Inside OS/2 Warp Server for e-business

Girish Basavalingaiah, Ron Bloor, Tonko De Rooy, Edgar Omar Gonzalez Espinosa, Roger Govind, Peter Marfatia, Oliver Mark, Frank Mueller, Indran Naick, Leon Van Der Linde, Frank Vanhulle

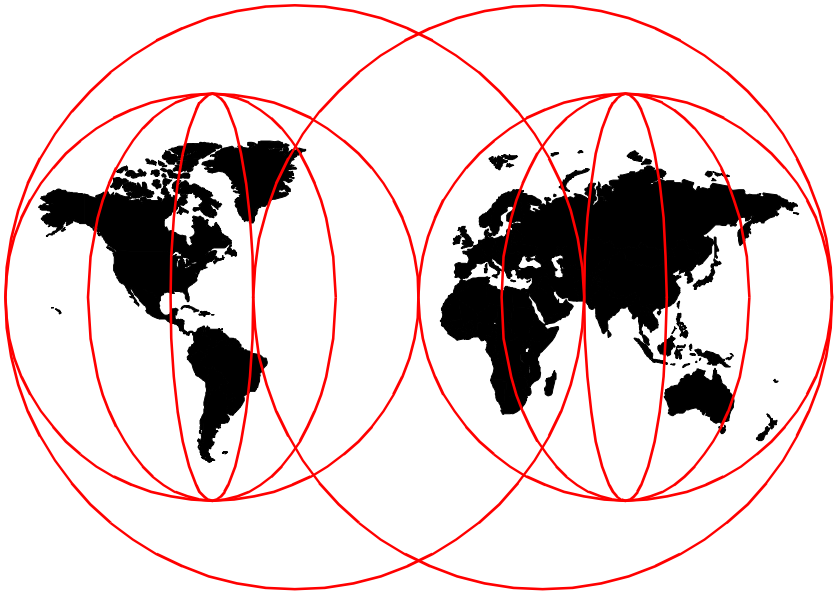

International Technical Support Organization

http://www.redbooks.ibm.com

SG24-5393-00

SG24-5393-00

International Technical Support Organization

Inside OS/2 Warp Server for e-business

July 1999

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix C, “Special notices” on page 429.

First Edition (July 1999)

This edition applies to OS/2 Warp Server for e-business for use on Intel server hardware.

Note

This book is based on a pre-GA version of a product and may not apply when the product becomes generally available. We recommend that you consult the product documentation or follow-on versions of this redbook for more current information.

Comments may be addressed to:
IBM Corporation, International Technical Support Organization
Dept. DHHB, Building 003 Internal Zip 2834
11400 Burnet Road
Austin, Texas 78758-3493

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.

Contents

Figures xi

Tables xvii

Preface xix
The team that wrote this redbook xix
Comments welcome xxiii

Chapter 1. Introduction and overview 1
1.1 Evolution 1
1.1.1 Host computing: 1964 1
1.1.2 Personal computing: 1981 2
1.1.3 Client/Server: 1988 3
1.1.4 Network computing: The world today 4
1.2 Overview 6
1.2.1 About e-business 6
1.2.2 Architecture and services 8
1.3 A versatile server 15
1.4 System requirements 17
1.4.1 Hardware requirements 18
1.4.2 Hard space disk requirements 19
1.5 Replaced or discontinued components 20
1.6 Server packaging and licensing 21
1.6.1 386 HPFS licensing 21
1.7 The installation process 22
1.7.1 The bootable CD 22

Chapter 2. Base operating system enhancements 23
2.1 Virtual memory support 23
2.2 Intelligent input and output I_2O 23
2.2.1 I_2O architecture 24
2.3 1024 cylinder restriction 32
2.3.1 The INT 13h interface 32
2.3.2 OS/2 boot process 34
2.4 Kernel enhancements and file subsystems 38
2.5 SMP support and microprocessor affinity 40
2.6 Large file support 42
2.6.1 SES and large file support 44
2.7 Java 1.1 45
2.8 Netscape 4.04 46
2.9 Graphical locale builder 47

2.9.1 Java locale 50
2.9.2 Euro currency support 50

Chapter 3. Adapters and protocol services 53
3.1 Overview of adapters and protocol services 53
3.1.1 Adapter and protocol support 53
3.1.2 Socket/Multiprotocol transport services 59
3.2 Installing adapters and protocol services 60
3.3 Additional configuration for adapters and protocol services 66
3.3.1 Configuring Socket/MPTS. 69
3.3.2 NETBEUI parameters 72
3.3.3 Configuring more than four LAN adapters. 72
3.3.4 DOS and Windows LAN applications on OS/2 79
3.3.5 NetBIOS over TCP/IP on OS/2 Warp Server for e-business 80
3.4 MPTS - Strong encryption 87
3.5 Removing adapters and protocol services 87

Chapter 4. Journaled file system and logical volume manager. 89
4.1 OS/2 file system support 89
4.1.1 File allocation table. 89
4.1.2 High performance file system 91
4.1.3 HPFS386 92
4.1.4 Journaled file system 92
4.1.5 HPFS386 features not available with JFS 93
4.1.6 Comparison of features of FAT, HPFS, and HPFS386 JFS. 93
4.1.7 Network file system 94
4.1.8 CD-ROM file system 94
4.2 Logical volume manager 95
4.2.1 LVM terminology. 95
4.2.2 Overview of LVM 96
4.2.3 LVM and FDISK 99
4.2.4 Key components of the LVM. 100
4.2.5 Bad-block relocation 112
4.2.6 LVM operation 112
4.2.7 Logical volume manager benefits 120
4.3 Journaled file system 121
4.3.1 JFS cache 121
4.3.2 JFS structure 123
4.3.3 JFS disk layout 123
4.3.4 JFS system structure 124
4.3.5 JFS utilities. 136
4.3.6 A closer look at chkdsk. 140
4.3.7 System limits 141

4.3.8 Performance considerations 142

Chapter 5. File and print services 145
5.1 Review 146
5.2 Installation 150
5.3 File and print services configuration 152
5.4 File and print services administration 156
5.5 386 HPFS 158
5.5.1 Attended installation of 386 HPFS with fault tolerance 158
5.5.2 Installing local security 160
5.5.3 Log files created during installation. 160
5.5.4 Installing 386 HPFS using CID 160
5.5.5 Existing ACLs 163
5.6 Functional rollups 163
5.6.1 IBM Neighborhood Browser Enabler 1.0 163
5.6.2 IBM Network Client 4.1 for Windows 95 164
5.6.3 IBM Network Client for Windows NT 165
5.7 Capacity enhancements 166
5.7.1 Overview 166
5.7.2 Maximum number of connections 167
5.7.3 Maximum open files 168
5.7.4 Maximum search value. 169
5.7.5 Maximum shares 171
5.7.6 Summary 172
5.8 Multiple server names 172
5.8.1 Configuration 173
5.8.2 Scenarios 174
5.9 Vinca StandbyServer 177
5.10 New NET USE switch 178

Chapter 6. Integrating Windows NT Servers 181
6.1 Overview 181
6.1.1 Key features 182
6.1.2 Benefits 182
6.1.3 Limitations 185
6.2 Architecture 185
6.2.1 User accounts 186
6.2.2 Windows NT Server alias support. 186
6.3 Concepts 187
6.3.1 Microsoft domain models 187
6.3.2 Microsoft workgroup model. 189
6.3.3 User authentication 189
6.3.4 The IBM Networks User Account Manager 191

6.4 Access control 192
6.4.1 An overview of access control on NT 192
6.4.2 Share permissions 192
6.4.3 Permissions on files and directories 194
6.4.4 ACLs, SIDs, and security related issues 199
6.4.5 Summary of domain security management 201
6.5 Installation of IBM Networks User Accounts Manager for NT 201
6.5.1 Prerequisites tasks 201
6.5.2 Installation on the NT server 202
6.5.3 Defining the NT server on the Warp Server domain 206
6.5.4 Post installation 210
6.5.5 Problem determination 211
6.5.6 Summary of installation 215
6.6 NT file and print servers 215
6.6.1 Defining an alias on an NT file server 215
6.6.2 Integrating NT print servers 222

Chapter 7. TCP/IP Version 4.21 new functions 225
7.1 TCP/IP 4.21 225
7.1.1 Changes in TCP/IP 4.21 227
7.2 TCP/IP 4.21 installation 229
7.2.1 Selective install for networking 229
7.2.2 Installation from the client CD-ROM 232
7.3 TCP/IP configuration notebook 236
7.3.1 Local configuration 236
7.3.2 Remote configuration 236
7.3.3 Network tab 238
7.3.4 Routing tab 244
7.3.5 Host names tab 246
7.3.6 Autostart tab 250
7.3.7 General tab 252
7.3.8 Security tab 254
7.3.9 SOCKS tab 268
7.3.10 Printing tab 274
7.3.11 NFS tab 276
7.4 DHCP/DDNS servers 279
7.5 Virtual private networks 279
7.5.1 Changes in VPN versions 281
7.5.2 Configuring IPSec clients 281
7.5.3 Configuring IPSec filters and tunnels 282
7.5.4 Creating a tunnel between two OS/2 machines 292
7.6 Creating a mini-firewall 293
7.6.1 Securing interfaces 293

7.6.2 Configuring filtering 293
7.7 Network file system 294
7.7.1 Introduction to NFS 294
7.7.2 Mounting an NFS share 297
7.7.3 NFS client utilities 299
7.7.4 Automatically mounting NFS shares on start-up 305
7.7.5 Sharing an NFS share using LAN Server 305
7.7.6 Creating your own NFS shares 306
7.8 TIMED 306
7.9 FTPD improvements 306
7.9.1 FTPD multithreading 307
7.9.2 Restarting broken connections 307
7.10 Line printer improvements 307
7.10.1 Streaming LPD 308
7.10.2 Streaming LPRPORTD 308
7.10.3 LPD security 309
7.11 TFTPD improvements 309
7.11.1 TFTPD multithreading 310
7.11.2 TFTPD blocksize 310
7.11.3 TFTPD security 310
7.12 IP aliasing 311
7.13 X Windows 312
7.13.1 Installing PMX over TCP/IP 4.21 312
7.13.2 Configuring PMX 313
7.13.3 Alternatives to PMX 315
7.14 TCP/IP development toolkit 316
7.15 Performance improvements 316
7.15.1 32-bit stack enhancements 316
7.15.2 inetcfg 317
7.15.3 SYN cookies 317
7.15.4 Reuse timewait state enhancements 317
7.15.5 HTTP fast path performance 318
7.15.6 SMP exploitation 318
7.15.7 New API calls 319
7.15.8 Variable cluster sizes 320

Chapter 8. Lotus Domino Go Web Server and WebSphere 321
8.1 Lotus Domino Go Webserver 321
8.2 WebSphere application server 321
8.3 Fastpath Install - Lotus Domino Go Webserver 322
8.4 Web server uninstall 329
8.5 Functional components 331
8.5.1 Tailoring the Web server 332

8.5.2 Using CGI programs, GWAPI programs, and Java servlets . . . 333
8.5.3 Managing your Web server with SNMP protocol 333
8.5.4 Restricting access 334
8.5.5 Mapping resources 334
8.5.6 Logging requests and errors 334
8.5.7 Running the server as a caching proxy 335
8.5.8 Running server with multiple IP addresses or virtual hosts 335
8.5.9 Using server-side includes 336
8.5.10 Customizing the server's error messages 336
8.5.11 Adding a search engine to your Web site 336
8.5.12 Using proxy authentication 336
8.6 Fastpath Install - WebSphere 336
8.6.1 Uninstall for WebSphere application server 341
8.7 WebSphere functional components 342
8.7.1 Accessing the product documentation 343
8.7.2 Starting the application server manager 344
8.7.3 Using the application server manager 345
8.7.4 Monitoring servlet activity 347
8.7.5 Establishing and maintaining security 348
8.7.6 Managing servlets 349
8.8 Developing and implementing servlets 349

Chapter 9. IBM remote access services 359
9.1 Overview 359
9.1.1 IBM remote access services environments 359
9.2 PPP support 361
9.3 Client support 361
9.4 System requirements 362
9.4.1 IBM remote access connection server requirements 362
9.4.2 Remote client system requirements 363
9.4.3 OS/2 RAS remote client system requirements 363
9.4.4 PPP client system requirements 364
9.4.5 Remote client support restrictions 364
9.5 Installing IBM remote access services 365
9.6 Configuring IBM remote access services 370
9.6.1 Configuring PPP support on the connection server 371
9.6.2 Reviewing IP address considerations for PPP clients 372
9.6.3 Configuring a TCP/IP protocol router 373
9.6.4 Binding the IBM TCP/IP protocol to the adapters 373
9.6.5 Setting TCP/IP configuration for PPP 375
9.6.6 Specifying PPP parameters in \WAL\WCLLOCAL.INI 378
9.6.7 IP addresses in WCLIPADR.INI 381
9.6.8 Using DHCP servers 381

9.7 Remote client 382
9.8 Windows 95 PPP clients 382
9.8.1 Modem installation and configuration 382
9.8.2 Windows 95 network configuration 383
9.8.3 Dial-up networking configuration 384
9.9 Windows NT Version 4 PPP clients 386
9.10 OS/2 Warp PPP clients 386
9.10.1 IBM dial-up for TCP/IP configuration 386
9.10.2 Configuring IBM 8235 DIALs client for OS/2 Version 4.52 392
9.11 Administration 394
9.11.1 PPP security 394
9.11.2 Address configuration 396
9.11.3 Base functions of the open as menu 396
9.11.4 Advanced functions of the open as menu 398
9.11.5 Allowing PPP-clients to change their PPP passphrase 398
9.12 IBM remote access services internal architecture 403
9.13 Remove IBM remote access services 406

Appendix A. More on Windows NT administration 407
A.1 Administration tools and utilities 407
A.1.1 Server management 407
A.1.2 Netfinity client for NT 4.0 408
A.1.3 Web administration of Microsoft Windows NT servers 415
A.1.4 Other administration utilities from NT Resource Kit 422
A.1.5 Summary of administration tools and utilities 424

Appendix B. CD-ROM contents 427

Appendix C. Special notices 429

Appendix D. Related publications 433
D.1 International Technical Support Organization publications 433
D.2 Redbooks on CD-ROMs 434
D.3 Product documentation 434
D.4 Other publications 434

How to Get ITSO Redbooks 437
IBM Redbook Fax Order Form 438

List of Abbreviations 439

Index 441

ITSO Redbook Evaluation 445

Figures

1. Centralized host computing 2
2. Personal computing 2
3. Client/Server computing, connected but proprietary 3
4. Network computing, open connectivity 5
5. OS/2 Warp Server for e-business, high-level architecture 8
6. OS/2 Warp Server, deployment into a number of distinct environments . . 15
7. I20 architecture. 25
8. Flow of information with an I2O adapter. 27
9. Adapter and protocol configuration, selecting the I2O LAN driver. 29
10. Adapter and protocol, I20 configuration 30
11. Adapter and protocol configuration, CONFIG.SYS update 31
12. FAT boot process. 35
13. HPFS boot process 36
14. OS/2 Warp Server, selective install, SMP support. 41
15. Locale icon. 48
16. New locale name 48
17. Settings for new locale. 49
18. Two defined locales. 50
19. MPTS components 54
20. NDIS - multiple protocols 57
21. Overview of Socket/MPTS 60
22. Adapter and protocol services configuration during the installation. 62
23. Add adapter driver to Adapters and protocol services 64
24. Add protocol driver to Adapters and protocol services 65
25. Change settings in Adapters and protocol services. 66
26. Additional configuration of Adapters and protocol services 67
27. Adapter and protocol configuration 68
28. NetBIOS socket access configuration 70
29. NetBIOS configuration for eight adapters 74
30. NetBIOS, NetBIOS over TCP/IP and TCP/IP structure 84
31. Coexistence TCPBEUI. 86
32. Overview of LVM disk management for a JFS volume. 97
33. Expanding a volume, initial 98
34. Expanding a volume, adding space 98
35. Expanding a volume, completed 99
36. Comparison chart of FDISK and LVM concepts 100
37. Block diagram summarizing DASD I/O path 103
38. Physical view via LVM.EXE 104
39. Physical view via LVMGUI.CMD 105
40. Logical view via LVM.EXE 105

41. Logical view via LVMGUI.CMD 106
42. LVM command syntax 106
43. LVM parameter options (1 of 3) 108
44. LVM parameter options (2 of 3) 109
45. LVM parameter options (3 of 3) 110
46. Partitioning used in the CID environment 111
47. LVM command line commands to partition the disk 111
48. Syntax for JFS initialization 122
49. JFS physical organization in a logical volume. 126
50. Anatomy of an i-node. 133
51. Aggregate details of a JFS partition. 135
52. File system utility framework. 136
53. SMB system architecture. 144
54. File and print service, architecture. 148
55. OS/2 client and server architecture 150
56. Selective install for networking. 151
57. File and pint services features 151
58. File and print services, configuration. 152
59. File and print services, selecting a network adapter. 153
60. Autostart configuration. 154
61. File and print services icon 154
62. LAN services, file and print 155
63. Domain administration. 157
64. User's window 157
65. 386 HPFS installation panel. 159
66. 386 HPFS configuration panel. 160
67. Seamless access to the integrated server domain 183
68. Single point administration. 184
69. The IBM user accounts manager for Windows NT Server. 186
70. Microsoft NT domain models 188
71. Warp Server accounts database synchronization 190
72. NT Member server's access to the domain SAM. 191
73. NTFS access control example. 197
74. Windows NT installation, select network services 203
75. Windows NT installation, insert disk 203
76. Windows NT installation: Select OEM option. 204
77. IBM Networks user account manager properties. 204
78. Windows NT installation, setup messages 205
79. Network identification. 205
80. Identification changes 206
81. LAN server administration 207
82. Domain view 208
83. Defined servers 208

84. Defined server - create 209
85. Defined servers 209
86. Event Viewer - system log 211
87. Event Detail - IBMLOGON service started successfully 212
88. Event Detail - account synchronization 213
89. Event Detail - synchronization failure. 214
90. Directory alias for NT resource 217
91. Manage the access of the NT resource 218
92. Setting the access permissions 219
93. TCP/IP stack structure. 226
94. TCP/IP, DLL interface to stack. 227
95. Selective install for networking: Selection screen 229
96. Selective install for networking: TCP/IP 4.21 options 230
97. Selective install for networking: TCP/IP configuration 231
98. TCP/IP 4.21 advanced path selection screen 233
99. TCP/IP configuration during advanced installation 234
100.TCP/IP configuration (local) 236
101.Create TCP/IP administrator password. 236
102.Allow remote configuration 237
103.Configure remote system 237
104.TCP/IP configuration authorization 237
105.TCP/IP configuration notebook: Network, basic 238
106.TCP/IP configuration notebook: Network, advanced 1 240
107.TCP/IP configuration notebook: Network, advanced 2 242
108.TCP/IP configuration notebook: Routing 244
109.TCP/IP configuration notebook: Routing, route entry 245
110.TCP/IP configuration notebook: Host names, name resolution. 246
111.TCP/IP configuration notebook: Host names, hosts 248
112.TCP/IP configuration notebook: Autostart, inetd 250
113.TCP/IP configuration notebook: General. 252
114.TCP/IP configuration notebook: Security, user access 254
115.TCP/IP security, add user 256
116.Configure FTP access. 257
117.Configure TELNET access 259
118.Configure REXEC access 260
119.Configure NFS access 261
120.TCP/IP configuration notebook: Security, RSHD 262
121.TCP/IP configuration notebook: Security, TFTPD. 264
122.TCP/IP configuration notebook: Security, Admin PW 266
123.TCP/IP configuration notebook: SOCKS, defaults 268
124.TCP/IP configuration notebook: SOCKS, direct routes 270
125.TCP/IP configuration notebook: SOCKS, direct routes, add entry 271
126.TCP/IP configuration notebook: SOCKS, servers 272

