination of SCSI-2 Fast and Wide offers substantial performance increases over previous SCSI solutions. It is a very cost effective solution for the server environment. The dual SCSI buses of the IBM SCSI-2 Fast/Wide Adapter/A prevents access to internal DASD (Direct Access Storage Device) from the external port and also allows the maximum cable lengths to be calculated individually for each bus. This allows for additional capability externally.

The IBM SCSI-2 Fast/Wide Adapter/A supports up to fifteen IBM SCSI option devices per adapter and as many as eight adapters per machine may be installed.

This non-array adapter (used in the PC Server 500 and 720 only) supports a total of fifteen SCSI devices. The internal channel supports up to seven SCSI devices, and the external channel supports from one to fourteen SCSI devices. However, the total number of SCSI devices attached to a single SCSI adapter cannot exceed fifteen. Please keep in mind that wide devices (16-bit) support SCSI IDs 0 to 15, but narrow (8-bit) devices support only IDs from 0 to 7.

Note: Some older systems may support only seven devices per adapter and four adapters per system.

This soft copy for use by IBM employees only.




Figure 6. SCSI-2 Fast/Wide Adapter/A

Figure 6 shows the layout of the SCSI-2 Fast/Wide Adapter/A.

External cable length is defined dependent on the device types attached as follows:

Device	Distance
Asynchronous SCSI devices	Up to 6 meters
Synchronous SCSI devices	Up to 6 meters
Synchronous SCSI-2 Fast devices	Up to 4 meters

This adapter has a dedicated 80C186 local processor on board. It supports SCSI Tagged Command Queuing (TCQ) which increases performance in DASD-intensive server environments. With SCSI-1 systems, only one command could be sent to a hard disk. With TCQ it is possible to send multiple commands to the hard disk, and the disk will store the commands

This soft copy for use by IBM employees only.

and execute each command in the sequence which will give optimal performance.

2.1.2 IBM PC Server SCSI-2 Fast/Wide PCI Adapter

The SCSI-2 Fast/Wide PCI adapter can be installed in the PC Servers 300, 320 and 720. This adapter is an intelligent PCI to an SCSI host adapter, providing a powerful multitasking interface between the PCI bus and the SCSI bus.

Multiple installed adapters provide additional SCSI channels which can be used to increase both the capacity and performance of your server's disks subsystem by adding drives, and distributing workload over multiple SCSI channels.

An on board RISC sequencer automates SCSI command processing, and reduces interrupts and command processing overhead. The 32-bit busmastering minimizes your computer's microprocessor overhead by transferring data into system memory directly. The adapter negotiates with devices on the SCSI bus to ensure that data transfer rates are acceptable by both the host computer and the SCSI devices. This prevents data from being lost and ensures error-free data transmissions.

This soft copy for use by IBM employees only.




Figure 7. SCSI-2 Fast/Wide PCI Adapter

Figure 7 shows the layout of the SCSI-2 F/W PCI adapter and the three connectors to attach SCSI devices, one internal and one external 16-bit connector, and one 8-bit internal SCSI connector. It can control up to fifteen 16-bit devices (SCSI-2 F/W) and up to eight 8-bit SCSI devices (SCSI-1 or SCSI-2 fast).

Adapter device drivers are available to support the most commonly used operating and network-operating systems such as OS/2, NetWare, Novell UnixWare, SCO-Unix, and Windows NT.

The adapter supports:

- Concurrent support of 8 and 16-bit SCSI-devices
- Single ended device support only (majority of current SCSI devices)
- Disk drives with capacities up to 8GB
- Up to 10 MBps data rate (SCSI-2 fast) for devices attached to the 8-bit SCSI connector

This soft copy for use by IBM employees only.

- Up to 20 MBps data rate (SCSI-2 fast/wide) for devices attached to the internal or external 16-bit SCSI connector
- 32-bit bus master transfers for effective and fast data-transfer to/from the host computer
- Up to 132 MBps data transfer rate on the PCI-bus
- Scatter/Gather to support data transfer to/from non-contiguous memory
- Onboard RISC sequencer for low SCSI-processing overhead
- SCSI-bus parity control for data integrity
- Tagged Command Queueing (TCQ)
- Disconnect/reconnect
- Multiple adapters to increase capacity and performance
 - PC Server 300: two adapters
 - PC Server 320: three adapters

SCSI devices can be attached to the adapter as follows:

- Up to seven 8-bit devices (SCSI-1, SCSI-2 and SCSI-2 fast) connected to the 8-bit internal connector.
 - Up to fifteen 16-bit (SCSI-2 fast/wide) devices connected to the internal 16-bit SCSI-connector
- or
- Up to fifteen 16-bit (SCSI-2 fast/wide) devices connected to the external 16-bit SCSI-connector

The internal and the external 16-bit SCSI connectors *must* not be used concurrently, that is, you can attach only fifteen 16-bit SCSI devices internally *or* externally but *concurrent* attachment of SCSI devices to both 16-bit connectors is *not* supported!

2.1.3 IBM Streaming RAID Adapter/A

The IBM Streaming-RAID Adapter/A is a high performance SCSI-2 fast and wide (16-bit SCSI bus) RAID adapter capable of 20MBps SCSI bus transfer rates. The combination of SCSI-2 Fast/Wide offers substantial performance increases over current SCSI solutions. RAID offers the additional data protection security inherent in RAID configurations.

This adapter offers functional advantages by supporting a wide range of IBM options, including the hot swappable fast and wide hard drives. The SCSI-2

This soft copy for use by IBM employees only.

Fast/Wide Streaming-RAID Adapter/A supports a 40MBps Micro Channel streaming data rate. Firmware 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 adapter also supports 8-bit SCSI devices using either asynchronous, synchronous, or fast synchronous (10MBps) SCSI data transfer rates per ANSI Small Computer System Interface 2 (X3T9.2/375R Rev. 10K) for SCSI-2 features. The dual bus of the adapter allows for maximum connection of up to fourteen drives, seven on each individual bus. For example, one bus cannot support internal and external devices simultaneously.




Figure 8. SCSI-2 Fast/Wide RAID Adapter Micro Channel

This soft copy for use by IBM employees only.

Figure 8 shows the Streaming RAID Adapter/A and the optional daughter card that allows the connection of both Fast/Wide channels externally. This daughter card connects with an internal cable to the second internal Fast/Wide port and occupies the next Micro Channel slot providing an external SCSI connector. This daughter card is only included, when the Streaming RAID Adapter/A is ordered as an additional option.

Key Features

- 32-bit, 40MBps Data Streaming Micro Channel Adapter
- Dual SCSI-2 16-bit Fast and Wide Channels, 20MBps, external and internal ports
- Supports SCSI Tagged Command Queuing Protocol (TCQ)
- i960 RISC processor
- Maximum connection of up to 14 devices, seven on each bus
- RAID 0, 1, hybrid 1, and 5 support
- 4MB of RAID HW disk cache, 60ns speed, configurable as Write-Thru (default) or Write-Back

2.1.4 IBM PC Server SCSI-2 Fast/Wide PCI-Bus RAID Adapter

This high-performance adapter has the same capabilities as the RAID adapter for Micro Channel systems with the only difference being this adapter is connected to the PCI bus with a 132MBps of throughput.

Figure 9 on page 53 shows the layout of the RAID PCI adapter and the position of the i960 RISC processor and the 4MB memory chip.

This soft copy for use by IBM employees only.




Figure 9. SCSI-2 Fast/Wide RAID PCI-Bus Adapter

2.1.5 RAID Adapters Features

Features which both PCI and Micro Channel RAID adapters have in common:

- **Support for Hybrid RAID 1, in addition to RAID-0,1 and 5**

With this capability you can configure RAID-1 arrays with odd or even number of drives.

- **Overlapped Input/Output Operation**

Because the RAID adapter provides multiple data paths to and from arrayed drives, your server can respond to requests from several users simultaneously. With its overlapped input/output operation, if one user requests data that resides on the first drive of the array, and a second user requests data that resides on the second drive, the adapter can simultaneously deliver both pieces of information.

This soft copy for use by IBM employees only.

- **Interleave Depth/Stripe Unit**

The granularity at which data from one file is stored on one drive of the array before subsequent data is stored on the next drive of the array is called the *interleave depth*. The interleave depth can be set for optimum system performance to stripe unit size from 8K to 64K bytes of data.

The collection, in logical order of these units, from the first drive of the array to the last drive of the array, is called a *stripe*.

- **Queue Depth**

Commands are queued in the adapter with a queue depth of 61. To obtain better performance, the commands in the queue will be *reordered* and *coalesced* on a hard disk drive basis. That is, the adapter organizes the commands according to which drive will be responding, and then orders and combines two or more commands, when possible, before sending them to the drives.

Figure 10 shows the configuration screen used to configure the stripe size.




Figure 10. RAID Advanced Functions Configuration Screen

This soft copy for use by IBM employees only.

- **Cache**

The RAID Adapter has 4MB of memory, of which more than 3MB is used as a data cache memory. This cache memory can be configured to operate in a *write-through* or *write-back* mode on a logical drive basis. Cache memory has parity to detect memory errors and retry algorithms to recover from errors that appear sporadically.

- **Hot Spare Drive Replacement**

A *hot spare drive* is a hard disk drive in your server that is defined for automatic use in the event of a drive failure. The hot spare drive must be of equal or greater capacity than the drives in the array that it is intended to replace. You can define as many hot spare drives as you want.

If a drive fails, the system automatically switches to the hot spare drive, and the data from the de-functional drive is reconstructed on the hot spare drive. The system automatically defines the replacement drive as a hot spare.

Only one hot spare drive is needed for both SCSI-channels.

- **One Array can span both SCSI-Channels**

This doubles the SCSI bus bandwidth to 40MBps and provides maximum performance.

- **Data Protection**

This Adapter provides the server with a data redundancy technique (RAID-1 or RAID-5) that distributes data across all the drives in the array. In the event of a single drive failure, read and write requests are serviced by the remaining disks in the array, without loss of data. The defunct hard disk drive *appears* to be operating correctly, if at reduced performance, despite its total failure.

When you replace the defunct drive, the data in logical drives assigned RAID level 1 and 5 is rebuilt on the new hard disk drive using the parity information stored on the remaining hard disk drives of the array. It is important for performance and reliability reasons, that the defunct drive be replaced as soon as possible. Once the drive contents are rebuilt, full performance and fault-tolerance functions are restored.

This soft copy for use by IBM employees only.

2.1.6 SCSI-2 Fast Adapters

Factory installed with the entry systems PC Server 300 and 310 family, are adapters that support the SCSI-2 Fast specification. These SCSI-2 Fast solutions support the type of workload that these servers are planned for.

IBM SCSI-2 Fast solutions:

- PC Server 300

This machine is equipped with the IBM SCSI-2 Fast PCI adapter. This adapter and a supporting internal SCSI 4-drop cable comes standard on the PC Server 8640-0PT.

The IBM SCSI-2 Fast PCI-Bus Adapter is a high-performance PCI SCSI-2 adapter for internal and external attachments of up to seven SCSI or SCSI-2 Fast devices. To attach external SCSI devices, use the PCI SCSI-2 Fast card (50 pin) to SCSI/SCSI-2 Fast (50-pin) cable . If you are adding this adapter to the PC Server and you want internal drive support, order PC Server Cable Kit. This kit includes a 7-drop SCSI cable for support of the PCI SCSI-2 Fast card, and a 2-drop IDE cable for support of two additional drives off the second IDE port on the 486 models.

- PC Server 310

The PC Server 310 is using an Adaptec 2940 compatible PCI Fast controller. The features of the adapter are identical with the adapter used for the PC Server 300.

Both these adapters support 10MBps data transfer rates, and provide an (Advanced SCSI Programming Interface) ASPI manager and features Tag Command Queuing support (TCQ) for popular network operating systems.

2.1.7 Summary of SCSI Adapters

Table 12 summarizes the features of the IBM SCSI Fast adapters used for the IBM PC Servers.

Attribute	SCSI-2 Fast PCI Adapter	SCSI-2 Fast PCI Adapter
IBM PC Server	PC Server 300	PC Server 310
U.S Part Number	32G3100	N/A 1
EMEA Part Number	SSPC01Y	N/A 1
SCSI Type	SCSI-2 Fast	SCSI-2 Fast

This soft copy for use by IBM employees only.

<i>Table 12 (Page 2 of 2). IBM SCSI-2 Fast Adapter Summary</i>		
Attribute	SCSI-2 Fast PCI Adapter	SCSI-2 Fast PCI Adapter
HW RAID Level Supported	None	None
I/O Bus	PCI	PCI
I/O Bus Speed	132MBps	132MBps
SCSI Channels	1	1
Maximum # Connectable Devices	7 (8-bit)	7 (8-bit)
On-board Processor	Yes	Yes
SCSI-1 & SCSI-2 Compatibility	Yes	Yes
Note:		
1 The PCI Bus SCSI-2 Fast Adapter for the PC Server 310 is not available as option.		

Table 13 summarizes the features of the SCSCI Fast/Wide adapters used for the IBM PC Servers.

<i>Table 13 (Page 1 of 2). IBM SCSI Fast/Wide Adapter Summary</i>				
Attribute	SCSI-2 Fast/Wide Adapter/A	SCSI-2 Fast/Wide PCI Adapter	SCSI-2 Fast/Wide Streaming RAID Adapter/A	SCSI-2 Fast/Wide RAID PCI-Bus Adapter
IBM PC Server	PC Server 500	PC Server 320, PC Server 520 1	PC Server 500, PC Server 720	PC Server 320, PC Server 520
U.S Part Number	70G8498	94G3771	70G9263	94G2764
EMEA Part Number	6451280	SPSW01Y	PRADADX	SPSR01Y
SCSI Type	SCSI-2 F/W	SCSI-2 F/W	SCSI-2 F/W	SCSI-2 F/W
RAID Level Supported	None	0,1,5	None	0,1,5
I/O Bus	Micro Channel	PCI	Micro Channel	PCI

This soft copy for use by IBM employees only.

Table 13 (Page 2 of 2). IBM SCSI Fast/Wide Adapter Summary

Attribute	SCSI-2 Fast/Wide Adapter/A	SCSI-2 Fast/Wide PCI Adapter	SCSI-2 Fast/Wide Streaming RAID Adapter/A	SCSI-2 Fast/Wide RAID PCI-Bus Adapter
I/O Bus Speed	40MBps	132MBps	40MBps	132MBps
SCSI Channels	2 2	1 3	2 4	2 4
Maximum # Connectable Devices per Channel	7 5	15 (16-bit) 7 (8-bit)	7 6	7 6
On-board Processor	Yes	Yes	Yes	Yes
SCSI-1 & SCSI-2 Fast Compatibility	Yes	Yes	Yes	Yes

Note:

- 1** Onboard PCI adapter in the Micro Channel/PCI models of PC Server 320 and 520
- 2** Two independent SCSI channels:
 - One internal channel with a 16-bit and a 8-bit connector (only one connector can be used)
 - One external channel with a 16-bit connector
- 3** One shared SCSI channel is provided:
 - Both internal 8-bit and internal 16-bit connectors can be used
 - or
 - Either internal connector (8 or 16-bit) and external 16-bit connector can be used
- 4** Two independent SCSI Channels are provided:
 - One shared between internal 16-bit and external 16-bit connectors
 - One dedicated channel with the internal 16-bit connector
- 5** If system BIOS level is less or equal to Version 3 then up to seven devices per channel are supported, if system BIOS level is greater than Version 3 up to fifteen devices per channel are supported.
- 6** Up to fourteen 16-bit devices are supported; up to 7 devices per channel, all of them can be internal (using the two internal connectors) or external (using the special daughter card to route the internal cable to an external connector).

This soft copy for use by IBM employees only.

2.2 Hard Disk Drives

All the PC Servers use hard disk drives to store data and perform the server function. The user relies on the quality and design of the hard disk drives to insure the integrity of the information that is stored on the drive.

The hard disk drive is the only component that when non-operational will significantly impact the use of the computer. Other components can be repaired at the cost of a few hours delay. However, without the ability to retrieve the data, the user is unable to use his/her computer.

Even when the data has been backed up, it requires a significant amount of time to rebuild the new hard disk drive and in most likelihood the backed up data is not recent. The worst feeling that computer users experience is to be confronted with a dead drive!

The purchaser of a PC server should consider not only the cost of the hardware but also the value and replacement cost of the data that will be stored in it. Only the hard disk drive can store and retrieve massive amounts of data. Thus, when making a PC server purchase, particular attention needs to be placed on the quality and reliability of the hard disk drive.

2.2.1 The Ultrastar Family

These drives are ideal for:

- Critical control
- Communications applications
- High-end technical and scientific workstations
- High-speed data storage in minicomputers
- Host-based systems
- High-performance storage in network servers
- Multimedia and digital video applications

The Ultrastar XP SCSI drives adhere to American National Standards Institute (ANSI) standard X3T9.2/375R, Revision 10K, dated March 1993, and are designed to be compatible with systems configured with SCSI-2 controllers that support this ANSI specification.

The Ultrastar XP SCSI Hard Disk Drives are based on the industry's most advanced recording technologies and meet the speed and storage

This soft copy for use by IBM employees only.

requirements demanded by applications where high, sustained transfer rates must be maintained and throughput optimized.




Figure 11. Ultrastar Hard Disk Drive Mechanical Design

Figure 11 shows the structure of an Ultrastar hard disk drive. To archive high capacity, many plates are included.

The Fast SCSI-2 interface is best suited in systems where one or two Ultrastar XP drives share a common SCSI controller. When three or more Ultrastar XP drives must share a common SCSI controller, the Fast/Wide SCSI-2 interface provides increased data path capacity.

Table 14 on page 61 summarizes the technical specifications for the entire Ultrastar Family. The values shown are important to be considered for performance and reliability evaluations and configurations.

This soft copy for use by IBM employees only.

<i>Table 14. Technical Characteristics of UltraStar Disk</i>					
Formatted capacity (GB)		1,12	2,25	4,51	5,31
Model Number	US, LA	94G2649 94G3052	94G3054 94G3055 94G2650	94G3196 94G3057 94G2651	94G3059 94G3787
	EMEA	OSW1GBM S1GHS1Y	OSW2GBM OSC2GBM S2GHS1Y	OSW4GBM S4GHS1Y	N/A
SCSI-2 Interface 1	Fast	N/A	OSC2GBM 94G3054	94G3196	94G3787
	F/W	OSW1GBM 94G3052	OSW2GBM 94G3055	94G3057 OSW4GBM	94G3059
	F/W Hot-Swap	94G2649 S1GHS1Y	94G2650 S2GHS1Y	94G2651 S4GHS1Y	N/A
Rotational speed (RPM)		7200	7200	7200	5400
Typical average read seek time (ms)		6.9	7.5	8.0	8.0
Number of disks		2	4	8	8
Average latency (ms)		4.17	4.17	4.17	5.55
Data buffer (KB)		512	512	512	512
MTBF - Projected power on hrs (Khs) 2		1000	1000	1000	1000
Typical sustained transfer rate (MBps) 3		5.5-7.4	5.5-7.4	5.5-7.4	4.4-7.6
Bay required	Non Hot-Swap	1/2 H 4	1/2 H	1/2 H	1/2 H
	Hot-Swap	1/3 H 5	1/3 H	2 x 1/3 H	N/A
<p>Note:</p> <p>1 For the <i>Fast</i> Interface 10MBps of maximum synchronous data transfer rate can be obtained, with the <i>Fast/Wide</i> Interface the maximum is 20MBps.</p> <p>2 MTBF includes early life and end of life failures. Product description data represents IBM's design objectives and is provided for comparative purposes; actual results may vary depending on a variety of factors. This product data does not constitute a warranty. Questions regarding IBM's warranty terms or methodology used to derive data should be referred to your IBM representative.</p> <p>3 Sustained transfer rate is a measure of the typical time used to transfer data between the media and the drive's data buffer. The rate is given as a range because it is affected by the number of sectors per track, which are different in each zone or notch, and by the bytes per sector, track and cylinder skew, error recovery, and command overhead.</p> <p>4 Half high 5.25" bay, also called 1.5" bay.</p> <p>5 Hot Swap 1/3 high, also called 1" bay.</p>					

This soft copy for use by IBM employees only.

2.2.2 The Deskstar Family

The target market for this family of hard disk drives is the workgroup or small enterprise network. These drives adhere to the SCSI-2 Fast specification and have a rotational speed of up to 5400 rpm, with good access times. They are the best cost/performance solution for small networks. For high-performance servers, the Ultrastar family should be considered. The 1.08GB and 540MB hard disk drives are a reduced cost solution to increase the hard disk capacity for the PC Server 3XX family.

Figure 12 shows the mechanical design of the Deskstar hard disk drive.



Figure 12. Deskstar Hard Disk Drive Mechanical Design

Table 15 on page 63 shows the technical specifications for the Deskstar hard disk drive family.

This soft copy for use by IBM employees only.

<i>Table 15. Technical Characteristics of Deskstar Disk</i>			
Formatted capacity		540MB	1,08GB
Model Number	US, LA	06H8558	94G3187
	EMEA	N/A	OSC1GBM
SCSI-2 Interface 1	Fast	06H8558	94G3187, OSC1GBM
	F/W 1	N/A	N/A
	F/W hot swap 2	N/A	N/A
Rotational speed (RPM)		4500	5400
Typical average read seek time (ms)		12.0	10.5
Number of disks		1	2
Average latency (ms)		6.67	5.56
Data buffer (KB)		128	512
MTBF - Projected power on hrs (Khs) 3		350	500
Typical Sustained transfer rate (MBps) 4		2.5-3.7	3.2-4.0
Bay required	Non hot swap	1/2 H	1/2 H
	hot swap	N/A	N/A
<p>Note:</p> <p>1 If the Fast/Wide is needed, consider de Ultrastar Family.</p> <p>2 If the hot swap capabilities are needed, consider the Ultrastar Family.</p> <p>3 MTBF includes early life and end of life failures. Product description data represents IBM's design objectives and is provided for comparative purposes. Actual results may vary depending on a variety of factors. This product data does not constitute a warranty. Questions regarding IBM's warranty terms or methodology used to derive data should be referred to your IBM representative.</p> <p>4 Sustained transfer rate is a measure of the typical time used to transfer data between the media and the drive's data buffer. The rate is given as a range because it is affected by the number of sectors per track, which are different in each zone or notch, and by the bytes per sector, track and cylinder skew, error recovery, and command overhead.</p>			

2.3 External Disk Subsystems

Sometimes the internal disk subsystem expansion possibilities are not enough for certain configurations. For fault tolerance or reliability purposes, external enclosures can be used to increase the amount of storage or add new functionality with different types of devices needed for a customized PC Server solutions. This section gives an overview of the different SCSI enclosures available for IBM PC Server products.

2.3.1 3516 Hot Swap Expansion Enclosure

Vital data to your LAN, and therefore for your business, is stored on your server. Protecting it from disk failure is essential. The 3516 Hot Swap Expansion Enclosure and the IBM SCSI-2 Fast/Wide Streaming-RAID Adapter/A are designed for that task. The 3516 Hot Swap Expansion Enclosure brings the safety of Redundant Array of Inexpensive Disks (RAID) technology to protect vital data. To reduce costly server downtime this IBM RAID solution allows you to install a spare drive, which can automatically take the place of the failed disk and even replace the defective drive without stopping the server.

The IBM SCSI-2 Fast/Wide Streaming-RAID Adapter/A is designed using the latest technology to minimize any performance overhead incurred by RAID. It has two SCSI-2 Fast/Wide channels that can transfer data at 20MBps. It uses the powerful Intel i960 RISC processor to perform the RAID calculations. To boost performance and minimize accesses to disks, 4MB of disk cache is provided. The adapter is a full 32-bit Micro Channel bus master which is capable of streaming data to and from system memory at 40MBps without assistance from the main CPU. The result of this extensive range of performance features is that the IBM RAID solution out performs many competitive non-RAID solutions.

This soft copy for use by IBM employees only.




Figure 13. 3516 Hot Swap SCSI Enclosure

Figure 13 shows an external view of the 3516 Hot Swap Enclosure.

The compact 3516 Hot Swap Expansion Enclosure attaches to the IBM SCSI-2 Fast/Wide Streaming-RAID Adapter/A and provides space for seven hot swap

This soft copy for use by IBM employees only.

disks. Two enclosures can be attached to an adapter and two adapters to a server. This enables the attachment of twenty eight disks to a single server.

2.3.1.1 Summary

Table 16 provides an overview of the key benefits of the 3516 Hot Swap Expansion Enclosure.

<i>Table 16. Key Benefits of 3516 Hot Swap Expansion Enclosure</i>		
3516 Hot Swap Expansion Enclosure	Benefit	Reason
Availability	Excellent	RAID protection against disk failure. On line option to install a spare drive that will automatically take over from the failed drive, without halting the server. Hot swap ability to replace failed drives without halting the server. LAN Server or NetWare.
Performance and Capacity	Excellent	Two SCSI 2 FAST/WIDE Channels. Data can be transferred to the disks at 20MBps. Ability to support fourteen drives per adapter using two 3516 Hot Swap Expansion Enclosure. 4MB disk cache improves performance by minimizing disk accesses. Most frequently used data is stored in fast cache memory. 32-bit streaming bus master adapter improves performance by enabling data to be transferred to and from memory at 40MBps without processor intervention.
Server Implementation	Excellent	Hot swappable disks and spare disks.
Use of Architectures	Excellent	SCSI-2, streaming data, Raid 0,1,5.

This soft copy for use by IBM employees only.

2.3.2 3511 External Enclosure for SCSI Devices

The 3511 External Enclosure for SCSI Devices enables up to seven additional SCSI devices to be housed in an external tower, and attached to the SCSI adapter in the PC Server. Since up to four SCSI adapters can be installed in a PC Server, the 3511 External Enclosure would typically be attached to its own SCSI adapter in order to spread the load on the adapters and SCSI buses.

Installing the 3511 External Enclosure enables additional disk storage to be attached to a PC Server (thus increasing the overall capacity). Also, it provides extra flexibility in systems implementation, since the 3511 External Enclosure could be easily unplugged from a failing PC Server and attached to a replacement PC Server.

2.3.2.1 Key Features

The 3511 External Enclosure is attached by a SCSI connecting cable to the SCSI adapter in the PC Server and consists of the following features:

- An external tower, similar in appearance to the PC Server Server 95, with its own power supply
- Up to seven bays for SCSI devices (two half-high 3.5" bays, one half-high 5.25" bay, and two full-height 5.25" bays, which can each be also used for two half-high devices instead)
- The ability to install fixed disks or other SCSI devices, such as tape backup, CD-ROM or re-writable optical disks in these bays

The 3511 External Enclosure could be used to implement a high-availability disk subsystem in conjunction with disk mirroring (using LAN Server or NetWare). Alternatively, it could be used to implement a high-performance disk subsystem in conjunction with NetWare data spanning (RAID-0). In both cases, the use of NetWare data spanning, combined with multiple 3511s attached to multiple SCSI adapters in the PC server, would increase performance still further.

2.3.2.2 Summary

Table 17 provides an overview of the key benefits of the 3511 External Enclosure.

3511 External Enclosure	Benefit	Reason
Availability	Good	Separate power supply; RAID-1 Disk Mirroring in conjunction with LAN Server or NetWare.

This soft copy for use by IBM employees only.

<i>Table 17 (Page 2 of 2). Key Benefits of 3511 External Enclosure for SCSI Devices</i>		
3511 External Enclosure	Benefit	Reason
Performance and Capacity	Very good	RAID-0 Data Striping using NetWare
System Security	Good	External disk subsystem can be located separately from the PC Server, lockable cover
Server Implementation	Good	Can be easily unplugged from a failing PC Server and reconnected to the replacement PC Server
Use of Architectures	Very good	SCSI

This soft copy for use by IBM employees only.

2.3.3 3510 SCSI Storage Enclosure

The 3510 SCSI Storage Enclosure enables one SCSI device to be housed in an external box, and attached to the SCSI adapter in the PC Server. Multiple 3510s can be attached to either the same or different SCSI adapters within the PC server.

Installing the 3510 Storage Enclosure both enables additional disk storage to be attached to a PC server (thus increasing the overall capacity), and also provides extra flexibility in systems implementation, since the 3510 Storage Enclosures can be easily unplugged from a failing PC Server and attached to a replacement PC Server. In addition, if the disk in the 3510 Storage Enclosure fails but it is part of a RAID-5 array, then a new disk can easily be plugged into the PC Server and the array rebuilt.




Figure 14. 3510 SCSI Enclosure

Figure 14 shows the external view of the 3510 SCSI Enclosure.

In addition to have hard disk configurations with the 3510 enclosure, this is a very flexible solution to provide specialized devices easy external installation (that is, backup tapes, optical drives, CD-ROMs).

This soft copy for use by IBM employees only.

2.3.3.1 Key Features

The 3510 Storage Enclosure consists of the following features:

- A housing unit, similar in appearance to an external CD-ROM, with its own power supply
- One half-height bay for a SCSI disk, either 3.5" or 5.25"

It is attached by a SCSI connecting cable to the SCSI adapter in the PC Server.

Multiple 3510s could be used to implement a high-availability disk subsystem in conjunction with disk mirroring (using LAN Server or NetWare).

Alternatively, multiple 3510s could be used to implement a high-performance disk subsystem in conjunction with NetWare data spanning (RAID-0).

In both cases, spreading the attachment of the 3510 Storage Enclosures across multiple SCSI adapters in the PC Server would increase performance still further.

2.3.3.2 Summary

Table 18 provides an overview of the key benefits of the 3510 SCSI Storage Enclosure.

3510 STORAGE ENCLOSURE	Benefit	Reason
Availability	Good	Separate power supply; RAID-1 Disk Mirroring in conjunction with LAN Server or NetWare
Performance and Capacity	Good	RAID-0 Data Striping using NetWare
System Security	Good	3510 can be located separately from the PC Server
Server Implementation	Good	Flexibility of re-plugging disk drives if either a PC Server or one drive in a RAID-5 disk array fails
Use of Architectures	Very good	SCSI

This soft copy for use by IBM employees only.

2.3.4 SCSI Storage Enclosure Summary

Table 19 gives an overview of the most important features for each of the previously discussed external SCSI enclosures.

US, LA Part Number	EMEA Part Number	Description
35167V0	SDHSEE1	DASD Hot Swap Expansion Enclosure (Array or Hot Swap Tower): Provides 7 Hot Swap drive bays for SCSI-2 Fast or SCSI-2 Fast/Wide Hot Plug drives. Includes a SCSI-2 Fast/Wide hot swap cable internally. Connects to either the SCSI-2 Fast/Wide Adapter/A (#70G8498) or SCSI-2 Fast/Wide RAID Adapter (#70G9263) using cable(#06H3231).
3510AV0	6451375	SCSI Enclosure: external enclosure for a single SCSI device (either SCSI-1, SCSI-2 Fast) or SCSI-2 Fast/Wide if you use the 3510-0V0 SCSI-2 Fast/Wide Upgrade Kit (#32G3920), 32 watts power supply.
35110V0	32G2964	SCSI Tower, external tower enclosure holding up to seven half-high bays, non hot swappable. Standard Tower configuration uses SCSI-1 or SCSI-2 Fast devices. If you want SCSI-2 Fast/Wide support, order the 3511-0V0 SCSI-2 Fast/Wide Upgrade Kit (#06H3153), 400 watts power supply.

2.4 Other IBM Options for the IBM PC Server

This section summarizes some other options that should be considered when configuring PC Server solutions.

2.4.1 CD-ROM Drives

CD-ROM has typically been used in environments where applications use large quantities of static or periodic data. Some examples are dictionaries, encyclopedias, online sales manuals, and parts books. In many LAN environments it has become necessary to share this type of media.

With the size of software ever increasing, it is becoming cumbersome to distribute on the diskette. CD-ROM (Compact Disc - Read Only Memory) drives and media are becoming more popular as they are able to offer large capacities and they are affordable. The following drives are available from IBM as options to be used with the IBM PC Server:

- IBM Internal SCSI 2X CD-ROM Drive
- IBM External SCSI 2X CD-ROM Drive

This soft copy for use by IBM employees only.

- IBM Internal SCSI 4X CD-ROM Drive
- IBM External SCSI 4X CD-ROM Drive

2.4.1.1 IBM External and Internal SCSI CD-ROM 2X Drives

Since the external CD-ROM drive is functionally equivalent to the internal version they will be described together.

The SCSI 2X CD-ROM Drive shown in Figure 15 is a synchronous SCSI high-performance, high-reliability compact disc read-only memory drive. The drive functions with compact discs compatible with: ISO 9660, High Sierra, Kodak** Photo CD (Orange Book), and CD-Digital Audio (Red Book) formats. The drive performs at a maximum sustained data rate of 372KBps and has an average access of 205ms, among the fastest available.




Figure 15. Internal SCSI 2X CD-ROM Drive

Key Features:

- High performance synchronous SCSI-2 Fast CD-ROM

This soft copy for use by IBM employees only.

- Industry standard SCSI-2 Fast Interface
- Industry standard compact disc (ISO 9660, High Sierra, and others)
- 600MB typical formatted capacity (disc dependent)
- Average Access Time
 - 1 x speed maximum sustained 330 ms
 - 2 x speed maximum sustained 205 ms
- Data Rate
 - 1 x speed maximum sustained 150 KBps
 - 2 x speed maximum sustained 327 KBps
 - Burst 3.0MBps
- 256KB Data Buffer
- CD-XA capable
- Universal 32 Watt power supply (applies to external drive only)
- Electronic eject
- Software volume control
- Independently certified as compatible with a broad mix of adapters by XXCAL Testing Laboratories
- Standard 5-1/4" half-high form factor
- Analog audio output
- Security, auditability and control
 - A mechanical lock to allow the user to lock the system cover.
 - A user-invoked password function to prevent the use of the system by unauthorized persons.
 - Enhanced security feature (Privileged-Access Password) when used with some systems.
 - The drive is enabled for software controlled eject and the ability to lock media in or out of the drive.

This soft copy for use by IBM employees only.

<i>Table 20. CD-ROM Feature Summary</i>	
Feature	IBM Internal and External SCSI CD-ROM 2X Drives
Performance and Capacity	Maximum Sustained Data Rate of 372KBps Average access of 205ms 256KB Data Buffer 600MB typical formatted capacity (disc dependent)
Compatibility	Accepts Industry Standard Compact Disc ISO 9660, High Sierra Compatible with a broad range of adapters
Security, Auditability and Control	Inherits System Security when installed internally Ability to lock media in or out of the drive

2.4.1.2 IBM External and Internal SCSI CD-ROM 4X Drives

The SCSI CD-ROM 4X Drive is the fastest available CD-ROM Drive from IBM to be used at the IBM PC Servers.

With 4X performance (600KB/sec data transfer rate) and MPC-2 compliance, these drives are the ideal solution for running today's demanding multimedia applications. This is the perfect solution to archive the better performance when you need to share the CD-ROM in a LAN.

A unique tray design allows the Internal SCSI drives to be mounted horizontally or vertically. These drives provide access to the growing list of software and databases now available on CD, and play audio CDs through the front panel headphone jack or through an optional sound card installed in the system

The technical specifications of this device are the same as described at section 2.4.1.1, "IBM External and Internal SCSI CD-ROM 2X Drives" on page 72, but the performance values are different as follows:

<i>Table 21 (Page 1 of 2). CD-ROM 2X and 4X Comparison Summary</i>		
Feature	SCSI CD-ROM 2X Drives	SCSI CD-ROM 4X Drives
Maximum sustained data rate	372KBps	600KBps
Burst data rate(MBps)	3MBps	5MBps
Data Buffer	256KB	256KB
Access time	300ms	200ms

This soft copy for use by IBM employees only.

<i>Table 21 (Page 2 of 2). CD-ROM 2X and 4X Comparison Summary</i>		
Feature	SCSI CD-ROM 2X Drives	SCSI CD-ROM 4X Drives
Standards	ISO 9660, High Sierra Group, Kodak Photo CD (Orange Book), CD-ROM XA (Yellow Book) and CD-Digital Audio (Red Book)	ISO 9660, High Sierra Group, Kodak Photo CD (Orange Book), CD-ROM XA (Yellow Book) and CD-Digital Audio (Red Book), MPC-2 compliant

2.4.2 Tape Drives

This section describes the tape options available for the IBM PC Server. With the large capacities of disk storage that servers are required to maintain, it has become mandatory that the server is equipped with a backup device. It is imperative that you maintain good backup procedures together with reliable media and backup devices.

The choice of device or media type may depend on a number of factors. For example, an unattended backup would mean that your device and media would have to support capacities greater than your current disk capacity. The following tape drives are available from IBM as options.

2.4.2.1 IBM 3450 1.2GB 1/4-Inch Cartridge Tape Drive Model 001

The IBM 3450 Model 001 is an external, self-contained 1.2GB 1/4-inch tape drive. The 3450 is a low cost 1/4-inch cartridge solution for larger capacity higher performance tape drives for personal computers. The 3450 is a 1.2GB 1/4-inch tape drive that attaches to a SCSI adapter in the host personal system. It supports many quarter-inch cartridge (QIC) formats. Software support for the 3450 comes in a wide variety of tape backup/restore utilities that offer support in DOS, OS/2, SCO-UNIX, and Novell NetWare environments.

Key Features:

- 1/4-inch tape drive.
- Provides up to 1.2GB of data per cartridge.
- Enhanced performance with a data transfer rate of 300 KBps.

This soft copy for use by IBM employees only.

- The 3450 can read and write tapes in QIC standard QIC-120, QIC-150, QIC-525, and QIC-1000 formats.
- SCSI I command support with a single-ended interface.
- 210KB data buffer.

2.4.2.2 IBM 3445 5.0GB 8mm Tape Drive Model 001

The IBM 3445 Model 001 is an external, self-contained 5.0GB/10.0GB 8mm tape drive. It is ideally suited for unattended backup in LAN servers or any other large capacity disk storage applications. Software support for the 3445 comes in a wide variety of tape backup/restore utilities that offer support in DOS, OS/2, SCO-UNIX, and Novell NetWare environments.

2.4.2.3 Key Features

- An 8mm helical scan tape drive.
- Capacity of 5GB of data per cartridge.
- Capacity can be extended to 10GB of data by enabling its data compression hardware.
- The data transfer rate of the 3445 is 500 KBps and can be increased to 1MBps (typical) when using data compression.
- 1/2-high design.
- Fewer mechanical parts.
- Gentler media handling.
- SCSI II command support with a single-ended interface.
- IDRC Compression.
- Downloadable firmware.
- 1MB data buffer.
- Clean Drive Indicator.
- Host independent diagnostics using IBM diagnostics cartridge.
 - Automatic self diagnostics upon insertion
 - No host intervention required
 - Failure data collection

This soft copy for use by IBM employees only.

2.4.2.4 IBM 3440 2.0GB 4mm Tape Drive Model 001

The IBM 3440 Model 001 is an external, self-contained 2.0GB 4mm tape drive. The 3440 offers an industry leading state-of-the-art technology, low cost-per-megabyte 4mm tape solution for the IBM PC Server. It addresses the growing requirements for unattended backup of LAN servers or any other large capacity disk storage applications. The small media size is easy to handle and even easier to store. The 3440 contains onboard data compression, which increases its capacity to 4-8GB with 5 GB being typical. There is a wide variety of tape backup/restore software available for DOS, OS/2, SCO-UNIX, and Novell NetWare environments.

Key Features

- 2.0GB capacity per cartridge
- Onboard data compression increases the capacity to 4-8GB of data with 5GB being typical
- 30 second average file search speed simplifies data retrieval
- Features non data grade media rejection for write operations (will not allow writing on audio grade media)
- A data transfer rate of 183 KBps up to 732 KBps depending on the type of data transferred and use of data compression
- DDS-Data Compression Format
- SCSI II command support with a single ended interface
- Synchronous and Asynchronous SCSI single ended interface
- 5MBps SCSI burst rate
- Three status indicators
 - Ready
 - Read-Write (media activity)
 - Disturbance (cleaning required)
- Diagnostics Cartridge
 - Automatic self diagnostics upon insertion
 - No host intervention required
 - Failure data collection

This soft copy for use by IBM employees only.

2.4.2.5 IBM 4/10GB 4mm DAT Tape Drive Option Kits

The new 4/10GB 4mm Digital Audio Tape (DAT) option kit is an industry leading state-of-the-art technology, low cost-per-megabyte 4mm tape solution for the IBM PC Server family. With its high capacity, it is the ideal solution for unattended backup of LAN Servers. The small, easy-to-handle media is even easier to store. The 4/10GB 4mm tape option contains onboard data compression that increases its capacity to 8-16GB with 10GB being typical.

There are two versions of the drive available. It is available as the 4mm 3.5" Tape Option Kit for the Array models of the Server 95 (Part# 74G8632). It is also available as the 4mm 5.25" Tape Option Kit for all other Server Models, that is, the 8640 and 8641 ³. The 5.25" drive can also be used in the 3510 and 3511 SCSI Enclosures.

NovaBack with Quick Start Reference, from NovaStor Corporation, is provided with the 5.25-inch option kit. NovaBack provides a full featured backup and restore solution for DOS, OS/2, and Windows. A Quick Start Installation and Reference Guide is included and full User Guide documentation is available through NovaStor. Backup software needed to support the 3.5-inch option kit has to be obtained separately.

Figure 16 on page 79 shows the layout of the 4/10GB 4mm DAT Tape Drive.

³ (Part# 74G68631). This 5.25" drive (#74G3961) was used as the example for almost all the tape backup solutions at Chapter 3, "Disk Subsystem - Setup and Configuration" on page 83

This soft copy for use by IBM employees only.




Figure 16. 4/10GB 4mm DAT Tape Drive

There is a wide variety of tape backup/restore software available for DOS, OS/2, SCO-UNIX, and Novell NetWare environments.

Key Features:

- SCSI-2 Fast interface
- 4/10GB capacity per cartridge, 8-16 (10GB typical) with data compression and with DDS2 data cartridges
- A data transfer rate of 400 to 1600 KBps with 1000 KBps typical in data compression mode.
- A 30-second average file search
- Uses the new DDS2 data grade media for higher reliability
- Features non data grade media rejection for write operations (will not allow writing on audio grade media)

This soft copy for use by IBM employees only.

- Features enhanced self-diagnostics and cleaning-required indicator
- Read and write compatibility with DDSIII tape cartridges in 2.0GB format
- Designed using Digital Audio Tape (DAT) technology
- Uses reliable Digital Data Storage (DDS) tape media

2.4.2.6 IBM 24/48GB Internal Tape Autoloader

This six cartridge autoloader (Part# 9462725) fits in the 5.25 full-high bay on the PC Server 5XX or 720. Unlike many autoloaders the cartridge is contained in the 5.25 full-high form factor allowing the system door to close while the autoloader is in use. This maintains security not only of the autoloader but of the diskette drive and CD ROM drive as well.

This is a DDS-2 DAT SCSI-2 Fast drive capable of backing up 48GB of data using compression on six cartridges. It has a typical sustained transfer rate of 1MBps.

This autoloader can be used with the PC Server 3XX non-array and the 5XX and 720 models.

The Array models of the PC Server 5XX and 720 require the IBM SCSI-2 Fast/Wide Adapter/A (#70G8498) for Micro Channel slots, and the IBM SCSI-2 Fast/Wide PCI adapter (#9463771) for PCI slots.

The following internal cables can be used:

- For Micro Channel adapters:
 - Part# 32G3921
 - Part# 32G3922
 - Part# 32G3098
- For PCI SCSI-2 F/W adapters:
 - Part# 70G9859
 - Part# 70G9264

The PC Server 500 requires BIOS level 4 or higher.

The Autoloader option kit contains:

- 24/48GB Autoloader
- SCSI wide to narrow converter
- Power Cable
- One 6 cartridge magazine

This soft copy for use by IBM employees only.

- Five data cartridges
- One cleaning cartridge
- User's manual

2.4.2.7 Available Software

A software tape backup/restore utility is required for the use of the above tape drives. The vendors listed below have certified that their indicated tape backup/restore software will support all the above products:

- IBM DualStor Backup Program

This product is offered with the following versions:

- IBM DualStor Backup Program for DOS and Windows Version 3.0
- IBM DualStor Backup Program for OS/2 Version 3.0

This product offers support for DOS, Windows and OS/2 environments.

- SYTOS-PLUS from Sytron Corporation supports DOS, OS/2, IBM LAN Server for OS/2, and client backup under Novell NetWare Software. Device drivers may not necessarily be included for all the above drives and can be obtained by contacting Sytron Technical Support.
 - SYTOS-PLUS Version 1.42 for DOS
 - SYTOS-PLUS Version 1.37 for OS/2
- ARCserve from Cheyenne Software, Inc. provides fully automated data management services for Novell Networks. It is available in a DOS edition and a Windows edition.
- ARCsolo from Cheyenne Software, Inc. provides high-performance backup and restore for DOS operating systems.
- CPbackup from Central Point Software, Inc.
- Novaback for DOS, OS/2, and Windows from NovaStor Corporation provides high-performance backup and restore.
- CorelSCSI four from Corel.
- Novell NetWare provides support in the NetWare environment.
- ARCHIVE Backup Exec from Archive Corporation.
- AIX PS/2 provides tape backup support in the AIX PS/2. environment
- SCO-UNIX provides support in the SCO-UNIX environment.

This soft copy for use by IBM employees only.

2.4.2.8 Tape Drives Summary

Table 22 summarizes the characteristics of the different tape solutions for the IBM PC Server.

<i>Table 22. IBM Tape Options - Summary</i>					
Attribute	3450-001	3445-001	3440-001	IBM 4/10GB 4mm DAT Tape	24/48GB Autoloader
Media	1/4 Inch Tape	8mm Tape	4mm Tape	4mm Tape	4mm Tape
Capacity	1.2GB	2.3GB	2.0GB	4.0GB	24.0GB
With Compression	N/A	5.0/10.0GB	4.0GB	8-16GB (10GB typical)	up to 48GB
Data Transfer Rate	300 KBps	245 KBps 500/1000 KBps	183/366 KBps	400/1600 KBps	400/1600 KBps
Access Time	30.0 sec Avg	2.6m sec (max) write 2.7m sec (max) read	N/A	N/A	N/A
Reposition Time	3.0 sec Avg	1.115 sec(max)	1.5 sec write 2.2 sec read		
Load Time	25 seconds maximum	N/A	N/A		

This soft copy for use by IBM employees only.

Chapter 3. Disk Subsystem - Setup and Configuration

This chapter helps you to define the SCSI options you need to make the different disk subsystems configurations operational.

We will explain how to configure the IBM PC Servers with the different SCSI options, both for internal and/or external SCSI device connection solutions.

The chapter includes only U.S. models and part numbers. A conversion table for other geographic regions is provided at Appendix A, "Part Number Equivalences for Europe" on page 207.

3.1 Internal SCSI Cabling

We will now look at how devices are connected internally to the host controller inside the PC Server.

3.1.1 PC Server 300 Internal Cabling

There is a PCI SCSI-2 Fast adapter standard on the 8640-0PT providing an external and internal port capable of handling up to seven devices in total. This Adapter is optional on PC Server 300 8640-0P0.

3.1.1.1 PC Server 300 Internal SCSI Cables

On the 8640-0P0 when you order the PCI SCSI-2 Fast Adapter you need to order the PC Server Cable Kit Part # 70G9264. It provides you with an 8-bit 7-drop SCSI cable and a 2-drop IDE cable. On the 8640-0PT, an 8-bit 4-drop SCSI cable comes standard. If you want to have more than four internal SCSI devices on PC Server 8640-0PT, you need to order the Cable Kit (#70G9264).

This soft copy for use by IBM employees only.

<i>Table 23. Options Installation Considerations for PC Server 300</i>		
Part #	Description	How it attaches to PC Server 300
06H8558	540MB Deskstar SCSI-2 Fast	Directly to 8-bit SCSI cable 1
94G3187	1.08GB Deskstar SCSI-2 Fast	Directly to 8-bit SCSI cable 1
94G3054	2.25GB SCSI-2 Fast	Directly to 8-bit SCSI cable 1
70G8503	Enhanced Internal CD-ROM II Drive	Directly to 8-bit SCSI cable 1
74G8631	4/10GB 4mm DAT Tape Drive (5.25")	Directly to 8-bit SCSI cable 1 Don't forget to add the tape backup software
Note:		
1 - For PC Server 8640-0P0, you need to add the PCI SCSI-2 Fast Adapter (#32G3100) and PC Server Cable Kit (#70G9264)		

Example 1 describes what you need in order to have a PC Server 300 8640-0PT with a 2GB SCSI-2 Fast disk drive, CD-ROM support and a 4/10GB Internal Tape Drive.

Example 1

PC Server 8640-0PT plus 1.08GB SCSI-2 Fast, CD-ROM, 4/10GB Internal Tape

- 86400PT PC Server 300, 16MB, 256KB L2, 1GB SCSI-2 Fast, PCI SCSI-2 Fast Adapter
- 94G3187 1.08GB SCSI-2 Fast Deskstar Drive
- 70G8503 Enhanced Internal CD-ROM II Drive
- 74G8631 4/10GB 4mm DAT Internal Tape Drive

Note: The Server 8640-0PT comes standard with an 8-bit 4-drop cable and will use up four out of four drops. Any additional internal SCSI devices would require the ordering of #70G9264

PC Server 300 Hot Swap Bay Option:

You can add a PC Server 300 Hot Swap Bay Option to the PC Server 300 to obtain up to six 1" high hot swap bays instead of the standard five 1.6" high non hot swap bays. This option, #94G2763, includes a 2-drop cable connecting the hot swap bank and CD-ROM connection to a SCSI-2 F/W PCI card, #94G3771. Since PC Server 300 models do not come standard with this adapter, the adapter must be ordered with the Hot Swap Bay Option.

Figure 17 on page 85 shows the connections layout for the PC Server 300.

This soft copy for use by IBM employees only.




Figure 17. PC Server 300 Internal SCSI Cables

This soft copy for use by IBM employees only.

3.1.2 PC Server 310 MCA or ISA Models Internal Cabling

One of the PCI card slot is occupied by a SCSI-2 Fast Controller connecting via a four drop (50-pin) SCSI cable to the standard 1.08GB hard disk drive and a Quad Speed CD-ROM, leaving 2 SCSI internal connectors for future expansions.

3.1.3 PC Server 320 Internal Cabling

This section describes the internal cabling for the PC Server 320 family, both non hot swap and hot swap connections are explained.

3.1.3.1 Non Hot Swap Models

This section describes the PC Server 320 non hot swap, non-array models.

PC Server 320 PCI/EISA (8640-0Y0, 0YT):

There is an SCSI-2 Fast/Wide PCI Adapter standard on the 8640-0Y0 and -0YT, providing one SCSI-2 F/W channel via one external port and 2 internal ports (1 Fast, 1 Fast/Wide). The adapter can handle up to 15 devices on the one channel via two ports maximum. On the 8640-0Y0 there is a PCI 8-bit 4-drop SCSI-2 Fast cable. If you require greater than 4 SCSI-1/SCSI-2 Fast devices, internally order the cable 16bit 7-drop cable #70G9859 and one 16/8-bit converter #32G3925 per SCSI-1/SCSI-2 Fast (8-bit) device along the cable. On the 8640-0YT there is a PCI 16bit 7-drop SCSI-2 Fast/Wide cable.

PC Server 320 PCI/MCA (8640-MX0, MXT):

An integrated PCI SCSI-2 Fast/Wide controller on the mother board provides one SCSI-2 F/W channel via one external port and 2 internal ports(1 Fast, 1 Fast/Wide). The controller can handle up to 15 devices on the one channel via two ports maximum. On the 8640-MX0 and -MXT there is a PCI 8-bit 4-drop SCSI-2 Fast cable. If you require greater than 4 SCSI-1/SCSI-2 Fast devices internally, order the cable 16bit 7-drop cable #70G9859 and one 16/8-bit converter #32G3925 per SCSI-1/SCSI-2 Fast (8-bit) device along the cable.