127.TCP/IP configuration notebook: SOCKS, servers, add entry 273
128.TCP/IP configuration notebook: Printing 274
129.TCP/IP configuration notebook: NFS 276
130.TCP/IP configuration notebook: NFS, add directory 277
131.TCP/IP configuration notebook: NFS, add host 278
132.Virtual private networks, implementation possibilities. 279
133.TCP/IP 4.21 configuration notebook: Autostart, nfsstart 298
134.\PIPE\LPD0 settings 308
135.Installing PMX from TCP/IP 2.0. 313
136.Old-style configuration notebook: PMX tab 314
137.PMX configuration: Initial settings, keyboard 315
138.Domino Go Webserver, Installation Panel with instructions 323
139.Installation screen, selecting the components and their locations 324
140.Domino Go Webserver, available disk space.. 326
141.Domino Go Webserver, configuration panel 327
142.Domino Go Webserver, installation status. 328
143.Domino Go Webserver, installation status 329
144.Domino Go Webserver, folder. 329
145.Domino Go Webserver, removing components 330
146.Domino Go Webserver, selecting components for removal 330
147.Domino Go Webserver, removal progress 331
148.Domino Go Webserver, status message. 331
149.WebSphere, installation 337
150.WebSphere, destination 337
151.WebSphere, selectable components. 338
152..WebSphere, selecting the Web server. 339
153.WebSphere, application folder selection 340
154.WebSphere, confirm options 340
155.WebSphere, Readme file 341
156.WebSphere, removal. 342
157.WebSphere, Log-In window 345
158.WebSphere, application server manager 346
159.WebSphere, set up Window for an applet. 347
160.WebSphere, first servlet output 351
161.Servlet, sample application 353
162.Servlet, application processing 357
163.Remote to LAN scenario 360
164.Setup and installation 366
165.IBM remote access services configuration 367
166.IBM remote access services configuration 368
167.Configure user ID and password. 369
168.IBM remote access services logon option. 371
169.Adapters and protocol services - Binding TCP/IP to LAN adapter 374

170.Adapters and protocol services - Binding TCP/IP to LAN adapter 375
171.Network section of the TCP/IP configuration notebook. 376
172.Network section of the TCP/IP configuration notebook. 377
173.Routing section of the TCP/IP configuration notebook 378
174.PPP section of the \WAL\WCLLOCAL.INI. 379
175.\WAL\WCLIPADR.INI file 381
176.Example environment 386
177.Dial-up configuration - Login info. 388
178.Dial-up configuration - Connect info 389
179.Dial-up configuration - Modem info 390
180.Closing dial configuration 390
181.IBM Dial-up for TCP/IP 391
182.PPP dialer connect info. 392
183.IBM remote access services folders 396
184.Change passphrase page. 399
185.RAS, protocol architecture 404
186.Warp Server alias creation flow. 408
187.Netfinity service manager 409
188.Security manager 409
189.Security manager: Add user ID and enable services 410
190.Security manager: Revoke services from public. 410
191.Directory alias - Create 411
192.Select the NT server for remote administration. 412
193.NT server manager - Share management. 413
194.NTFS security access management 414
195.Microsoft Web administration 415
196.Microsoft Internet service manager. 418
197.WWW service properties. 419
198.Web administration of NT server. 420
199.Web administration - File system management 421
200.Web Administration - Share management 421
201.Web Administration - NTFS management. 422

Tables

1. OS/2 Warp Server for e-business, disk space requirements 19
2. Limitations of various interfaces. 33
3. Adapters and protocol services installation 62
4. Configuration menu items 68
5. NetBIOS Interface configuration for socket access 71
6. Relationship between partition size and cluster size 90
7. Feature comparison between FAT, HPFS, HPFS386, and JFS. 93
8. Configuration variables for LVM. 106
9. JFS initialization parameters 122
10. Parameters for JFS format. 137
11. Parameters for CHKDSK 138
12. Parameters for DEFRAGFS. 139
13. Parameters for EXTENDFS 139
14. IBM Network Clients for Windows NT – Summary 165
15. Maxconnections, details. 168
16. Maxopens, details. 168
17. maxsearches, details. 169
18. maxshares, details. 171
19. Summary of new capacity enhancements. 172
20. External Resources on an OS/2 Warp Server domain 179
21. Sharing a directory. 193
22. File and directory permissions on Windows NT Server 194
23. Differences between TCP/IP versions 227
24. Selective install for networking: TCP/IP services. 230
25. Selective install for networking: TCP/IP configuration 231
26. TCP/IP installation: Select and configure components 233
27. Configuration during advanced installation 235
28. TCP/IP 4.21 settings: Network, basic 239
29. TCP/IP 4.21 settings: Network, advanced 1 241
30. TCP/IP 4.21 settings: Network, advanced 2 243
31. TCP/IP configuration notebook: Routing 244
32. TCP/IP configuration notebook: Add routing entry. 245
33. TCP/IP configuration notebook: Host names, name resolution 247
34. TCP/IP configuration notebook: Host names, hosts 249
35. TCP/IP configuration notebook: Autostart, inetd 250
36. TCP/IP configuration notebook: General settings 253
37. Sample timezones. 253
38. TCP/IP configuration notebook: Security, user access 255
39. TCP/IP security, add user 256
40. Configure FTP access 257

41. Configure TELNET access 259
42. Configure REXEC access 260
43. Configure NFS access 261
44. TCP/IP configuration notebook: Security, RSHD 263
45. TCP/IP configuration notebook: Security, TFTPD 265
46. TCP/IP configuration notebook: Security, Admin PW 267
47. TCP/IP configuration notebook: SOCKS, defaults 269
48. TCP/IP configuration notebook: SOCKS, direct routes 271
49. TCP/IP configuration notebook: SOCKS, direct routes, add entry 271
50. TCP/IP configuration notebook: SOCKS, servers 273
51. TCP/IP configuration notebook: Printing 275
52. TCP/IP configuration notebook: NFS 277
53. TCP/IP configuration notebook: NFS, add directory 278
54. TCP/IP configuration notebook: NFS, add host 278
55. Configure filter, parameter description 283
56. IP Sec policy file, parameter description 286
57. Admin command, parameter description 291
58. UNIX file system access controls 295
59. NFSSTART parameters 298
60. Mount options 299
61. TIMED parameters 306
62. Domino Go Webserver, selectable components 325
63. Domino Go Webserver, configuration values 327
64. WebSphere, understanding the components 338
65. Servlets Monitor, tasks and their uses 348
66. Servlet security, tasks 348
67. Servlet management, tasks 349
68. Modification of CONIFIG.SYS and PROTOCOL.INI 370
69. PPP section of the WCLLOCAL.INI 379

Preface

This redbook describes some of the core functions of the IBM OS/2 Warp Server for e-business product based on the experience of IT Specialists who participated in the ITSO, Austin Center project.

The purpose of this redbook is to provide information and guidance on new features that have been introduced into OS/2 Warp Server for e-business. This document does not describe the systems management, backup and restore, software distribution and advanced print services components of IBM OS/2 Warp Server for e-business. A later redbook is planned to cover those components.

This redbook is not intended to provide information from the ground up. Instead, knowledge of IBM LAN Server 4.0 or OS/2 Warp Server 4.0, as well as an understanding of TCP/IP, is assumed.

The team that wrote this redbook

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

Girish Basavalingiah is currently a Software Analyst in IBM Global Services India. He has a Bachelor's Degree in Computer Science and Engineering. He has six years of experience in system software development and support and is a member of the OS/2 LAN SERVER change team providing Level 2.5/3 support from India.

Ronald Bloor is a Warp Server Instructor and Network Consultant. He is President of Connection Business Solutions Inc., which is in Toronto, Ontario Canada. He holds a number of IBM and Microsoft certifications that include IBM Network Communications Engineer and Microsoft Certified Systems Engineer. For the past seven years, he has been a trainer and consultant in the areas of communications networks, IBM OS/2, IBM Communication Server, IBM Warp Server, and Microsoft NT. Prior to this, he worked for IBM Canada Ltd.

Tonko De Rooy works for IBM EMEA Netfinity Servers Level 2 support in the United Kingdom and has worked at IBM for four years. His areas of expertise include IBM OS/2, IBM LAN Server, networking, IBM PC servers, and IBM network hardware. He holds a number of IBM and Novell certifications including Certified OS/2 Warp Engineer and Certified OS/2 Warp Server Engineer. He has written extensively on TCP/IP 4.21.

Edgar Omar Gonzalez Espinosa is an Electronic Engineer and is working as an I/T Specialist for Operative Systems in Mexico. He has three years of experience in OS/2 Warp and LAN Server. He had worked at an IBM-affiliated company for three years. Currently, he is an IBM employee. His areas of expertise include six Microsoft courses for Microsoft Certified Systems Engineer. He has taught IBM courses about Warp Server and Voice Type. He has written extensively on File and Print Services.

Roger Govind is a Desktop/Server Specialist for SONZ in New Zealand. He has seven years of experience with OS/2. He has worked at IBM for two years. His areas of expertise include OS/2, OS/2 Warp server, and PC/Server hardware. He has written extensively on Logical Volume Manager (LVM) and the Journaled File System (JFS).

Oliver Mark is a Senior I/T Specialist for Client/Server at Global Services in IBM Germany. He has eight years of experience in OS/2 and Warp environments as well as experience with OS/2 Warp Server, CID Installation, NetFinity, and Remote Access Services. He holds certifications as an IBM OS/2 Warp Engineer, Warp Server Engineer, Warp Instructor, Warp Server Instructor, IBM Professional Server Expert, and Microsoft Certified Professional. He has written extensively on IBM NetFinity.

Frank Mueller is a System Administrator at the Zevener Volksbank eG, Germany. He has five years of experience in OS/2 Client/Server configuration, administration, and internal user support. His areas of expertise include the OS/2 family of products and PC hardware.

Peter Marfatia is the Senior Technical On-Site Support Consultant in IBM Global Services at IBM in Melbourne, Australia He holds a bachelor of science in computer science from Deakin University and has been an IBM Professional Server Specialist since 1996. Peter has worked for IBM Global Services Australia since 1996. He has six years of experience in a variety of areas related to OS/2 including troubleshooting, technical support, developing and supporting OS/2 images for IBM Asia-Pacific, LAN Systems, and NVDM/2.

Indran Naick is a Senior IT Specialist at the International Technical Support Organization, Austin Center. He has 10 years of experience with IBM and writes extensively on OS/2, Warp Server, and WorkSpace On-Demand. Before joining the ITSO in 1999, Indran worked in IBM South Africa as a Software Solutions Architect.

Leon Van Der Linde is a Consultant Systems Engineer in IBM South Africa. He joined IBM 27 years ago and spent 19 years on the mainframe specializing in VM but also supporting VSE, MVS/SP, and Storage Products.

He was country system specialist for VM/SP for several years. He started supporting OS/2 in 1994 and is OS/2 Warp Certified. Leon has written extensively on VM, MVS, and AIX.

Frank Vanhulle is the manager of the personal computer development and support team of Anhyp NV, a banking company in Belgium that is a member of the AXA group. This team supports over 2000 personal computers running OS/2 and Windows operating systems totally controlled via Netfinity. Frank started his career at Anhyp in 1995 and holds a degree in Electrical Engineering specializing in computer science from the Vrije Universiteit Brussel (VUB). His primary role is to define the strategic direction of personal computers for Anhyp.

.

The first team that helped write the book:

From the top left: Oscar, Oliver, Frank and Indran
From the bottom left: Ronald, Tonko, Roger and Edgar

The second team that wrote the book:

From the top left: Girish, Frank and Peter
From the bottom left: Indran and Leon

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

Temi Rose
International Technical Support Organization, Austin Center

Starwalker Jj
Graphic Illustrations

David Medina
Lead Developer, IBM Austin

Steven French
Lead Developer, IBM Austin

IBM Austin Development Leads:

Ron Aguirre
Barry Arndt
Chandra Bagchi
Randy Baker
Jim Brisson
Oscar Cepeda

David Dutcher
Peter Greulich
Frank Grubbs
Bill Hartner
Mark Johannsen
David Kleikamp
David Klein
Telford Knox
Sharon Lucas
Wade Mahan
Cristi Nesbitt
Velma Pavlasek
Steven Pratt
Ben Rafanello
Bill Sinclair
John Stiles
Paul Thayer
Steve Tipton
Andreas Tuerk
Allen Wissinger

Other IBM Locations:
Timothy F. Sipples
Peter Degotardi

IBM India:
Amol Mahamuni
Shamarukh Mehra

Comments welcome

Your comments are important to us!

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

- Fax the evaluation form found in “ITSO Redbook Evaluation” on page 445 to the fax number shown on the form.
- Use the online evaluation form found at: `http://www.redbooks.ibm.com`

  For IBM Intranet users `http://w3.itso.ibm.com`

- Send us a note at the following Web address:

  redbook@us.ibm.com

Chapter 1. Introduction and overview

OS/2 Warp Server for e-business is the entry-level member of the IBM family of scalable server platforms and builds on OS/2 Warp Server's 10-year track record for mission-critical applications in providing a reliable foundation for network computing and complex e-business solutions.

OS/2 Warp Server for e-business is Year 2000 ready, supports eurocurrency, and, when packaged with WebSphere and Lotus Domino Go, is able to provide content-rich Web sites with strong enterprise host systems interoperability.

This chapter will describe the transformation of OS/2 Warp Server for e-business, a summary of the features, the systems requirements, and the installation process.

1.1 Evolution

Over time, computing trends change. Technology choices made by individuals and corporations are based on potential immediate and long-term benefits. These choices are influenced by a number of factors. Historically, it appears that the economics of any technology is the most significant factor. Both the related costs and the return on investment are important factors.

This section will describe the major trends in computing and describe how OS/2 Warp Server for e-business fits into the current computing paradigm.

1.1.1 Host computing: 1964

At one time, computing was defined by the host. These machines were controlled by specialists in large businesses and institutions who understood the emerging field of electronic computing. Host computing, in its heyday, was all about centralized everything: Centralized computing power, centralized administration, centralized decision-making. The high cost of computing at the time did not allow for a distributed architecture.

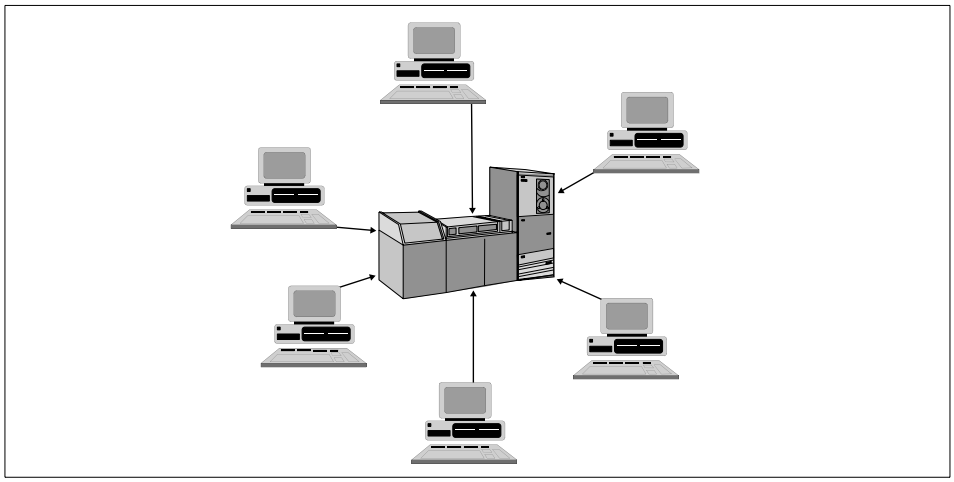

Figure 1. Centralized host computing

Centralization has some significant benefits: With everything centralized, it was easier to build, deploy, and manage new applications. The IBM mainframe became the de-facto work horse of the airline industry, banking industry, retail industry, and dozens of others that harnessed enterprise computing for data processing and transactions. The minicomputers that became popular in the '70s and UNIX systems are both variations of host computing.

1.1.2 Personal computing: 1981

Personal Computing was made possible by the advent of the microprocessor, a computer on a single chip. The personal computer radically changed the face of information technology not by marginalizing host systems but by opening up new uses and new markets for computing.

Figure 2. Personal computing

At the price point, although much higher than it is today, the PC era put information technology into the hands of millions of individuals. Computing was suddenly affordable. It could be bought from retail outlets, over the phone, or online. Applications came in a box. What this era accomplished, more than anything else, was to make consumers and end users more savvy about and more open to what computing could help them achieve. But, this model was not without its flaws.

1.1.3 Client/Server: 1988

Client/server attempted to bridge the gap. The imperfection of the PC era was that PCs were islands of computing, disconnected from mainframes and other enterprise systems. Valuable data, applications, and services on these larger systems were inaccessible to users on PCs, and the valuable information stored on PCs could not be shared with other users. Client/server computing attempted to rectify this. It envisioned a world of enterprise systems serving data and applications to client PCs, and it promised universal connectivity between any client and any server. Client/server gave us powerful new applications like groupware, which supports collaboration.

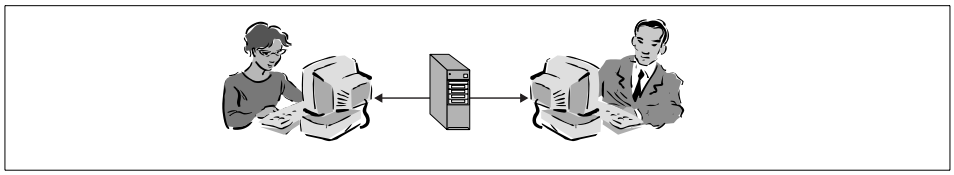

Figure 3. Client/Server computing, connected but proprietary

But, there are disadvantages.

- A number of client side platforms exist, each with its own set of APIs. This means that client side applications written for one platform will not work on

other platforms. To support multiple clients requires multiple copies of the client side application.
- The cost of managing the clients are high. Both the operating system and the application needs to be maintained. Also, as new releases of the operating system and applications become available, more powerful hardware has to be purchased and deployed.
- The client side also requires complex administration with much of the functionality of the solution down at the client. This problem is not so evident with a small number of localized machines but becomes a major problem when a large number of machines are distributed.
- Designing client/server applications is complex.

1.1.4 Network computing: The world today

Network Computing is a new model of computing that is taking the world by storm. By definition, it is a network-based style of computing based on seamless access to information and applications.

It is different from traditional computing in that the applications and data the user accesses reside on servers rather than PCs. The connection is provided by a network. Because network computing is based on standards like the Internet protocols and Java, most of the world's computers can, for the first time, communicate and share applications and data even though they are made by different companies and run different operating systems. And, because the heavy workload of computing shifts to servers, entirely new end-user access devices are now possible: Everything from Web-enabled TVs and personal digital assistants to screenphones and smartcards.

This computing model brings significant relief to the problems experienced with client/server implementations. This is mostly due to the high rate of

adoption of Internet technologies. Having large numbers of people adopt a technology brings the costs down making it accessible to even more.

Seemingly overnight, the Internet and the World Wide Web have delivered on the promise of *any client to any server*. Now, anyone using a PC who wants access to information on another computer can get it as long as the computing systems adhere to the technical standards of the Internet (and the vast majority of computers today do). Not only can the world's computers talk to each other and share data, they can support transactions of all kinds effortlessly and inexpensively. And, with the Java programming language, they can share applications as well.