You can add a PC Server 300 Hot Swap Bay Option to the PC Server 320 to obtain up to six 1"high hot swap bays instead of the standard five 1.6"high non hot swap bays. This option, #94G2763, includes a 2-drop cable connecting the hot swap bank and CD-ROM connection to the standard SCSI-2 F/W PCI card. The PC Server 300 Hot Swap Bay Option is standard on PC Server 320 models 8641-MY0, MYT, 1Y0, 1YT.

This soft copy for use by IBM employees only.

Table 24 on page 87 shows a summary of the options that can be used for the PC Server 320 non hot swap, and how the options are attached.

<i>Table 24. Options Installation Considerations for PC Server 320 Non Array</i>		
Part #	Description	How it attaches to PC Server 320 (Non Hot Swap)
94G3187	1.08GB SCSI-2 Fast	0YT-Use 16/8-bit converter, Others-Directly
94G3052	1.12GB SCSI-2 F/W	0YT-Directly, Others-Cable #70G9859
94G3054	2.25GB SCSI-2 Fast	0YT-Use 16/8-bit converter, Others-Directly
94G3055	2.25GB SCSI-2 F/W	0YT-Directly, Others-Cable #70G9859
94G3196	4.51GB SCSI-2 Fast	0YT-Use 16/8-bit converter, Others-Directly
94G3057	4.51GB SCSI-2 F/W	0YT-Directly, Others-Cable #70G9859
94G3787	5.31GB SCSI-2 Fast	0YT-Use 16/8-bit converter, Others-Directly
94G3059	5.31GB SCSI-2 F/W	0YT-Directly, Others-Cable #70G9859
94G2686	SCSI CD-ROM drive (std.)	Additional: 0YT-Use 16/8-bit converter, Others-Directly
74G8631	4/10GB 4mm DAT Tape Drive (5.25")	0YT-Use 16/8-bit converter, Others-Directly. Don't forget to add the tape backup software

The following three examples show typical configurations for PC Servers 320.

Example 2

PC Server 320-90MHz PCI/EISA with 2x2.25GB SCSI-2 Fast,CD-ROM and 4/10GB Tape Drive

86400Y0 PC Server 320-90, 16MB, 256KB, Diskless, CD-ROM,
 PCI SCSI-2 F/W Adapter, 4-drop/8-bit
 94G3054 2.25GB SCSI-2 Fast (quantity 2)
 74G8631 4/10GB 4mm DAT Internal Tape Drive, 5.25"

Example 3

PC Server 320-90MHz PCI/EISA with 2x1.12GB SCSI-2 F/W with Disk Duplexing, CD-ROM, 4/10GB Tape

86400YT PC Server 320-90, 16MB, 256KB, 1x1.12GB SCSI-2 F/W,
 CDROM, PCI SCSI-2 F/W controller, 7-drop/16-bit.
 94G3771 SCSI-2 Fast/Wide PCI Adapter (additional card)
 70G9859 PCI SCSI-2 Fast/Wide Card 7-drop Internal Cable
 94G3052 1.12GB SCSI-2 F/W Disk
 74G8631 4/10GB 4mm DAT Internal Tape Drive, 5.25"

This soft copy for use by IBM employees only.

Example 4

PC Server 320-75Mhz PCI/MCA with 2x2.25GB SCSI-2 F/W Hot Swap, CD-ROM, 4/10GB Tape

8640MX0 PC Server 320-75, 16MB, 256KB, Diskless, CD-ROM,
PCI SCSI-2 F/W controller, 4-drop/8-bit
94G2763 PC Server 300 Hot Swap Bay Option
94G2650 2.25GB SCSI-2 F/W Hot Swap (quantity 2)
35100V0 External Single SCSI enclosure (*)
70G9858 PCI SCSI-2 F/W Card to SCSI-1/SCSI-2 Fast Option Cable (*)
32G3919 8-bit Active Terminator (External) (*)
74G8631 4/10GB 4mm DAT Internal Tape Drive 5.25"

Note (*): When Hot Swap Bay Option, #94G2763, is added to a PC Server 320 or comes standard, there's no suitable bay internally to hold a 4/10GB DAT Tape Drive, so it must be added externally inside a 3510 SCSI enclosure.

This soft copy for use by IBM employees only.

Figure 18 and Figure 19 on page 90 show the internal organization of the different components of the PC Server 320 with PCI/EISA bus (Figure 18) and PC Server 320 with PCI/MCA bus (Figure 19 on page 90). The bays, I/O slots, and CPU disposition are shown.




Figure 18. PC Server 320 PCI/EISA Non Hot Swap Internal SCSI

This soft copy for use by IBM employees only.




Figure 19. PC Server 320 PCI/MCA Non Hot Swap Internal SCSI

3.1.3.2 Hot Swap Models

The PC Server 320 Array models (8640-1Y0,1YT, MYR) have a PCI SCSI-2 Fast/Wide RAID Adapter with two channels (one usable internally, one externally) each capable of handling up to seven SCSI devices. There is a PCI 16-bit 2-drop SCSI cable to connect a six bay (1") hot swap enclosure and a double-speed CD-ROM drive. The 8640-MYR also comes with an integrated SCSI-2 F/W controller (Non RAID). PC Server 320 Non-Array Hot Swap models (8640-MY0, MYT) come standard with an integrated SCSI-2 F/W (non-RAID) controller with a 16-bit 2-drop SCSI cable to attach the six bay hot swap enclosure and CD-ROM.

This soft copy for use by IBM employees only.

Table 25 on page 91 shows a summary of the options that can be used for the PC Server 320 hot swap array, and how the options are attached.

<i>Table 25. Options Installation Considerations for PC Server 320 Array</i>		
Part #	Description	How it attaches to PC Server 320 Array
94G2649	1.12GB SCSI-2 F/W Hot Swap	Drive slides into a hot swap bay.
94G2650	2.25GB SCSI-2 F/W Hot Swap	Drive slides into a hot swap bay.
94G2651	4.51GB SCSI-2 F/W Hot Swap	Drive slides into two hot swap bays.
70G9851	SCSI-2 Fast Hot Swap Drive Tray	Wraps NON hot swap drive, slides in enclosure.
70G9741	SCSI-2 F/W Hot Swap Drive Tray	Wraps NON hot swap drive, slides in enclosure.
74G8631	4/10GB 4mm DAT Tape Drive (5.25")	Use external 3510 enclosure with cable. Don't forget to add tape backup software.

Figure 20 on page 92 shows the position of the bays and planar building blocks.

This soft copy for use by IBM employees only.




Figure 20. IBM PC Server 320 Array Mechanical Design

Figure 21 on page 93 shows the cabling scheme for the PC Server 320 with hot swap capabilities (array or non-array models).

This soft copy for use by IBM employees only.




Figure 21. PC Server 320 Array or 320 with Hot Swap (SCSI-2 F/W PCI Adapter or SCSI-2 F/W PCI RAID Adapter)

Example 5 explains what options you need to consider to have PC Server 320 Array with three 1GB hard drives with CD-ROM and a 4/10GB tape drive.

This soft copy for use by IBM employees only.

Example 5

PC Server 320 PCI/EISA Array (8640-1YT) with 3x1GB SCSI-2 F/W, CDROM, 4/10GB Tape Drive

86401YT	PC Server 320, 16MB, 256KB, 2x1.12GB F/W, CD-ROM, SCSI-2 F/W RAID PCI Adapter
94G2649	1.12GB SCSI-2 F/W Hot Swap (qty.1)
35100V0	External Single SCSI enclosure (*)
70G9858	PCI SCSI-2 F/W Adapter to SCSI-1/SCSI-2 Fast Option Cable (*)
32G3919	8-bit Active Terminator (External) (*)
74G8631	4/10GB 4mm DAT Internal Tape Drive 5.25"

Note (*): These options are included because there is no suitable internal 1/2 high non hot swap bay for the tape drive.

3.1.4 PC Server 500, 520 and 720 Internal Cabling

This section describes the internal cabling for the PC Server 500, 520 and 720 families. All three server families share the same basic frame and internal bay structure.

- The array models of the PC Server 500 and 720 both use the Micro Channel based Streaming RAID Adapter/A.
- The PC Server 520 PCI/EISA Non-Array models come with a PCI SCSI-2 Fast/Wide adapter, while the the Array models use a PCI SCSI-2 Fast/Wide RAID adapter.
- The PC Server 520 PCI/Micro Channel Non-Array models come with a standard onboard PCI SCSI-2 Fast/Wide controller. The PCI/Micro Channel Array models include this same controller, plus a PCI SCSI-2 Fast/Wide RAID Adapter.

Table 26 on page 95 shows a summary of the options that can be used for the PC Server 500, 520 and 720 and how the options are connected to each server.

This soft copy for use by IBM employees only.

<i>Table 26. Options Installation Considerations for PC Server 500, 520 and 720</i>		
Part #	Description	How it attaches to the PC Server
94G2649	1.12GB SCSI-2 F/W Hot Swap	Tray with drive slides into Bank
94G2650	2.25GB SCSI-2 F/W Hot Swap	Tray with drive slides into Bank
94G2651	4.51GB SCSI-2 F/W Hot Swap	Tray with drive slides into Bank
70G9851	SCSI-2 F Hot Swap Drive Tray	Slides into Bank (holds non hot swap drive)
70G9741	SCSI-2 F/W Hot Swap Drive Tray	Slides into Bank(holds non hot swap drive)
70G9740	PC Server Hot Swap Backplane	Provides hot swap support Bank D or E
70G9739	PC Server Power Supply Upgrade	Provides 220W extra for Bank D & E (required)
74G8631	4/10GB 4mm DAT Tape Drive (5.25")	Directly into 16/8-bit converter (#32G3925) and then into Bay B1 or Bay B2. Don't forget to add tape backup software

Figure 22 on page 96 shows the bay structure for the PC Server 500 and 720. The different bay banks are shown. It's very important to get familiar with the bank names and the positioning inside the case.

This soft copy for use by IBM employees only.




Figure 22. PC Server 500, 520, and 720 Internal Bay Structure

3.1.4.1 Array Models

Figure 23 on page 97 shows the internal structure for the PC Server 720. The front view helps to find the external bay access, and the side view shows the position of the main components on the planar board of the server.

This soft copy for use by IBM employees only.




Figure 23. PC Server 720 and 720 Array Mechanical Design

Figure 24 on page 98 shows the internal structure for the PC Server 500. The front view helps to find the external bay access, and the side view the position of the main components on the planar board of the server.

This soft copy for use by IBM employees only.




Figure 24. IBM PC Server 500 - Mechanical Design

PC Server 500/720 Array (8641-1Y0,-0YR,-0YS, 8642-1Z0,-2ZS,-4ZS) Cables:

There is a SCSI-2 Streaming Fast/Wide RAID Adapter in a Micro Channel slot with two SCSI-2 Fast/Wide 16-bit 20MBps channels that can be used internally, externally or both. The controller handles 40MBps data streaming along the MCA bus on the PC Server 500 Arrays and 80MBps on the PC Server 720 Arrays. This RAID card physically supports up to 14 SCSI devices, 7 per channel. Within the PC Server 500, you could have up to 12 hot swap SCSI-2 F/W drives in Bank C and D plus the standard CD-ROM drive or 11 hot swap drives (5 on Bank C and 6 in Bank D) plus the standard CD-ROM drive AND an optional 4/10GB DAT Tape Drive, all connected to the one RAID card.

This soft copy for use by IBM employees only.




Figure 25. PC Server 500/720 Array Models — SCSI-2 F/W RAID Adapter/A Connections

Figure 25 shows the internal SCSI cabling for the PC Server 500 or 720 with one RAID controller and two backplanes.

If an additional RAID adapter is added you could add six more hot swap drives (in Bank E), as shown is Figure 26 on page 100.




Figure 26. PC Server 500/720 Array Models — Two SCSI-2 F/W RAID Adapters/A Connections to Three Backplanes.

RAID function allows you to create up to four physical arrays with one physical array defined as seven drives maximum across the two channels. You can also define up to eight logical arrays within the four physical arrays on the RAID adapter.

The RAID adapter comes pre-cabled with a 16-bit SCSI-2 Fast/Wide cable connecting one of its channels to the CD-ROM drive in Bay A2. One of the half-high bays (Bay B1 or B2) is for a removable media (tape drive, etc) or non hot swap drives, in addition to connecting to Bank C (Bays 1-6) for up to six 1/3 high hot swap bays or three 1/2 high hot swap bays, or two full-size hot plug bays. The SCSI cable, therefore, has two 16-bit non hot swap drops (connections) and 1 drop to the first backplane (which has six hot swap

This soft copy for use by IBM employees only.

connections). Keep in mind that you can have a maximum of 7 SCSI devices per channel, so if you are using the CD-ROM and six 1/3 high hot swap drives in Bank C, a tape drive installed in Bay B1 or B2 could not be connected to the same SCSI cable. Although there are nine connections on the cable and backplane, a maximum of seven devices can be attached.

You can mix SCSI-2 Fast hot swap drives with SCSI-2 Fast/Wide hot swap drives within the same physical array as long as they're the same size (that is, 1GB).

To attach a SCSI-2 Fast or SCSI-1 non hot swap device in Bay B1 or B2 to the standard 16-bit cable, you need to add a 16/8-bit converter for each device. If more hot swap drives are required, then an additional backplane and power supply should be ordered. This will give an additional six 1/3 high hot plug bays (or three 1/2 high, or two full high) that can be installed in Bank D. The 16-bit SCSI cable that comes with the extra backplane will plug into the second internal channel on the SCSI-2 Fast/Wide Streaming RAID Adapter. The cable only has one drop which connects to the extra backplane containing six hot swap connections. More hot swap bays beyond this require ordering a second additional backplane (a second extra power supply is not required for this third backplane for Bank E). The 16-bit SCSI cable that comes with this final backplane demands a SCSI-2 Fast/Wide Streaming RAID Adapter/A be ordered to provide a third SCSI-2 Fast/Wide channel to connect to.

PC Server 520 Array (8641-EZ1,EZS,EZE,MZ1,MZS,MZE) Cables:

All the array models of the 520 family use a SCSI-2 Fast/Wide PCI-Bus RAID Adapter. This adapter has the same functionality as the Micro Channel adapter. All the considerations made for the 500 and 720 can be used for these systems as well.

Figure 27 on page 102 shows the PC Server 520 EISA model planar position and its components.

This soft copy for use by IBM employees only.




Figure 27. IBM PC Server 520 EISA/PCI - Mechanical Design

3.1.4.2 Non-Array Models

In this section we will describe the two main options for the SCSI controller, the PCI version used at the PC Server 720 and 520 and the Micro Channel used for the PC Server 500.

PC Server 500 (8641-0Y0,-0YT,-0YV) SCSI Channels and Cables:

These models come standard with a SCSI-2 Fast/Wide Adapter/A that has two SCSI-2 Fast/Wide 20MBps channels (one internal, one external). The internal channel comes standard with an 8-drop (two non hot swap and six hot plug) 16-bit cable. The same cable is used for the array version of the PC

This soft copy for use by IBM employees only.

Server 500. To attach a SCSI-2 Fast or SCSI-1 non hot swap drive in bay B1, or B2, you need to add a 16/8-bit converter for each device or drive (same as on the PC Server 500 Array). If you need additional hot swap drive bays beyond the standard six 1/3 high hot plug bays found in Bank C, the same requirements for ordering of additional backplanes and power supply apply for the non-array models as on the array models. Keep in mind that if you continue to use the SCSI-2 Fast/ Wide Adapter/A, you'll need to order an additional adapter per extra bank. Each additional bank, for example Bank D or E, requires a separate SCSI-2 Fast/Wide channel for the 16-bit hot swap cable that comes with each additional backplane ordered. The SCSI-2 Fast/Wide Adapter/A only comes with one internal channel.

This non-array adapter (used in the PC Server 500 and 720 only) supports a total of fifteen SCSI devices. The internal channel supports up to seven SCSI devices, and the external channel supports from one to fourteen SCSI devices. However, the total number of SCSI devices attached to a single SCSI adapter cannot exceed fifteen. Please keep in mind that wide devices (16-bit) support SCSI IDs 0 to 15, but narrow (8-bit) devices support only IDs from 0 to 7.

Figure 28 on page 104 shows a typical configuration of the internal SCSI cabling and the SCSI Fast/Wide Controller/A for the PC Server 500 non-array models.

This soft copy for use by IBM employees only.




Figure 28. PC Server 500 Non-Array with SCSI-2 F/W Adapter Connections

Example 6 and 7 show typical configurations for the PC Server 500 Array.

This soft copy for use by IBM employees only.

Example 6

PC Server 500 with 5x2.25GB SCSI-2 F/W hot swap, RAID support, CD-ROM, Tape Drive

86410YS	PC Server 500, 32MB, 90/60Mhz, 3x2.25GB SCSI-2 F/W, CD-ROM, F/W RAID card
94G2650	2.25GB SCSI-2 F/W hot swap (quantity 2)
74G8631	4/10GB 4mm DAT Int tape drive, 5.25"
32G3925	16/8 bit converter (for tape drive connection to 16-bit cable)

Note:

Please note that we connected as many SCSI devices as allowed on one SCSI-2 Fast/Wide channel of the RAID Adapter (7 SCSI devices, five 2.25GB drives plus a CD-ROM drive plus a 4/10GB DAT tape drive). If we want a sixth hot swap drive in the first Bank, we must ensure the tape drive is connected to the second channel (use cable# 32G3922 and 16/8 bit converter #32G3925) or put on a separate additional SCSI adapter (for example, MC Adapter #70G8498) so that we don't exceed 7 SCSI devices per RAID channel. If the above steps are not taken, you would need to move one of the drives down into the second backplane (#94G9740) and also order the extra power supply (#70G9739) required when using the second or third backplane.

This soft copy for use by IBM employees only.

Example 7

PC Server 500 with 8 x 2.25GB SCSI-2 Fast/Wide Hot Swap, 2 x 4.51GB SCSI-2 F/W Hot Swap, RAID Support, CD-ROM, 24/48GB DAT Tape Drive

86410YS	PC Server 500, 32MB, 90/60Mhz, 3x2.25GB SCSI-2 F/W, CD-ROM, F/W RAID MCA
94G2650	2.25GB SCSI-2 F/W hot swap 1" (quantity 5)
94G2651	4.51GB SCSI-2 F/W hot swap 1.6" (quantity 2)
94G2725	24/48GB Internal SCSI DAT Tape Auto-loader 5.25" Full high (Bay B1&B2)
32G3925	16/8 bit converter (for tape drive connection to 16-bit cable)
70G9740	PC Server Backplane (quantity 1)
70G9739	PC Server Second Power Supply (220 watts)

Note:

The SCSI configuration for this example implies that you must use a second backplane with an additional power supply. The first SCSI channel will have connected the six 2.25GB drives only and the second SCSI channel will have the CD-ROM, the tape, two 2.25GB drives and two 4.51GB drives (six devices). In order to make this configuration, you need to change the cable that connects the CD-ROM to the first backplane and connect it to the second backplane, otherwise you won't have enough SCSI IDs at the first SCSI channel.

PC Server 520 (8641-MZ0, MZV) SCSI channels and cables:

There is a SCSI-2 Fast/Wide PCI controller onboard these machines, For further information about this SCSI support, refer to the "PC Server 320 PCI/MCA (8640-MX0, MXT)" on page 86. The internal cabling is the same as the one described in the following section.

Figure 29 on page 107 shows the PC Server 520 MCA model planar position and its components.

This soft copy for use by IBM employees only.




Figure 29. IBM PC Server 520 MCA/PCI - Mechanical Design

PC Server 720 (8642-0Z0) and 520 (8641-EZ0, EZV) SCSI channels and cables:

This model comes standard with a SCSI-2 Fast/Wide PCI Adapter (Non Array) and the same hot swap bays as on the 720 Array models. Since this adapter can support up to fifteen SCSI devices off of its one channel, it can connect the standard CD-ROM drive, an optional internal tape drive and six hot swap drives in Bank C, in addition to seven external SCSI devices.

This soft copy for use by IBM employees only.




Figure 30. PC Server 720 Non-Array with SCSI-2 F/W PCI Adapter Connections

Figure 28 on page 104 shows the configuration used with the SCSI-2 Fast/Wide Controller/A for the PC Server 500 non-array models. Figure 30 shows the same configuration used with the SCSI-2 Fast/Wide PCI Controller for the PC Server 720.

3.1.5 Internal Bays/Cabling Upgrades

Figure 31 on page 109 shows what you need in order to upgrade your internal SCSI cabling at your PC Server. First of all, identify the disk subsystem controller the system has installed and then the device type you want to connect to it. Follow the lines that connect them (start at the controller and follow the lines to the devices) and you can find the identification (part number) of the upgrade you need in order to make additional internal connections possible.

This soft copy for use by IBM employees only.




Figure 31. Internal Cabling Upgrade Options Summary

3.1.6 Upgrade Trays and Hot Swap Bay Options

In addition to 3.1.5, "Internal Bays/Cabling Upgrades" on page 108, the options in Table 27 on page 110 covers the following additional situations:

- Upgrades from Fast to Fast/Wide cabling for SCSI Enclosures
- Drive upgrades for hot swap capabilities
- System upgrades for hot swap capabilities
- Kits for non hot swap hard drive installation

This soft copy for use by IBM employees only.

<i>Table 27. Trays for Disk Drives and Bay Options for Hot Swap Upgrade</i>	
Part Number	Description
32G3920	SCSI-2 F/W (16-bit) Upgrade Kit for 3510-0V0.
06H3153	SCSI-2 F/W (16-bit) Upgrade Kit for 3511-0V0.
70G8506	PS/2 Hard Drive Install Kit includes rails for up to five drives within a server. Order one kit per five non hot swap hard drives within a server. Used for 3511-0V0.
06H7939	PS/2 Floor standing Hard Disk Drive Kit includes rails for up to one drive within a server or 3511. Order one kit per non hot swap hard drive.Used for 3511-0V0.
94G2648	3510 SCSI Drive Install Kit Includes rails for one drive inside a 3510 SCSI enclosure.
70G9851	SCSI-2 Fast hot swap drive tray - PC Server. Upgrades a non hot swap SCSI-2 Fast drive with hot plug tray support for use in a PC Server 300/320, 500, and 720 (for the 300/320 non-array array models #94G2763 is needed.)
70G9741	SCSI-2 Fast/Wide hot swap drive tray - PC Server. Upgrades a non hot swap SCSI-2 Fast/Wide drive with hot plug tray support for use in a PC Server 320 Array, 500, and 720.
94G3346	SCSI-2 Fast/Wide hot swap drive tray - 95 Array. Upgrades a non hot swap SCSI-2 Fast/Wide drive with hot plug tray support for use in a Server 95 Array, 3516-7V0 Tower.
94G2763	PC Server 300/320 hot swap Bay Option. Upgrades a non hot swap PC Server 300 or 320 (8640-0PO,-OPT,-0Y0,-0YT) <i>only</i> providing up to six 1/3 hot swap bays (1") or three half-high bays (1.5") using the space of five original 1/2 high non hot plug bays.

This soft copy for use by IBM employees only.

3.2 SCSI External Cabling

Sometimes when we are configuring a PC Server solution, we need to include external SCSI devices.

When should you consider an external SCSI device configuration?

- When there is no suitable bay in the PC Server case for the new SCSI device.
- When there is not enough bays in the PC Server case for all the SCSI devices.
- When you need more flexibility with the SCSI devices. For example, an expensive external tape drive can be easily moved from one server to another.

3.2.1 SCSI Adapter Connectors

Different SCSI connectors are used on the external port of the SCSI controllers, each of the different connectors need a special Card to Option Cable to configure the external chain of SCSI devices. The SCSI connectors used at the IBM SCSI Adapters are:

- Centronics 50-pin connector
- IBM 50-pin I/O card interface
- IBM 60-pin I/O card interface
- 50-pin high density interface
- 68-pin high density interface

Figure 32 shows the Centronics 50-pin connector.




Figure 32. Centronics 50-pin Connector

This soft copy for use by IBM employees only.

Figure 33 on page 112 describes the high density connectors used at the PCI controllers.




Figure 33. High Density PCI Controller SCSI Connectors

Figure 34 on page 113 describes the IBM SCSI connectors used on the Micro Channel controllers.

This soft copy for use by IBM employees only.




Figure 34. IBM SCSI I/O External Interface Connectors on Micro Channel Adapters

3.2.2 Card to Option Cables

Figure 35 on page 114 shows the different options available to connect the SCSI subsystem controller to the first external device. First of all you need to identify the adapter that your system has installed and then the type of SCSI optional device you want to connect as the first one. The line between them represents the cable needed for the connection. The circle in the line shows the part number of the cable and the cable length.

This soft copy for use by IBM employees only.




Figure 35. External Card to Option Cables Summary

The information in Figure 35 can also be used for the onboard SCSI support found at the PC Server 320/520 Micro Channel. You must assume that this onboard adapter has the same SCSI external connector as the SCSI-2 F/W PCI Adapter (#94G3771).

The *Fast Single Port Device (50 pin)* referred in Figure 35 are the popular Tape Backup Drives IBM 3431, 3440, 3445 and 3450.

3.2.3 Option to Option Cables and Terminators

Figure 36 on page 115 shows the options that allows the connection of more than one device to an external SCSI chain.

This soft copy for use by IBM employees only.




Figure 36. Option to Option Cables and Terminator Summary

In Figure 36, when connecting a Fast Single Port Device (50 pin) in a daisy-chain, #32G3099 Card to Option Cable must be included in the connection to the host adapter.

Whenever you need to connect 16-bit and 8-bit external devices, first connect all the 16-bit devices closer to the host controller, and then the 8-bit devices.

This soft copy for use by IBM employees only.

<i>Table 28. IBM External Micro Channel SCSI Cabling Summary - (IBM 50 or 68 pin external connector)</i>		
Description	PC Server 500/500A: 8641-0Y0,1Y0,0YT, 0YV,0YR,0YS	PC Server 720/720A: 8642-0Z0,1Z0,2ZS, 4ZS
<i>SCSI-2 Fast/Wide RAID Adapter, #70G9263 (two channels, External or Internal)</i>	Standard Adapter on 8641-1Y0,0YR,0YS, Optional Adapter on 8641-0Y0,0YT,0YV. Use cable (#06H3231) between card & tower	Standard Adapter on 8642-1Z0,2ZS,4ZS, Use cable (#06H3231) between card and 3516/3511 Tower
<i>SCSI-2 Fast/Wide Adapter/A, #70G8498, has a 16-bit SCSI-2 Fast/Wide external port</i>	Standard Port on 8641-0Y0,0YT,0YV, Optional Adapter/Port on -1Y0,0YR,0YS	Optional External Port via Optional adapter (#70G8498)
<i>Using "SCSI-2 F/W Adapter/A" (#70G8498) and first external SCSI device is 16-bit (SCSI-2 F/W)</i>	Use Cable #06H3231 (ideal) or #32G3913	Use Cable #06H3231 (ideal) or #32G3913
<i>Using "SCSI-2 F/W Adapter/A" (#70G8498) & 1st external SCSI device is 8-bit (SCSI-2 F, SCSI-1)</i>	Use Cable #32G3915 and (Optional Adapter) #70G8498	Use Cable #32G3915 and (Optional Adapter) #70G8498
<i>Using "SCSI Adapter MCA 32bit" (#70G8497) & 1st external SCSI device is 8-bit (SCSI-2 F, SCSI-1)</i>	Use Cable #32G4143 and Optional Adapter #70G8497	Use Cable #32G4143 and Optional Adapter #70G8497
<i>Using "SCSI-2 F/W Adapter/A" (#70G8498) and Add'l SCSI device is 16-bit (SCSI-2 F/W) with a 16-bit device connected before it</i>	Use Cable #32G3914 and Optional Adapter #70G8498	Use Cable #32G3914 and Optional Adapter #70G8498
<i>Using "SCSI-2 F/W Adapter/A" (#70G8498) and Add'l SCSI device is 8-bit (SCSI-2 F, SCSI-1) with a 16-bit device connected before it</i>	Use Cable #32G3916 and Optional Adapter #70G8498	Use Cable #32G3916 and Optional Adapter #70G8498
<i>If Additional SCSI device is 8-bit (SCSI-2 Fast or SCSI-1) with an 8-bit device connected before it</i>	Use Cable #32G3917 and Optional Adapter #70G8498	Use Cable #32G3917 and Optional Adapter #70G8498
<i>If you want a single SCSI-2 Fast/Wide device enclosure (external)</i>	Use 3510-0V0 SCSI Enclosure with SCSI-2 Fast/Wide Upgrade Kit(#32G3920) and appropriate external cable	Use 3510-0V0 SCSI Enclosure with SCSI-2 Fast/Wide Upgrade Kit(#32G3920) and appropriate external cable
<i>If you want a single SCSI-2 Fast or SCSI-1 device enclosure (external)</i>	Use 3510-0V0 SCSI Enclosure with appropriate cable	Use 3510-0V0 SCSI Enclosure with appropriate cable
<i>For last SCSI device on the chain (Cable)</i>	Add Terminator (#32G3918) if it's a 16-bit device, #32G3919 for 8-bit	Add Terminator (#32G3918) if it's a 16-bit device, #32G3919 for 8-bit

This soft copy for use by IBM employees only.

Table 29 (Page 1 of 2). IBM External PCI SCSI Cabling Summary - (50 or 68 pin high density external connector)						
Description	PC Server 300: 8640-0P0,0PT and PC Server 310 8639-0X0,0XT	PC Server 320: 8640-MX0,MXT, 0Y0,0YT and PC Server 520: 8641-EZ0,EZV, MZ0,MZV	PC Server 320/300 with hot swap Bay Option or 8640-MY0,MYT	PC Server 320 Array: 8640-1Y0,1YT, MYR PC Server 520 Array: 8641-EZ1,EZS, EZE,MZS,MZE	PC Server 720: 8642-0Z0	PC Server 720 Array: 8642-1Z0,2ZS, 4ZS
SCSI-2 Fast PCI Adapter, #32G3100 (2 ports: 1 External, 1 Internal) One channel, max 7 devices	0P0,0X0,0XT: Optional adapter 0PT: Standard adapter (use cable #70G9874 between adapter & 1st 8-bit option)	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported
SCSI-2 F/W PCI Adapter (#94G3771) or Integrated PCI SCSI-2 F/W Controller: 1 external port	Optional adapter	Standard PCI F/W external port	Standard PCI F/W external port	Optional adapter (#94G3771) required on 8640-1Y0,1YT	Standard PCI F/W external port	Optional adapter (#94G3771) required
SCSI-2 F/W RAID PCI Adapter, #94G2764 (2 channels: 1 internal, 1 external) Up to 7 devices per channel	Optional External Port via Optional adapter (#94G2764)	Optional External Port via Optional adapter (#94G2764)	Optional External Port via Optional adapter (#94G2764)	Standard external port via adapter	Not Supported	Not Supported
If first external SCSI device is 16-bit (SCSI-2 Fast/Wide)	Use Cable #70G9857 and optional PCI F/W adapter	Use Cable #70G9857 with standard PCI adapter	Use Cable #70G9857 with standard PCI adapter	Use Cable #70G9857 with standard PCI adapter	Use Cable #70G9857 with standard PCI adapter	Use Cable #70G9857 and optional PCI adapter #94G3771
If first external SCSI device is 8-bit (SCSI-2 Fast or SCSI-1)	Use Cable #70G9874 with standard PCI Fast adapter	Use Cable #70G9858 with standard PCI adapter	Use Cable #70G9858 with standard PCI adapter	Use Cable #70G9858 with standard PCI adapter	Use Cable #70G9858 with standard PCI adapter	Use Cable #70G9858 and optional PCI F/W adapter #94G3771
If Additional SCSI device is 16-bit (SCSI-2 Fast/Wide) with a 16-bit device connected before it	Use Cable #32G3914, #70G9857 & opt'l PCI F/W adapter	Use Cable #32G3914	Use Cable #32G3914	Use Cable #32G3914	Use Cable #32G3914	Use Cable #32G3914 with optional PCI Adapter #94G3771

This soft copy for use by IBM employees only.

Table 29 (Page 2 of 2). IBM External PCI SCSI Cabling Summary - (50 or 68 pin high density external connector)						
Description	PC Server 300: 8640-0P0,0PT and PC Server 310 8639-0X0,0XT	PC Server 320: 8640-0X0,MXT, 0Y0,0YT and PC Server 520: 8641-EZ0,EZV, MZ0,MZV	PC Server 320/300 with hot swap Bay Option or 8640-MY0,MYT	PC Server 320 Array: 8640-1Y0,1YT, MYR PC Server 520 Array: 8641-EZ1,EZS, EZE,MZS,MZE	PC Server 720: 8642-0Z0	PC Server 720 Array: 8642-1Z0,2ZS, 4ZS
<i>If Additional SCSI device is 8-bit (SCSI-2 Fast or SCSI-1) with a 16-bit device connected before it</i>	Use Cable #32G3916, #70G9857 & opt'l PCI F/W adapter	Use Cable #32G3916	Use Cable #32G3916	Use Cable #32G3916	Use Cable #32G3916	Use Cable #32G3914 with optional PCI Adapter #94G3771
<i>If Additional SCSI device is 8-bit (SCSI-2 Fast or SCSI-1) with an 8-bit device connected before it</i>	Use Cable #32G3917, #70G9874, and std PCI Fast adapter	Use Cable #32G3917	Use Cable #32G3917	Use Cable #32G3917	Use Cable #32G3917	Use Cable #32G3917 with optional PCI Adapter #94G3771
<i>If you want a single SCSI-2 Fast/Wide device enclosure (external)</i>	Use 3510-0V0 SCSI Enclosure with SCSI-2 Fast/Wide Upgrade Kit (#32G3920) and appropriate external cable	Use 3510-0V0 SCSI Enclosure with SCSI-2 Fast/Wide Upgrade Kit (#32G3920) and appropriate external cable	Use 3510-0V0 SCSI Enclosure with SCSI-2 Fast/Wide Upgrade Kit (#32G3920) and appropriate external cable	Use 3510-0V0 SCSI Enclosure with SCSI-2 Fast/Wide Upgrade Kit (#32G3920) and appropriate external cable	Use 3510-0V0 SCSI Enclosure with SCSI-2 Fast/Wide Upgrade Kit (#32G3920) and appropriate external cable	Use 3510-0V0 SCSI Enclosure with SCSI-2 Fast/Wide Upgrade Kit (#32G3920) and appropriate external cable
<i>If you want a single SCSI-2 Fast or SCSI-1 device enclosure (external)</i>	Use 3510-0V0 SCSI Enclosure with appropriate cable					
<i>For Last SCSI device on the chain (Cable)</i>	Add Terminator (#32G3918) if it's a 16-bit device, #32G3919 for 8-bit	Add Terminator (#32G3918) if it's a 16-bit device, #32G3919 for 8-bit	Add Terminator (#32G3918) if it's a 16-bit device, #32G3919 for 8-bit	Add Terminator (#32G3918) if it's a 16-bit device, #32G3919 for 8-bit	Add Terminator (#32G3918) if it's a 16-bit device, #32G3919 for 8-bit	Add Terminator (#32G3918) if it's a 16-bit device, #32G3919 for 8-bit

This soft copy for use by IBM employees only.

3.2.4 PC Server 300 and 310 External Cabling

The PCI SCSI-2 Fast adapter used on the PC Server 300 and 310 are physically different cards, but use the same interfaces and cables.

The external port on the SCSI-2 Fast PCI Adapter is a standard 50-pin 8-bit interface supporting SCSI-1/SCSI-2 Fast (8-bit) devices. Ensure that there is an 8-bit active terminator on the final SCSI device.

Example 8

Attaching one 2GB SCSI-2 Fast Drive & 4/10GB DAT Tape Drive Externally

70G9874	SCSI-2 Fast Card (50-pin) to SCSI (8-bit) Cable
35100V0	SCSI External enclosure for one SCSI device (8-bit std)
94G3055	2.25GB SCSI-2 Fast Drive (8 bit)
32G3917	SCSI (8-bit) Option to SCSI (8-bit) Option (Short) Cable, 0.3m or
6451042	SCSI(8-bit) Option to SCSI(8-bit) Option Cable, 0.6 meters
35100V0	SCSI External enclosure for one SCSI device (8-bit std)
74G8631	4/10GB 4mm DAT Int Tape Drive (8-bit)
32G3919	8-bit Active Terminator (attaches to 3510 containing Tape)

Note: You can have maximum of 7 SCSI-1/SCSI-2 Fast devices connected to the SCSI-2 Fast PCI Adapter internal and external combined

Figure 37 on page 120 shows the connection layout for Example 8.

This soft copy for use by IBM employees only.




Figure 37. PC Server 300 External Cabling Example (PCI SCSI-2 Fast)

3.2.5 PC Server 320, 320 Array, 520, 520 Array and 720 Non-Array External Cabling

The external port on the SCSI-2 Fast/Wide PCI Adapter (RAID or non-RAID) is a high density 68pin (16-bit) interface supporting SCSI-1, SCSI-2 Fast, or SCSI-2 Fast/Wide devices. Ensure that there is an 8-bit active terminator on the final SCSI device if it is an 8-bit device (for example, SCSI-1 or SCSI-2 Fast) and a 16-bit active terminator if the last device on the external SCSI cable is Fast/Wide.

The PC Server 320, 320 Array, 520, 520 Array and 720, all of them have the same external connector at the disk subsystem controller, therefore the following examples are valid for all of these systems.

This soft copy for use by IBM employees only.

Example 9

Attaching a SCSI-2 Fast 4/10GB DAT Tape Drive

- 70G9858 SCSI-2 Fast/Wide PCI Adapter to SCSI (8-bit) Cable, 1 meter
- 35100V0 SCSI External enclosure for one SCSI device (8-bit std)
- 74G8631 4/10GB 4mm DAT Internal Tape Drive (8-bit)
- 32G3919 8-bit Active Terminator (attaches to 3510 containing tape)

Example 9 and Figure 38 shows the SCSI-2 Fast 4/10GB DAT tape drive external connection when using a PCI host adapter.




Figure 38. PCI SCSI-2 Fast/Wide External SCSI (8-bit External Device)

This soft copy for use by IBM employees only.

Example 10

Attaching a SCSI-2 F/W External Device (2.25G F/W)

70G9857 SCSI-2 Fast/Wide PCI Adapter to SCSI F/W Cable, 1 meter
35100V0 SCSI External enclosure for one SCSI device (8-bit std)
32G3920 SCSI-2 F/W (16-bit) Upgrade Kit for 3510-0V0
94G3055 2.25GB F/W non hot swap drive (16-bit)
32G3918 16-bit Active Terminator (attaches to 3510 F/W)

Note: You can have maximum of 15 SCSI-1/SCSI-2 Fast devices connected to SCSI-2 Fast/Wide PCI Adapter internal and external combined and 7 devices per channel (one channel could be external) on the PCI SCSI-2 F/W RAID Adapter

Example 10 and Figure 39 show a SCSI-2 Fast/Wide hard drive external connection.




Figure 39. PCI SCSI-2 Fast/Wide External SCSI (16-bit External device)

Example 11 shows the connection of the 3516 hot swap external enclosure connection.

This soft copy for use by IBM employees only.

Example 11

7GB SCSI-2 F/W External Attachment for the PC Server 320 Array

70G9857	SCSI-2 F/W Card to SCSI-2 Fast/Wide Option Cable (1.0m)
35167V0	3516 hot swap Enclosure (holds up to 1/3 high bay hot swap drives)
70G9743	1GB SCSI-2 Fast/Wide hot swap Drive (quantity 7 in one 3516)

3.2.5.1 PC Server 500, 500 Array and 720 Array, External Cabling

The external SCSI support is provided at the PC Server 500 and 720 Array by the factory installed SCSI-2 Fast/Wide Adapter/A. One of the channels of the controller can be configured for external SCSI devices connections. The external channel comes *standard* with 16-bit SCSI-2 FAST/Wide interface.

If you have a mixture of external 8-bit (SCSI-1 or SCSI-2 Fast) and 16-bit SCSI-2 Fast/Wide devices, start by first connecting your SCSI-2 Fast/Wide devices to the external SCSI-2 Fast/Wide external port on the Server. Your SCSI-1 devices should be the farthest along the SCSI cable chain.

This soft copy for use by IBM employees only.

Example 12

Attaching External 2x2GB Fast/Wide drives and External 4/10GB DAT Tape

32G3913 SCSI2 F/W Card to SCSI2 F/W (16-bit) Option Cable, 1.5m
32G3920 SCSI2 F/W (16-bit) Upgrade Kit for 3510-0V0 (8-bit std)
35100V0 SCSI External enclosure for one SCSI device (8-bit std)
70G8494 2GB SCSI-2 Fast/Wide drive (16 bit)
32G3914 SCSI2 F/W (16-bit) Option to SCSI2 F/W (16-bit)
Option Cable, 0.3m
32G3920 SCSI2 F/W (16-bit) Upgrade Kit for 3510-0V0 (8-bit std)
35100V0 SCSI External enclosure for one SCSI device (8-bit std)
70G8494 2GB SCSI-2 Fast/Wide drive
32G3916 SCSI2 F/W (16-bit) Option to SCSI (8-bit) Option Cable, 0.3m
35100V0 SCSI External enclosure for one SCSI device (8-bit std)
74G8631 4/10GB 4mm DAT Int Tape Drive (8-bit)
32G3919 8-bit Active Terminator (attaches to 3510 containing Tape)

Note: We have now used 3 out of 7 drops possible along the external port and we have included an Active Terminator at the end of the SCSI chain since it's an external chain. Please note that if your last device on the external chain is a 16-bit (SCSI-2 Fast/Wide) device you need to use a 16-bit Active Terminator (#32G3918), not an 8-bit terminator. Keep in mind SCSI distance limitations. Assume a 0.3m distance for each 3510 (external SCSI device) used in addition to cable lengths.

Example 12 and Figure 40 on page 125 show the configuration of an external chain of F/W and Fast SCSI devices.

This soft copy for use by IBM employees only.




Figure 40. External 16 and 8-Bit SCSI Configuration

Example 13

Attaching an External 8-bit (SCSI-1,SCSI-2 Fast) device, (for example 4/10GB Tape)

- 32G3915 SCSI2 F/W Card to SCSI (8-bit) Option Cable
- 35100V0 SCSI External enclosure for one SCSI device (8-bit std)
- 74G8631 4/10GB 4mm DAT Int Tape Drive (8-bit) Inside the 35100V0
- 32G3919 8-bit Active Terminator (attaches to 3510 containing Tape)

Note: We have now used 1 out of 7 drops possible along the external port and we have included an Active Terminator at the end of the SCSI chain since it's an external chain.

Example 13 and Figure 41 on page 126 show the SCSI-2 Fast 4/10GB DAT tape drive external connection when using a MCA host adapter.

This soft copy for use by IBM employees only.




Figure 41. External 8-Bit SCSI Configuration

This soft copy for use by IBM employees only.

Chapter 4. Disk Subsystem Sizing and Performance

This chapter contains two parts. In the first part you will find considerations on how to define the most appropriate disk subsystem configuration that fits your business needs and requirements.

The second part contains disk subsystem performance considerations, which will assist you in optimizing the disk subsystem performance.

4.1 Disk Subsystem Sizing

This section describes criteria for defining and sizing your disk subsystem. Examples provide detailed information about the various dependencies and eventual trade-offs you may make while trying to find the best implementation.

4.1.1 Introduction

When deciding on a disk subsystem configuration that can best meet your needs, you need first to prioritize your key business requirements. Because there is no perfect solution, you may need to compromise on performance, capacity, fault tolerance and cost. These characteristics are often contradictory. For example, when cost is a major concern and criterion, you may be forced to abandon or reduce your requirements for performance, capacity or fault tolerance. If, on the other hand, performance and fault tolerance are your highest priority, then you may need to compromise on price.

4.1.2 Configuration Considerations

Common requirements for disk subsystems are:

- *Capacity*

As mentioned in Chapter 1, determining a server's disk subsystem storage capacity depends on the number of users and the amount of disk space they need, as well as on the amount of data and the number of programs to be stored. Other factors, such as cost, performance and fault tolerance, can make this simple task rather complex.

This soft copy for use by IBM employees only.

- *Cost*

Often a limited budget prevents you from purchasing a server with all the functions you need to meet your business requirements.

If cost is the major criterion, then you need to compromise on capacity or on some other disk subsystem requirement.

- *Performance*

If there is a need for almost instant data availability at the workstation, as for example in a manufacturing environment to control various and complex processes or in radio/TV stations, where digitized voice, music and image information need to be available without delays and interruptions, then performance becomes your major concern. This requires more powerful components and a more sophisticated implementation such as an intelligent RAID controller and multiple SCSI-2 Fast/Wide disk drives defined as a RAID-0 array. Again, you may be forced to compromise on capacity to reduce cost, or you may need to accept a higher price tag if you need both capacity and performance.

- *Fault Tolerance*

Increasingly, mission- and business-critical data and programs are stored on your server's disk subsystem. Multiple disk drives are combined in arrays to create large-capacity volumes or logical drives to satisfy the growing demand for disk storage capacity.

This disk storage is used to hold your most valuable asset: the data and programs that run and execute the tasks required by your business.

In addition, a failing server system will not only affect one workstation and its user, but also all other connected workstations. Besides the lost productivity and the efforts to recover, you may lose valuable data. Even after a server problem is fixed and the server is up and running again, the loss of data may seriously impact your business operation.

To avoid such events and their aftermath, protection against server failures and data loss caused by disk failures is essential.

Protection against loss of data and productivity is achieved by implementing a disk subsystem that remains operational even if there is a failure in one of its major components. Such disk subsystems are called fault tolerant, because they tolerate a component failure but remain operational.

RAID technology can help to protect you against the effects of a disk failure.

This soft copy for use by IBM employees only.

Implementing a fault tolerant disk subsystem may result in a considerably higher cost due to redundant components required and also impact the disk subsystem's performance and capacity depending on the RAID implementation used.

However, if you take into consideration that a server failure may affect hundreds of workstation users with a resulting loss in productivity, the additional investment in fault tolerance may well pay off.

4.1.2.1 Capacity

Hard disks of different capacities are available. You may decide to obtain the total disk capacity required by using either one (or more) large-capacity drives. Since one large drive is usually lower-priced than two smaller drives with the same capacity, you may find purchasing the single large disk to be the better solution.

If your major criterion is cost, and fault tolerance and high performance are less important for your business, selecting large-capacity drives (for example 4.5GB) may be the best implementation for your requirements.

However, other factors may become important. Then you need to carefully evaluate their impact to your disk subsystem.

Other factors that are related to capacity include:

Cost:

When trying to realize a disk subsystem at the lowest cost for a given capacity, you may choose the following solution:

Use large capacity drives and a standard hard disk controller.

- Advantages
 - A large-capacity drive is less expensive than multiple, smaller drives with the same total capacity.
 - Low-cost hard disk controllers can be used.
- Disadvantages
 - Installing more than one drive results in multiple drive letters (for example, C:, D:), which may become a problem depending on your application program's disk space requirements.
 - Software is required to create a single multidisk volume.

This soft copy for use by IBM employees only.

- Seek times of a single large-capacity drive are longer than for an array of the same capacity built from smaller hard disks.

Performance:

To achieve higher performance, consider the following suggestions:

1. Use RAID technology (RAID-0 or RAID-1) to distribute the load over multiple drives, to shorten the seek time per drive, and to allow multiple concurrent I/O operations.
2. Use intelligent RAID controllers to off-load RAID management and I/O request processing tasks from the system processor.
3. Use fast hard disks with short access times and a high data transfer rate.

Fault Tolerance:

To improve your disk subsystem's reliability, you can implement fault tolerance, which prevents data losses and server downtime due to failing components in the disk subsystem. Fault tolerance requires redundant components to take over and continue operation when a failure occurs.

Advantages

- Protection against disk subsystem component failures.
- Performance improvement over single drives, when using RAID-1 (mirroring).
- Intelligent hard disk controllers reduce CPU load by off-loading disk subsystem-related processes from the server's main processor to the hard disk controller.

Disadvantages

- Higher cost for the same disk subsystem net capacity. (Using RAID-1 requires twice as many drives as does using either single drives or RAID-0. Therefore, the drive cost is doubled.)
- Performance is degraded when using RAID-5 (in comparison to RAID-0 and RAID-1, but only one additional drive per array, resulting in only a moderately higher price tag).
- Higher cost for an intelligent hard disk controller (approximately three times higher).

This soft copy for use by IBM employees only.

Cost versus Capacity:

The following section describes the positive and the negative aspects of this approach.

The following examples are based on an IBM PC Server 500. Very similar calculations can be done for the other members of the IBM PC Server family by simply substituting the IBM SCSI-2 Fast/Wide Adapter/A with an equivalent adapter for the other IBM PC Server systems.

The components included in this calculation are:

- IBM SCSI/2 Fast/Wide adapter
- IBM 4.5GB SCSI-2 Fast/Wide hard disks
- IBM 2.25GB SCSI-2 Fast/Wide hard disks
- IBM 1.12GB SCSI-2 Fast/Wide hard disks
- Additional power supply

Required when more than one row of hot-plug drives will be installed.

- Additional hot-plug backplane

For each additional row, you need to add one backplane to your calculation. Therefore, a maximum of two additional backplanes may be required when you plan to populate rows 2 and 3 in the PC Server 500.

Figure 42 on page 132 illustrates the difference when using, for a given disk capacity (here 4.5GB), either one drive or up to four individual drives. Using a smaller number of large-capacity drives reduces the disk subsystem's price tag, because as mentioned above, a single 4.5GB drive is priced lower than four 1.12GB drives required to achieve the same disk storage capacity of 4.5GB. For the same disk subsystem capacity, the price is:

- US\$ 3,264.- when using a single 4.5GB disk drive
- US\$ 4,751.- when using four 1.12GB disk drives

This soft copy for use by IBM employees only.




Figure 42. Disk Subsystem Cost for 4.5GB Total Capacity

Figure 43 on page 133 shows four different disk subsystem capacities (4.5GB - 18GB) and the cost when using either 1.12GB, 2.25GB or 4.5GB drives. Be aware that you need to add the price for up to two more backplanes and the additional power supply to install more than one six hard disks in this server.

These parts are required on the PC Server 500 for installation of more than one bank of hard disks.

This soft copy for use by IBM employees only.




Figure 43. Disk Subsystem Cost (1.12GB, 2.25GB and 4.5GB Drives)

Figure 44 on page 134 illustrates the trade-off between cost and performance. It compares the throughput of a disk subsystem with a total net capacity of 4.5GB.

It shows the negative impact on performance when using one large 4.5GB hard disk instead of four 1.12GB drives, providing the same capacity and considerably better performance.

At this point you need to decide what is more important: a lower price with lower performance, or significantly higher performance at a moderately higher price.

This soft copy for use by IBM employees only.




Figure 44. IBM PC Server 500 with a Single 4.5GB Drive versus Four 1.12GB Drives

Test environment:

- PC Server 500, 32MB RAM
- IBM SCSI-2 Fast/Wide Adapter/A
- One 4.5GB SCSI-2 Fast drive for the first test
- Four 1.12GB SCSI-2 Fast/Wide drives
- Multidisk volume defined using NetWare 3.12
- Net hard disk capacity of 4.5GB
- NOS: NetWare 3.12
- Test: NetBench 3.01, Disk_Mix, 4 - 60 stations

This soft copy for use by IBM employees only.

Conclusion

Although a single large-capacity disk drive provides the lowest priced solution, you have to accept considerably lower disk subsystem performance with higher disk subsystem loads.