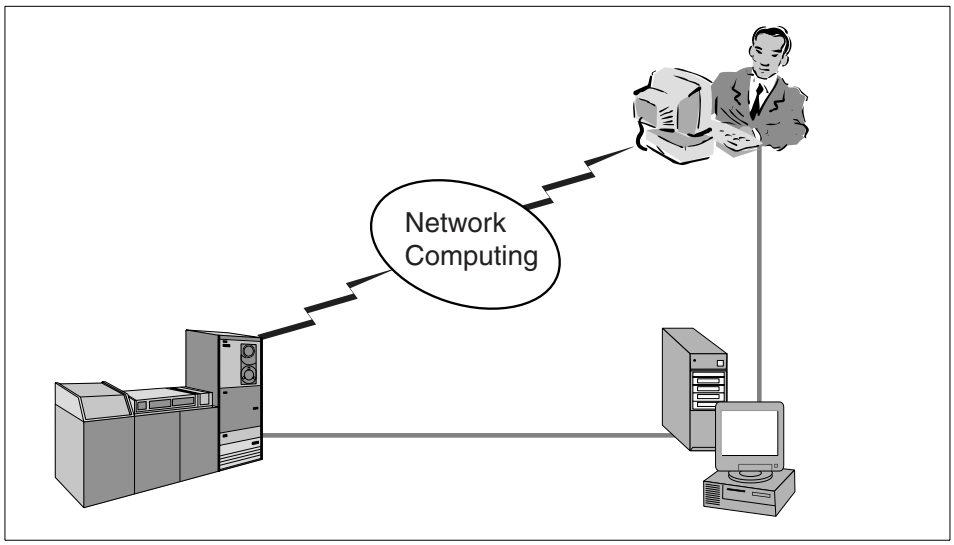

Figure 4. Network computing, open connectivity

The reason we believe this era of network computing is going to last for a long time is that it builds upon the advantages of all the eras that have come before while eliminating virtually all of their downsides. It recentralizes the development and deployment of applications. As in the host computing era, applications are centrally deployed and managed. At the same time, network computing actually expands the freedom individuals gained during the PC era by conceptually giving them access to every individual, business, and institution in the world. Finally, network computing offers the collaboration and connectivity that the client/server era promised but never delivered.

1.1.4.1 Pervasive computing

By the year 2000, it is estimated that more than 40 percent of the access to the Net will be via devices that are not PCs-PDAs, web-based televisions, kiosks, or small devices. Truly, pervasive computing is what will happen when devices disappear and chips and tiny disk drives become embedded in things, such as our cars, machine tools, appliances, houses, even clothing. This will make possible a new class of application: online services that are interwoven into the fabric of people's lives. The opportunity is enormous for consumers and for the businesses that will provide those applications.

1.1.4.2 Deep computing

The story gets more profound. All of these devices will enable companies to amass valuable data. Every one of these things connected to the Net will be able to transmit information. So, for instance, a car manufacturer will be able to gather data about how consumers use their vehicles. Here is where deep computing comes in. It is the astonishing combination of ultrafast computing combined with sophisticated analytical software. It will enable companies to analyze, find patterns, and take action on all the data they have gathered.

1.2 Overview

When the first PC's were introduced, they were stand-alone. The high cost of disk space and printers introduced the concept of a server. One of the most popular network services was file and print sharing. The server's hardware architecture was similar to that of the PC client, but the software was not. Network Operating Systems, as they were known, were available from vendors, such as IBM, Novell, Microsoft, and Banyan Vines.

In 1984, just a few years after the introduction of the IBM Personal Computer, IBM announced the IBM PC Network Program software designed to allow peer-to-peer communications and resource sharing in a DOS environment among IBM PCs. The product evolved into OS/2 LAN Server, and then into OS/2 Warp Server. OS/2 Warp Server was the first to introduce an extensive range of application services. The newest member of that family, OS/2 Warp Server for e-business, further extends the set of application services that are provided.

1.2.1 About e-business

The current computing model, namely, network computing, as we defined it, is a network-based style of computing based on seamless access to information and applications.

e-business describes the customer benefits of network computing. It is about transforming key business processes with Internet technologies.

e-business embraces e-commerce (buying and selling over the Net), but it is more than that. e-business includes intranet applications that let employees better manage their knowledge and operations, extranets that transform the way an organization works with its suppliers, distributors and partners, and the very important noncommercial Web applications in education, healthcare, and government. In sum, an e-business is any company or institution that conducts its core business via the Internet.

e-business is any activity that connects critical business systems directly to their critical constituencies (customers, employees, vendors and suppliers) via intranets, extranets and over the World Wide Web

In general, most people understand e-business as doing business electronically. There has been a lot of talk about network computing, e-business and e-commerce, and, while the terms are not synonymous, there is overlap between them. e-business is a subset of network computing that is a broader concept than e-commerce.

This is e-business – where the strength and reliability of traditional information technology meet the Internet. This new Web + IT paradigm merges the standards, simplicity, and connectivity of the Internet with the core processes that are the foundation of business. The new killer apps are interactive, transaction intensive, and let people do business in more meaningful ways.

In order to support e-business, a server must possess the necessary application-level software to allow a customer to leverage the network computing model. These layers have now been added to OS/2 Warp Server, thus the name OS/2 Warp Server for e-business.

OS/2 Warp Server for e-business consolidates previous OS/2 Warp Server V4 releases into a single CD-media package. Current OS/2 Warp Server V4.0 FixPaks and current OS/2 Warp Server V4.0 Software Choice features are rolled into the product including Year 2000 and euro currency readiness, Java and Netscape, as well as improved TCP/IP. Also included are Netfinity V5.2, LDAP Client, Dynamic IP Client for Win95 and NT, and OS/2 Warp Server Backup/Restore V6.0. And, the new JFS delivers increased server reliability. Seamless Windows NT Server Management, which allows Windows NT 4.0 servers to be seamlessly integrated into a network, fortifies OS/2 Warp Server for e-business's capability in mixed environments. Industry standard support for Intelligent Input/Output (I2O) adapters for SCSI disks and LAN

Ethernet and token ring is also planned. What this adds up to is a powerful mission-critical foundation for the e-business transformation. The Server's features are designed to meet the more rigorous performance and availability requirements of universal access. It can facilitate the transition to network computing and the Java Application Model with the Java and LDAP toolkits. Yet, its support of existing OS/2 applications and eased administration of Windows platforms can preserve existing investments. And, most definitely, it is the optimum platform for WorkSpace On-Demand, which is a proven IBM product for reducing cost of ownership, speeding deployment of new applications, and transitioning to network computing. OS/2 Warp Server for e-business is the definitive Intel server in a comprehensive computing infrastructure that meets today's and tomorrow's demanding requirements.

1.2.2 Architecture and services

The figure below is a high-level view of the architecture of OS/2 Warp Server for e-business. Built into the base operating system is support for a number of file systems. Communication support is also tightly integrated into the base. Each of the services that the server can provide runs as an application on top of the base operating system. Services can be selected during installation depending on the tasks you require the server to perform.

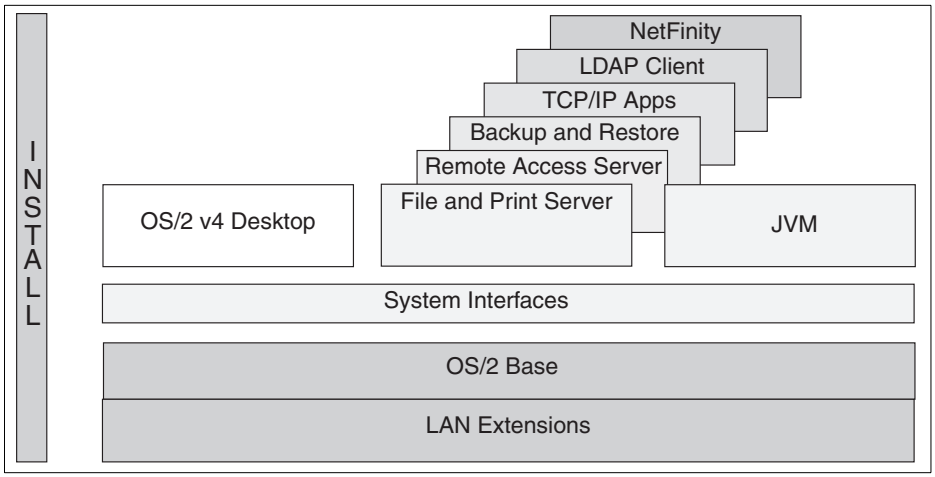

Figure 5. OS/2 Warp Server for e-business, high-level architecture

1.2.2.1 Base operating system services

The base operating system for OS/2 Warp Server for e-business is an advanced, multitasking, 32-bit operating system that runs DOS, Windows,

and OS/2 16 and 32 bit applications and utilizes SMP hardware configurations.

The base operating system interface is based on the intuitive OS/2 Warp 4.0 interface making it easy for a system administrator to perform server tasks.

Memory support has been enhanced. Applications running on the server can access a virtual memory address space of up to 3GB.

The base operating system also boasts some significant new server level features.

- A new 32-bit SMP-enabled Kernel Execution Environment (KEE) significantly improves the SMP scalability of the product.
- The server includes support for both the uniprocessor and Symmetric Mutiprocessor (SMP) support. SMP support is optimized for 8-way CPUs with support within the architecture for up to 64-way systems.
- Two new OS/2 kernel APIs that allow OS/2 applications to specify which processor they would like to run on in a multiprocessor environment.
- Security Enablement Services (SES) and Kernel Program Interfaces (KPIs) previously available in FixPaks are included. Both the uniprocessor and SMP versions of SES are available in the base. In addition, there are modified KPI's for supporting files greater than 2GB in size.

Large files greater than 2GB in size are now supported.

The base operating systems include support for Intel's intelligent I/O (I_2O) ethernet, token-ring, and SCSI specifications. The I_2O architecture is an emerging standard for the development of device drivers. The architecture enables adapters to run driver functions that process transactions that are normally handled by the CPU, thus, decreasing the CPU utilization.

These and other enhancements are discussed in more detail in Chapter 2, "Base operating system enhancements" on page 23.

1.2.2.2 Adapter and protocol services

OS/2 Warp Server for e-business provides you with a wide rage of supported networking protocols and communication adapters that you may use in many combinations to suit your requirements. This is for a server system in a LAN environment as well as for wide area networking. Adapters and Protocol Services may be called the communications engine of OS/2 Warp Server for e-business since it provides the communication support for all the other components of this product.

This is discussed in more detail in Chapter 3, “Adapters and protocol services” on page 53.

1.2.2.3 File and print

The File and Print Sharing services component of OS/2 Warp Server for e-business is a local area network (LAN) application used to share hardware and software resources that are located on the server machine.

Despite the increasing popularity of web-based applications, file and print services are still commonplace in many of today’s work environments. Some of the enhancements in File and Print Services include:

- Windows NT User Account manager: This enables administrators to manage user IDs on Windows NT Server systems that are defined as additional servers in the OS/2 Warp Server domain. The User Account Manager function lets administrators consolidate the administration of both NT Server-based additional servers and OS/2 Warp Server-based systems into a single interface, thus, saving time.
- Windows 95/NT client support: The popularity of clients running Windows95 and Windows NT has increased greatly throughout the 1990s. Although these clients are compatible with OS/2 Warp Server, IBM has released applications to enable these clients to become full participants in the domain, which includes support for IBM-based enhancements, such as aliases, logon assignments, and disk limits.

IBM has also increased the maximum numbers allowable for some resources within the File and Print Services component, such as MAXSEARCHES, MAXSHARES, MAXCONNECTIONS, and MAXOPENS. This enables OS/2 Warp Server for e-business to continue to scale even higher than it could before making it the ideal choice for enterprise customers.

1.2.2.4 TCP/IP services

TCP/IP services have been part of OS/2 Warp Server since the Version 4 release in 1996. Since this time, many enhancements in function and usability have been incorporated into subsequent TCP/IP Services releases. OS/2 Warp Server for e-business includes the latest version of TCP/IP, Version 4.2.1. Some of the highlights and improvements are:

- 32-bit enablement: Since Version 4.1, IBM has enhanced TCP/IP for 32-bit APIs. These can provide better throughput because of 32-bit data grouping.
- New file APIs: Applications that manipulate, send, and receive files, such as web servers, will see better performance through new send_file() and accept_and_recv() APIs.

- Enhanced daemons: The File Transfer Protocol Daemon (FTPD) is enhanced for multi-threading for better performance. It also supports the restart of broken connections. The Trivial File Transfer Protocol Daemon (TFTPD) is also multi-threaded and supports block sizes of up to 8 KB. There is also a security mechanism that enables the administrator to restrict access to certain directories for TFTP clients. The line printer daemons, LPD and LPRPORTD, now have Streaming support to enable IBM Network Station clients to use OS/2-based print servers.
- NFS client and server support: The Network File System, previously available in TCP/IP Version 2.0, is enhanced and included. This NFS is a 32-bit implementation that is ready for supporting the IBM Network Station. In addition, the administrator can mount NFS shares on UNIX-based systems and redirect these shares locally using NETBIOS, thus acting as an NFS gateway for OS/2 Warp Server clients that do not have NFS client support installed.
- DHCP and DDNS support: These services have been enhanced significantly over the past few years. The inclusion of a BINL server that supports Intel Wired for Management specifications is new. For a good description of the architecture and implementation of DHCP and DDNS in IBM products. Refer to the redbook *Beyond DHCP - Work Your TCP/IP Internetwork with Dynamic IP*, SG24-5280.
- VPN support: This function allows the administrator to create a Virtual Private Network (VPN) between two OS/2 systems running TCP/IP Version 4.1 or higher. The IP-filtering function of VPN also enables the administrator to create a mini-firewall. For more exhaustive information about VPN, refer to the redbook *A Comprehensive Guide to Virtual Private Networks, Volume 1: IBM Firewall, Server and Client Solutions*, SG24-5201.

This is discussed in more detail in Chapter 7, “TCP/IP Version 4.21 new functions” on page 225.

1.2.2.5 Remote access services

In OS/2 Warp Server Version 4, IBM included the ability for OS/2 and Windows systems to dial into a Remote Access Server to access resources as if the remote clients were actually on the LAN. This function was provided by the LAN Distance program. Since this time, IBM has also added a Point-to-Point Protocol (PPP) server and the ability for any PPP client to connect to it. PPP clients include the IBM 8235 Dialer, Windows 95/98, Windows NT Remote Access, and Shiva PPP.

This PPP server and client support has been included in OS/2 Warp Server for e-business. Remote Access Services has also been enhanced to provide client support for the dynamic assignment of IP address known as dynamic IP.

This is discussed in more detail in Chapter 9, “IBM remote access services” on page 359.

1.2.2.6 Systems management

IBM has a long history in the Systems Management area on many computing platforms. On the Intel platform, the original systems management function in OS/2 Warp Server Version 4 was provided by OS/2 SystemView. With the release of the OS/2 Warp Server SMP Feature, the Netfinity server function was introduced. OS/2 Warp Server for e-business now includes Netfinity Version 5.2.

Also new in OS/2 Warp Server for e-business is the inclusion of a Tivoli Management Agent (TMA). This TMA enables the server to become a managed object in the Tivoli Managed Environment, a very popular cross-platform systems management framework.

1.2.2.7 Backup and recovery

Warp Server for e-business includes PSnS Version 6.01. This component allows the backup of applications, data, and access controls to diskette, tape, local and remote hard disks, optical disks, and the ADSTAR Distributed Storage Manager (ADSM) server on any platform. This version has more hardware device support and includes the following major enhancements:

- Usability: IBM has enabled PSNS to support backup and restore procedures written in C and REXX through the corresponding APIs (see the online documentation for more details).
- Additional media: PSnS now supports backup to removable partitioned media, such as the Iomega JAZ and ZIP drives.
- Dual device backup sets: PSnS allows the first backup to be on one type of media and subsequent incremental backups to be stored on a different type of media. This enables administrators to select the best backup solutions for their specific environment.
- Support for files greater than 2GB in size.
- Support for restoring backups taken with older versions of PSnS.

1.2.2.8 Advanced print services

The Advanced Print Services of OS/2 Warp Server for e-business are still based on Print Services Facility/2 (PSF/2). PSF/2 allows the administrator to define print conversion transforms for HP-PCL, PPDS, Postscript and 3270 host data streams. One major feature of Advanced Print Services is the ability

to print postscript output on non-postscript printers. With Advanced Print Services, it is also possible to print large PC-based jobs onto host printers with the Upload and Print function.

1.2.2.9 File systems

As you already know, OS/2 Warp Server Version 4 supports the FAT, HPFS and HPFS386 file systems, which have existed for several years now. OS/2 Warp Server for e-business supports these as well as a new file system called the Journaled File System (JFS). Let us describe each one briefly to refresh your memory, and then we will discuss JFS in more detail.

FAT

The File Allocation Table used a linear search mechanism to find files, which became very slow as hard disks grew in size and the number of files in a directory increased. FAT dates back to the early days of DOS and is still supported as a bootable file system primarily for backwards compatibility and data interchange. Installation of OS/2 Warp Server for e-business on a FAT partition is, generally, not recommended.

HPFS

The High Performance File System is another bootable file system that is supported in OS/2 Warp Server for e-business. HPFS is preferred over FAT because of the many enhancements in speed and reliability that were made. JFS (see Section 1.3.2.4 below) is faster than HPFS and offers better scalability, caching and shorter recovery times. Thus, it can replace HPFS in all applications except the boot partition.

HPFS386

HPFS386 contains a Ring 0 SMB server. This means that the transfer of data from the HPFS386 cache to the network adapter driver occurs much more quickly. HPFS386 is best used for file serving and for 802.2 Remote Initial Program Load (RIPL) of DOS, OS/2, and WorkSpace On-Demand RIPL clients. As of this writing, it is also the only choice if you need DASD limits or the Fault Tolerance Feature since (as of the writing of this redbook) JFS does not implement these. Before installing over a machine with HPFS386, you have to remove the HPFS386 DASD limits and access controls (ACLs) as well as the Fault Tolerance from the target volume.

JFS

JFS, a 32-bit file system, is especially suited for application servers, such as hosting the data of a Web server or a Lotus Notes server. JFS can be used to replace HPFS in most cases. It offers larger and faster caching capabilities and improves performance over HPFS. Currently, JFS is the only file system in OS/2 Warp Server for e-business that can be extended by adding more

partitions to a volume, thus, increasing available file space. JFS also offers better performance and scalability on SMP machines due to the changes in the I/O subsystem and features special optimizations for IP-based services. Similar to 386HPFS cache, JFS cache does not have a specific maximum; it is dependent on the amount of real memory installed on the system. The default is set to 1/8 of memory, but this default may be overridden by the Tuning Assistant (in the case of an attended installation). The memory used by JFS cache is allocated from the system arena and is non-swappable. JFS also includes the following capabilities:

- Unicode support
- Variable block size (512 - 4096 bytes)
- Extended filename support (254 characters)
- Online volume expansion
- Online defragmentation
- Sparse file support

1.2.2.10 Lightweight Directory Access Protocol (LDAP)

LDAP is a mechanism used for communicating with servers that provide global directory functions, sometimes called Yellow Pages after the telephone books that list all business in a particular area in alphabetical order. OS/2 Warp Server for e-business includes an LDAP client, which allows it to send and receive directory information from an LDAP server.

1.2.2.11 Web server

OS/2 Warp Server for e-business includes Lotus Domino Go Webserver 4.6.2.6. Lotus Domino Go Webserver is a scalable high-performance Web server that is easy to install and maintain. It includes state-of-the-art security, site indexing capabilities, and support for JDK 1.1.x. Lotus Domino Go Webserver makes it possible to maintain a productive Web presence in a diverse and dynamic environment.