Using large-capacity drives is acceptable when:

- Cost is your major criterion
- User access frequency is low
- Amount of data per transaction is low

On the other hand, when many users frequently access your server's disk subsystem, considerably more time is required to complete a request. This may lead to lower productivity and can justify the implementation of a disk subsystem that provides better response time and throughput.

4.1.2.2 Performance

To achieve maximum performance, you need intelligent hard disk controllers to off-load as many disk-related tasks as possible from the system's main processor. Fast hard disks with high data transfer rates are required as well as a high-bandwidth SCSI bus to efficiently handle the data and command flow between the controller and the drives. Multiple SCSI channels increase the bandwidth, provide higher performance, and further improve configuration flexibility and fault tolerance.

To obtain maximum performance, you should do the following:

- *Use RAID technology (RAID-0 or RAID-1).*

Both RAID-0 and IBM's RAID-1 implementation enable you to combine several physical drives into a single logical drive. Data will be distributed equally over all affected drives in stripes. For a given file size, this results in less disk space used on each individual drive, thus reducing the time required for seek operations to locate the data. In addition, because I/O operations can be initiated to all drives in the array simultaneously, the completion of these I/O requests takes less time than it would take on a single drive. As a result, a RAID-0 disk array provides higher performance than a single drive of the same capacity as the array.

- *Use hard disks with a short seek time and latency.*

A short seek time results in faster head positioning, thus reducing the time required to reach the desired track. High rotation-per-minute (rpm) drives (drive spindle turns at 7200 rpm) have a shorter latency; thus,

This soft copy for use by IBM employees only.

there is a shorter delay until a disk sector passes the drive's read/write heads, and can be accessed for read or write operations.

This soft copy for use by IBM employees only.

- *Use RAID-1 for maximum performance when using mainly read I/O operations (eg. executable program loads, video data, archive data etc.)..*

RAID-1 uses a similar concept of data striping but also offers fault tolerance (mirroring). For more details, see section 4.1.2.3, "Fault Tolerance" on page 139. RAID-1 can also offer significant performance improvements because read requests are always directed to both the original drive and the drive holding the mirrored data.

If the majority of all I/O requests on a RAID-1 array are read requests, then a RAID-1 array can be at least as fast as a RAID-0 array.

If the majority of all I/O operations are write requests, then a RAID-1 array will be slower than a RAID-0 array, because each write operation needs to be executed twice, once to the original stripe *and* once to the mirrored stripe. Therefore, the performance obtained using a RAID-1 array depends on whether the I/O load is write- or read-intensive.

To achieve the best performance using RAID-0 configuration, follow these principles:

- The combination of an intelligent RAID controller and multiple drives provides the best performance and off-loads any disk subsystem management tasks from the server system's processor. In addition, your configuration is operating system-independent, because the array handling and control are executed by the RAID controller itself. The RAID controller will present a disk array (independent of the number of its physical drives) as a *single* drive to the operating system.
- Another method for creating a logical drive that consists of multiple physical drives is to use NetWare as your network operating system (NOS). NetWare can create a disk volume (the equivalent of a logical drive) consisting of multiple disk drives. This configuration does not require a RAID controller; a normal disk controller can be used. The operating system handles the management of individual drives as one large multidisk volume. This appears to be a RAID-0-like configuration.

Using the RAID controller and a RAID-0 array will provide a high-performance and operating system-independent solution. The software solution is always operating system-dependent.

Figure 45 on page 138 shows the performance of both implementations, the hardware implementation using a RAID controller handling a RAID-0 disk array (here called HW RAID-0) and the software implementation, a standard SCSI-2 Fast/Wide disk controller, where the software (NetWare 3.12) handles and combines multiple drives into one (RAID-0-like) multidisk volume.

This soft copy for use by IBM employees only.

The hardware RAID-0 implementation shows higher performance under a medium and an especially heavy load than does the software RAID-0 solution. If many users concurrently access your server's disk subsystem, the hardware RAID-0 implementation will provide better performance and thus higher end-user productivity.




Figure 45. IBM PC Server 500 with RAID-0 versus NetWare Multidisk Volume

Test environment:

- PC Server 500, 64MB RAM
- Streaming RAID Adapter/A (HW RAID-0)
- IBM SCSI-2 Fast/Wide Adapter/A (Multidisk volume)
- Six 1.12 SCSI-2 Fast/Wide drives
- RAID-0 array with six drives
- Hard disk net capacity 6.7GB
- Multidisk volume defined using NetWare 3.12, drive spanning feature
- NOS: NetWare 3.12

This soft copy for use by IBM employees only.

- Test: NetBench 3.01, Disk_Mix, 4-60 stations

4.1.2.3 Fault Tolerance

Fault tolerance can be provided by redundant components or other methods to keep your server's disk subsystem operational even if one component fails. It protects you against the loss of data and programs in case of a disk failure, reduces or eliminates server downtime, and allows maintenance and replacement of failing components at a later and scheduled time instead of at the time of failure. This also applies for remote locations and branch offices where no on site support is available.

RAID - Redundant Array of Inexpensive Disks:

This section briefly discusses RAID technology and the RAID levels used in IBM PC Server systems. For details, see the ITSO redbook titled *IBM Personal Computer Disk Subsystem Considerations*.

A disk array consists of individual and independent disk drives, which are grouped together as an array and controlled by array management software. Array management software controls the operation of the disk drives and presents them as one or more virtual disks or logical drives to the host operating system.

Array management software can be part of an operating system. In this case, it is executed by the host system's processor. This method allows RAID implementation using standard disk controllers, but results in considerably higher host system CPU utilization.

To off-load RAID management tasks from the CPU and reduce its load, intelligent disk controllers equipped with dedicated and fast processors can be used to execute all RAID management tasks, thus reducing the main processor's load and providing operating system independence.

Disk arrays may improve some or all characteristics compared to an individual disk drive:

- *Capacity*

Multiple, independent drives are combined by array management software into a single logical drive and presented as one disk drive to the host operating system. This principle permits the creation of disk volumes of very large capacity using standard physical disk drives.

- *I/O performance*

This soft copy for use by IBM employees only.

Data is evenly distributed across the disks in the array. This is also called striping. As multiple disk drives are concurrently involved in writing and reading data, the I/O request processing runs on all drives in parallel and reduces the time required for seek and data transfer compared to a single disk drive.

- *Fault tolerance by redundancy*

When one disk drive in the array fails, the array remains operational. This redundancy can be accomplished in several ways. Following is a short explanation of the two most frequently used principles:

- Disk mirroring. All data is written to two drives in parallel, the original drive and a second (mirrored) drive. If one drive fails, data is still available on the other drive and your array becomes fault-tolerant.
- Check sum (parity) generation. For each block of data written to the disk array, parity information is created and written as a separate block to one of the member disks in the array. To create this parity information, the Boolean Exclusive Or function is used. Parity blocks use the capacity of one drive in the array and can be either all written to a separate drive or to parity blocks, distributed across all drives in the array.

When one of the drives in the array fails, the missing data can be reconstructed by using the remaining data and the parity block information.

RAID Levels:

Of the six different disk array models or RAID levels (RAID-0 through RAID-5), three have proved commercially attractive and are used by IBM for further improvement of performance, capacity and reliability in their PC Servers' disk subsystems.

These RAID levels are:

- *RAID level 0 or RAID-0*

As explained above, with RAID-0, data is distributed evenly across all disk drives in the array. This is also called disk striping.

Although the absence of redundancy in a striped array makes the term RAID-0 a misnomer, this definition is used to describe disk striping as RAID level 0.

This soft copy for use by IBM employees only.

RAID-0 disk arrays provide excellent performance and the highest net capacity of all RAID implementations because its total capacity is calculated as follows:

$$\text{RAID-0 net capacity} = \text{Drive capacity} \times N$$

(Where N is the number of drives in the array)

Therefore, RAID-0 provides RAID technology at the lowest cost compared to the other commonly used RAID implementations. However, it is important to mention that RAID-0 does *not* offer fault tolerance.

Figure 46 illustrates the RAID-0 implementation.




Figure 46. RAID-0, Disk Striping

This soft copy for use by IBM employees only.

- *RAID level 1 or RAID-1*

Fault tolerance is accomplished by disk mirroring. To improve performance over which one of two individual drives, IBM's RAID-1 implementation combines disk striping with disk mirroring.

For read-intensive applications, this results in excellent performance close to or even exceeding RAID-0 because seek and read operations are executed on both drives simultaneously, and the drive that reads the the data first due to a shorter seek time and latency will deliver them to the controller.

However, for applications issuing predominantly write data requests, the performance will be lower due to the fact that data must be written to both disks.

The combination of both disk striping and mirroring also permits the use of either an odd or an even number of disks in the array, providing greater flexibility in defining disk arrays.

RAID-1 offers only 50% net capacity because for each drive to be mirrored, a second drive is required. Its total capacity is calculated as follows:

$$\text{RAID-1 net capacity} = (\text{Drive capacity} \times N) / 2$$

(Where N is the number of drives in the array)

Therefore, RAID-1 provides fault tolerance and excellent performance but at a higher cost than the other commonly used RAID implementations.

Figure 47 on page 143 illustrates the RAID-1 implementation.

This soft copy for use by IBM employees only.




Figure 47. RAID-1, Disk Mirroring

- RAID level 5 or RAID-5

RAID-5 offers fault tolerance comparable to that of RAID-1 with read performance similar to that of disk striping. Data and computed parity information (blocks or stripes) are stored in the same disk stripe, that is, distributed across all disks in the array. This implementation provides some relief from the write bottleneck that characterizes RAID level 4, which uses a dedicated drive to store the computed parity information.

As explained above, parity blocks are generated from the data blocks, and both are written to same disk stripe. When one drive in the array fails, all data on this drive are lost but can be reconstructed using the

This soft copy for use by IBM employees only.

remaining data blocks *and* the parity blocks, thus keeping the array operational and providing fault tolerance.

The parity blocks *always* require the capacity of one drive in the array as shown in the following example.

The net capacity of a RAID-5 array is calculated as follows:

RAID-5 net capacity = Drive capacity * (N-1)

(Where N is the number of drives in the array)

To illustrate the higher net capacity provided by a RAID-5 disk array, assume the following:

- Your disk array consists of five 1GB drives providing a total array capacity of 5GB.
- The capacity of one drive is used to hold the parity blocks.
- This leaves you with a net capacity of: $5 - 1 = 4\text{GB}$.
- When using RAID-1, the net capacity would be: $5/2 = 2.5\text{GB}$ or 50% of the total capacity.

The overall performance of a RAID-5 array is generally lower than the performance of a RAID-0 and RAID-1 array because of the overhead in processing when writing a block of data. First the target data block and its corresponding parity block are read, the old data block's parity information is stripped from the parity block, then the new parity block is calculated, and finally, both the parity and the data block are written to the disk array.

Therefore, RAID-5 provides lower performance than RAID-0 and RAID-1, but it provides fault tolerance at a considerably lower cost compared to RAID-1.

Figure 48 on page 145 illustrates the RAID-5 implementation.

This soft copy for use by IBM employees only.




Figure 48. RAID-5, Striping Plus Parity Blocks

Cost:

When your major criterion is fault tolerance due to mission-critical applications and data, you still have the choice between two implementations offering fault tolerance. Again, you need to prioritize your requirements.

When your application requires high performance combined with fault tolerance, you should choose RAID-1 as the most suitable solution. However, be aware that RAID-1 delivers only 50% disk subsystem net capacity, which is a drawback because of its inherent cost.

You may try to reduce the cost for a RAID-1 disk array by using significantly lower-priced, normal disk controllers instead of intelligent RAID controllers,

This soft copy for use by IBM employees only.

and let the network operating system perform the RAID-1 management (mirroring or duplexing) tasks, but again there is a trade-off between cost and performance. Also, if you calculate the cost for a disk subsystem built, for example, from six individual drives, the difference between a normal disk controller and an intelligent RAID controller becomes less significant, because the disks account for most of the cost.

Assume the following:

- Standard SCSI controller: US\$ 500.-
- RAID controller: US\$ 1500.-
- Disk drives: US\$ 1300.- per drive
- Total cost using a standard disk controller: US\$ 8300.-
- Total cost when using a RAID controller: US\$ 9300.-
- Difference between both implementations: US\$ 1000.- or only 11%

Figure 49 on page 147 shows the difference in performance you can expect from either solution.

As the figures above demonstrate, a hardware-based RAID-1 implementation using an intelligent RAID controller provides superior performance over the software-based implementation, especially when a large number of attached workstations are accessing the server's disk subsystem.

Therefore, considering the higher performance of the HW RAID-1 solution, we recommended that you use the HW RAID-1 implementation on networks with many attached users making frequent accesses to the server's disk subsystem.

The SW RAID-1 implementation can be used for smaller networks or when the frequency of user accesses is low.

This soft copy for use by IBM employees only.




Figure 49. Software versus Hardware Mirroring

Test environment:

- PC Server 500, 64MB RAM
- Streaming RAID Adapter/A (HW RAID-1)
- IBM SCSI-2 Fast/Wide Adapter/A (SW RAID-1)
- 6 x 1.12 SCSI-2 Fast/Wide drives
- RAID-1 array with six drives
- Net capacity 3.35GB
- NOS: NetWare 3.12
- Test: NetBench 3.01, Disk_Mix 4-60 stations

This soft copy for use by IBM employees only.

When fault tolerance is required, but cost is also a major concern, then you need to compromise on performance and choose a RAID-5 implementation.

RAID-5 keeps the overall cost for your disk subsystem considerably lower than a RAID-1 implementation at the penalty of lower performance.

Figure 50 on page 149 illustrates the relationship between price and net capacity for all three RAID implementations, RAID 0, 1, and 5.

It shows on the Y-axis (left) the maximum disk subsystem capacity for each RAID implementation and on the Y2-axis (right) the cost in US Dollars per MB of disk subsystem storage capacity. The calculation is based on a disk subsystem, built from 12 x 1.12 GB disk drives. The maximum net capacity is as follows:

- RAID-0
13.5GB net capacity, sum of all individual drives (100%)
- RAID-1
6.75GB net capacity due to mirroring (50%)
- RAID-5
11.2GB net capacity. As the number of drives in one array is limited to eight drives, two arrays are defined, each requiring the capacity of one drive in the array to store the parity blocks. Therefore, 2 x 1.12GB must be deducted from the total disk capacity of 13.5GB.

The chart also shows the cost in US Dollars/MB of the RAID-1 and RAID-5 implementations, identifying RAID-5 as providing the best cost/capacity ratio in conjunction with fault tolerance. However, the trade-off is the lower performance of RAID-5 compared to disk subsystems using RAID-1.

This soft copy for use by IBM employees only.




Figure 50. Disk Subsystem Cost (US Dollars/MB) for RAID-0, 1 and 5

4.1.3 Summary

In sizing your server's disk subsystem, consider the following:

- What are the most important criteria? Are they:
 - Cost
 - Capacity
 - Performance
 - Fault tolerance
- If your major criterion is cost and you do not need fault tolerance, then you should use a standard disk controller and high-capacity disk drives. Remember that a single 4.5GB drive costs less than four 1.12GB or two 2.25GB drives.

This soft copy for use by IBM employees only.

Based on the current list prices for 4.5GB, 2.25GB and 1.12GB IBM SCSI-2 Fast/Wide Ultrastar drives, this will result in the following cost per GB storage capacity:

- 4.5GB drives = US\$ 611
 - 2.25GB drives = US\$ 720
 - 1.12GB drives = US\$ 941
- When your major criterion is performance, then you should use a RAID controller, define RAID-0 arrays and use more drives of either 1.12GB or 2.25GB capacity. This will provide better performance than the 4.5GB drives as shown in Figure 44 on page 134.

The cost per GB of disk storage capacity is identical to the examples shown above.

- When your major criterion is fault tolerance, then you should use an intelligent RAID controller and choose either RAID-1 or RAID-5.

The next factor to consider is the cost for a given disk subsystem storage. If cost is the main criterion for your disk subsystem, then you should choose the RAID-5 implementation because it provides the highest storage capacity.

If, on the other hand, performance is your main criterion, then you should use the RAID-1 implementation because it delivers considerably higher performance than RAID-5.

Based on the current list prices for 4.5GB, 2.25GB and 1.12GB IBM SCSI-2 Fast/Wide Ultrastar drives, this will result in the following cost per GB of storage capacity, depending on the RAID level (1 or 5) you choose and assuming a net capacity of 9GB:

- RAID-1
- For 9GB net capacity, you need (US\$/GB):
- 4 x 4.5GB drives = US\$ 1,222
 - 8 x 2.25GB drives = US\$ 1,439
 - 16 x 1.12GB drives = US\$ 1,883

As with RAID-0, you may need to compromise on performance when choosing a smaller number of large-capacity drives (for example, 4.5GB drives instead of 1.12GB drives).

- RAID-5
- For 9GB net capacity, you need (US\$/GB):

This soft copy for use by IBM employees only.

- 3 x 4.5GB drives = US\$ 916
- 5 x 2.25GB drives = US\$ 899
- 9 x 1.12GB drives = US\$ 1059

Here, for the first time, the use of 2.25GB drives carries a lower price tag than the implementation using 4.5GB drives. Remember that RAID-5 requires the capacity of *one drive* to store the parity blocks. A net capacity of 9GB requires three x 4.5GB drives. In other words, you lose the 4.5GB required for storing the parity blocks.

When using, for example, 2.25GB drives, you need five drives (4 + 1 for the parity blocks), which means you lose only 2.25GB instead of the 4.5GB in the previous example.

Therefore, from all aspects, performance and cost for a given net capacity, the use of lower-capacity drives on RAID-5 is more efficient.

Disk Subsystem Selection: Figure 51 on page 152 provides a flow chart-like approach for selecting the disk subsystem implementation according to your business needs.

The first question you need to ask is "How much net capacity do you need for the disk subsystem?" The second row shows the selection you need to make, depending on whether your major requirement is performance, cost or fault tolerance. The chart also provides recommendations about the disk subsystem implementations.

This soft copy for use by IBM employees only.




Figure 51. Disk Selection Guide

Remark:

The reason for recommending a software-based RAID-1 (mirroring) solution to meet the requirement for fault tolerance at the lowest possible cost is

This soft copy for use by IBM employees only.

simply the fact that the minimum requirements for RAID-5 are three disk drives and a RAID controller. When you intend to use up to four drives, then network operating system-based (IBM LAN Server 4.0 Advanced or NetWare 3.12 or 4.x) mirroring using standard disk controllers provides fault tolerance at a lower cost than the hardware RAID-5 implementation.

The break-even point on cost between a hardware- and a software-based RAID implementation is approximately at four drives as shown in Figure 52. However, be aware that four 1.12GB disk drives provide 2.24GB of disk space when using RAID-1 and 3.36GB of disk space when using RAID-5.




Figure 52. Cost Software RAID-1 versus Hardware RAID-5

The software RAID-1 solution requires a net capacity of 3.3GB six 1.12GB disk drives. Meanwhile, the hardware RAID-5 solution requires only four disk drives. Therefore, when you need to use more than four disk drives to achieve the required disk subsystem net capacity, you should use hardware

This soft copy for use by IBM employees only.

RAID-5 to implement a fault-tolerant disk subsystem at the lowest possible cost.

If you compare the performance of both the software and hardware RAID implementation, there is a difference of approximately 10% between the software RAID-1 and the hardware RAID-5 implementation, which should be acceptable for these implementations, where besides fault tolerance, cost is your major criterion.




Figure 53. Performance Software RAID-1 versus Hardware RAID-5

Test environment:

- PC Server 500, 64MB RAM
- Streaming RAID Adapter/A (HW RAID-5)
- SCSI-2 Fast/Wide Adapter/A (SW RAID-1)
- 6 x 1.12 SCSI-2 Fast/Wide drives
- RAID-1 and RAID-5 array with six drives

This soft copy for use by IBM employees only.

- Net capacity 5.6GB (RAID-5) and 3.36GB (RAID-1)
- NOS: NetWare 3.12
- Test: NetBench 3.01, Disk_Mix 4-60 stations

Conclusion:

Except for small disk subsystems with two to four drives, the hardware RAID-5 solution provides a fault-tolerant disk subsystem at the lowest cost for a given disk subsystem capacity. From a performance standpoint, there is only a small difference, for medium and heavy loads, between the software RAID-1 and the hardware RAID-5 solution as shown in Figure 53 on page 154.

4.2 Disk Subsystem Performance

The following section addresses frequently asked questions about the disk subsystems used in the IBM PC Server family. Its findings are based on numerous test runs using two members of the IBM PC Server family, the IBM PC Server 500 array and non-array and the IBM PC Server 320 array.

The benchmark program used is Ziff-Davis' NetBench Version 3.01 from Ziff-Davis.

See Appendix B, "Description - Test Environment" on page 211 for a detailed description of the test environment (e.g., servers, memory, network operating system, hardware and software configuration, driver levels). If the test environment differs from the specifications listed in Appendix B, the differences will be clearly described for each affected example.

The charts shown in this chapter are not intended to show a server's maximum throughput rate; their purpose is to illustrate the effect of configuration modifications and variations, and to assist you in optimizing your disk subsystem.

4.2.1 RAID Performance

This paragraph discusses the differences in performance when your disk subsystem is configured for RAID-0, RAID-1 or RAID-5.

As you can see from Figure 54 on page 158, RAID-0 provides the best performance, but no fault tolerance. A disk drive's reliability is defined as *Mean Time Between Failures* (MTBF), which is the average number of operating hours before a disk failure occurs.

When combining several independent disk drives into one array, the MTBF figure for the entire array is different and lower than that for an individual drive.

The formula for calculating an array's MTBF is as follows:

$$\text{Array MTBF} = \text{drive MTBF} / \text{number of drives}$$

The more drives you use to create a disk array, the higher the probability of a non-recoverable disk failure. Therefore, for improved disk subsystem reliability, you should use the RAID levels that provide fault tolerance.

This soft copy for use by IBM employees only.

RAID-1 (disk mirroring) provides fault tolerance and excellent performance as shown in Figure 54 on page 158. However, be aware that, especially on RAID-1, the maximum obtainable performance depends to a large extent on the mix of read and write requests issued to the disk subsystem.

If the requests are more write-bound, the RAID-1 performance will be lower than that shown in Figure 13 because all data needs to be written to *both* sides of the mirror before the next write request can be processed. If, for example, only read requests will be issued, then the performance of a RAID-1 array may reach or exceed that of a RAID-0 array.

It should be understood that throughput figures are always request-dependent, and even when using exactly the same hardware configuration, you will see different figures (for example, using other test suites or application programs having different I/O request patterns).

RAID-5 shows the lowest performance, but as mentioned in the previous chapter, it combines the advantages of fault tolerance with a high net capacity and therefore shows the best price per usable MB of disk storage.

Figure 54 on page 158 shows and summarizes the differences in performance between the three RAID implementations.




Figure 54. IBM PC Server 500, RAID-0, 1 and 5 Throughput

Test environment:

- PC Server 500, 64MB RAM
- Streaming RAID Adapter/A
- 6 x 1.12 SCSI-2 Fast/Wide drives
- RAID-0, 1 and 5 array with six drives
- Net capacity 6.7GB (RAID-0), 5.6GB (RAID-5) and 3.36GB (RAID-1)
- NOS: NetWare 3.12
- Test: NetBench 3.01, Disk_Mix 4-60 stations

This soft copy for use by IBM employees only.

4.2.2 Hardware versus Software RAID

Today's network operating systems, such as IBM OS/2 LAN Server 4.0 Advanced or Novell NetWare 3.12 or 4.x, offer disk mirroring or duplexing capabilities to improve your disk subsystem's reliability.

The following chart compares the two possible implementations, either using a standard disk controller (here the IBM SCSI-2 Fast/Wide Adapter/A) and letting the network operating system perform the task of disk mirroring (RAID-1), or using an intelligent RAID controller (here the IBM Streaming RAID Adapter/A) to off-load the disk subsystem management tasks from the system's main processor to the RAID controller. Figure 55 on page 160 illustrates the difference between both implementations and shows significantly higher performance when using the Streaming RAID Adapter/A.




Figure 55. Software versus Hardware Mirroring

Test environment:

- PC Server 500, 64MB RAM
- Streaming RAID Adapter/A (HW RAID-1)
- IBM SCSI-2 Fast/Wide Adapter/A (SW RAID-1)
- 6 x 1.12 SCSI-2 Fast/Wide drives
- RAID-1 array with six drives
- Net capacity 3.35GB
- NOS: NetWare 3.12
- Test: NetBench 3.01, Disk_Mix 4-60 stations

This soft copy for use by IBM employees only.

4.2.3 Number of Drives

When adding more drives to an array, there is often the question of whether the higher demand for capacity will result in lower disk subsystem performance due to bandwidth limitations of the SCSI channel the drives are attached to.

Remember, when using the SCSI-2 Fast/Wide implementation, the bandwidth of one SCSI channel is 20MBps. Since each IBM Ultrastar disk drive offers a maximum transfer rate up to 7.4MBps,⁴ the concern is:

- When connecting three drives to the same SCSI-channel, do we reach or exceed the SCSI channel's bandwidth and does adding more drives result in a performance degradation? The answer is No, there is no performance degradation caused by more than three, up to a maximum of six, drives per SCSI channel on the IBM RAID adapters.

The maximum transfer rate that each drive can handle is 7.4MBps, but there is no continuous data stream to or from a drive. The transfer is performed in bursts, each with a transfer rate of up to 7.4MBps. Therefore, the overall transfer rate is lower. Even with six drives installed and connected to the same SCSI-channel, you don't usually reach the maximum bandwidth of 20MBps on a SCSI-2 Fast/Wide channel.

Today, the maximum number of drives that can be attached to one SCSI-channel is limited to seven drives when using the IBM 3516 external storage enclosure and to six drives per channel due to the backplane used in the following IBM PC Server models offering only six connectors:

- PC Server 320, array models
- PC Server 500
- PC Server 520
- PC Server 720

Figure 56 on page 162 illustrates that, in addition to a higher disk subsystem capacity, you may even gain higher performance due to distributing the load to more drives in the array and reduced seek times.

⁴ For more information on the technical specification for the IBM Ultrastar disk see Table 14 on page 61.

This soft copy for use by IBM employees only.




Figure 56. IBM PC Server 320, Throughput versus Number of Drives

Test environment:

- PC Server 320, 32MB RAM
- IBM PCI RAID Adapter
- 2-6 x 1.12GB SCSI-2 Fast/Wide drives
- RAID-0 array with two to six drives
- Net capacity 2.2 - 6.7GB
- NOS: NetWare 3.12
- Test: NetBench 3.01, Disk_Mix, 60 stations

This soft copy for use by IBM employees only.

4.2.4 SCSI Channel Bandwidth

Due to the limitation caused by the backplane used to connect disk drives on the SCSI bus, you can connect up to six disk drives to each of the two SCSI channels provided by the IBM RAID controllers.

The IBM RAID controllers also allow disk drives, connected to both SCSI channels, to be combined as one array that can span both SCSI channels.

However, you cannot combine drives as one array, which are connected to SCSI channels of *different* RAID controllers.

The advantages of this implementation are:

- Two SCSI-2 Fast/Wide channels each with a bandwidth of 20MBps increases the total SCSI bandwidth to 40MBps per RAID controller, permitting the connection of more or even faster future devices without exceeding the combined bandwidth provided by the IBM RAID adapters.
- A maximum of 12 drives can be connected to one RAID controller, providing further disk subsystem expandability as shown in Figure 57 on page 164.
- A single disk array may contain up to eight drives, exceeding the capabilities of a single SCSI channel, but here the drives can be evenly distributed over both SCSI channels as shown in Figure 58 on page 165.
- In addition, IBM RAID controllers provide the option to define a spare drive that will automatically replace a failing drive without external intervention.

The advantages are clear:

- No immediate human intervention is required because the disk subsystem automatically replaces a failing drive.
- Protection against the failure of an additional drive is provided. Only one drive can fail. If two or more drives fail, the array becomes defunct, and all your data and programs are lost and need to be restored after the failing disks have been replaced.
- Reduced performance will be limited to the time required to rebuild the array, that is, reconstructing the lost data on the new drive, automatically inserted into the array.
- Only *one spare drive* is required for both SCSI channels. The spare drive can be connected to either SCSI channel.

This soft copy for use by IBM employees only.




Figure 57. Disk Arrays on Separate SCSI Channels

This soft copy for use by IBM employees only.




Figure 58. Disk Array Spanning Two SCSI Channels

4.2.5 Stripe Size

Another factor impacting your disk subsystem's performance is your array's stripe size. For best performance, it should match your network operating system's volume block size when defining NetWare volumes or the *CRECL*= definition in the *IFS*= statement used by OS/2. These values should match in order to optimize your disk subsystem's performance. It is recommended that the parameter *set cache buffer size* = match the block size chosen during the NetWare volume definition (this is valid for NetWare 3.12 but not for NetWare 4.x).

The disk array's stripe or block size is defined using the following procedure:

- Start the RAID configuration utility by booting from the RAID option diskette.
- Select Create/Delete Array.

This soft copy for use by IBM employees only.

- Define your disk array and define your logical drive(s). When defining logical drives, you also select the RAID level to be used for this logical drive.
- Exit and select Advanced Functions.
- Select the option Define Stripe Unit Size.
- Select the required stripe size and save the configuration on the RAID adapter.
- Exit and select Initialize Array.

Don't forget this step; otherwise, you may encounter problems when trying to install the network operating system.

A stripe unit size different from the NetWare volume block size may result in more I/O requests being issued to the individual disk drives in the array, to more seek operations, and thus result in a performance degradation.

Figure 59 on page 167 illustrates the impact of selecting a stripe unit size different from the NetWare volume block size.

This soft copy for use by IBM employees only.




Figure 59. IBM PC Server 320, RAID Controller Stripe Unit Size RAID-0

Test environment:

- PC Server 320, 32MB RAM
- IBM PCI RAID Adapter
- 6 x 1.12GB SCSI-2 Fast/Wide drives
- RAID-0 array with six drives
- Net capacity: 6.7GB
- Three test runs, one for each selected stripe unit size
- Stripe unit size: 8KB, 16KB, 32KB
- NOS: NetWare 3.12
- Test: NetBench 3.01, Disk_Mix, 60 stations

This soft copy for use by IBM employees only.

4.2.6 RAM Disk Cache

To achieve higher throughput, server systems frequently use memory to cache data read from the disk subsystem as well as data that should be written to the disk subsystem.

When an application asks a server for data, the server first looks for the data already in cache memory (located in the server's main memory). When the data is found in cache, it is called a cache-hit. When data is not found, it is called a cache-miss and the data must be retrieved from the disk subsystem. Cache-hits usually take only a few microseconds to complete; they involve the server's CPU and the network interface. When disk accesses are required due to a cache-miss, considerably more time is required to complete the operation. A higher cache-hit ratio means better performance, although it strains the capacity of the LAN subsystem.

Because a larger cache memory can hold more data and eventually helps to reduce disk accesses due to a higher cache-hit rate, it usually improves the overall server performance.

A larger disk cache should improve the cache-hit ratio, reduce the frequency of disk accesses due to cache-misses, and thus lead to better server performance.

4.2.7 Hardware Disk Cache

As mentioned above, disk caching can help to improve the performance of your server's disk subsystem. Caching can be done on read data requests to reduce the number of disk accesses by trying to get the data from cache first. If there is a cache-miss, that is, data is not found in the cache memory, then the data is retrieved from the disk subsystem, which takes considerably more time to complete due to the physical disk access required.

Caching is also used to reduce the time it takes to write data to the disk drives.

Assume the following:

- While the disk subsystem is executing a write request and is currently too busy to write the data to the disk drives, a second write request is issued.
- Until all data is written to the disk drives and the first write request is completed, the second write request will remain queued, waiting for execution. If other write requests are issued, they also will become queued, until all previous write-request processing completes. This

This soft copy for use by IBM employees only.

serializes the execution of write requests and can reduce the overall performance of the server and its disk subsystem. This may become a serious bottleneck, because a file server environment is always I/O-request-intensive; that is, for maximum performance the disk subsystem must be able to handle multiple I/O requests concurrently.

- Caching write requests helps to reduce the write-request serialization bottleneck. The write request is stored in cache memory and then immediately posted as complete to the server's NOS file system. However, the actual disk write operation is executed at a later moment when the disk subsystem completes the previous request and is ready for the next write operation. This type of operation is called *write-back caching*.

But what happens if there is a serious problem, such as a power loss, which causes the server to crash?

- Some write requests that are still in cache awaiting execution are lost. Because they are already posted as completed, the application issuing the write request assumes that the write request has been completed and the data has been saved on your disk drives. Therefore, when using write-back, you may compromise on data integrity.
- The risk of a total power loss can be eliminated by attaching an Uninterruptible Power Supply (UPS) to your server, which automatically takes over in case of a total loss of power, but even more important, it protects your server against impulse noise, spikes and voltage changes (brownouts). These highly intermittent events can cause server problems, which are difficult to diagnose.
- In addition, the use of ECC memory and highly reliable components reduces considerably the risk of losing data caused by a server failure. Therefore, the default definition cache-write policy for the RAID controller's hardware disk cache is write-back.

If you want to further reduce the risk of lost data, then you should change the cache-write policy to write-through at the penalty of lower write performance, which may impact the overall performance of your server's disk subsystem.

- The cache-write policy can be changed using the RAID configuration utility, which is started when you boot your server system from the RAID option diskette.

Select "Advanced functions" and then "Change cache write policy." You may toggle between write-back (WB) and write-through (WT).

This soft copy for use by IBM employees only.

The cache-write policy can be changed without redefining the array and the loss of data already stored on your disk array. You may change the cache-write policy at any time.

Figure 60 shows the difference for a given disk load (NetBench Version 3.01, Disk Mix) when using either write-back or write-through as cache-write policy for the RAID controller's hardware disk cache.

The differences shown in Figure 19 are caused by the fact that on RAID-1, and especially on RAID-5, you need more I/O operations to complete a write request. On RAID-5, first the target data block is read, parity information is read and updated, and both the new data and the updated parity block are written to the disk array.

This processing overhead explains why, on RAID-5, you see the largest difference in performance when using either write-through or write-back cache-write policy.




Figure 60. IBM PC Server 320, Write-Back versus Write-Through

This soft copy for use by IBM employees only.

Test environment:

- PC Server 320, 32MB RAM
- IBM PCI RAID Adapter
- 6 x 1.12GB SCSI-2 Fast/Wide drives
- RAID-0, 1, 5 array with six drives
- Net capacity: 6.7GB, 3.3GB, 5.6GB
- NOS: NetWare 3.12
- Test: NetBench 3.01, Disk_Mix for 60 users only

4.2.8 RAID-5 Degraded

If one drive in a RAID-5 array fails, the data lost needs to be reconstructed from the remaining drives' data and parity blocks. This provides some processing overhead, reducing the entire RAID-5 array's performance.

To answer the frequently asked question about the performance degradation caused by a drive failure in a RAID-5 array, we compared the results of two different tests:

- One test environment shows the performance of a RAID-5 array with all drives operational.
- In the second test environment, we used exactly the same configuration but pulled one drive to simulate a failing drive.

Figure 61 on page 172 shows the performance degradation on a RAID-5 array with one defunctional drive. The test results can vary widely if, for example, your test array is built from a different number of drives, or you are using drive types. Therefore, the results shown in Figure 20 should be interpreted as:

- Yes, there is a considerable performance degradation, and one should consider including a spare drive, which will at least limit the duration of the performance degradation.
- No, these throughput figures should not be used to define the exact throughput in MBps to show the maximum performance of a fully operational array versus an array with a defunctional drive.

This soft copy for use by IBM employees only.




Figure 61. IBM PC Server 320, RAID-5 with One Defunctional Drive

Test environment:

- PC Server 320, 32MB RAM
- IBM PCI RAID Adapter
- 6 x 1.12GB SCSI-2 Fast/Wide drives
- RAID-5 array with six drives, one drive defunctional
- Net capacity: 5.6GB
- NOS: NetWare 3.12
- Test: NetBench 3.01, Disk_Mix, 4 - 60 stations

Chapter 5. IBM PC Server High Availability Examples

The purpose of this chapter is to explain sample solutions that will provide high availability in a LAN environment where access to a server is critical.

The importance of data and the networks and servers that provide it continue to increase. More and more mission critical data is stored on PC server systems. Keeping a server system running and its data available to all attached workstations is essential for uninterrupted data access.

New technologies such as error correcting code (ECC) memory, fault tolerant disk subsystems, and protection against power failures are implemented to improve a server's reliability. Although servers contain fault tolerant components to reduce the server downtime, there are still server components whose fault may result in a server failure. In addition, a server may also become in-operational due to problems caused by software.

Protection against the situation, where the entire server becomes unavailable to the network can be prevented in several ways.

- A standby server replaces automatically and without any interruption a failing primary server.

Here server redundancy is accomplished continuously mirroring data from the primary server to a standby server to maintain data integrity.

A separate high speed communication link between the primary and the standby server ensures instant mirroring and avoids additional network load caused by the continuous mirroring process.

Permanent monitoring of the primary server's status ensures the immediate detection of a server fault and instant activation of the standby server as primary server for uninterrupted data access.

The standby server location can be different from that of the primary server. This is an important consideration, because it allows you to physically separate the two servers and install them in either different rooms, buildings or even at different sites, provided the high speed server to server link connection is available at both locations.

When for example one server is destroyed by fire, the second server placed at a different location will be saved and the network remains operational despite a totally destroyed and unrecoverable server. This concept fits well into disaster avoidance plans, which outlines methods

This soft copy for use by IBM employees only.

and implementations to keep a mission critical network operational even when severe events occur, such as fire or flooding.

- Switching an external disk subsystem using a SCSI-switch from the primary to the standby server when the primary server fails. The external disk subsystem holds all, operating system, application programs and data.

After switching from the primary to the standby server, the standby server re-boots from the external disk subsystem as primary server, and after a short interruption the server is again available to all network users.

5.1 NetWare SFT III by Novell

NetWare SFT III by Novell provides complete server fault tolerance by redundancy. Full redundancy is accomplished by mirroring any server I/O activity over a separate high speed link to a standby or secondary server. To the network itself only that server is visible, which currently acts as the primary server.

When the primary server fails, its mirrored partner automatically takes over. This process is fully transparent to the attached workstations and guarantees uninterrupted server availability.

After the failing server is repaired and activated again, it will announce itself to the only (or primary) server as secondary server. Both servers synchronize themselves by re-mirroring only information that has changed while one server was down.

In addition to the mirroring function the secondary server provides split seek and read operations. Multiple read requests are split between the primary and the secondary server for simultaneous processing to improve the servers performance on disk read requests.

A separate and dedicated high speed link, called the **MIRRORED SERVER LINK (MSL)** connects the primary and the secondary server for mirroring.

Applications and **NETWARE LOADABLE MODULES (NLMs)** that do access hardware cannot be mirrored and are thus unavailable to servers running SFT III.

This soft copy for use by IBM employees only.




Figure 62. SFT III Concept

5.1.1 SFT III Components

The following provides a short overview about the major components used in server systems using SFT III.

- MSL (Mirrored Server Link)

This is a dedicated high speed link connecting both servers. It acts like a bus extension from the primary to the secondary server. One additional high speed adapter is required in each server, directly connecting the two servers (for example by coaxial cable or fiber optics).

- IOEngine (Input/Output Engine)

This soft copy for use by IBM employees only.

This component handles physical processes (for example network and disk I/O), and is not mirrored and therefore each server has its own IOEngine. The primary server's IOEngine controls the entire logical server (both the primary and the secondary server). It converts all network requests, disk I/O, etc. into SFT III events, which then are sent to the MSEngines of both servers.

- MEngine (Mirrored Server Engine)

Both servers share the same MEngine. The MEngine's mirror applications, data and non-hardware related NLM's.

When a workstation wants to access a resource on the primary server, it sends a request to this server. The primary server's IOEngine sends a copy of this request over the mirrored server link to the secondary server's IOEngine. Both IOEngines now send the request to their part of the shared MEngine. Each part of the MEngine triggers processing after receiving the request. When the process is completed, a reply is sent to each IOEngine. However, only the primary IOEngine sends a response to the client.

This soft copy for use by IBM employees only.




Figure 63. SFT III IO- and MEngines

NetWare SFT III requires two identical servers (same make and model), equipped at least with Intel 80386 processors, equal amounts of memory and disk storage. Also identical monitors and monitor adapters are required.

The two servers are not required to be on the same LAN segment or even required to have the same type of LAN adapters installed.

Both servers need to be at the same inter-network and all workstations can send/receive packets to both of them.

5.1.2 Server Failure Handling

Server failure handling prevents network down time due to a server problem such as hardware, software, LAN adapter, and cable failure.

Possible failure situations are:

- Primary Server Failure
 - The secondary server detects a timeout on the MSL and no responses over the LAN. This indicates, that the primary server is down.
 - The secondary server displays an error message, indicating that the primary server is down and that the secondary server now takes over the primary server's role.
 - A message is sent to all attached clients that this server is now the primary server.
 - Client requests are now rerouted to the (new) primary server.
 - Any disk changes are tracked following the server failure for later re-synchronization.
 - After the failed servers problem is solved this server is restarted.
 - Both servers, the (new) primary server and the newly started server re-synchronize themselves.
 - The primary server sends all disk changes over the mirrored server link to the repaired server until both disk subsystems hold again identical data.
 - The repaired server becomes now the secondary server.
- Secondary Server Failure
 - Network service continues without interruption.
 - The primary server informs the network administrator, that the secondary server is down.
 - The primary server now tracks all disk changes for later re-synchronization.
 - After the failing (secondary server) is repaired, it is brought up again and announces its availability to the primary server.
 - Both servers synchronize themselves and exchange all information updated while the server was down.
- Mirrored Server Link Failure
 - Network service continues without interruption.
 - The secondary server detects the MSL failure. Server state packets still received across the network indicate that the primary server is still operational and active.
 - The secondary server shuts down because no more mirroring is possible.

This soft copy for use by IBM employees only.

- The primary server now tracks all disk changes for later re-synchronization.
- After the failing mirrored server link is repaired, the secondary server is started.
- Both servers synchronize themselves and exchange all information updated in the primary server while the secondary server was down.
- LAN Adapter failure in the primary server
 - The primary server's operating system detects a failing LAN adapter.
 - The condition of the secondary server's LAN adapter is checked. If the secondary server's LAN adapters are all functional and the server is still mirrored, the secondary server takes over as primary server and continues to serve the workstations.
 - To repair the failing LAN adapter, the administrator shuts down the former primary server.
 - The primary (former secondary) server now tracks all disk changes for later re-synchronization.
 - After the failing LAN adapter is repaired, the server is started and will become the secondary server.
 - Both servers synchronize themselves and exchange all information updated in the primary server while the secondary server was down.

5.2 StandbyServer 32 by Vinca

Vinca's StandbyServer 32 for LAN Server provides for the implementation of a hot spare file server capable of automatically taking over the functions of the failing server and becoming the primary server. By using StandbyServer 32 for LAN Server, data is continuously mirrored to a second computer over a dedicated high-speed link included with the product. When the primary server fails, either from hardware failure or software corruption, the standby system takes over server function automatically. Users can continue to access data with only a short pause as the switch over happens. The failing server then can be diagnosed and repaired while users continue working.

StandbyServer 32 for LAN Server consists out of two high-speed link cards, an interconnect cable, Vinca software and IBM's NetFinity system management software. The link cards and the cable provide the high-speed dedicated link from the primary to the standby server. The Vinca software gives IBM LAN Server the ability to mirror data from the primary server to the standby machine across the high-speed link.

By using IBM's NetFinity software, switch over from the primary to the standby machine is fully automated. NetFinity also allows a system

This soft copy for use by IBM employees only.

administrator to customize alerts which can be sent to a local client, a remote console or a digital pager.

5.2.1 Benefits

Vinca's StandbyServer 32 for LAN Server mirroring system can virtually eliminate server downtime, and offers the following benefits for a LAN Server network.

- **Full server fault tolerance**

Because data is being mirrored to a completely separate computer, all server components are redundant. RAID can protect disk drives and redundant power supplies can ensure continuous power, but the Vinca solution replicates all hardware components. In addition, since data is mirrored to a separate computer, the Vinca solution can even recover from a software failure on the primary server.

- **Primary server and standby machine do not need to be identical**

This allows a system administrator to configure a less expensive computer as the standby machine; perhaps using last year's 486 servers as a standby machine with a more expensive Pentium based system as a primary server. When the primary server fails and the users are switched to the slower standby machine, they could notice a performance difference between the two servers but they continue to have access to their data.

- **Full automation of server switch over**

When the primary server fails and the standby machine is initialized as the main machine, no user intervention is required. A system administrator doesn't need to drop everything and rescue the downed network and the stranded data. At remote sites where a system administrator may not be available immediately, users can continue working.

- **Ease of installation and use**

Vinca's StandbyServer 32 for LAN Server uses standard LAN Server advanced fault tolerance. The product is a simple, yet elegant and powerful solution to get maximum fault tolerance and data accessibility. No complex software needs to be installed or maintained. Standard OS/2 compatible Vinca device drivers are installed on both the primary and standby machines. A system administrator can install the StandbyServer 32 for LAN Server easily, and it requires no additional training past an understanding of LAN Server fault tolerance.

This soft copy for use by IBM employees only.

- **Remote notification**

With the use of NetFinity, Vinca's StandbyServer 32 for LAN Server can notify a system administrator of a server failure at a remote workstation or even on a digital pager. If a system administrator is offsite or is responsible for several networks at different locations he can always be aware of the status of the Vinca system.

- **Enhanced diagnosis of server failure**

StandbyServer 32 for LAN Server gives system administrators the ability to diagnose system failures while users continue to have access to their data and network services. Further, it allows an administrator the time to fully analyze why the server failed and how to prevent the problem in the future.

- **Does not impact network traffic**

Data is mirrored through Vinca's high-speed dedicated cable rather than over the network wiring. Data is also transmitted at a very high speed (160MBps). Therefore, StandbyServer 32 for LAN Server does not affect network traffic, and in some configurations it can actually enhance the performance of the network.

5.2.2 System Configuration

The Vinca server mirroring solution is designed to work on a wide variety of Intel based computers. The machines to be used as both the primary server and the standby machine must have the following general characteristics. Most of these considerations are requirements for standard LAN Server mirroring.

- Both the primary server and the standby machine must be capable of running OS/2, OS/2 LAN Server and IBM's NetFinity. This includes OS/2 WARP and OS/2 LAN Server 3.0 or 4.0 advanced.
- The primary server and the standby machine do not need to be identical in either bus configuration or processor speed and type. Vinca currently produces StandbyServer kits in EISA, Micro Channel or EISA/Micro Channel configurations. Both machines must support one of these bus types.
- The standby machine needs to be configured with network interface cards supporting the same protocol and LAN topology then the primary server for all LAN segments to be protected. The network interface cards in the standby machine must be attached to the LAN. The interface cards used in the primary server and the standby machine do not need to be identical.
- A Vinca server mirroring card must be installed in each machine.

This soft copy for use by IBM employees only.

- Both the primary server and the standby machine must have at least two drives installed or a RAID adapter installed, capable to present two independent drives to the system. StandbyServer 32 for LAN Server does not allow the first drive in each system to be mirrored. Mirroring begins at the second drive.
- Drives in a mirrored pair must report to OS/2 identical drive geometric. Drives must be either identical or have controller cards that support geometry translation.

5.2.3 Installation Overview

In a typical LAN Server configuration StandbyServer 32 for LAN Server is used to protect a mission-critical server data on a domain server or an additional server. One high speed Vinca board is installed in the primary server with the second Vinca board installed in a second computer which then becomes the standby machine. LAN Server's fault tolerance or disk mirroring uses this link to transmit all mirrored data from the primary server to the standby machine. To LAN Server, it appears as if normal local disk duplexing is occurring when in fact data is being transmitted to the standby machine. This transfer happens in real time. There is no latency or data buffering occurring on the high speed link or in the standby machine. Data is written to both drives in a mirrored pair just as they would be under local disk duplexing. There are several requirements which must be considered when installing and configuring both OS/2 and LAN Server for StandbyServer 32 for LAN Server. These include:

- The first drive in both the primary server and the standby machine must be configured with the same number and type of partitions. This insures a consistency in drive letter assignment before and after a server switch over. Please see attached description *Hard disk configuration requirements*
- Both OS/2 and LAN Server must be installed on the first drive in each system. Since the first drive is not mirrored, different hardware may be used and the necessary system configuration information is kept intact. On the primary server the fault tolerance option must be selected for installation. On both the primary server and the standby machine, the replicator must be selected for installation if the server is a domain controller. LAN Server may be configured for automatic loading on the primary server, but must not be configured for automatic load on the standby machine.
- Both servers are configured with the same server name. This does not pose a conflict since LAN Server is only running on one of the machines at any time.

This soft copy for use by IBM employees only.

- If the server being protected is a domain controller, the replicator must be configured to replicate network configuration information located on the first drive. This would be information such as user accounts which may change on a periodic basis.
- Vinca device drivers must be installed on each machine using the OS/2 program ddinstall. Both systems must then be re-booted to load the Vinca drivers.
- Drive mirroring is defined between the two systems starting with the second drive in each system. The Vinca software hides the first drive in the standby machine from the primary server. Please see attached Figure A and description. In establishing the mirrored pair of drive(s), the primary mirrored partition(s) must be located on the standby machine. This allows the switch over to occur without reconfiguring partitions.
- NetFinity must be installed on both machines and configured for autoloading: NetFinity services on the primary server and NetFinity services and manager on the standby machine. A NetFinity command file is included with the Vinca software which must be loaded also.

LAN Server is only active on the primary machine. While in the standby mode of operation, the standby machine is running OS/2, the OS/2 requester (if replication is being used), and StandbyServer 32 for LAN Server. In addition to OS/2 and LAN Server, the primary server is running Vinca device drivers as well. When all these software components are loaded and running, the primary server accesses and mirrors data to hard drives physically located in the standby machine.

When in standby mode, the standby machine is not accessed directly by users, they simply interact with the primary server as usual. Client workstations need no additional software or configuration.

Figure 64 on page 184 summarizes the concepts used by the StandbyServer 32 for LAN Server solution. Basically simply the path writing data to the primary partition is intercepted by the StandbyServer 32 for LAN Server, redirected via the high speed link to the standby machine, and there written to the primary partition. The OS/2 LAN Server itself only knows about the primary and the secondary partition and no changes are required in either its setup nor the definitions using the FTSETUP utility.




Figure 64. StandbyServer 32 for LAN Server Concept

When the primary server fails from either a hardware or software related failure the following steps are executed by the standby machine:

1. A failure is detected across the Vinca high-speed link.
2. The standby machine verifies the primary server failure by attempting to contact the primary server across the network link. When the standby machine is unable to contact the primary server across the network link, the primary server is reported as having failed.
3. The standby machine then reports to a system administrator workstation the failure status of the primary server (It could also broadcast to users that the primary server has failed and a switch over is in progress).

This soft copy for use by IBM employees only.

4. The standby machine then unloads the StandbyServer 32 for LAN Server bringing the standby machine back to OS/2.
5. LAN Server is then loaded automatically on the standby machine and the standby machine becomes the main server.

All data that was mirrored by the original server is now available for users to access on the standby machine after a brief pause. They may also see a message indicating a switch over from primary server to standby machine is in progress. After the switch over is complete users do not need to reattach to the standby machine, they simply continue working.

The system administrator can now diagnose and repair the problem on the original server without critical data being affected. As soon as the problem with the original server is found and corrected, the system administrator can restore the original server as the primary server and the standby machine as the standby server. This restoration process requires bring the LAN down to switch back server functions and verify the mirrored data. During the mirror verification any data that has changed on the standby machine while it was the main server is mirrored back to the primary server. Data is then current and accurate on both machines.

5.2.4 Hard Disk Configuration Requirements

In the following we will describe a minimum basic configuration for the StandbyServer 32 by Vinca solution.