1.2.2.12 Java application server

After you install Lotus Domino Go Webserver, you can add Java support by installing IBM WebSphere Application Server 1.1. IBM WebSphere Application Server is a plug-in for Lotus Domino Go Webserver and includes:

- A Java servlet engine that implements the JavaSoft Java Servlet API
- IBM additions and extensions to the Java Servlet API for enhanced session tracking and personalization
- Support for JavaServer Pages (JSP) is a powerful approach to dynamic Web pages
- A database connection manager for caching and reusing connections to JDBC-compliant databases

- Data access JavaBeans (additional Java classes for accessing JDBC-compliant databases)
- CORBA Support—an object request broker (ORB) and a set of services that are compliant with the Common Object Request Broker Architecture (CORBA)

1.3 A versatile server

As described in the previous section, OS/2 Warp Server for e-business supports a number of application services. This section describes just some of the more common environments where the server can be deployed:

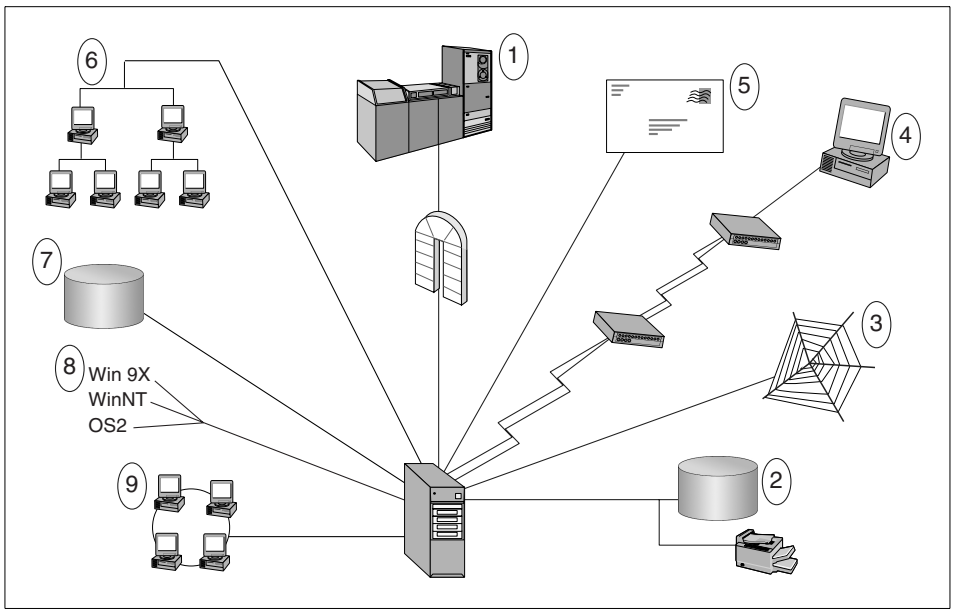

Figure 6. OS/2 Warp Server, deployment into a number of distinct environments

1. Communications Gateway

   Many enterprise environments have large mainframes, midrange computers, or minicomputers that store and process large amounts of mission-critical data. With additional Communications Server software, OS/2 Warp Server for e-business can act as a gateway to the many client machines that require access to corporate data-processing centers.

2. File and Print Server

In a more traditional role, OS/2 Warp Server for e-business is often used to share hardware resources, such as disk space and printing. Users can share data, applications, and network devices.

With the separately purchasable 386 HPFS, OS/2 Warp Server for e-business provides world-class file and print performance for large enterprise environments.

3. Web Servers

   OS/2 Warp Server for e-business comes with all of the necessary software to set up simple or highly complex web sites. It provides unparalleled performance and reliability at extremely low cost.

   The latest OS/2 Java implementation (Java 1.1.7 based) provides state-of-the-art Java performance integrated into OS/2. Significant enhancements to HTTP related APIs allow the server to handle a wider number of connections much faster.

   These improvements allow customers the potential to combine high-capacity, high-performance web-serving and Java serving.

4. Remote Access Server/ISP

   To enable mobile workers to gain access to corporate data or Internet Service Providers(ISP's) to set up a robust dial in infrastructure, OS/2 Warp Server for e-business comes with Remote Access Services. This feature offers users a wide variety of remote client access support and offers new levels of performance and security plus a low cost of ownership.

5. Mail Server

   With Lotus Notes, OS/2 Warp Server is able to support thousands of users without forsaking performance or reliability.

   Benchmarks indicate that OS/2 Warp Server with Notes provides a high performance server capability. It shows that this system provides outstanding price/performance and highly cost-effective solutions to customer's Notes server needs making it the optimum choice as a Notes Server platform for most customers.

6. Workgroup Systems Management

   As workers become more and more dependent on their computers, it is imperative that they are carefully managed. With Netfinity Manager and Client Services included, OS/2 Warp Server for e-business gives network administrators the visibility and capabilities they need to initiate and manage services for LAN-connected clients and servers. Netfinity for OS/2 is the ideal solution for departmental server/client management.

7. Database Server

   With DB2 Universal Database, OS/2 Warp Server for e-business extends into a powerful, reliable, Database server with universal client connectivity.

   DB2 Universal Database version 5.2 runs on Windows NT, Windows 95/98 (Personal Edition only), OS/2, AIX, Hewlett-Packard HP-UX, Sun Solaris, and Unixware 7. Client access is provided from these platforms plus DOS, Apple MacOS and Silicon Graphics IRIX, plus Web access with popular browsers and Java applications using DB2's native Java/JDBC support.

8. Server-Managed Clients

   OS/2 Warp Server for e-business is the premier platform for WorkSpace On-Demand, an IBM thin client offering. Customers who are deploying WorkSpace On-Demand have found that it reduced their client ownership costs, expedited deployment of new applications, and supported their transformation to the e-business application model while allowing them to maintain and use their existing OS/2 applications.

   WorkSpace On-Demand now supports a diverse set of clients including Windows NT workstation, Windows 98, Windows 95, and OS/2.

9. Collaborative Application Server

   Together with Domino, OS/2 Warp Server for e-business provides a complete and powerful platform for building secure collaborative Web applications. The Domino Application Server is an open secure platform optimized to deliver collaborative Web applications. These applications can be used to integrate your enterprise systems and allow for rapidly changing business processes.

1.4 System requirements

Prior to installing OS/2 Warp Server for e-business, make sure you have the required hardware and software and that your hardware is supported by OS/2 Warp Server for e-business.

You can find the most current information on supported hardware and additional device support on the Internet at the following Web site:

`http://www.software.ibm.com/os/warp/support`

The Personal Computer Manufacturer (PCM) table lists systems that have been tested and verified to work with OS/2. The PCM table is available at the following Web site:

`http://www.software.ibm.com/os/warp/hw-cert`

New hardware that becomes available after OS/2 Warp Server for e-business is installed can be supported with modifications to the diskettes used to start the installation process. You can find the most current information about additional device drivers on the Internet at the following Web address:

`http://service.software.ibm.com/os2ddpak/index.htm`

1.4.1 Hardware requirements

To install OS/2 Warp Server for e-business and use its services on the server, you need at least the following minimum hardware:

1. One or more Intel-compatible Pentium or higher processors with a speed of at least 133 MHz.

> **Note**
>
> A multiprocessor system must either comply with the Intel Multiprocessor Specification, Version 1.4 or 1.1, or it must be one of the following computers, each of which has its own proprietary SMP architecture:
>
> - Compaq Proliant 2000
> - Tricord PowerServer, models 30 and 40
> - IBM PC Server 720
>
> If this is not applicable to your system, your hardware manufacturer might have written its own *.PSD driver file to support OS/2 Warp Server for e-business SMP on his hardware. Check ahead of time to see whether such support is available.

2. You need a minimum of 32MB of random access memory (RAM), but 64MB or more provides better performance, depending on which services are installed.
3. A minimum of 120MB of available hard disk space for the base operating system. A minimum of 200MB is required for the base operating system and all default installation items. A total of 500MB is recommended for a typical installation depending on which services and components are installed. For installation requirements of services and components, see the table in section 1.4.2 “Hard Disk Space Requirements”.
4. A 1.44MB 3.5-inch diskette drive configured as drive A.
5. 640 x 480 (16-color) or higher resolution VGA display.
6. An IBM-compatible mouse.
7. A CD-ROM drive supported by OS/2.

8. A LAN adapter card supported by MPTS.
9. Remote Access Services requires the supported remote access adapters, which are a subset of the supported MPTS LAN adapters.
10. A modem that supports speeds of 9600bps or higher if you plan to use RemoteAccess Services.
11. An Internet-enabled LAN or a modem if you plan to use the Internet.

1.4.2 Hard space disk requirements

The requirements in this section are based on information available at publication time.

Table 1. OS/2 Warp Server for e-business, disk space requirements

Service	Hard Disk Space Requirements (MB)
OS/2 Base Operating System All optional OS/2 components	180.0
File and Print Sharing Services	15.0
TCP/IP Services	30.0
Remote Access Services	6.0
Netscape Communicator	11.0
Tivoli Management Agent	1.5
Personally Safe 'n' Sound	7.2
LDAP Services Toolkit	4.2
Advanced Print Services	54.0
First Failure Support Technology (FFST) 1.2	0.1
Online Books	10.0
Total (if all components and services are installed)	439.0

1.4.2.1 Keyboards supported during installation

For SBCS versions of OS/2 Warp Server for e-business, code page 850 is the only code page that is supported during the first phase of the installation process or if you start the system from utility diskettes. As a result, you must select one of the Latin-1 keyboards even if you normally select a non-Latin-1

keyboard. You can still set the country code to any valid country. This setting may affect the default country and keyboard settings used later in the installation process, including those used by the Logical Volume Management Tool (LVM), and for command line processing when the system is booted from utility diskettes. However, later in the installation, you can specify your preferred country and keyboard settings.

1.5 Replaced or discontinued components

The following components, previously installed with OS/2 Warp or OS/2 Warp Server, are replaced in OS/2 Warp Server for e-business:

1. The Fixed Disk Utility (FDISK.COM) is replaced by the Logical Volume Management Tool (LVM.EXE). The Fixed Disk Presentation Manager Utility (FDISKPM.EXE) is replaced by Logical Volume Manager Graphical User Interface (LVMGUI.CMD).
2. Java 1.0 is replaced by Java 1.1.7.
3. Pulse is replaced by CPU Monitor (CPUMON).
4. LAN Distance is superseded by Remote Access Services.
5. SystemView Agent (Netfinity TME 10) is replaced by Tivoli Management Agent (TMA).

The following components previously installed with OS/2 Warp or OS/2 Warp Server are not part of OS/2 Warp Server for e-business:

- Password Coordinator.
- Network Signon Coordinator.
- BonusPakOS/2 Warp Tutorial.
- OpenDoc.
- WarpGuide.
- VoiceType.
- Hibernate (Trapdoor).
- Novell NetWare Client for OS/2.
- Dual boot is not supported.
- Easy Path installation.
- Mobile File Sync.
- PCOMM Lite 4.1.
- Keyworks.

- Remote client installation is not supported.

1.6 Server packaging and licensing

As of this writing, OS/2 Warp Server for e-business is comprised of eight CD-ROMs. In addition to the CD-ROMs described below, the product ships with a bootable CD and a Quick Beginnings hard copy manual for the specific language version purchased. This section describes the contents of each CD-ROM.

CD-ROM #1 This is the bootable CD. Should your BIOS not support boot from CD-ROM, you can still create and use the diskettes to boot.

CD-ROM#2 This contains the server code including the OS/2 base, the uniprocessor and symmetric multiprocessor kernels, OS/2 LAN Server functions, MPTS, TCP/IP, Remote Access Services, Backup and Recovery Services, Tivoli Management Agent, Java Version 1.1.7, Netscape Communicator, and online documentation. This CD-ROM comes in the specific language version that was purchased.

CD-ROM #3-5 This contains the Client Connect Pak. This is the client function for OS/2 Warp 4, Windows 95/98, Windows NT Version 4 and Windows 3.1. The functions included (not necessarily for all platforms) are File and Print, TCP/IP, Remote Access, Tivoli Management Agent, and Java Runtime. These functions are provided in all supported languages.

CD-ROM #6 This contains the Netfinity Version 5.2 component. This function is provided in the specific language version that was purchased.

CD-ROM #7 This contains the Lotus Domino Go Web Server Version 4.6.2 and the IBM WebSphere Application Server Version 1.1. This CD-ROM is provided in the specific language version that was purchased.

CD-ROM #8 This is the security CD and contains either 56 bit or 128 bit encryption for use within the USA. This CD-ROM is separately purchasable.

1.6.1 386 HPFS licensing

386 HPFS is available as a separate product. For customers who currently have OS/2 Warp Server advanced, this is a no-charge upgrade. For

customers who have made a new purchase of OS/2 Warp Server for e-business or customers who previously had OS/2 Warp Server Entry and wish to use 386 HPFS, it will need to be purchased separately.

1.7 The installation process

OS/2 Warp Server for e-business and the many optional components can be installed using either an attended or an unattended installation procedure. If you are installing just a few servers, you will probably choose to use the attended installation process. This is covered in great detail in the online documentation, *Quick Beginnings: Installing OS/2 Warp Server for e-business*, and will not be covered here.

In each of the chapters that follow, a description of the installation of that particular component may be included.

Sample response files are also provided with the product and the optional components for unattended installation.

Two significant changes from the previous version are:

- The easy installation path has been removed. The new attended installation takes you through what was previously the advanced install path.
- The product now ships with a bootable CD, which can be used in place of the two or three boot diskettes that came with OS/2 Warp 4.0 and OS/2 Warp Server.

1.7.1 The bootable CD

OS/2 Warp Server now ships with a bootable CD. This speeds up the boot and install processes. The CD, however, may not work for some systems. This may be for one of the following reasons:

- Some older hardware does not support boot from CD-ROM.
- The CD contains support for a limited number of devices. If the controller to which either your Hard File or CD-ROM drive is attached is not recognized, you will need to create diskettes.

  These diskettes can then be modified to add the necessary device drivers.

Chapter 2. Base operating system enhancements

The base operating system for OS/2 Warp Server for e-business is an advanced multitasking, 32-bit operating system that runs DOS, Windows, and OS/2 16 and 32 bit applications and utilizes SMP hardware configurations. OS/2 Warp Server version 4.0 was based on the OS/2 Warp Connect operating system. OS/2 Warp Server for e-business is based on a modified OS/2 Warp 4.0 operating platform.

Many of the components that operate on upper layers make use of the essential services provided by the base operating system. Many changes have been made to make OS/2 Warp Server for e-business a premier operating platform. Some changes are discussed here, and others are covered in the chapters that follow together with the service that they enhance.

Many new features and enhancements have been incorporated into this version of the OS/2 base operating system. All of the service elements that have been changed since OS/2 Warp 4.0 was released as well as those incorporated into OS/2 Warp Server and OS/2 Warp Server SMP have been integrated into this version of the base operating system.

In this chapter, we will give a brief introduction to the new I_2O driver model, talk about the SMP enhancements, provide information about the large file support, and the 1024 cylinder boot restriction. We will also discuss the new version of Netscape, Java 1.1.7, the global locale builder, and Euro support.

2.1 Virtual memory support

The virtual memory map of the system has been changed, which may allow up to an additional 32MB of virtual memory. This OS/2 High Memory area support, introduced in OS/2 Warp Server SMP, is enhanced in this release and available on both the uniprocessor and SMP versions of the operating system and allows pure 32-bit OS/2 applications the potential to address up to 3GB of virtual memory.

2.2 Intelligent input and output I_2O

OS/2 Warp Server for e-business provides support for Intelligent Input/Output (I_2O) based LAN and SCSI adapters. I_2O is a specification that defines a standard architecture for intelligent I/O that is independent of the specific device and host operating system. It defines an approach where the interrupt

processing tasks are off-loaded from CPU to I/O processors designed specifically to handle I/O, thus, relieving the CPU of I/O tasks. I_2O makes it easier for developers to implement cross-platform intelligent I/O solutions.

2.2.1 I_2O architecture

In a networking environment, a significant amount of CPU time is used to perform I/O tasks. A significant amount of time, money and effort is also spent on developing device drivers on multiple platforms for hardware devices manufactured by different hardware vendors. The I_2O specification is an attempt to standardize the device driver development to make it easier to develop cross-platform intelligent I/O solutions.

Intelligent Input/Output (I_2O) is a specification that provides an open architecture for the development of device drivers for intelligent I/O. It is independent of the operating system and attempts to standardize the driver development by transferring the interrupt-intensive tasks from the CPU to I/O processors (IOPs). The IOPs are designed specifically to handle I/O transactions. The I_2O specification is intended to provide a standard approach to driver development that makes it easier for the developer to write a device driver for a class of adapters (like LAN disk adapters).

Off-loading the interrupt-intensive tasks from the CPU to the IOPs is achieved by implementing a split driver model. According to the I_2O specification, this split driver model consists of two parts: The first part interfaces with the host OS, and the second part resides and interfaces with the adapter. The two parts communicate through a communications layer. The first part is known as the Operating System Service Module (OSM) and the second part is known as the Hardware Device Module (HDM). The OSM and the HDM communicate through the communication system. This communication system consists of two layers. The Message layer, which sets up the communication and a Transport layer, which defines how information will be exchanged.

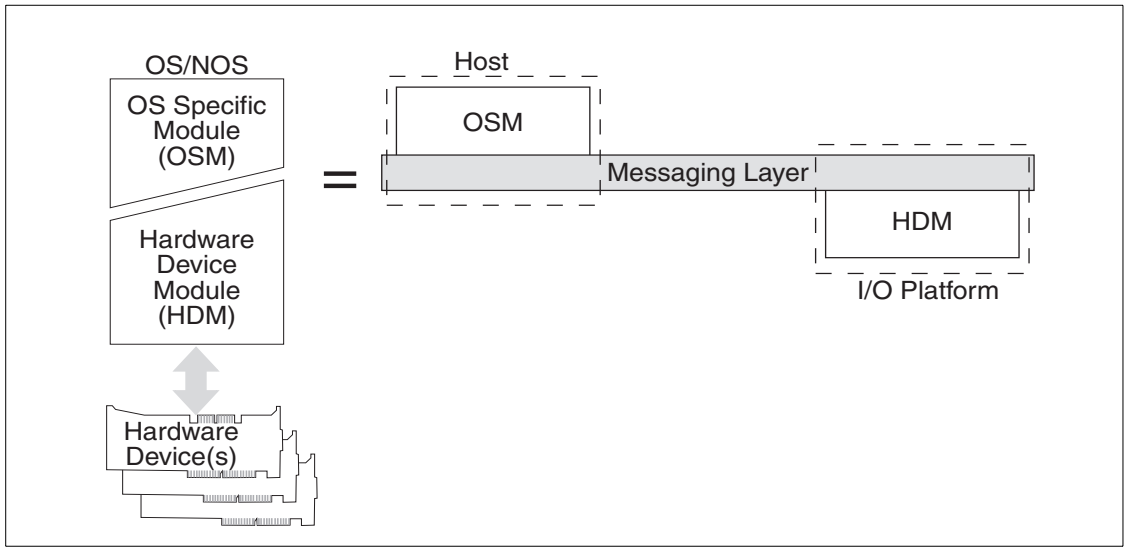

Figure 7. I_20 architecture.

2.2.1.1 Operating System Service Module

The OSM is the part of the driver that sits on the operating system and interfaces with the operating system. The OSM interfaces with the I_2O communication layer on the lower end. The requests from the OS are translated by OSM into messages and are sent to the HDM through the message layer. The HDM communicates the results back to the OSM through the message layer. The OSM then translates these messages and passes them to the operating system. The OSM is the interface between the operating system and the device. Operating system vendors develop OSMs.

2.2.1.2 Hardware Device Module

The Hardware Device Module (HDM) which resides on the adapter is the device-specific portion of the driver. It interacts with the OSM through the Message layer. HDMs receive I/O requests in the form of messages from the OSM through the Message layer. The HDM processes the request and communicates the results back to OSM using the Message layer. The HDMs are developed by the hardware device vendor.

2.2.1.3 Communications system

The communication system of the I_2O consists of two layers.

1. Transport layer:

   It defines how the OSM and HDM exchange information

2. Message layer:

   It sets up the communication between OSM and HDM. The Message layer is the backbone for I_2O. It handles all the communication that takes place between OSM and HDM. It handles the receiving and dispatching of messages and provides service routines for processing the messages. The messages consist of two parts: A header and a payload. The header consists of the type of service and the return address of the requestor. There are three components in the message layer:

   1. Message handle
   2. Message queue
   3. Message Service Routine

When a request is made, a message is deposited in the message queue and a message service routine is activated to process the request. A message handle, which is the address of the service routine, is returned to the requester.

2.2.1.4 I_2O LAN OSM implementation

The LAN I_2O OS Service Module (OSM) supports network adapters attached to the OS/2 system via I20 as specified in the Intelligent I/O Architecture Specification dated November 1996, Draft Revision 1.5a adopted by the SIG Membership March 21,1997.

The specific implementation for OS/2 is as a Network MAC driver conforming to the 3COM/Microsoft LAN Manager Network Driver Interface Specification (NDIS) version 2.0.1. It utilizes the operating system message services developed for the I_2O storage and SCSI peripheral (I_2OEXPORT.SYS).

The I_2O architecture specifies a split driver module that allows the operating system vendor to develop an Operating System Service Module(OSM).

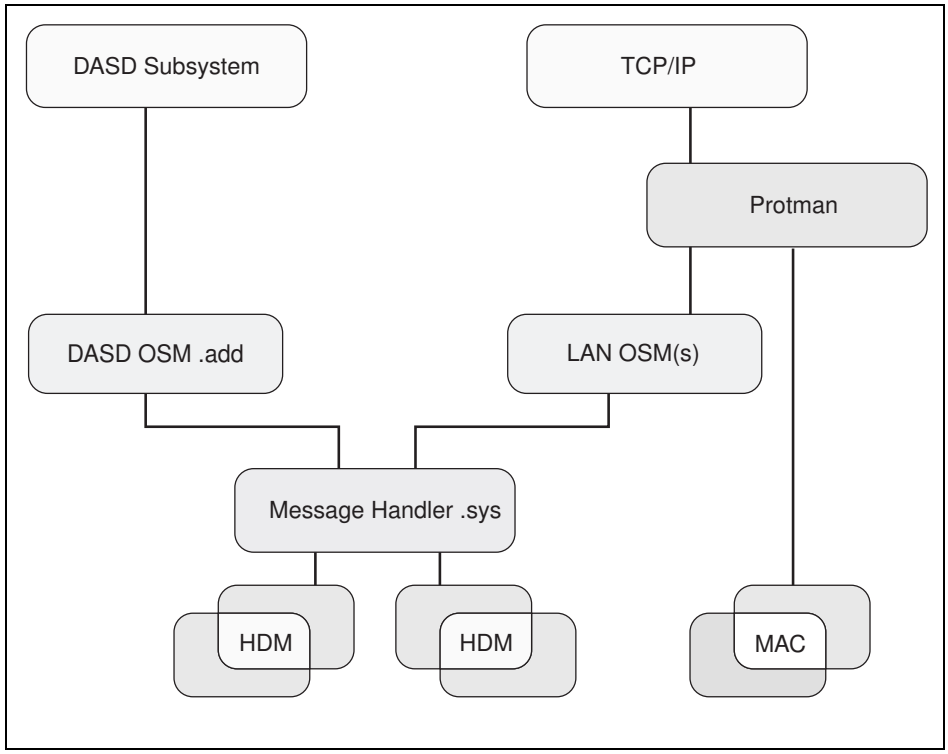

Figure 8. Flow of information with an I_2O adapter.

The normal flow of information from TCP/IP utilizing Protman remains in effect with the addition of the ability of Protman to recognize the OSM as a MAC driver.

The LAN OSM will communicate with the rest of the OS using the NDIS 2.0.1 Interface and will communicate with the HDM using the LAN I_2O interface through the Message Handler**.**

The I_2O LAN OSM utilizes the message handler developed for I_2O SCSI (I2OXPORT.SYS) and the I2OCALLS (I2OCALLS.LIB) developed for that subsystem. At system boot time, I2OXPORT.SYS detects IOPs on the system and initializes each IOP accordingly. The LAN OSM loads after I2OXPORT.SYS and calls it to find out each IOP's ID. The LAN OSM then calls each IOP to find out its configuration and registers with the kernel; so, it can send I/O Request Blocks. The LAN OSM allocates memory to contain the formatted I_2O message. The I_2O message handler copies the I_2O message to shared memory (shared between the processor and the IOP).

2.2.1.5 Hardware support

The I_2O LAN OSM(s) and Message Handler requires an I_2O LAN Adapter. This is a LAN adapter with an Intelligent I/O processor on the adapter which will host Hardware Device Module(HDM) software which will interface to the LAN OSM(s) via the Message Handler.

In general, four adapters per server is considered the maximum number supported. This is for practical reasons of system capability with reasonable performance, however, there is no actual system constraint which prohibits more than four.

2.2.1.6 Error logging

If the LAN SOM fails to load at boot time, it will log a message at boot time; it will log a message to the screen during processing of CONFIG.SYS, and the user must hit ENTER to continue.

The LAN OSM does not log messages to lantran.log; however if a protocol fails to bind with the I_2O LAN adapter, the protocol will log an error message to lantran.log.

2.2.1.7 Installation and configuration issues

Installation of the LAN I_2O OSM is the same as for any other LAN driver. In our example, we added a pre-production Soliton EtherPUMP adapter into a Netfinity 5500 server. To activate the adapter, we did the following:

1. Start MPTS by typing `MPTS` at an OS/2 command prompt and pressing **Enter**.
2. Select the configure button, and you will get the panel displayed in Figure 9 on page 29.

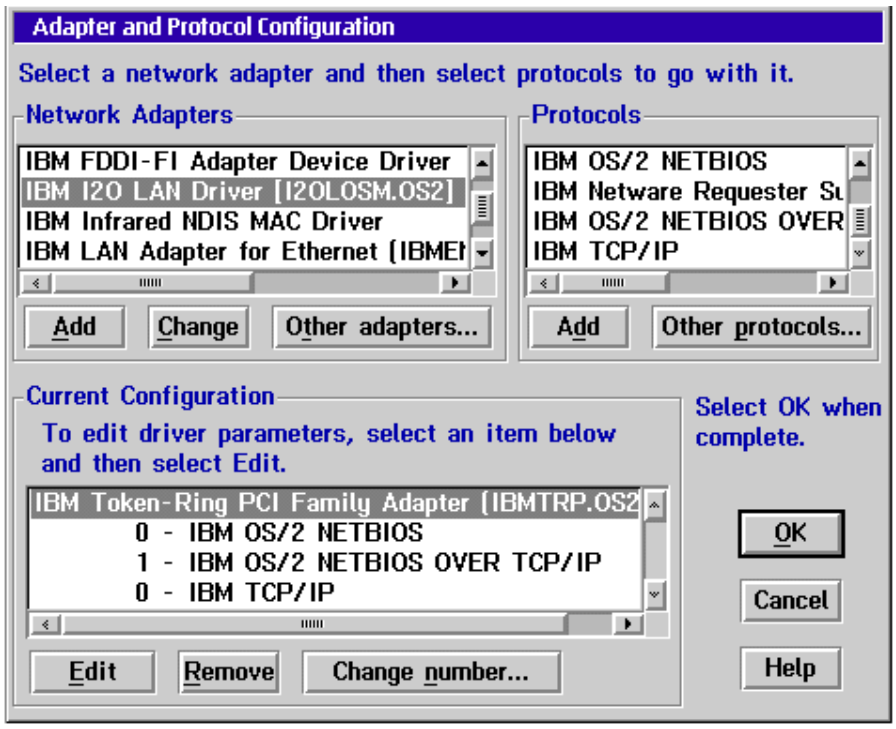

Figure 9. Adapter and protocol configuration, selecting the I_2O LAN driver.

3. In the Network Adapters window in the top left corner, select the **IBM I20 LAN Driver**; then, click on the **Add** button.
4. This will add it to the current configuration window in the bottom left corner of the panel. Select the **adapter** in the current configuration window and, then, select a **Protocol** from the top right window in the panel.
5. Select **Add** to add the protocol to the adapter. In the example, we added the **NetBIOS** protocol to the adapter. The configuration is shown in Figure 10 on page 30

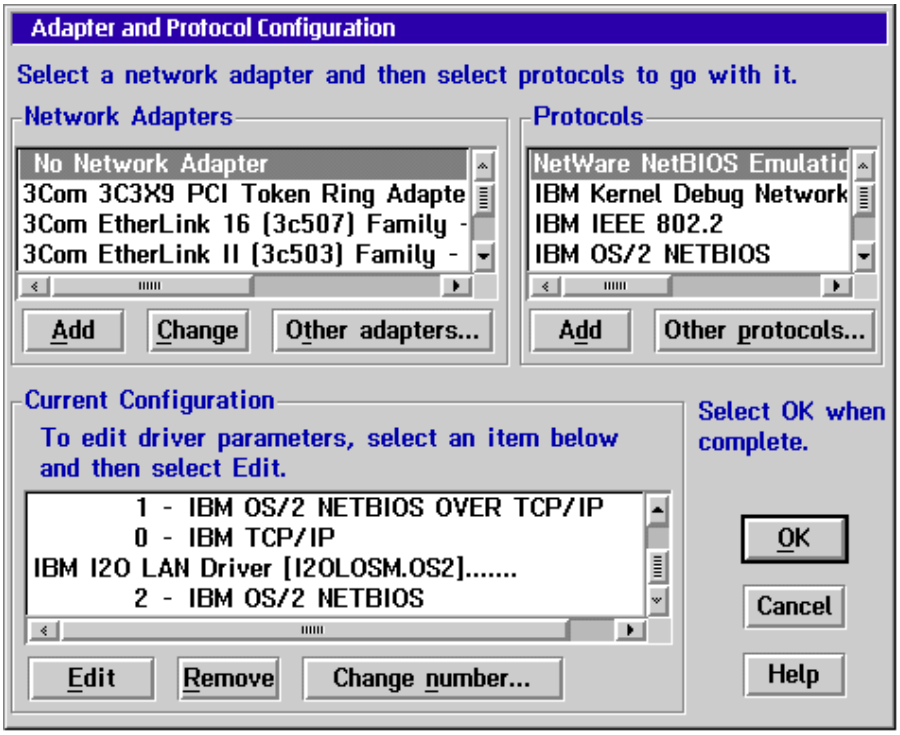

Figure 10. Adapter and protocol, I_2O configuration

6. Select **OK** to commit the configuration. A confirmation window appears allowing you to select which CONFIG.SYS files you wish to update. Note that if you have multiple partitions that have multiple installed copies, they will all be listed here if they are visible to this partition. Decide carefully if you do not want to unnecessarily update other partitions.

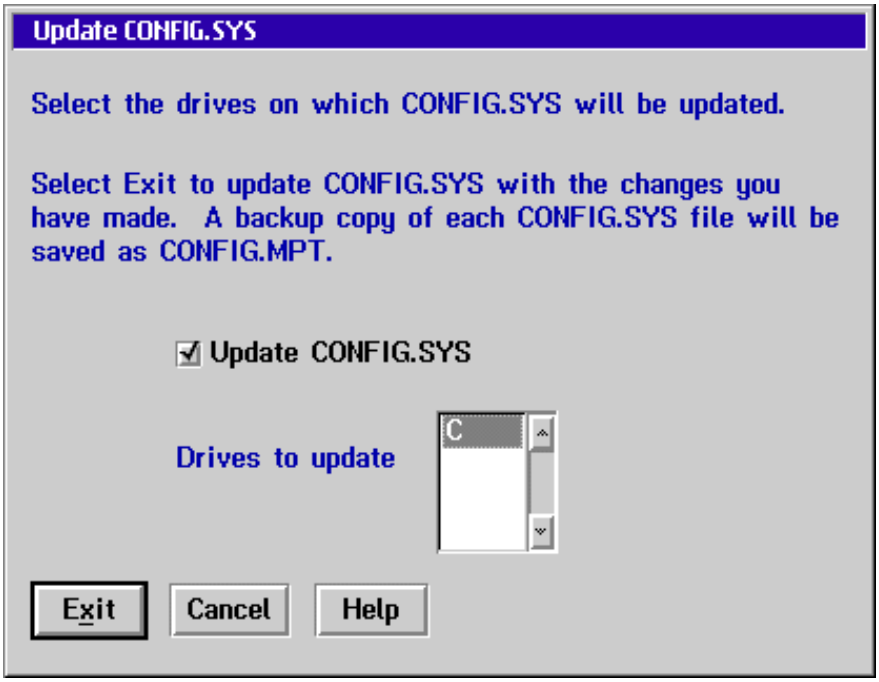

Figure 11. Adapter and protocol configuration, CONFIG.SYS update

7. Select **Exit** to finally update the configuration. You will need to reboot the machine to have the drivers take effect.

IBM Lock File Device Driver

MPTS configuration makes use of the lock file device driver. Files that need to be replaced and are loaded at system initialization time are replaced using the lock file device driver.

If, at any time after rebooting, you are unable to re-enter the MPTS configuration, remark out or delete the lock file device driver from your CONFIG.SYS and rename or delete the file x:\OS2\INSTALL\IBMLANLK.LST.

8. By adding the adapter, both your CONFIG.SYS and PROTOCOL.INI file have been changed. The LANTRAN.LOG file contains a listing for the LAN initialization. The LANTRAN.LOG reported the following after installing the I_2O adapter:

```
LT00073: FFST/2 is installed but is not started. LANTRAN.LOG is being created.
IBM OS/2 LANMSGDD [04/07/99] 5.05 is loaded and operational.
IBM OS/2 NETBEUI 5.50.0
NETBEUI: Using a 32-bit data segment.
IBM OS/2 TCPBEUI 5.50.0
```

```
TCPBEUI: Using a 32-bit data segment.
Installing NETWKSTA.200 Version 6.0. IBM LAN Redirector ( Apr 12, 1999)

IBM OS/2 NETBIOS 4.0
Adapter 0 has 30 NCBs, 152 sessions, and 27 names available to NETBIOS applications.
Adapter 1 has 123 NCBs, 28 sessions, and 6 names available to NETBIOS applications.
Adapter 2 has 252 NCBs, 254 sessions, and 41 names available to NETBIOS applications.
NETBIOS 4.0 is loaded and operational.
IBM OS/2 Warp 3.0 and above NDIS 2.0 IBM Token-Ring PCI Family Adapter, Version 2.00
IBM OS/2 LAN Netbind
Port 00000000: The IBM Token-Ring PCI Family Adapter UAA (BIA) is 00203503cb33.
Port 00000000: The IBM Token-Ring PCI Family Adapter opened at 16 Mbps, half duplex.
TCPBEUI: completed initialization and is bound to TCP/IP network interface lan0.
Adapter 0 is using node address 00203503CB33.
```

Adapter 2, in the above listing, is the I20 adapter.

2.3 1024 cylinder restriction

When the operating system or an application wants to access a hard disk in real-address mode, it normally does it via the BIOS (BASIC Input/Output System) via the INT 13h services. The BIOS provides a standard interface to the hard disk. This INT 13h interface allows the operating system and applications to access the hard disk without the need for the applications to understand the complexities of accessing the hard disk. However, the standard INT 13h interface does not allow the operating system to boot from any point on the hard disk and limits the capacity that can be addressed to 8.4GB.

OS/2 Warp Server for e-business overcomes this limitation. Thereby, allowing you to install and boot the operating software on any point on the hard drive and recognizing drives that are greater than 8.4GB in capacity.

2.3.1 The INT 13h interface

To understand the 1024 Cylinder restriction, it is important to know how the data is accessed on the hard disk drives. To read data from the hard disk drives, the location of the data on the disk platters has to be specified. The location on the disk is expressed in terms of drive geometry. The geometry of the drive is normally expressed in terms of the cylinder, head, and sector that the system wants the drive to read. A request is sent to the drive over the disk drive interface giving it this address and asking for the sector to be read. The geometry of the drive is specified using the INT 13h interface.

The INT 13h supports many different commands, such as disk read, write, and so on, that are then submitted by the BIOS to the hard disk directly. When a command is passed on to the hard disk, the head, cylinder, and

sector number have to be provided. The head, cylinder, and sector addressing is called drive geometry. The drive geometry is passed on in registers, and INT 13h provides 24 bits to specify the drive geometry. The 24 bits are broken as follows:

10 bits for the cylinder number (Maximum of 1024)

8 bits for the head number (Maximum of 256 heads)

6 bits for the sector number (Maximum of 63 sectors; sectors are numbered from 1, not 0)

The sector size is a standard 512 bytes.

The default mode used for disk access is the CHS mode. CHS stands for Cylinder, Head, Sector. This interface, therefore, restricts the number of sectors (per track) to 63, the number of heads to 256, and the number of cylinders to 1024 for a total of 8.4 gigabytes (63*256*1024*512=8.4GB).

But, the hardware specification of the ATA drives (IDE drives), limits the number of sectors/tracks to 255, heads to 16, and cylinders to 65536. This limits the device to 136.9 GB (255 * 16 * 65536 * 512 =136.9 GB). If a drive is to be accessed through BIOS, the smaller value must be used as the limit. 16 heads supported by ATA drives is a smaller value. Considering this, the access is now restricted to 528 MB (63 * 16 * 1024 * 512 = 528 MB).

Table 2. Limitations of various interfaces

	BIOS	ATA	Limit
Max Sectors/Track	63	255	63
Max Heads	256	16	16
Max Cylinders	1024	65,536	1024
Capacity	8.4 GB	136.9 GB	528 MB

The above limitations can be overcome in two ways:

1. Create a false geometry that fits within the Int 13h limitations and also uses the full capacity of the drive. This is called drive translation. There is a problem with drive translation. Since there is no single standard for drive translation, several independent translation method exists. If a drive is moved from one computer to a different computer, the data on the drive

will not be recognized unless both computers use the same translation algorithm.

2. Address the drive using the enhanced INT 13h extensions. The enhanced INT 13h interface, in addition to overcoming the 528MB limit, also overcomes the 8.4 GB limitation problem faced when a conventional INT 13h interface is used.

   The new enhanced INT 13h services are fundamentally different from the existing INT 13h services in the following ways:

   - They support the passing of data structures and are not register based
   - Buffers are used to pass on the media address information
   - Flags are used to identify optional capabilities

2.3.2 OS/2 boot process

The OS/2 boot process within OS/2 Warp Server for e-business now uses the enhanced INT 13h services to access the hard disk. With the 1024 cylinder limitation, the user was restricted to boot from partitions that were contained within the first 1024 cylinders. This limitation has now been removed. This section explains the boot process that gives an idea about where all the INT 13h service is employed during the boot process:

The OS/2 boot process uses the INT 13h services to access the hard disk before all of the base device drivers are loaded.

2.3.2.1 FAT and HPFS boot flow

The first six steps are common to both FAT and HPFS boot because they are file system independent flow.