5.2.4.1 StandbyServer 32 for LAN Server Minimum Basic Configuration

The following section describes the minimum configuration requirements for each, the primary and the standby machine.

Figure 65 on page 186 shows the basic configuration. The digits shown inside the disk drive symbols will refer to the text labels used in the description following this figure.

This soft copy for use by IBM employees only.




Figure 65. StandbyServer 32 for LAN Server Minimum Configuration

1. The first drive of each system is:
 - Not mirrored
 - Only large enough for OS/2 and OS/2 LAN Server (200MB is satisfactory)
 - Identical in number and type of partitions
2. The second drive on each system is:
 - The first mirrored drive in the system

This soft copy for use by IBM employees only.

- Identical in reported drive geometries
 - a. Identical in brand, model and size
 - or.....
 - b. Able to report identical geometries (number of heads and number of sectors per track) because they use drive controllers that support geometry translation.

Additional mirrored drives may be used provided they meet the identical geometry requirement.

RAID-5 subsystems may be mirrored if they meet the requirement of identical geometries, have OS/2 device drivers, and can be configured to report two separate drives to OS/2.

5.2.4.2 System Configuration in Normal Operating State

This paragraph explains, why the first drives in each system are not mirrored and only hold OS/2 and the OS/2 LAN Server as shown in Figure 66 on page 189.

- OS/2 LAN Server and OS/2 use system configuration files similar to DOS's CONFIG.SYS file. This system configuration files contain information specific to that computer the OS/2 is running on. These system configuration files are located on the drive (and partition) the computer boots from, which is typically the first drive in the system. By not mirroring this drive, different hardware can be used for the primary server and the standby machine because the configuration files on each system can remain different.
- OS/2 uses a static drive letter scheme. Partitions receive assignment letters at boot time and not again until the next boot. Many applications store information by assigning that data to a specific drive letter. Those drive letters must remain the same for the application to continue to find and access its data. Vinca has developed a way to keep drive letter assignments identical regardless of whether the primary server is running or the standby machine has taken over and becomes the primary server. This process is as follows:
 - As described in the hard disk requirements, the partitions on the first drive in each system must be identical in number and type. This allows consistency for OS/2 and all the boot information. If necessary, you need to define *dummy partitions* on the standby machines first disk drive to ensure that the same drive IDs will be assigned when the standby machine takes over in case of a primary server failure.

This soft copy for use by IBM employees only.

- The standby machine *must* be booted first and the Vinca software *must* be running before the primary server is booted. While in standby mode, the standby machine appears to the primary server as an additional disk subsystem. OS/2 only scans for devices when it is booted. This means the standby machine must be booted and already running by the time OS/2 starts on the primary server or the drives in the standby machine will not be recognized.
- During normal operation the first drive in the standby machine is hidden to OS/2 on the primary server by the Vinca software. Drive lettering starts with the first drive in the primary server, skips the hidden drive in the standby machine and proceeds to the first mirrored set of drives which receive the same drive letter. Figure Figure 66 on page 189 illustrates the drive assignments in both the primary server and in the standby machine.

This soft copy for use by IBM employees only.




Figure 66. Drive Assignments in Both the Primary Server and in the Standby Server

5.2.4.3 Primary Server Failure

When the primary server fails, and switch over to the standby machine occurs (please refer to Figure 67 on page 190), the hidden drive is already known by the OS/2 in the standby machine. The standby machine having the same drive assignments as they were used in the primary server and thus becomes now visible for the OS/2 LAN Server started at the standby machine. This occurs because these drives are configured with identical partition configurations. The mirrored drive(s) then keep the same drive letter as they originally had, thus allowing applications to continue and function normally.




Figure 67. Switch Over to the Standby Machine

5.2.4.4 Mirrored Partition Priority Designations

LAN Server disk mirroring assigns different partition types to each partition in a mirrored pair. This does not affect performance or information access during normal operation. However, once one of the disks in the mirrored pair fails, OS/2 treats the two partitions differently next time it boots. If the secondary partition is the partition that failed, the primary partition can be accessed directly without any reconfiguration. If the primary partition fails, the secondary partition requires some reconfiguration, before it can be accessed. For this reason, Vinca requires, that the primary mirrored partition(s) be located on mirrored drive(s) in the standby machine. When a

This soft copy for use by IBM employees only.

failure on the primary server occurs, the standby machine can initialize as the primary server without any user intervention.

Figure Figure 68 shows the partition assignments in both the primary server and the standby machine.




Figure 68. Partition Assignment in Both Primary Server and Standby Machine

This soft copy for use by IBM employees only.

5.2.4.5 Establishing a Mirrored Set of Disks

To make OS/2 LAN Server's fault tolerance easy to use and to administer, a number of processes are automated including responsibility for selecting a disk area to mirror a partition to. When a partition or disk is selected to be mirrored, OS/2 LAN Server scans for the first available area on a separate disk with a separate controller. The OS/2 LAN Server then reserves that area and mirrors the original partition to that area. In a multiple disk configuration be sure that the correct disks are mirrored together. The best way to do this is to begin to define the mirrored partitions from the largest down to the smallest.

5.3 Dual Servers for High Availability

In this example of server high availability we have two server systems that will share a disk subsystem that contains all the production files necessary. The control of the external disk subsystem is performed by a SCSI switching box which will be described later. In comparison to other fault tolerant system solutions on the market, this solution does not provide a 100% transparent switch to the backup server.

There will be a brief interruption in service during the backup process where users may have to re-log on or re-boot to establish their connection again. In a NetWare environment, the SFT III product for example is 100% fault tolerant, but much more expensive than the Dual Server for High Availability solution. The other issue, of course is that the SFT III solution only works in a NetWare environment. The VINCA fault tolerant solution offers also a 100% fault tolerant scenario for both, OS/2 LAN Server and NetWare, but again is much more expensive.

5.3.1 Objectives

The challenge in designing the Dual Server scenario had several criteria:

- Offer a degree of high availability.
 - Service will be maintained although a re-boot is acceptable.
- Allow more flexibility as to the backup configuration.
 - Allow the backup server to be used as a workstation or additional server.
- Support OS/2 and NetWare environments.
- Unattended switching to backup.
 - Requires for remote branch locations or area where no administrator was available.
- Minimal cost.

This soft copy for use by IBM employees only.

One of the benefits in the design is that the backup server can be an active device on the network while it is waiting for the primary server system to fail. In the first example documented, the backup server is actually participating on the LAN as an additional server as well as acting as the monitoring station on the network. This is possible due to the alternate boot environment the backup server has when it boots from its internal hard disk.

A more realistic scenario to this solution would have a different device other than the backup server acting as the monitoring station. In most cases, it is recommended that the workstation, (normally the LAN Administrator), act as the monitoring station for the entire network, not just the servers. The advantage to this is that the tasks of monitoring or performing LAN management are not lost when the server systems are in backup mode.

Although there are several ways of performing this monitoring role, by far the simplest is by using the NetFinity product by IBM. For those environments that are implementing this scenario as an enterprise-wide solution, there may be benefits in using other products like LMU/2, which is a member of the SystemView family of products. NetFinity is positioned as a complementary product to the NetView suite, or can be used as a standalone network management tool.

In the examples documented, we have assumed that the primary server has failed for one of several different reasons:

- Power loss to the primary server
- Internal hardware failure
- Loss of access to the LAN (that is, LAN adapter card)

In each case, the monitoring application has sensed a failed heartbeat from the primary server. Once detected, the backup server will perform different actions based on different environments.

There have been several other scenarios since the first venture of this project, but as you can easily tell from the first example, the basic design or concept of how the solution works remains the same. Each customer environment will have different conditions, different responses, and different reactions based on their unique situations. The limits are purely your own creativity on what the client wants to do, based on what has failed. The various scenarios documented are from the most basic, with a minimal amount of hardware, to a more complex situation which offers a slightly higher level of fault tolerance (that is, RAID disk protection). There will be a basic system prerequisites, growing from there to a more complex scenarios which offer greater flexibility.

This soft copy for use by IBM employees only.

5.3.2 Limitations

As far as limitations, the only two discovered are those which deal with cable length once you get into the SCSI-II fast and wide devices, and those that deal with dissimilar servers. The cable distances when these devices are employed using SCSI-II are extremely minimal. If greater distances between the two servers are required, there are SCSI Extenders available from the same company that manufactures the SCSI switch.

The second limitation has to deal with booting two dissimilar servers in an OS/2 environment. Because the OS/2 install process loads various device drivers and BIOS files that are unique to each system, there may be a necessity to have both the primary and backup server be of the same machine type (a model 77 for example, cannot backup a model 95 server because the BIOS files are different).

5.3.3 System Prerequisites

The solution has several key components as a minimum configuration. The following is a list of prerequisites which will provide fault tolerance at the processor level. Other scenarios which involve greater degrees of fault tolerance will be documented in follow-up examples.

- Primary server with SCSI adapter
 - SCSI adapter must have external connection capability
 - No prerequisite as to SCSI level (that is, can be SCSI-1 or SCSI-2)
 - If external RAID required, then RAID adapter is necessary
- Backup server (requires same type of external SCSI support)
- Internal drives
 - In our testing, we have found that in order to keep the proper drive letters in sync, it is necessary to use the same number of hard drives in both servers. If you do not, the location of files and/or applications will change and will cause problems under the OS/2 desktop icons, which launch the programs.
- NetFinity
 - LAN management software that provides both managing, as well as managed portions of this solution (LMU may be used if your environment is already implementing the NETVIEW suite of products).
 - Client managed code installed on primary server, while manager code installed on backup server or monitoring station.
 - ServerGuard
- ServerGuard adapter
 - Optional, but recommended for server maintenance function.
 - This adapter comes packaged with NetFinity software.

This soft copy for use by IBM employees only.

- External Disk Subsystem
 - All data is stored on an external hard disk (or tower).
- SCSI switch
 - Electronic switch device which evokes the physical connection of the disk subsystem from one server to the other, see 5.3.6, “SCSI Switch” on page 196.
 - Various types and sizes of switches available as shown in Figure 69 on page 197. Make sure you order the proper SCSI model with respect to SCSI-1 versus SCSI-2, as well as proper cable connections.

5.3.4 Design Specification

The design point of this alternative was to allow a backup server to assume the production role, when the primary server had failed. In a normal environment, this would be accomplished by simply moving the production data to the backup server. This usually involved physically moving disks from one server to another, or moving cables around to the backup server. In each case, there was human intervention required which posed a delay in getting to backup mode.

The automation of this task is accomplished through a series of events which are documented in 5.3.5, “Setup Procedure.” The key to the solution is the physical SCSI switching device which will perform this unattended switching of the disk subsystem from one server to another.

5.3.5 Setup Procedure

The primary and backup server are both configured to boot from the external DASD as the first boot location (after diskette drive). This is set up using the reference disk under the selectable boot sequence procedure, specifying the SCSI ID of the external drive (that is, SCSI ID 6). If we are using a RAID configuration, note that these would be referenced as hard drive 0, then 1, then 2 instead of using SCSI IDs.

The next boot location from selectable boot sequence would be the internal hard drive (SCSI ID 5, or hard drive 0,1 etc.).

Both servers are attached to the external DASD via the SCSI switch described in 5.3.6, “SCSI Switch” on page 196. The production server would be on port A, the backup server on port B. In normal production mode, the SCSI switch would have the DASD tower active on port A.

The SCSI switch has a 9-pin serial connection which attaches to a COM port on the backup server or the monitoring workstation (this is how you can send

This soft copy for use by IBM employees only.

commands to the switch via a batch file). The internal hard drive can be seen as drive D.

The backup server will attempt to boot from the external disk(s), but because the SCSI switch is active in port A, it will not see the drive(s).

You will then receive a I9990301 BOOT DEVICE NOT FOUND. The system is trying to boot from the external drive but discovers that it is not there so it proceeds to boot from the next designated location which is the internal drive. No intervention is required as this message will be cleared during the boot sequence.

Both servers are now booting with their respective hard drives.

- The primary server off the external disk, which has the production operating environment, LAN software and applications and the backup server.
- The backup server from its own internal hard drive with whatever you choose, a workstation, an additional server, and/or monitoring station.

As a monitoring station, the backup server is running NetFinity which will keep tabs on the heartbeat of NetFinity on the primary server.

When the monitoring server detects a failed heartbeat from the primary server, you may then choose to execute your backup scenario. To do this, a series of commands will carry out several steps that you code through an ASCII file.

In this scenario the batch file has three statements:

- Issue a SET MODE command to set up the COM port that the SCSI switch is attached to (via the serial cable that comes with the switch)
- Issue an ASCII command to the SCSI switch to activate the external DASD from port A to port B
- Re-boot the backup server (which will now boot from the external drives)

5.3.6 SCSI Switch

The following chapter will provide details on the SCSI switch used in the documented "High Availability" examples. The SCSI switch is a high performance electronic crosspoint switch for use with the SCSI bus. The switch is manufactured by Applied Concept and we used the model ACI 2014 shown in Figure 69 on page 197.

This soft copy for use by IBM employees only.




Figure 69. APPLIED Concepts (4x2 Electronic Crosspoint Switch)

This product enables six independent SCSI buses to be selected and electrically connected in any combination via internal electronic switching circuits, thus eliminating the need to swap and reconfigure SCSI cables and bus terminators each time a new system configuration is required.

The SCSI switch front panel can be used to select the switch configuration. Port A and B are the two primary SCSI ports. Ports 1,2,3 and 4 are secondary ports. Depressing a position on the front panel will connect and disconnect each port as indicated.

In addition to selected ports, the front panel indicators provide SCSI "busy" activity.

The switch can be operated and controlled via the rear panel RS-232 interface using supplied Windows, Macintosh or UNIX software. Standard modem software and stand-alone terminals can also be used to control the switch by simple ASCII commands. An additional hand-held remote control unit is also available. For the purpose of the IBM Dual Server Solution, the switch is connected directly to COM2 of the backup server. This will allow the ASCII codes to be issued directly to the switch from a backup file.

This soft copy for use by IBM employees only.

The rear panel of the SCSI switch, as shown in Figure 77 on page 205, identifies each of the six independent SCSI ports. Each port has two daisy-chained connectors for installation of SCSI devices or external terminators, and can be installed at any point on the SCSI bus. The SCSI switch does not require a device ID and is completely "transparent" to all computers and peripherals on the SCSI bus. No additional software is needed for installation and operation. A universal power supply auto-switches AC input voltages between 105 to 240VAC.

The SCSI switch incorporates proprietary high performance switching logic to isolate and regenerate single-ended SCSI bus signals. Each port of the switch supports single ended SCSI cable length to 6 meters (19.7 feet). Therefore, the maximum cable distance between computers and devices connected single ended ports on the SCSI switch can be as far as 12 meters (82 feet). See Figure 70 for an illustration of the Applied Concepts 4x2 electronic crosspoint switch.




Figure 70. APPLIED Concepts (4x2 Electronic Crosspoint Switch)

The differential version of the SCSI switch is available with 8-bit narrow or 16-bit wide capability. This switch will support SCSI devices over distances of 25 meters (82 feet). The 16-bit connectors conform to the industry standard 68-pin design. This allows standard IBM cables to be used:

This soft copy for use by IBM employees only.

- From the server to the SCSI switch (IBM 16-bit Fast/Wide card to option cable, Partnumber 06H3231).
- From the SCSI switch to the SCSI Fast/Wide devices (IBM 16-bit Fast/Wide option to option cable, Partnumber 32G3914)
- From the SCSI switch to the SCSI-2 Fast or SCSI-1 devices (IBM 16-bit Fast/Wide option to 8-bit SCSI-2 Fast or SCSI-1 Option Cable, Partnumber 32G3916)

The user has the option of selectively enabling each internal terminator on the single ended version of the SCSI switch. A rear panel switch controls internal "active" terminators for each individual port. By selecting terminator switch A,B,1,2,3, or 4, each internal terminators activated or deactivated.




Figure 71. APPLIED Concepts (Backing up Multiple Servers)

5.3.7 Configuration Examples

The following section provides configuration examples.

This soft copy for use by IBM employees only.

5.3.7.1 Minimum Configuration

Figure 72 shows a minimum configuration working successfully under OS/2 or NetWare.

- Both, primary and backup server have a SCSI adapter.
- Both have an internal hard drive (for an alternative boot environment).
- Primary accesses external disk via port A on switch.
- Backup server accesses external disk via port B.
- Switch is connected to serial port of backup server.
- ServerGuard is optional in this minimal configuration.
- Internal drive of backup server configured as an OS/2 workstation in the domain running NetFinity manager code.
- Monitor software set up to check the heartbeat of the primary server.
When the heartbeat is overdue, batch file is executed to switch to backup server. Also have a message sent out to the pager.




Figure 72. High Availability Server (Minimum Configuration)

This soft copy for use by IBM employees only.

5.3.7.2 External RAID Tower

The configuration shown in Figure 73 is an identical configuration to Figure 72 on page 200 with the following three changes:

- The SCSI adapter is one of the RAID adapters to allow RAID support of the external disk drives.
- This solution requires the internal hard drive on the backup server to be on a second SCSI adapter, but does not have to be a RAID adapter.
- The external DASD is enclosed in a hot plug 3516 tower.

This example provides a higher level of fault tolerance with the drives now protected via RAID-5 or RAID-1.




Figure 73. High Availability Server (External RAID Tower)

5.3.7.3 External Mirrored RAID-1 Towers

This configuration hot plug towers mirrored with RAID-1 will provide fault tolerance in the event of losing an entire tower.

Limitations:

IBM's RAID-1 definition when the environment has more than three drives, will use the enhanced RAID-1 scheme. This will not allow you to mirror tower

This soft copy for use by IBM employees only.

1 with tower # 2. What you will have to do is mirror each drive in tower # 1 with each drive in tower # 2, in effect giving you the same required level of fault tolerance. This does mean however, that you will be restricted to partition sizes that are the same size as the hard drives, which will allow partitions sizes of 1GB, 2GB and 4GB respectively.




Figure 74. High Availability Server (External Mirrored RAID-1 Towers)

5.3.7.4 External Duplexed RAID-1 Towers

This configuration is using hot plug towers duplexed with RAID-1 and will provide fault tolerance in the event of losing a channel on the RAID adapter.

Limitations:

IBM's RAID-1 definition when the environment has more than three drives, will use the enhanced RAID-1 scheme. This will not allow you to mirror tower # 1 with tower # 2. What you will have to do is mirror each drive in tower # 1 with each drive in tower # 2, in effect giving you the same required level of fault tolerance. This does mean however, that you will be restricted to partition sizes that are the same size as the hard drives, which will allow partitions sizes of 1GB, 2GB and 4GB respectively.

This soft copy for use by IBM employees only.




Figure 75. High Availability Server (External Duplexed RAID-1 Towers)

5.3.7.5 Sample NetFinity Environment

Figure 76 on page 204 illustrates how NetFinity is installed throughout the network. Although this document details that the backup server does the monitoring and switching, it can also be accomplished by a workstation on the network as well.

The 9-pin serial cable in this case would be attached to the monitoring workstation instead of the backup server.

In reality, this scenario is a more realistic approach, as the LAN administrators workstation would normally be monitoring the entire network as well as for failed servers.

This soft copy for use by IBM employees only.




Figure 76. High Availability Server (NetFinity Implementation)

Figure 77 on page 205 shows how two server systems and an external disk subsystem (DASD Tower) are connected to the SCSI switch.

This soft copy for use by IBM employees only.




Figure 77. Sample Cabling for two CPU's

5.3.7.6 Sample Programs

Following is a sample of the actual procedure used in configuration shown in Figure 72 on page 200 and Figure 73 on page 201. The instructions assume that you have a working knowledge of NetFinity.

Set Up NetFinity

1. Open the Remote System Manager panel on the monitoring system.
2. Open the group which contains the server you wish to monitor.
3. Click with the right mouse button on the primary server.
4. Click on System Notification, and you should see options to turn on Notify when the system is online and offline.
5. Click to turn on the Offline selection, select a severity code, then use the spin buttons to specify a timeout value. This parameter is the time NetFinity will use to assume a failed response from the primary server.
6. Return back to the services panel for the backup machine, and click on the Alert Manager.
7. From the options on the bottom of the screen, select Actions.
8. Now that you are on the Alert Actions panel, you must select what action you wish to perform based on the failure.
9. Click on New at the bottom of the screen.

This soft copy for use by IBM employees only.

10. When the Action Editor menu appears set the following parameters:
 - a. Alert Type is Any.
 - b. Severity - select the same severity code you set with the Notification panel in the previous step 3 (these severity range is from 1-7).
 - c. Application ID is Any.
 - d. Application alert type is Any.
 - e. Sender ID is Any.
 - f. Action definition - scroll through and select Execute Command < P 1 > .
 - g. In the box provided below, supply the name of the batch file you wish to execute upon the alert with the appropriate path. An example batch file is shown in 5.3.7.4, "External Duplexed RAID-1 Towers" on page 202.
 - h. Save the action parameter and you are ready to set up the hardware and test.

Sample Batch File:

This sample batch file can be modified to suit particular port setups. The name of the batch file is called SWITCHIT.CMD and it will execute two other files, one called B4A1 the other called REBOOT.CMD.

- Sample Batch File SWITCHIT.CMD
 - MODE COM2: 9600,N,8,1 .* sets up the COM2 port
- Sample File Called B4A1
 - Following is the B4A1 batch file. It is an ASCII control string that must be sent to switch via COM2 (as per the file SWITCHIT.CMD).
 - **NOTE:**. case sensitive! Use all upper case letters!
 - //B4K .* turns on port B to access the tower via port 4
 - //A1K .* terminates the SCSI cable for the server in port A
- Sample File called REBOOT.CMD
 - "SETBOOT /IBD:C" .* issues a reboot to the local machine

This soft copy for use by IBM employees only.

Appendix A. Part Number Equivalences for Europe

<i>Table 30 (Page 1 of 3). US, EMEA Part Number Equivalences</i>		
US, LA	EMEA	Product Description
SCSI Contollers		
32G3100	SSPC01Y	SCSI-2 Fast PCI-Bus Adapter
94G3771	SPSW01Y	SCSI-2 Fast/Wide PCI Adapter
94G2764	SPSR01Y	SCSI-2 Fast/Wide RAID PCI-Bus Adapter
70G8498	6451280	SCSI-2 Fast/Wide Adapter/A
70G9263	PRADADX	SCSI-2 Fast/Wide Streaming RAID Adapter/A
Deskstar Hard Disk Drives		
06H8558	N/A	540MB SCSI-2 Fast Disk
94G3187	OSC1GBM	1.08GB SCSI-2 Fast Disk
Ultrastar Hard Disk Drives		
94G3052	OSW1GBM	1.12GB SCSI-2 Fast/Wide Disk
94G2649	S1GHS1Y	1.12GB SCSI-2 Fast/Wide Hot Swap Disk
94G3054	OSC2GBM	2.25GB SCSI-2 Fast Disk
94G3055	OSW2GBM	2.25GB SCSI-2 Fast/Wide Disk
94G2650	S2GHS1Y	2.25GB SCSI-2 Fast/Wide Hot Swap Disk
94G3196	N/A	4.51GB SCSI-2 Fast Disk
94G3057	OSW4GBM	4.51GB SCSI-2 Fast/Wide Disk
94G2651	S4GHS1Y	4.51GB SCSI-2 Fast/Wide Hot Swap Disk
94G3787	N/A	5.31GB SCSI-2 Fast Disk
94G3059	N/A	5.31GB SCSI-2 Fast/Wide Disk
Additional Hard Disk Drives		
70G8495	32G4336	2GB Fast SCSI-2 Disk Drive 1
70G7164	70G7164	1GB Fast SCSI-2 Disk Drive for 3516
70G9743	S1GHSD1	1GB Fast/Wide SCSI-2 Disk Drive for 3516
External SCSI Enclosures		
35110V0	32G2964	3511 SCSI Enclosure
3510AV0	6451375	3510 SCSI Enclosure

This soft copy for use by IBM employees only.

<i>Table 30 (Page 2 of 3). US, EMEA Part Number Equivalences</i>		
US, LA	EMEA	Product Description
35167V0	SDHSEE1	3516 DASD Hot Swap Expansion Enclosure
SCSI-2 Fast Backup Tape Drives		
74G8632	8191386	Internal 3.5 4mm DAT 4/10GB SCSI Tape Drive
74G8631	8191441	Internal 5.25 4mm DAT 4/10GB SCSI Tape Drive
94G2725	STAL01Y	24/48 Internal SCSI DAT Tape AutoLoader
74G8630	46G4004	3450 External 1.2GB 1/4" Tape Drive
3440001	46G4016	3440 External 2.0GB 4mm Tape Drive
74G8629	46G2643	3445 External 5/10GB 8mm Tape Drive
55F9464	46G2365	PS/2 2.0GB Internal 4mm SCSI Tape Drive
SCSI CD-ROM Drives		
06H7648	OCS4XIM	Internal SCSI-2 Fast 4X CD-ROM Drive
3509707	OCS4XEM	External SCSI-2 Fast 4X CD-ROM Drive
94G2686	OCDIS6M	Internal SCSI-2 Fast 2X CD-ROM Drive
3509705	OCDES6M	External SCSI-2 Fast 2X CD-ROM Drive
Optical Drives		
05H2794	ORI230M	Imagestar 230MB Internal Optical Disk Drive
3432715	ORE230M	Imagestar 230MB External Optical Disk Drive
3431001	46G2685	PS/2 External Rewritable Optical Disk Drive
External SCSI Cables		
70G9874	PCSC01Y	PCI SCSI-2 Fast Card to SCSI-2 Fast (50 pin) Option Cable
70G9858	SS2C01Y	PCI SCSI-2 F/W Card to SCSI-2 Fast (50 pin) Option Cable
70G9857	SS2C02Y	PCI SCSI-2 F/W Card to SCSI-2 F/W (68 pin) Option Cable
32G3915	32G3915	MCA SCSI-2 F/W Card to SCSI-2 Fast (50 pin) Option Cable
06H3231	PRADC1M	MCA SCSI-2 F/W Card to SCSI-2 F/W (68 pin) Option Cable
32G3099	32G3099	MCA SCSI-2 F/W Card to SCSI-2 Fast Single Port Option Cable
N/A	SSSPC1Y	PCI SCSI-2 F/W Card to SCSI-2 Fast Single Port Option Cable
6451042	6451042	SCSI-2 Fast Option to SCSI-2 Fast Option (0.6m)
32G3914	32G3914	SCSI-2 F/W Option to SCSI-2 F/W Option
32G3916	32G3916	SCSI-2 F/W Option to SCSI-2 Fast Option
32G3925	32G3925	SCSI Converter (68-pin to 50-pin)

This soft copy for use by IBM employees only.