2.3.2.2 FAT boot flow

The following figure illustrates the FAT boot flow:

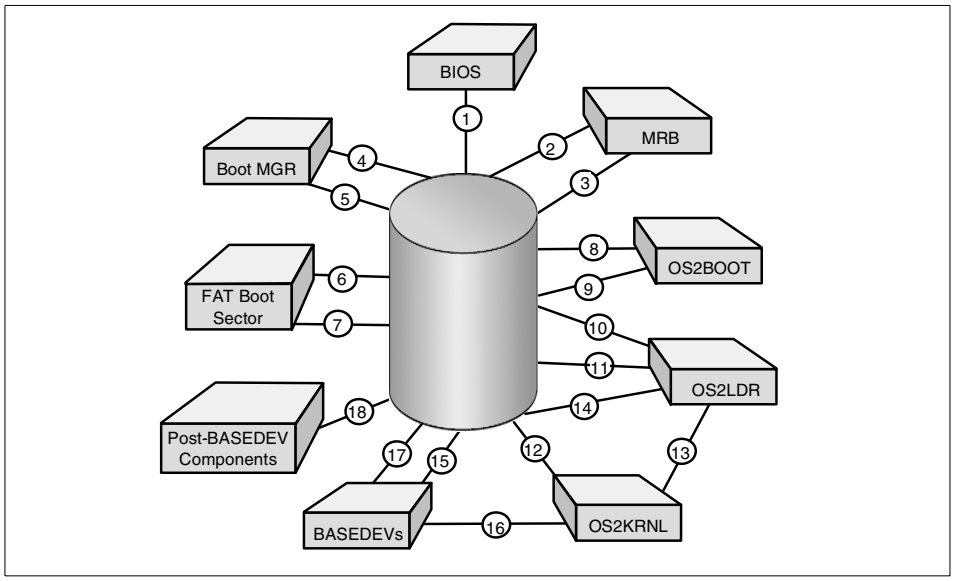

Figure 12. FAT boot process

The steps involved in the FAT boot process are listed below:

1. System BIOS issues a request via BIOS INT 13h services to the hard disk to read the Master Boot Record (MBR).
2. The MBR is loaded into the memory. The MBR inspects the partition table to determine whether Boot Manager or an active bootable partition should be loaded.
3. The MBR issues a request to the hard disk via the INT 13h services to read the boot sector for the selected partition into memory.
4. If the selected partition was Boot Manager, the Boot Manager boot sector is loaded into memory.
5. If the selected partition was Boot Manager, Boot Manager will determine which partition should be booted and will issue a request to a hard disk via BIOS INT 13h services to read the boot sector for the bootable partition into memory.
6. The boot partition for the bootable partition is loaded into memory.
7. The FAT Boot sector determines the location of OS2BOOT and issues a request to the hard disk via BIOS Int 13h services to read OS2BOOT into memory.
8. OS2BOOT is loaded into memory and determines the location of OS2LDR.

9. OS2BOOT issues a request to the hard disk via BIOS INT 13h services to read OS2KRNL into memory.
10. OS2LDR is loaded into memory and determines the location of OS2KRNL.
11. OS2LDR issues requests to the hard disk via BIOS INT 13h services to read OS2KRNL into memory.
12. OS2KRNL is loaded into memory.
13. OS2KRNL issues requests to OS2LDR to read various BASEDEVs into memory.
14. OS2LDR issues requests to the hard disk via BIOS INT 13h services to read various BASEDEVs into memory.
15. BASEDEVs are loaded into memory.
16. After all BASEDEVs have been loaded and initialized, any further read requests are issued to the BASEDEVs.
17. The BASEDEVs issue read requests to the hard disk via direct hardware access.
18. The requested information is loaded into memory.

2.3.2.3 HPFS boot flow

The following figure illustrates HPFS boot flow:

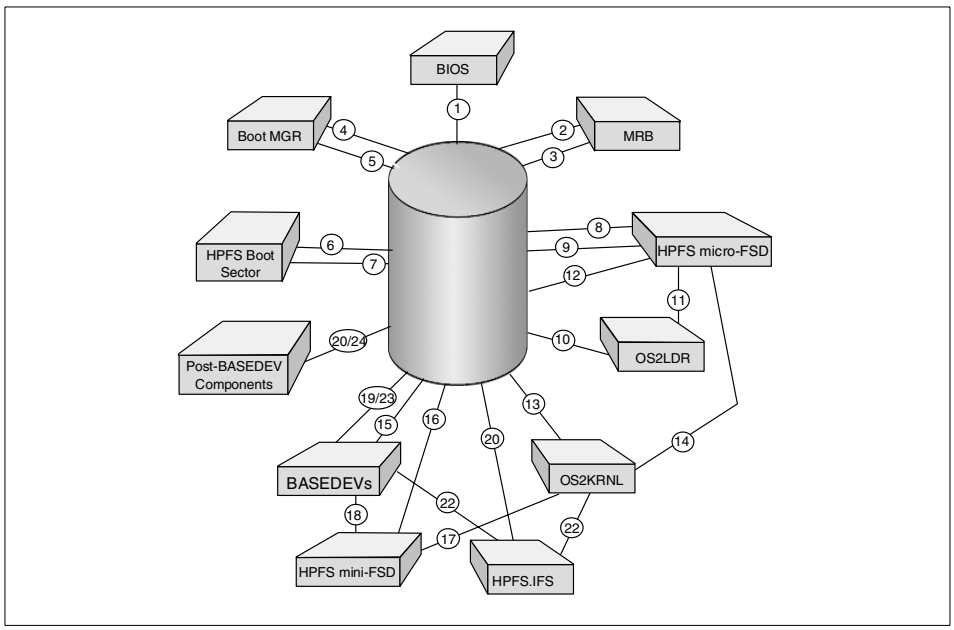

Figure 13. HPFS boot process

The steps involved in the HPFS boot process are listed below:

1. System BIOS issues a request to the hard disk via BIOS INT 13h services to read the Master Boot Record (MBR).
2. The MBR is loaded into the memory. The MBR inspects the partition table to determine whether Boot Manager or an active bootable partition should be loaded.
3. The MBR issues a request to the hard disk via INT 13h services to read the boot sector for the selected partition into memory.
4. If the selected partition was Boot Manager, the Boot Manager boot sector is loaded into memory.
5. If the selected partition was Boot Manager, Boot Manager will determine which partition should be booted and will issue a request to a hard disk via BIOS INT 13h services to read the boot sector for the bootable partition into memory.
6. The boot partition for the bootable partition is loaded into memory.
7. The HPFS Boot sector determines the location of the HPFS micro-FSD and issues a request to the hard disk via BIOS INT 13h services to read the HPFS micro-FSD into memory.
8. The HPFS micro-FSD is loaded into memory and determines the location of OS2LDR.
9. The HPFS micro-FSD issues a request to the hard disk via BIOS INT 13h services to read OS2LDR into memory.
10. OS2LDR is loaded into memory.
11. OS2LDR issues a request to the HPFS micro-FSD to read OS2KRNL into memory.
12. The HPFS micro-FSD determines the location of the OS2KRNL and issues a request to the hard disk via the INT 13h services to read OS2KRNL into memory.
13. OS2KRNL is loaded into memory.
14. OS2KRNL issues a request to HPFS micro-FSD to read BASEDEVs into memory.
15. The HPFS micro-FSD issues requests to the hard disk via the INT 13h services to read various BASEDEVs and the HPFS mini-FSD into memory.
16. BASEDEVs and HPFS mini-FSD are loaded into memory.
17. After all BASEDEVs and HPFS mini-FSD are loaded and initialized, read requests are issued to the HPFS mini-FSD to load additional system components, including HPFS.IFS, into memory.
18. The HPFS mini-FSD determines the location of these components and issues read requests to the BASEDEVs.
19. The BASEDEVs issue read requests to the hard disk through direct hardware access.

20. The additional system components and HPFS.IFS are loaded into memory.
21. After HPFS.IFS is loaded and initialized, any further read requests will come from OS2KRNL to HPFS.IFS
22. HPFS.IFS issues a read request to the BASEDEVs.
23. The BASEDEVs issue read requests to the hard disks through direct hardware access.
24. The requested information is loaded into memory.

FDISK currently prohibits a user from setting a partition as installable or startable if it is not wholly contained within the first 1024 cylinders. LVM will allow installable and startable partitions of up to 1024 cylinders unless the BIOS supports large hardfiles. If BIOS supports large hardfiles, a partition in any location may be set installable or startable and added to the Boot Manager menu.

2.3.2.4 Dependencies

In order for this support to be activated, the machine BIOS must support the Enhanced INT 13h interface. Also, the BIOS must support drives with more than 1024 cylinders (capacity greater than 8.4 GB) if this support is to be activated above 1024 cylinders.

2.4 Kernel enhancements and file subsystems

This section describes enhancements that have been made to kernel and file subsystem environment. Changes have been made to the Installable File System Manager (IFSM) and the Kernel Execution Environment(KEE) in OS/2 Warp Server for e-business to improve the performance and SMP scalability of the file system (JFS/LVM/OS2DASD) and the communications stack (TCP/IP).

An installable File System (IFS) is a dynamically loadable module. It is loaded by using a CONFIG.SYS command, *IFS=...* HPFS, HPFS386, and JFS are examples of an IFS. The IFS handles all file system requests for its file system. The IFS is also referred to as a File System Driver (FSD). The Installable File System Mechanism (IFSM) defines the relationship between an IFS and the OS/2 kernel File System Router. The File System Router is responsible for routing file system requests to the appropriate entry point in the IFS.

The Kernel Execution Environment is a set of new, 32-bit, SMP-enabled, subsystem API's.

The performance enhancement is achieved by converting part of the existing kernel execution environment from 16 bit to 32 bit.

Poor SMP scaling is also due to the fact that the OS/2 kernel must provide serialization for the legacy device drivers on SMP since they are not SMP enabled. The OS/2 SMP kernel provides the serialization that protects the device driver's data structures from data corruption (concurrent access by multiple threads on SMP) by using a single system-wide spinlock. This spin lock is known as the SUBSYS spin lock. The SUBSYS spinlock is acquired by the kernel on behalf of the device driver every time the kernel calls into the device driver. Because a single spin lock is employed, only one processor can be executing code in a device driver at any one point in time. This leads to poor SMP scaling when there are multiple threads needing to execute in the device driver. The other threads must wait (spin) until the owning processor releases ownership of the SUBSYS spin lock.

To correct the scaling problem, the device drivers will use the new 32-bit SMP-enabled Kernel Execution Environment (KEE). KEE is a set of 32-bit SMP-enabled APIs. KEE provides a set of system service APIs that will enable data structures with more granular spin locks instead of using the system-wide SUBSYS spin lock. This improves the SMP scalability. In addition, the KEE APIs require fewer CPU cycles than the 16-bit device helper APIs; so, performance improves on both UNI and SMP.

Some disadvantages of a 16-bit code path are:

- 16 bit code runs slower than 32 bit code.
- The file system router is a 16-bit router that limits the size of a DosRead and DosWrite to a value less than 64K. Therefore, an application request larger than 64K is broken into multiple calls to the IFS.
- The OS/2 SMP kernel will serialize the execution of all code in an installable file system (HPFS) and device driver (OS2DASD, *.ADD) using a single system-wide spinlock (SUBSYS spinlock). In the 16-bit code path, this spin lock is held over most of the code path and results in poor SMP scalability.
- The kernel stack selector is 16-bit. Therefore, overhead is incurred when an address of a local variable is required to be passed to a function.
- The system services, device helper, and file system helper are 16-bit. The device helper services use a router mechanism. The overhead of the device help router is significant.

The new 32-bit code path provides the following improvements:

- Portions of the kernel File System Router layer, JFS, LVM, and OS2DASD (IO path) are 32-bit resulting in improved execution time.

- Portions of the File System Router are 32-bit; there is no limit on the 64K size for an I/O call like DosRead or DosWrite.
- The kernel File System Router will not acquire the SUBSYS spinlock prior to calling the 32-bit JFS entry points for DosRead, DosWrite, and DosChgFilePtr. JFS/LVM/OS2DASD use the new KEE spinlock APIs to provide more granular locking of its data structures resulting in improved SMP scalability. The SUBSYS spinlock will be acquired prior to calling the SCSI ADD and represents only a small percentage of the total I/O path.
- The system services (KEE calls) are 32-bit and do not use a router. The code path for a call is significantly reduced.

2.5 SMP support and microprocessor affinity

Previously, SMP support was provided in a separate OS/2 Warp Server Advanced feature. Now, support for up to 64 processors is included in one package. The OS/2 Warp kernel in OS/2 Warp Server for e-business as well as the file systems have been further tuned to improve processing scalability.

Since OS/2 Warp Server for e-business comes with both a kernel for machines with one processor and an SMP enhanced Version, you now have the option to start on a Server with only one processor installed. If your requirements outgrow this, you can add the SMP Support later using selective install. A multiprocessor system must either comply with the Intel Multiprocessor Specification, Version 1.4 or 1.1, or it must be one of the following computers, each of which has its own proprietary SMP architecture:

- Compaq Proliant 2000
- Tricord PowerServer, models 30 and 40
- IBM PC Server 720

During installation, you will be allowed to select Multiprocessor Support. This is shown in Figure 14 on page 41.

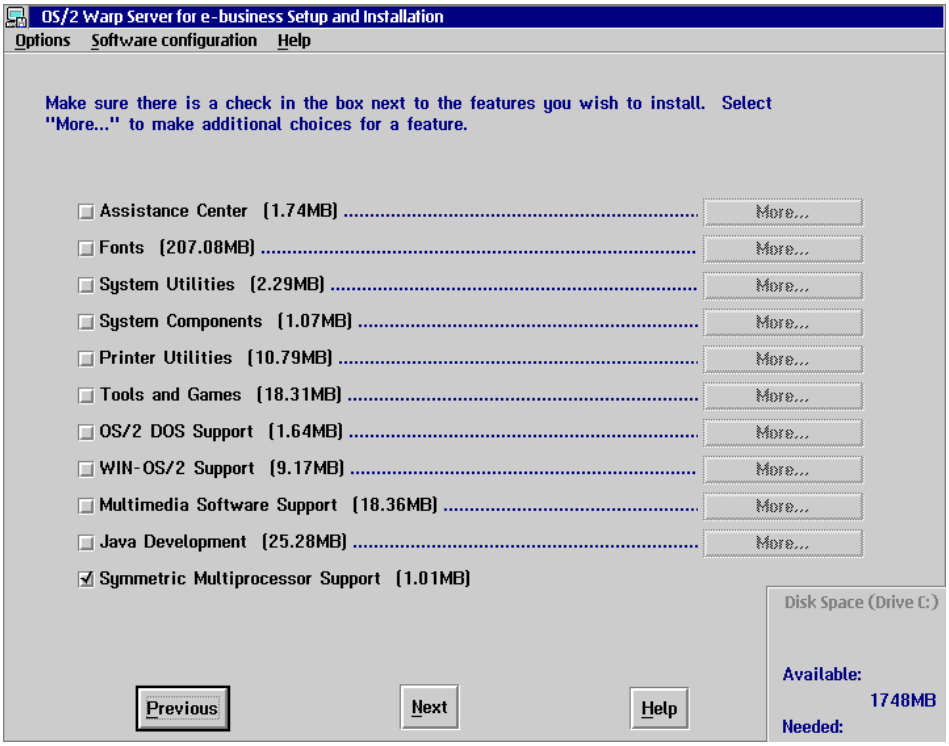

Figure 14. OS/2 Warp Server, selective install, SMP support

If you do not meet the requirements above, your hardware manufacturer might have written its own *.PSD driver file to support OS/2 Warp Server for e-business SMP on his hardware. In this case, when you select SMP Support as shown in Figure 14 on page 41, you will need to provide the location of a valid PSD file. It would be best to check this in advance whether such support is available.

Some applications running on Multi-Processor systems require the capability of designating which processor or group of processors a thread is eligible to run on. This designation is called processor affinity. There are two new OS/2 kernel APIs which allow OS/2 applications to specify Processor Affinity in a multiprocessor environment. These APIs are

1. DosSetThreadAffinity ()
2. DosQueryThreadAffinity()

These APIs will have no effect on UNI Processor environment.

Support for OS/2 processor affinity is controlled by two new kernel data fields:

qwCapableAffinity - This is the mask of processors available to OS/2 in the system; it is new to the kernel's global data.

qwThreadAffinity - This is the mask of processors available to the thread; it is new to each thread's internal data.

The processor affinity mask fields are 64bit wide fields and use 1 bit to represent each processor in the system. This support is therefore limited to multiprocessor systems with 64 or fewer processors. If an affinity mask bit is set, the OS/2 scheduler will make the thread eligible to be run on the corresponding processor(s). If the affinity mask bit is not set, the scheduler will never run the thread on the corresponding processor(s). Non-existent processors will always be represented by reset bits in any affinity mask.

The system initializes the qwCapableAffinity mask with one bit set on for each processor available to OS/2. The bits will be set contiguously starting with bit zero.

The kernel maintains each thread's current affinity, qwThreadAffinity, in the Thread Control Block (TCB.) Any attempt to change a thread's qwCurrentAffinity (via DosSetThreadAffinity) must be attempting to redefine Processor Affinity as a subset of qwCapableAffinity. Any attempt to set an affinity mask to zero will be rejected.

All new thread created will inherit the current thread's qwThreadAffinity mask. The new thread will then be capable of setting its current affinity, via DosSetThreadAffinity, to any subset of the system's qwCapableAffinity mask.

Application developers may use Processor Affinity APIs to force a thread's execution to one or more specific processors. This is useful for collecting diagnostic information from a specific processor. Secondarily, the ability to control which processor set an application runs on is useful as a performance tuning mechanism.

2.6 Large file support

Today’s data storage requirements have exploded with new data types, such as video and the need to house large databases. Currently, OS/2 APIs and filesystem structures use datastructures that impose a limit of 2 GB for file sizes. New APIs and data structures with 64 bit fields to hold file sizes/offsets have been added to overcome this restriction. For the rest of this section, Large File will refer to files that are greater than 2 GB in size.

JFS was designed with large files in mind; so, it already has 64 bit values internally. Changes have been made in the Installable File System Manager (IFSM) and the DOSCALLS interface to support large files.

Since JFS is the only IFS that will support Large Files, the IFSM will hide this new support from FAT, HPFS, and HPFS386 (and all other legacy 16bit IFSs, such as NFS, SRVIFS,...) and enforce the old 2GB limit to the same extent it does today even when an application uses the new APIs.

Note

Support for large files is NOT transparent to applications

Support for large files is NOT transparent to applications. In order to create and use large files, applications have to be modified to take advantage of these new APIs. This support is limited to partitions using the JFS.

A number of system utilities, such as XCOPY, are modified to take advantage of the new APIs and support Large Files.

The Filesystem will handle access to Large Files in two ways. First, on a DosOpen, the new DosOpenL API will be used. If an application attempts to open a Large File using the old DosOpen, that application will be denied with a return code of ERROR_ACCESS_DENIED. The application will, therefore, not be able to read or write to the file. The second method of determining if an application is Large File Aware is on the data structures used. Applications which are Large File Aware should always use the new types even if they do not expect to use Large Files.

The IFSM will handle some of the cross conversion between application and IFS. FS APIs have been modified to inform JFS as to when an application is not allowed to open Large Files.

The rules for reading and writing Large Files are as follows. Applications using the new APIs on JFS will have complete access to all files at all times following the current access control. Applications using the old DosOpen will not be allowed to open files larger than 2GB. Applications using the old APIs (specifically the DosOpen API) will be allowed normal access to files less than 2GB in size regardless of whether or not the file is opened by a process using the new DosOpenL. In order to prevent legacy applications from taking an unexpected error in the middle of processing a file, a limitation is placed on the new API. The limitation is that a file may not be grown past the 2GB limit while a legacy open is in progress. If this condition occurs, a sharing violation error will be returned to the application. Since this error can only

occur if the file is opened by multiple applications, the new application can prevent this by opening with the OPEN_SHARE_DENYREADWRITE mode flag. If the application requires shared access to a file, a new sharing flag OPEN_SHARE_DENYLEGACY may be specified. This flag is used in addition to the other sharing flags and only affects opens through the legacy DosOpen. It has no effect on DosOpenL calls. Setting this flag on an open will have the same effect on legacy applications as OPEN_SHARE_DENYREADWRITE. This means that the only error relative to large file that a legacy application should see is an ERROR_ACCESS_DENIED on a DosOpen.

Applications that are required to work on versions of OS/2 prior to OS/2 Warp Server for e-business should not issue the new APIs directly. They should use DosQueryProcAddr to determine if the new API exists, and, if not, issue the old API. For the APIs that use new infolevels but do not have a new API name, the application should issue the command with the new Infolevel, and, if an ERROR_INVAID_PARAMETER is received, the application should then reissue the command using the old Infolevel.

All old APIs and data structures will continue to be supported; so, all current applications will continue to work with the one exception that they will not be able to see or open files greater than 2GB in size that have been created by new applications.

2.6.1 SES and large file support

There are new SES KPIs for use with large files within JFS volumes. Third party security products implement Installable Security Subsystems (ISS) that make use of SES APIs. These vendors and/or customers will need to modify their Installable Security Subsystems to be aware of large files.

Existing ISSs using the old SES API may prohibit access to data in large files (beyond 2GB). This may cause the system to halt or hang because the ISS does not know how to deal with large files. It is not possible to determine the exact behavior of an ISS using the old SES API on a JFS file that grows beyond 2GB.

Note

Running an ISS implemented with the old SES API on a JFS file that grows beyond 2GB is not supported. This can result in data loss, a system crash, or some other unpredictable behavior.

A customer *should not use* an old API ISS with large files. A customer should get an upgraded ISS from their ISS vendor that has been developed and tested to support large files.

SES support in OS/2 Warp Server for e-business only provides the ISS vendors with an API extension to support large files. It is the ISS vendor's responsibility to implement that API.

Tests indicate that legacy ISSs on OS/2 Warp Server for e-business with JFS files smaller than 2GB are acceptable. But, as soon as a file on JFS grows to more than 2GB in size running a legacy ISS, there could be a system failure.

2.7 Java 1.1

The OS/2 Warp family has consistently provided highly-competitive Java performance. The latest OS/2 Java implementation (Java 1.1.7 based) provides state-of-the-art Java performance integrated into OS/2. This provides customers the potential to combine high-capacity, high-performance web-serving, Java serving, and traditional network file and print serving, all using one platform.

IBM OS/2 Warp Developer Kit, Java(TM) Edition, Version 1.1.7 (Developer Kit) is based on Sun Microsystems' Java 1.1.7A maintenance level.

This kit is also supported in the following environments:

- OS/2 Warp 4
- OS/2 Warp Server Version 4
- OS/2 Warp Server Advanced Version 4
- OS/2 Warp Server Advanced Version 4 SMP Feature
- WorkSpace On-Demand clients and servers
- OS/2 Warp 3
- OS/2 Warp Connect

Each of the above platforms may have individual fixpack requirements. Please check the online readme file for more information. OS/2 Warp Server automatically installs this component if it is selected. All prerequisite fixes are included in the package.

IBM OS/2 Warp Developer Kit, Java(TM) Edition, Version 1.1.7 replaces all previous versions of Java 1.1.x for OS/2 Warp. If you do not uninstall the previous version, the installation program will replace it with the Developer Kit. If a previous version of a Java component was installed and you have not selected to reinstall that component, the installation program displays a window warning you that this component will be downlevel and then lets you choose to upgrade the component.

2.8 Netscape 4.04

Netscape 4.04 is available for installation. The version that is shipped is enabled to work with just 16 colors.

Communicator builds on the existing Navigator 2.02 function, which is still supported. You can connect to the Internet and send e-mail worldwide, establish group communications with friends and colleagues, search for data, browse, view documents including those with video and audio, explore Virtual Reality Modeling Language (VRML) 3-D, and even publish on the World Wide Web.

Some of the technical details that help enhance the new features include:

- Dynamic HTML - technologies that transform static HTML pages into dynamic animated pages
- JavaScript 1.2 - Dynamic HTML uses new functions of JavaScript:
  - Style Sheets - both Cascading Style Sheets (CSS) and JavaScript-accessible Style Sheets (JASS)
  - HTML layering and positioning tags - layer elements such as graphics and text that can be positioned and made visible or invisible under the control of JavaScript.
- AutoInstall - Automatically download a plug-in when a page needs it.
- Native OS/2 and Microsoft Windows 3.1 Plug-ins - Provided with the Communicator 4.04 for OS/2 Warp Plug-In Pack
- MPEG - Support for MPEG movies
- Invoke plug-ins using the Object tag in addition to the embed tag.
- Dynamic Fonts - Send a font along with a document
- Java 1.1.6 - Support for the latest Java 1.1 releases, including Java 1.1.4 for OS/2 and OS/2 Warp Developer's Toolkit for Java 1.1.6.
- Security - Communicator added security features to help create a secure environment for conducting transactions over the Internet. Easing people’s security fears is essential to increasing business over the Internet. The new security features help make Communicator 4.04 a good choice to utilize the latest Web technologies in your transition to e-business.

Communicator 4.04 delivers:

- A padlock icon that calls the Security Advisor. The details of your current connection's security status help you better determine what you can safely do or, perhaps, not do.
- A 128 bit-strong encryption version for U.S. and Canada that uses strong cryptology to protect sensitive data is available separately.
- Secure access controls when two or more people need to share the same computer.
- PCMCIA-based security services.
- Secure, encrypted, mail and news Certificates for additional identity verification.

2.9 Graphical locale builder

An e-business efficiently connects valuable information to the people who need it unecumbered by time zones and national borders. When customers, suppliers, or business partners are able to interact with Internet applications in their native language, more meaningful communications occur.

OS/2 Warp Server for e-business transverses language barriers as it extends the existing I/T infrastructure to the internet. The IBM Graphical Locale Builder (GLB) help make business applications more intelligible to global customers and business partners who use different language, currency, and date formats. With the Graphical Locale Builder, companies can use a Graphical User Interface to quickly modify internationalization settings.

A company can develop applications in their native language and then use any of the 14 languages supported by OS/2 Warp Server for e-business. From the Country Palette's GUI interface, developers or administrators create a locale. Once locales are created for the various languages and countries, a user can drag a locale from the Country Palette and drop it onto applications that are fully enabled for international support. The international settings of that application then change to reflect the parameters of the locale. For example, a multinational bank can have its teller program running in German, and the language can be switched to French or even Japanese (double-byte character set) without having to reboot the system. An application that was originally set up for use in the U.S can be modified for use in France by using the GLB to change the currency to French Francs and the date to dd/mm/yy, even if the dates are in the year 2000 or beyond.

The system locale defines a user environment. The locale specifies a language, currency, date format, and how the system handles formatting, conversions, sorts, and text processing procedures.

Applications will need to be written to take advantage of this feature. Here we will describe how to create a new locale and use it within some of the default applications. To create a new locale, you need to do the following:

1. Double click on the Locale icon within the OS/2 System Setup folder. This icon is shown in Figure 15 on page 48.

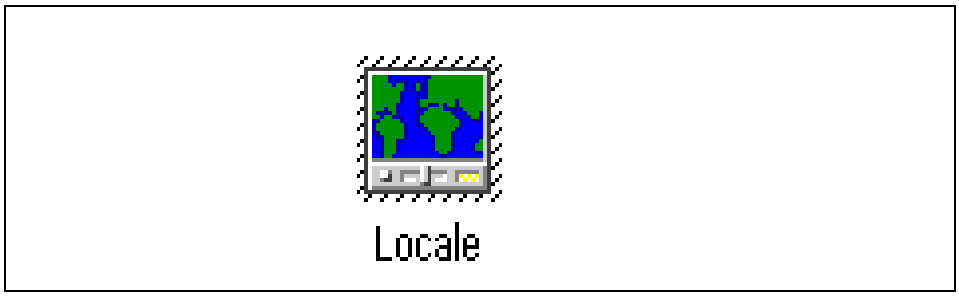

Figure 15. Locale icon.

2. This will bring up the Locale Folder. Select the **Locale** item of the action bar and then select **new Locale** to create a new Locale. This will bring up a dialog box as shown in the following figure:

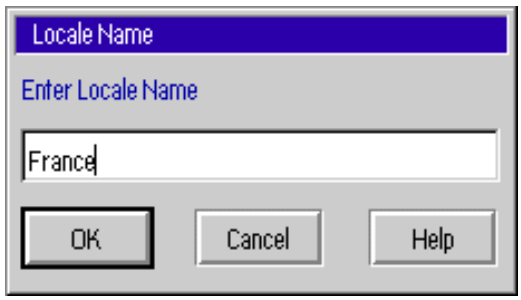

Figure 16. New locale name

3. Enter a descriptive name for the new Locale. Select **OK**. This will bring up a settings notebook for your new Locale. The settings notebook is shown in Figure 17 on page 49.

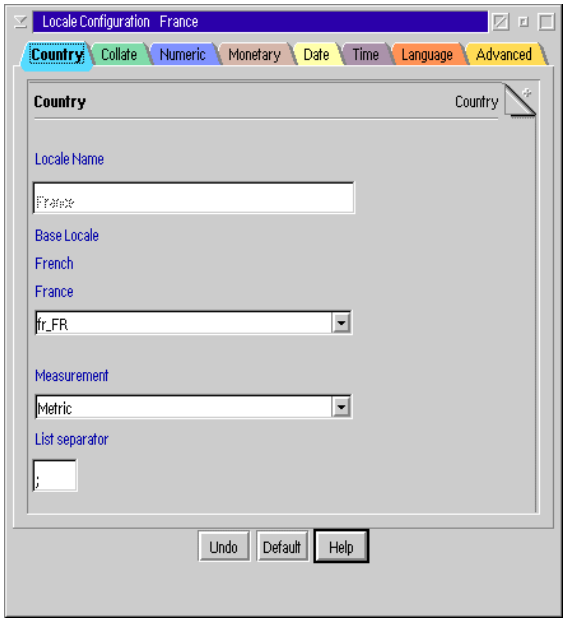

Figure 17. Settings for new locale

4. The settings notebook allows you to set all the variables that would be associated with that region. The application that is *locale enabled* will gather its settings from the values entered here. The first page allows you to set the base Locale from a drop down list. The syntax for base Locale is as follows:

   xx_yy

   Where xx represents the two-character ISO language identifier, and yy represents the two-character ISO country identifier. On this page, you can also select the Measurement to be Metric or English as well as the seperator.

5. Use the help button for more information on each of the options. The other settings pages allow you to define the Numeric, Date, and other formats. When you are done, close the settings notebook and you will have a new Locale. Our new Locale called France is shown in Figure 18 on page 50.

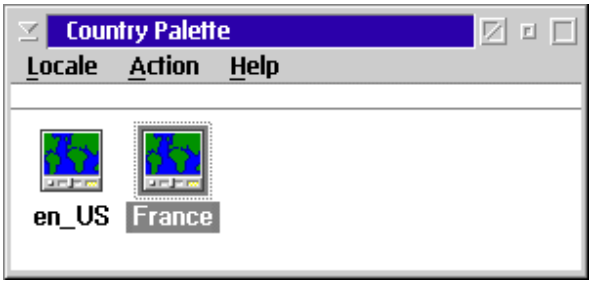

Figure 18. Two defined locales.

To use the new locale, simply drag and drop the **Locale** icon onto the application that is Locale enabled. You can try this out by dropping the new icon onto the System Clock application within the System Setup folder or the OS/2 System Editor which is contained within the Utilities folder. You should notice language and date/time format changes.

2.9.1 Java locale

To create a Java Locale, you will need to have the JDK 1.1.7 or later installed. If this is not installed, you will get a cryptic error message. To create a corresponding Java Locale, simply:

1. Select the **Locale icon** you created earlier.
2. Then select the **Action** item off the Menu.
3. Then select **Create Java Locale**.

2.9.2 Euro currency support

Euro currency support is essential in some countries and industries today and will become more critical as the currency goes into circulation. Therefore, OS/2 Warp Server for e-business supports the euro currency character and enables using it in applications and Java applets. The euro is not supported in Win OS2.

OS/2 support for the euro consists of adding euro character fonts, code pages, and keyboards. The following code pages are supported:

- 850 PC Latin 1
- 857 PC Turkish
- 1004 Windows Extended
- 1250 Windows Latin 2
- 1251 Windows Cyrillic
- 1252 Windows Latin 2
- 1254 Windows Turkish

- 1257 Windows Baltic PM

The euro character has been added to the following fonts:

- Courier
- Helv
- Helvetica
- System Monospace
- System Proportional
- System VIO
- Times New Roman
- Times New Roman MT 30
- Tms Rmn
- WarpSans

The euro has been added to the following keyboards:

Belgium, Germany Latin, Sweden, Swiss French, Canada, America, Germany, Swiss German, Canadian French, Netherlands, Iceland, Turkey, Denmark, Norway, Iceland 101, Turkey 440, Finland, Portugal, Italy, UK, France, Spain, Italy Extended, UK 168, US International.

Chapter 3. Adapters and protocol services

OS/2 Warp Server for e-business provides you with a wide rage of supported networking protocols and communication adapters that you may use in many combinations to suit your requirements for a server system in a LAN environment as well as for wide-area networking. Adapters and protocol services may be called the communications engine of OS/2 Warp Server for e-business since it provides communication support for all the other components of this product.

This chapter describes Adapters and protocol services and its component parts, new features in this release of OS/2 Warp Server for e-business, and some configuration examples.

OS/2 Warp Server for e-business includes support for I_20 LAN Adapters. For more information on I_20 and setting up an I_20 LAN adapter, see Chapter 2.2, “Intelligent input and output I_2O” on page 23.

3.1 Overview of adapters and protocol services

Adapters and protocol services can be divided into two parts:

- LAN-Adapter and Protocol Support (LAPS)
- Multiprotocol Transport Services (MPTS)

3.1.1 Adapter and protocol support

Originally, LAPS was called LAN Adapter and Protocol Support, but, since it also includes drivers for wide area networking adapters, we skipped the word LAN and kept the familiar acronym LAPS. LAPS includes the following networking protocols based on the Network Driver Interface Specifications (NDIS) standard:

- IBM OS/2 NetBIOS
- IBM TCP/IP
- IBM IEEE 802.2
- Netware NetBIOS Emulation over IPX

  This protocol ships with the Netware Requester for OS/2. LAPS simply provides an NDIS NIF file; so, a user who has installed the NetWare Requester (along with the NetWare NetBIOS protocol that comes with it) can configure it with LAPS.
- IBM OS/2 NetBIOS over TCP/IP

- NetWare Requester support

The files that make up LAPS and its configuration are placed under the \IBMCOM directory tree on the OS/2 boot drive.

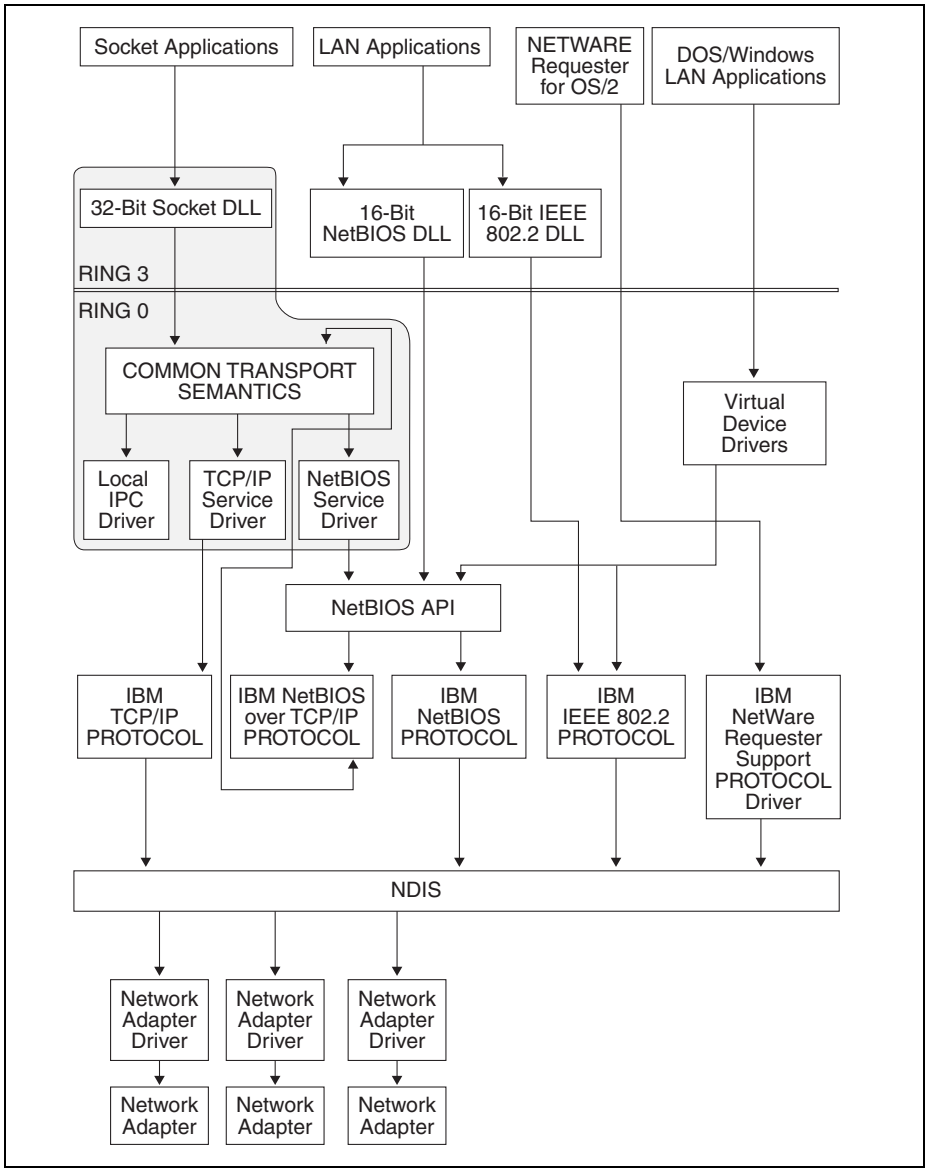

Figure 19. MPTS components

A wide range of LAN adapters for token-ring, Ethernet, and FDDI are included as well as WAN adapters and serial and parallel communications support for NDIS. A complete list of adapters supported by Adapters and protocol services can be found in the \IBMCOM\MACS\REDMAC.TXT file.

3.1.1.1 Network Driver Interface Specifications (NDIS)

IBM transport strategy is based on the Network Driver Interface Specifications (NDIS) version 2.0.1- a standard jointly developed by 3COM and Microsoft Corporation. NDIS allows different network protocols to operate over the same LAN interface at the same time.

NDIS is a standardized Medium Access Control (MAC) interface for network adapter and protocol drivers. It has become a de facto industry standard providing a common open interface that enables different manufacturers of network adapters and LAN software developers to produce products that communicate with each other.

NDIS separates protocol handling from hardware manipulation by defining functions that protocol drivers and network adapter drivers must provide to each other.

NDIS defines:

- Specifications for network protocol drivers
- Specifications for network adapter drivers
- Interface between the above two layers
- Binding process to link these protocol and adapter drivers

A *network protocol driver* provides the communication between an application and a network adapter driver.