<i>Table 30 (Page 3 of 3). US, EMEA Part Number Equivalences</i>		
US, LA	EMEA	Product Description
32G3919	32G3919	8-bit Active Terminator
32G3918	32G3918	16-bit Active Terminator
SCSI Kits		
94G2763	SHSB01Y	PC Server 300/320 Hot Swap Bay Option
70G9264	S86C01Y	PC Server 300 Internal Cable Kit
70G9740	SHSBP1Y	PC Server Hot Swap Backplane
70G9739	SAPS501	PC Server Second Power Supply (220 watts)
70G9851	SHST02Y	SCSI-2 Fast Hot Swap Drive Tray - PC Server
70G9741	SHST03Y	SCSI-2 Fast/Wide Hot Swap Drive Tray - PC Server
94G2763	SHSB01Y	PC Server 300/320 Hot Swap Bay Option
94G2648	OHKITCY	Model 3510 Enclosure Drive Install Kit
6451145	6451145	Enhanced CD-ROM Drive Installation Kit A
6451146	6451146	Enhanced CD-ROM Drive Installation Kit B (used for 3510)
32G3920	32G3920	SCSI-2 F/W Upgrade for 3510-0V0
06H3153	PRADWCK	SCSI-2 F/W Upgrade for 3511-0V0
Note:		
1 These disk drives do not belong to the Deskstar or the Ultrastar families.		

This soft copy for use by IBM employees only.

This soft copy for use by IBM employees only.

Appendix B. Description - Test Environment

This appendix provides details on the test environment used for the tests performed for Chapter 4, "Disk Subsystem Sizing and Performance" on page 127.

B.1 Test Environment Used for the IBM PC Server 320

- Benchmark program used
 - Ziff-Davis' NetBench version 3.01
 - Test suite
 - Disk_Mix for 4, 12, 20, 28, 36, 44, 52, and 60 clients (DM_4.TST DM_60.TST)
- Server model = IBM PC Server 320
- Processor = 1 x Pentium 60/90 MHz
- RAM = 32MB
- L2 cache = 256KB
- I/O Bus = EISA/PCI
- Disk adapter = IBM PCI RAID Adapter
- Disk device driver = IBMRAID.DSK, IBM RAID Driver For NetWare 3.1x, Version 3.07, February 9, 1995
- Disk drives = IBM Ultrastar SCSI-2 Fast/Wide 1.12GB
- Number of drives = 2, 3, 4, 5, 6
- RAID-level = HW RAID level 0, 1, and 5
- LAN Topology = IBM Token-Ring
- LAN speed = 16Mbps
- LAN adapter = 4 x Madge Smart 16/4 Ringnode
- LAN adapter device driver = Madge Fastmac Plus Token-Ring MLID, Version 4.55, October 26, 1994
- NOS = NetWare 3.12
 - NOS parameters:
 - set maximum packet receive buffer size = NetWare default

This soft copy for use by IBM employees only.

- set enable disk read after write verify = OFF
- set immediate purge of deleted files = ON
- set maximum packet receive buffers = 1000
- set minimum packet receive buffers = 500
- set maximum concurrent disk cache writes= NetWare default
- set enable disk read after write verify = off
- set NCP packet signature option = 0
- set cache buffer size = 16384
- set new service process wait time = NetWare default
- set maximum service processes = NetWare default
- NetWare Volume block size =16KB
- Network type = 16 Mbps Token-Ring, 4 segments, 15 stations per segment
- Client type = IBM P/2 model 70, 1 x Intel 80386, 20MHz
- Client RAM = 2MB
- Client LAN adapter = IBM 16/4 Token-Ring Adapter/A
- Client DOS = IBM PC DOS 6.3
- Client network software = NetWare 3.12 DOS ODI client
- Control terminal operating system = IBM PC DOS 6.3, Windows for Workgroups 3.11

All products used in the tests are shipping versions available to the general public.

The test was performed without independent verification by Ziff-Davis.

B.2 Test Environment Used for the IBM PC Server 500

- Benchmark program used
 - Ziff-Davis' NetBench version 3.01
 - Test suite
 - Disk_Mix for 4, 12, 20, 28, 36, 44, 52, and 60 clients (DM_4.TST DM_60.TST)
- Server model = IBM PC Server 500

This soft copy for use by IBM employees only.

- Processor = 1 x Pentium 60/90 MHz
- RAM = 32MB and 64MB
- L2 cache = 256KB
- I/O Bus = IBM Micro Channel
- Disk adapter = IBM Streaming RAID Adapter/A
- Disk device driver = IBMRAID.DSK, IBM RAID Driver for NetWare 3.1x, Version 3.07, February 9, 1995
- Disk adapter = IBM SCSI-2 Fast/Wide Adapter/A
- Disk device driver = PS2SCSIA.DSK Device Driver for IBM PS/2 SCSI Adapters with ASPI support Version 2.01 December 12, 1994, Copyright 1994 IBM, inc. All rights reserved.
- Disk drives = IBM Ultrastar SCSI-2 Fast/Wide 1.12GB
- Number of drives = 2, 3, 4, 5, 6
- RAID-level = HW RAID level 0, 1, and 5
- LAN Topology = IBM Token-Ring
- LAN speed = 16Mbps
- LAN adapter = 4 x IBM Lanstreamer MC32
- LAN adapter device driver = IBMMPCO.LAN, IBM Streamer Family ODI NetWare LAN Driver Version 1.38, June 27, 1995
- NOS = NetWare 3.12
 - NOS parameters:
 - set maximum packet receive buffer size = NetWare default
 - set enable disk read after write verify = OFF
 - set immediate purge of deleted files = ON
 - set maximum packet receive buffers = 1000
 - set minimum packet receive buffers = 500
 - set maximum concurrent disk cache writes= NetWare default
 - set enable disk read after write verify = off
 - set NCP packet signature option = 0
 - set cache buffer size = 16384
 - set new service process wait time = NetWare default

This soft copy for use by IBM employees only.

- set maximum service processes = NetWare default
- NetWare Volume block size =16KB
- Network type = 16 Mbps Token-Ring, 4 segments, 15 stations per segment
- Client type = IBM P/2 model 70, 1 x Intel 80386, 20MHz
- Client RAM = 2MB
- Client LAN adapter = IBM 16/4 Token-Ring Adapter/A
- Client DOS = IBM PC DOS 6.3
- Client network software = NetWare 3.12 DOS ODI client
- Control terminal operating system = IBM PC DOS 6.3, Windows for Workgroups 3.11

All products used in the tests are shipping versions available to the general public.

The test was performed without independent verification by Ziff-Davis.

This soft copy for use by IBM employees only.

Glossary

A

arbitration. A method with which multiple devices attached to a single bus can bid to get control of that bus.

asynchronous. A mode of data transfer across the SCSI Bus where each byte of data transferred must be acknowledged as received by the target before the next byte can be sent. Maximum transfer rate possible 2Mbps.

B

buffer fill rate. Determines how fast the device can fill the buffer on a controller. The buffer fill rate is measured in Mbps.

bus master. An intelligent device that when attached to the Micro Channel bus can bid for and gain control of the Micro Channel bus to perform its specific task.

C

CCS. The SCSI Common Command Set. A set of SCSI commands that is specified in the ANSI standard that all SCSI device must be able to use in order to be fully compatible with the ANSI standard.

CD-ROM. Compact Disk Read Only Media is a disc that you can only read data from. Data cannot be written to CD-ROM.

D

device level copying. When two devices attached to the SCSI Bus perform data transfers between each other across the SCSI Bus without using the attachment feature.

disconnect. When a device has received a

command and disconnects from the SCSI Bus to enable other device to use the SCSI Bus while it processes its command.

DMA. Direct Memory Access. A method used to transfer data directly from device to system memory without using the main system processor.

I

initiator. A device attached to the SCSI Bus that sends a command to another device on the SCSI Bus. The device that receives that command is a Target.

L

latency. The time interval between the instant at which an instruction control SCSI Bus. unit initiates a call for data and the instant at which the actual transfer of the data starts. Synonymous with waiting time.

logical unit. A device attached to a SCSI device and is NOT attached directly to the SCSI Bus.

LUN. Logical Unit Number. A number given to a device that is attached to a SCSI device and not directly to the SCSI Bus. The device is known as a Logical Unit.

M

mean time between failure (MTBF). Is the stated period in the life of a functional unit, the mean value SCSI of the length of time between consecutive failures under stated conditions.

This soft copy for use by IBM employees only.

P

PUN. Physical Unit Number. Another term used to describe a device attached directly to the SCSI Bus. Also known as a SCSI Device or SCSI ID.

R

reconnect. When a device that has finished processing a command arbitrates for the SCSI Bus in order to reconnect to it to perform its data transfer.

rotation speed. Rotation speed of the disk drive, measured in revolutions-per-minute SCSI Bus (rpm).

S

SCSI. Stands for Small Computer Systems Interface. Defines the interface between an attachment feature and intelligent devices.

SCSI attachment feature. The feature that attaches to the main system unit and the SCSI Bus. It is the controlling feature of the SCSI subsystem.

SCSI bus. A term used to describe the 50-Conductor cable that attaches intelligent devices to the SCSI attachment feature.

SCSI device. An intelligent device that is directly attached to the SCSI Bus. It conforms to

the ANSI Standard X3.131-1986 for attached SCSI devices.

SCSI ID. A number configured on a SCSI device so that it can be accessed on the SCSI Bus. Each SCSI device has a unique SCSI ID number. It is in the range 0 to 7.

seek. To selectively position the access mechanism of a direct access device.

seek time. Time required to selectively position the access mechanism of a direct access device.

synchronous. A mode of data transfer across the SCSI Bus where each byte of data transferred does not have to be acknowledged as received by the target device before the next byte can be sent. Maximum transfer rate possible 5 Mbps.

T

target. A device attached to the SCSI Bus that receives and processes commands sent from another device on the SCSI Bus. The device that sends the command is known as an Initiator.

terminator. A piece of hardware that must be attached to both ends of the 50-Conductor SCSI attachment cable (commonly known as the SCSI Bus).

This soft copy for use by IBM employees only.

List of Abbreviations

APA	All Points Addressable	LAN	Local Area Network
ANSI	American National Standards Institute.	MCA	Micro Channel Architecture
CPU	Central Processing Unit	MFM	Modified Frequency Modulation
DASD	Direct Access Storage Device	PCI	Peripheral Component Interconnect
ESDI	Enhanced Small Device Interface	PROFS	Professional Office System
Hz	Herz	RLL	Run Length Limited
IBM	International Business Machines Corporation	SCSI	Small Computer System Interface
ICA	Industry Standard Architecture	ST506	Seagate Technology ST506
ITSO	International Technical Support Organization		

This soft copy for use by IBM employees only.

This soft copy for use by IBM employees only.

Index

A

abbreviations 217
acronyms 217
Action Editor 206
Adaptec 2940 56
Adapter
 Device Drivers 49
 IBM PC Server SCSI-2 Fast/Wide PCI-Bus
 RAID Adapter 52
 Non-array 46
 PCI SCSI-2 Fast Adapter 83
 PCI SCSI-2 Fast/Wide RAID Adapter 90
 SCSI-2 Fast Adapters 56
 SCSI-2 Fast/Wide Adapter 90
 SCSI-2 Fast/Wide Streaming RAID
 Adapter 101
 Summary of SCSI-2 Adapters 56
Adapter/A
 IBM PC Server SCSI-2 Fast/Wide 45
 IBM SCSI-2 Fast/Wide Streaming-RAID
 Adapter/A 64
 IBM Streaming RAID 50
 SCSI-2 Fast/Wide Adapter/A 102
 SCSI-2 Fast/Wide Streaming RAID
 Adapter/A 101
 SCSI-2 Fast/Wide Streaming-RAID 51
 SCSI-2 Streaming Fast/Wide RAID
 Adapter/A 98
 Summary of SCSI-2 Adapters 56
Alert Actions 205
Alert Manager 205
application server 5, 6
Arbitration
 Definition 215
Array 55
Array Configurations 51
Array Models 96
ASPI, Advanced SCSI Programming
 Interface 56

Asynchronous Mode
 Definition 215
Audio Grade Media 79
Audio, CD-Digital 72
Automation 195

B

BBS, Bulletin Board Systems 4
Buffer Fill Rate
 Definition 215
Bus Master
 Definition 215
Busmastering 48

C

Cable Kit, PC Server 56
Cable Length, External 47
Cabling
 Card to Option 113
 Internal Bays Upgrades 108
 Option to Option 114
 PC Server 300 83
 PC Server 300/310 External 119
 PC Server 310 MCA or ISA Models 86
 PC Server 320 86
 PC Server 320 PCI/EISA 86
 PC Server 320 PCI/MCA 86
 PC Server 320/320 Array External 120
 PC Server 500 SCSI 102
 PC Server 500, 520, 720 Internal 94
 PC Server 500/720 Array 98
 PC Server 520 Array 101
 PC Server 520 SCSI 106
 PC Server 520/520 Array External 120
 PC Server 720 and 520 SCSI 107
 PC Server 720 Non-Array External 120
 SCSI External 111
 SCSI Internal 83

- CCS (Common Command Set)
 - Definition 215
- CD ROM
 - Definition 215
- CD-Digital Audio 72
- CD-ROM
 - 2X speed 16
 - 4X speed 13, 29
 - Drives 71
 - IBM External and Internal SCSI CD-ROM 2X Drives 72
 - IBM External and Internal SCSI CD-ROM 4X Drives 74
 - IBM External SCSI 2X CD-ROM Drive 71
 - IBM External SCSI 4X CD-ROM Drive 72
 - IBM Internal SCSI 2X CD-ROM Drive 71
 - IBM Internal SCSI 4X CD-ROM Drive 72
- client/server 5
- Coalesced 54
- communications server 4
- Compression
 - DDS-Data 77
- Configurations, Array 51
- Controllers 45
- Controllers, Disk Subsystem 45
- CPU 5

D

- DASD, External 195
- DAT, Digital Audio Tape 80
- data integrity 3
- Data Paths, Multiple 53
- Data Protection 55
- data security 3
- Data Streaming 45
- database server 5
- Daughter Card 52
- ddinstall 183
- DDS-Data Compression 77
- DDS, Digital Data Storage 80
- DDSIIII 80
- departmental LANs 10
- Depth/Stripe Unit 54

This soft copy for use by IBM employees only.

- Deskstar Family, Drives 62
- Device drivers 49
- Device-Level Copying
 - Definition 215
- Devices
 - SCSI 50
- Disaster Avoidance Plans 174
- Disconnect
 - Definition 215
- Disk Subsystem
 - External 64
 - Setup and Configuration 83
- Disk Sybsystem
- distributed applications 5
- DMA
 - Definition 215
- Drive Replacement, Hot Spare 55
- Drives
 - 2GB SCSI-2 Fast Disk Drive 84
 - CD-ROM 71
 - Hard Disk 59
 - IBM External and Internal SCSI CD-ROM 2X Drives 72
 - IBM External and Internal SCSI CD-ROM 4X Drives 74
 - The Deskstar Family 62
 - The Ultrastar Family 59
- Dual Server 192
- Dual Servers
 - Configuration Examples 199
 - Design Specifications 195
 - External Duplexed RAID-1 Towers 202
 - External Mirrored RAID-1 Towers 201
 - External RAID Tower 201
 - Setup Procedure 195
 - System Prerequisites 194

E

- ECC - Error Correcting Code 173
- ECC (Error Correcting Code) memory 8
- ECC memory 24
- ECP/EPP 16
- EISA 11, 181

This soft copy for use by IBM employees only.

EISA/Micro Channel 181
Enclosure
 3510 SCSI Storage Enclosure 69
 3511 External Enclosure for SCSI
 Devices 67
 3516 Hot Swap Expansion Enclosure 64
 External Duplexed RAID-1 Towers 202
 External Mirrored RAID-1 Towers 201
 External RAID Tower 201
ExpressPrint 24
External Cable Length 47
External DASD 195

F

Failure, Single Drive 55
Fast Single Port Device 114
Fault Tolerant 173, 174
FAX server 5
file server 3
FloThru cooling system 25, 29
FloThru, cooling system 17
FTSETUP 183

G

gateway 4
groupware applications 5
GUI 5

H

Hard Disk Drives 59
Heartbeat 193, 196, 200
High Sierra 72
Hot Plug 201
Hot Spare Drive Replacement 55
Hot Swap Bay Option 84
Hot Swap Bay Options 109
Hot Swap Models 90
Hybrid RAID 1 53

I

I/O bus 6

IBM 24/48GB Internal Tape Autoloader 80
IBM 3440 2.0GB 4mm Tape Drive Model 001 77
IBM 4/10GB 4mm DAT Tape Drive Option
 Kits 79
IDRC 76
IDRC Compression 76
Initiator
 Definition 215
Input/Output Operation, Overlapped 53
Interleave 54
Interleave Depth/Stripe Unit 54
Internet 4
IOEngine - Input/Output Engine 175
ISDN 5
ISO 9660 72

K

Key Features
 3510 SCSI Storage Enclosure 70
 3511 External Enclosure for SCSI
 Devices 67
 IBM 3445 5.0GB 8mm Tape Drive Model
 001 76
 IBM 3450 1.2GB 1/4-Inch Cartridge Tape
 Drive Model 001 75
 Internal SCSI 2X CD-ROM Drive 72
 Streaming RAID Adapter/A 52
Kodak** Photo CD 72

L

LAN subsystem 8
Latency
 Definition 215
LMU/2 193
Logical units
 Definition 215
LogicLock security 35
Lotus Notes 5
LUNs
 Definition 215

This soft copy for use by IBM employees only.

M

Macintosh 197
mainframe access 4
Mean Time Between Failure (MTBF)
 Definition 215
Media, Audio Grade 79
Media, Non Data Grade 77, 79
Micro Channel 181
Mirrored Partition 190
modem pool 4
Monitoring 193
MPC-2 74
MSEngine - Mirrored Server Engine 176
MSL - Mirrored Server Link 174, 175
multi-function server 3
multifunction server 6
multimedia server 3
Multiple Data Paths 53

N

NetFinity 14, 16, 179, 181, 183, 193, 194, 196,
 203
NLMs - Netware Loadable Modules 174
Non Data Grade Media 77, 79
Non-Array Models 102
Notification 206
NovaBack 78
NovaStor 78

O

online services access 4
Operation, Overlapped Input/Output 53
Option Kits, IBM 4/10GB 4mm DAT Tape
 Drive 78
Options 45
OS/2 181, 187
OS/2 LAN Server 181, 187, 192
OS/2 WARP 181
Overlapped Input/Output Operation 53
Overview 45

P

Parity 55
PC Server 300 56
 description 10
 specifications 12
PC Server 310 56
PC Server 310 (ISA or MCA)
 description 13
 specifications 14
PC Server 320 (EISA/PCI)
 description 16
 entry models specifications 19
 specifications 17
PC Server 320 (MCA/PCI)
 description 21
 specifications 21
PC Server 500
 description 23
 highlights 25
 specifications 26
PC Server 520
 description 28
 highlights 29
PC Server 520 (EISA/PCI)
 specifications 29
PC Server 520 (MCA/PCI)
 specifications 31
PC Server 720
 description 34
 highlights 36
 L3 Victim cache 36
 SMP 36
PC Server Cable Kit 56
PCI 11
PCI Adapter
 IBM PC Server SCSI-2 Fast/Wide 48
 IBM PC Server SCSI-2 Fast/Wide PCI-Bus
 RAID Adapter 52
 IBM SCSI-2 Fast PCI-Bus Adapter 56
 PCI SCSI-2 Fast Adapter 83
 PCI SCSI-2 Fast/Wide RAID Adapter 90
 PCI to SCSI Host 48
 SCSI-2 Fast/Wide PCI Adapter 107
 Summary of SCSI-2 Adapters 56

This soft copy for use by IBM employees only.

PCI SCSI-2 Fast Adapter 119
Pentium
 100MHz 29
 100MHz speed 36
 75MHz speed 13
 90MHz speed 16, 24
PFA (Predictive Failure Analysis) 24
print queue 4
print server 4
Protection, Data 55
PUNs
 Definition 216

Q

QIC, Quarter-Inch Cartridge 75
Queue Depth 54

R

RAID 100
RAID Adapter
 Cache 55
 Features 53
 IBM PC Server SCSI-2 Fast/Wide PCI-Bus
 RAID Adapter 52
 IBM SCSI-2 Fast/Wide Streaming-RAID
 Adapter/A 64
 PCI SCSI-2 Fast/Wide RAID Adapter 90
 SCSI-2 Fast/Wide Streaming RAID
 Adapter 101
 SCSI-2 Fast/Wide Streaming RAID
 Adapter/A 101
 SCSI-2 Streaming Fast/Wide RAID
 Adapter/A 98
RAID-5 187
Real Time 182
Reconnect
 Definition 216
Redundancy 174
remote LAN access 4
Remote System Manager 205
Reordered 54
RISC
 Sequencer 48

Rotation Speed
 Definition 216
RS-232 197

S

SCSI
 Adapter Connectors 111
 Definition 216
 IDs 46
 Internal Cabling 83
 PCI to SCSI Host Adapter 48
 Summary of SCSI Adapters 56
 Switch 196
 Ultrastar XP SCSI Drives 59
SCSI Attachment Feature
 Definition 216
SCSI bus
 Definition 216
SCSI device
 Definition 216
SCSI ID
 Definition 216
SCSI-1 101, 103
SCSI-2 101, 103
 IBM PC Server SCSI-2 Fast/Wide PCI-Bus
 RAID Adapter 52
 IBM SCSI-2 Fast PCI-Bus Adapter 56
 IBM SCSI-2 Fast Solutions 56
 IBM SCSI-2 Fast/Wide Streaming-RAID
 Adapter/A 64
 PCI SCSI-2 Fast Adapter 83
 PCI SCSI-2 Fast/Wide RAID Adapter 90
 SCSI-2 Fast Adapters 56
 SCSI-2 Fast/Wide Adapter 90
 SCSI-2 Fast/Wide Adapter/A 102
 SCSI-2 Fast/Wide PCI Adapter 107
 SCSI-2 Fast/Wide Streaming RAID
 Adapter 101
 SCSI-2 Fast/Wide Streaming RAID
 Adapter/A 101
 SCSI-2 Fast/Wide Streaming-RAID
 Adapter/A 51
 SCSI-2 Streaming Fast/Wide RAID
 Adapter/A 98

This soft copy for use by IBM employees only.

SCSI-2 (*continued*)

- Summary of SCSI-2 Adapters 56
- SCSI-2 Fast PCI Adapter
 - PC Server 310 (ISA or MCA) 13
- SCSI-2 Fast/Wide PCI Adapter 120
 - &320e. 16
 - &520e. 29
 - PC Server 720 36
- SCSI-2 Fast/Wide RAID Adapter/A
 - PC Server 720 36
- SCSI-2 Fast/Wide RAID PCI-Bus Adapter
 - &320e. 16
 - PC Server 520 29
- Seek
 - Definition 216
- Seek Time
 - Definition 216
- Selectable Boot Sequence 195
- server
 - bottlenecks 6
 - downtime 9
 - load 6
 - location 8
 - number of users 8
 - positioning 2
 - purpose 8
 - selection 2, 9
 - type 6
 - types 3
- Server Failure Handling
 - LAN Adapter Failure 179
 - Mirrored Server Link Failure 178
 - Primary Server Failure 178, 189
 - Secondary Server Failure 178
- ServerGuide 14, 16, 26
- SFT III 192
- Single Drive Failure 55
- small business sever 10
- Software
 - AIX PS/2 81
 - ARCHIVE 81
 - ARCserve 81
 - ARCsolo 81
 - CorelSCSI 81
 - CPbackup 81

Software (*continued*)

- IBM DualStor Backup Program 81
- NetWare SFT III 174
- Novaback 81
- Novell NetWare 81
- SCO-UNIX 81
- SYTOS-PLUS 81
- Standby Machine 190
- StandbyServer 32
 - Benefits 180
 - Hard Disk Configuration Requirements 185
 - Hot Spare 179
 - Installation Overview 182
 - System Configuration 181
- Stripe 54
- Summary
 - 3510 SCSI Storage Enclosure 70
 - 3511 External Enclosure for SCSI Devices 67
 - 3516 Hot Swap Expansion Enclosure 66
 - Tape Drives 82
- Synchronous mode
 - Definition 216
- SynchroStream technology 24
- System Notification 205
- SYSTEMVIEW 193

T

- T1 lines 5
- Tape Drives
 - 3450 1/4-Inch Drive 75
 - 4/10GB Internal Tape Drive 84
 - IBM 3440 2.0GB 4mm Tape Drive Model 001 77
 - IBM 3445 5.0GB 8mm Tape Drive Model 001 76
 - IBM 3450 1.2GB 1/4-Inch Cartridge Tape Drive Model 001 75
 - IBM 4/10GB 4mm DAT Tape Drive Option Kits 78
 - Software 81
 - Summary 82
- Target
 - Definition 216

This soft copy for use by IBM employees only.

TCQ, Tagged Command Queuing 45

Terminator

 Definition 216

Tower, External 67

transaction

 frequency 8

 size 8

U

Ultrastar Family, Drives 59

UNIX 197

UPS (Uninterrupted Power Supply) 8

V

VPD (Vital Product Data) 24

W

Windows 197

Write-back 55

Write-through 55



This soft copy for use by IBM employees only.

Printed in U.S.A.

SG24-4525-00