A network adapter driver or MAC driver provides the communication between a network adapter and a protocol. The main function of the network adapter driver is to support network packet reception and transmission.

Each driver has an upper and a lower boundary. The drivers are linked together to form a stack by binding the lower boundary of one driver to the upper boundary of another driver. The MAC driver at the bottom of the stack always has its lower boundary connected to the physical layer - the network adapter hardware.

The NDIS specification defines the binding process of the drivers. Three components are used to form and manage the protocol stack from individual drivers. These are:

PROTOCOL.INI

This is an ASCII file that defines the protocol drivers and adapter drivers in use and their binding information.

PROTMAN.OS2

This is a Protocol Manager.

NETBIND.EXE

This initiates the final binding process.

The LAPS component of Adapter and Protocol Services contains the above three files, the protocol and adapter drivers, and a utility for easy installation and configuration of the required drivers. LAPS also contains Virtual Device Drivers which make the installed protocols available to DOS and Windows sessions under OS/2 without the need for specific DOS protocol drivers.

3.1.1.2 Multiple Protocol Support

NDIS allows multiple protocols to be bound to a single MAC driver, that is, to share a network adapter. Figure 20 shows the NDIS protocol stacks when NetBIOS, IEEE 802.2, and TCP/IP are loaded together. In this example, two LAN adapters are in use. NetBIOS and IEEE 802.2 are bound to one of the adapters, and the other adapter is dedicated to the TCP/IP protocol (although there is no reason why all three protocols could not have been bound to both adapters). The configuration information defining which protocol(s) is/are bound to which adapter(s) is contained in the PROTOCOL.INI file.

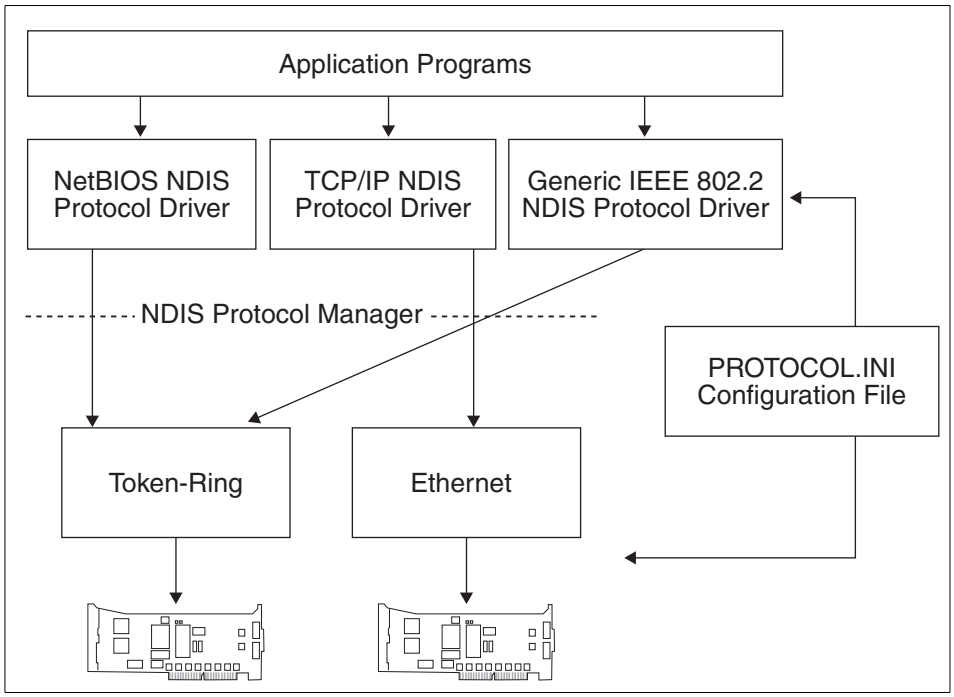

Figure 20. NDIS - multiple protocols

3.1.1.3 PROTOCOL.INI

PROTOCOL.INI contains the NDIS configuration information for network adapter drivers and protocol drivers for a workstation. PROTOCOL.INI is an ASCII file that can be edited manually, but this is generally not recommended. We recommend that you always use the Adapter and Protocol Services configuration utility to ensure the creation of valid PROTOCOL.INI and CONFIG.SYS files. The PROTOCOL.INI file consists of four sections:

- Protocol Manager

- Configuration
- Protocol drivers
- MAC (network adapter) drivers

All these sections have the following structure:

```
[module name]
    parameter=value
```

The following is an example of a PROTOCOL.INI file configured with both NetBIOS and TCP/IP protocol stacks (similar to Figure 20). The first entry is the protocol manager, which is the driver that controls the binding process.

```
[PROT_MAN]
   DRIVERNAME = PROTMAN$
```

The configuration section defines which protocols are used and what types of adapters are configured. In the following example, netbeui_nif and tcpbeui_nif are the protocol drivers, and IBMTOK_nif is the adapter configuration (in this case an IBM Token-Ring adapter).

```
[IBMLXCFG]
   netbeui_nif = netbeui.nif
   tcpip_nif = tcpip.nif
   IBMTOK_nif = IBMTOK.NIF
```

The Bindings = statements under the various protocol drivers specify the module name of the MAC driver to which the protocol driver will bind to form a protocol stack or stacks. In this example, NetBIOS, the NetBIOS API is using the NetBEUI protocol driver, which itself is bound to the token-ring MAC driver. TCP/IP is also bound to the token-ring MAC driver.

```
[NETBIOS]
   DriverName = netbios$
   ADAPTER0 = netbeui$,0
[netbeui_nif]
   DriverName = netbeui$
   Bindings = IBMTOK_nif
[tcpip_nif]
   DriverName = TCPIP$
   Bindings = IBMTOK_nif
[IBMTOK_nif]
   DriverName = IBMTOK$
```

More statements of the kind of parameter = value may appear under each protocol and MAC section. The meanings of parameter and the allowed ranges and types for value are contained in network information (.NIF) files, which exist for each protocol and MAC driver. The configuration program parses those NIF files to check what can be configured for any given section in PROTOCOL.INI and whether a configuration item is valid. If no additional

parameters are specified in the PROTOCOL.INI sections, default values will be used as defined in the .NIF files.

3.1.2 Socket/Multiprotocol transport services

The Sockets interface allows you to develop distributed or client/server applications using various transport protocols. The application can select the transport protocol or request that the Socket/MPTS layer determine the protocol. Most socket applications available today communicate with either TCP or UDP.

Sockets are duplex, which means that data can be transmitted and received simultaneously. Sockets allow you to send to and receive from the socket as if you are writing to and reading from any other network device.

Socket/MPTS provides the support for three kinds of address families for the Sockets Application Programming Interface (API):

TCP/IP address family (AF_INET)

NetBIOS address family (AF_NB)

OS/2 address family (AF_OS2)

It also provides a local IPC transport for Sockets applications (inter-process communications support that does not issue any calls to the network).

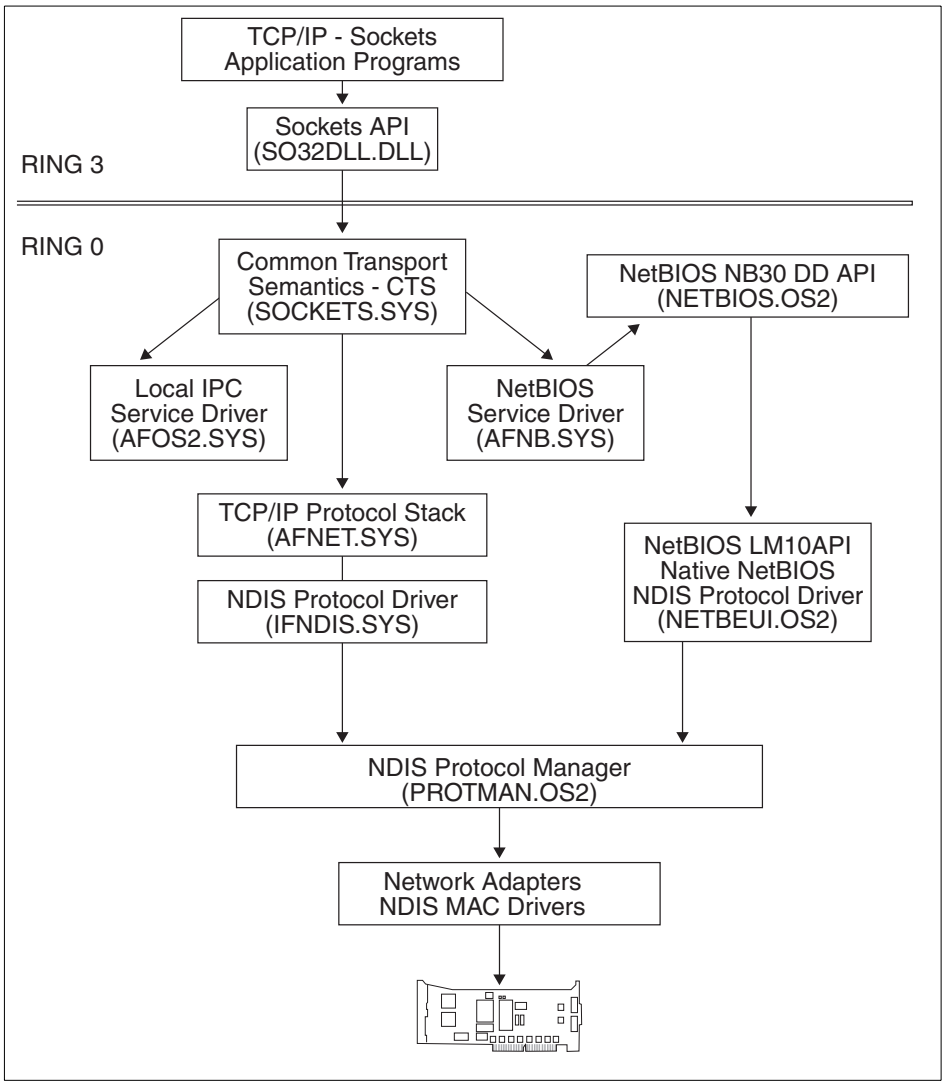

Figure 21. Overview of Socket/MPTS

3.2 Installing adapters and protocol services

Since this component of OS/2 Warp Server for e-business provides communication support to all other parts of the product, it will always be installed when you install an OS/2 Warp Server for e-business. Therefore, it

cannot be explicitly selected in the OS/2 Warp Server for e-business Setup and Installation Menu.

Note

If you are installing OS/2 Warp Server remotely in unattended mode (response file installation method using CID), you must include the Adapter and Protocol Services component in the installation procedures. Otherwise, your system will not be able to access the LAN.

Adapters and protocol services will attempt to detect and identify any LAN adapters that are installed in your system. See the \IBMCOM\MACS\READMAC.TXT file for a list of adapter drivers that are supplied with OS/2 Warp Server for e-business. However, only the first adapter that can be found will be automatically included in the configuration. You may add more adapters, and you may add additional adapter drivers that are not included in OS/2 Warp Server for e-business, which we will describe later in this chapter.

As far as network protocols are concerned, the initial configuration of Adapter and Protocol Services depends on the features that you have selected to be installed.

The following Figure shows a possible configuration menu during the installation of OS/2 Warp Server for e-business.

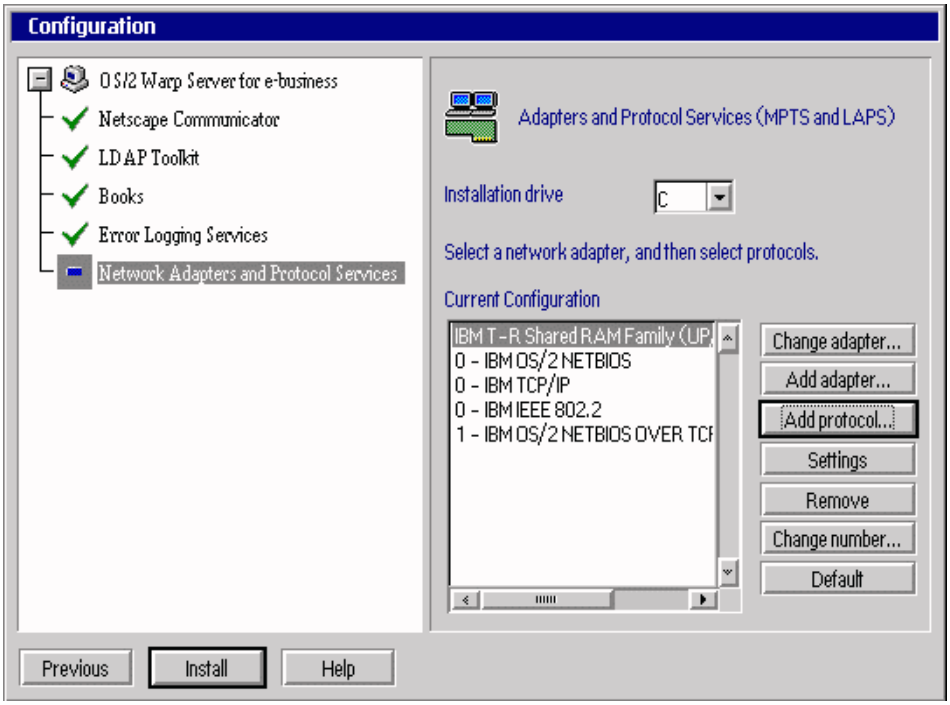

Figure 22. Adapter and protocol services configuration during the installation

The following table summarizes the configuration parameters of this page and describes their purposes.

Table 3. Adapters and protocol services installation

Configuration Item	Configuration Data
Change Adapter	Press this button if you want to change a network adapter for the current configuration. This may be necessary if the installation program could not detect your adapter configuration properly.
Add adapter	Press this button if you: 1. Want to add an adapter to your configuration 2. Want to add a new adapter driver that is not supplied with Adapters and protocol services This will bring up the menu that is shown in Figure 23 on page 64.

Configuration Item	Configuration Data
Add protocol	Press this button if you: 1. Want to add a protocol to your configuration 2. Want to add a new protocol driver that is not supplied with Adapters and protocol services. This will bring up the menu that is shown in Figure 24 on page 65.
Settings	Press this button if you want to change the parameters for any item in the configuration list. This will bring up the menu that is shown in Figure 25 on page 66.
Remove	Press this button if you want to remove an item from the configuration list. You can only remove one item at a time. You can only remove an adapter if you have previously removed all protocols that have been associated with that adapter.
Change number	Press this button if you need to change the logical sequence in which a protocol driver will address adapters if this protocol has been associated with more than one adapter. Application programs will use these numbers when they issue calls to the network interface.

If the LAN adapter in your system is not supplied with OS/2 Warp Server for e-business, but you have an NDIS compliant OS/2 device driver for that adapter, click on the **Add adapter** button. The following menu will appear:

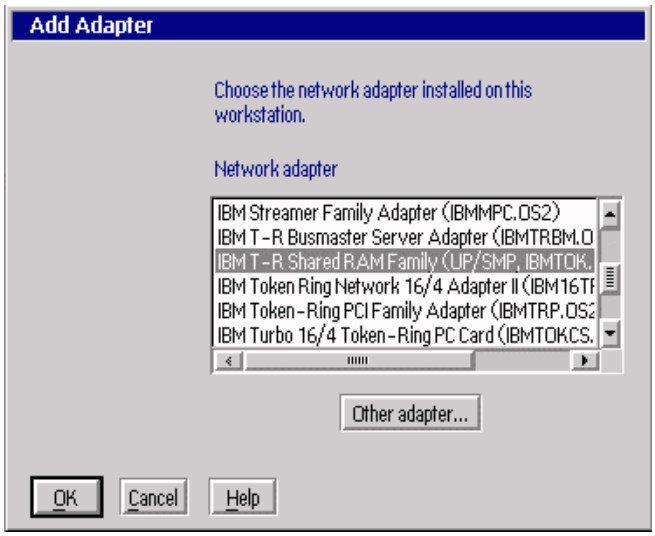

Figure 23. Add adapter driver to Adapters and protocol services

To add a supplied adapter to your configuration, select the appropriate driver from the list and press **OK**. To add a new adapter driver, select **Other adapter...**, then specify the source drive from where the new driver will be copied to your system.

If you want to use a protocol that is not supplied with OS/2 Warp Server, but you have an NDIS-compliant OS/2 device driver for that protocol, click on the **Add protocol** button. The following menu will be shown:

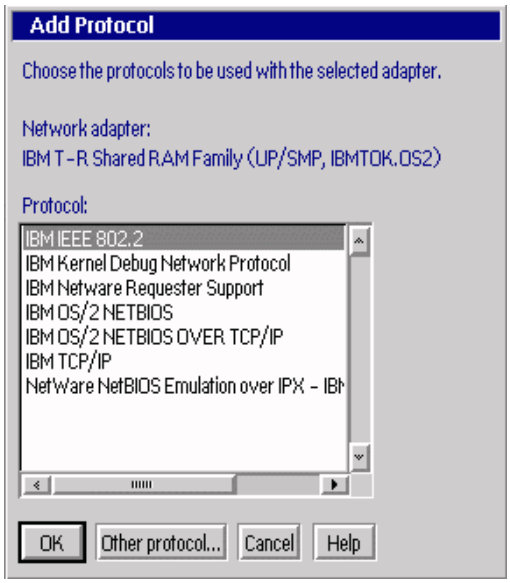

Figure 24. Add protocol driver to Adapters and protocol services

To add a supplied protocol to your configuration, select the appropriate driver from the list, then press **OK**. To add a new protocol driver, select **Other protocol...**, then specify the source drive from where the new driver will be copied to your system.

Adding an adapter or protocol may affect the configuration of other OS/2 Warp Server components; so, you may want to check the items on the configuration tree again. For instance, adding the NetBIOS over TCP/IP protocol will result in another LAN adapter, which can be selected for File and Print Sharing Services.

If you want to view or change the configuration of any adapter or protocol driver, select the appropriate item on the list, then click on the **Settings** button. The following menu will be shown:

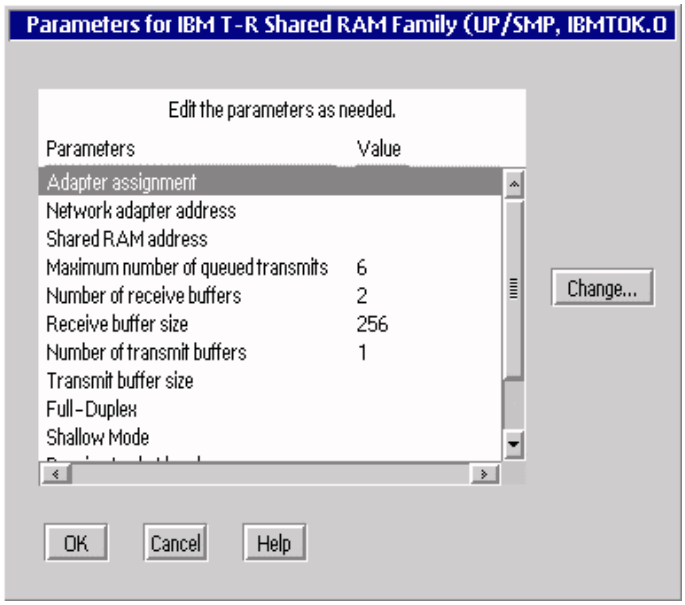

Figure 25. Change settings in Adapters and protocol services

3.3 Additional configuration for adapters and protocol services

If you want to make changes to the configuration of Adapters and protocol services after OS/2 Warp Server for e-business has been installed, you can easily do so by using the Adapters and protocol services configuration program. To start this program, either click on the **MPTS Network Adapters and protocol services icon** in the System Setup folder or type **MPTS** at an OS/2 command prompt.

Then, select **Configure** to bring up the following figure:

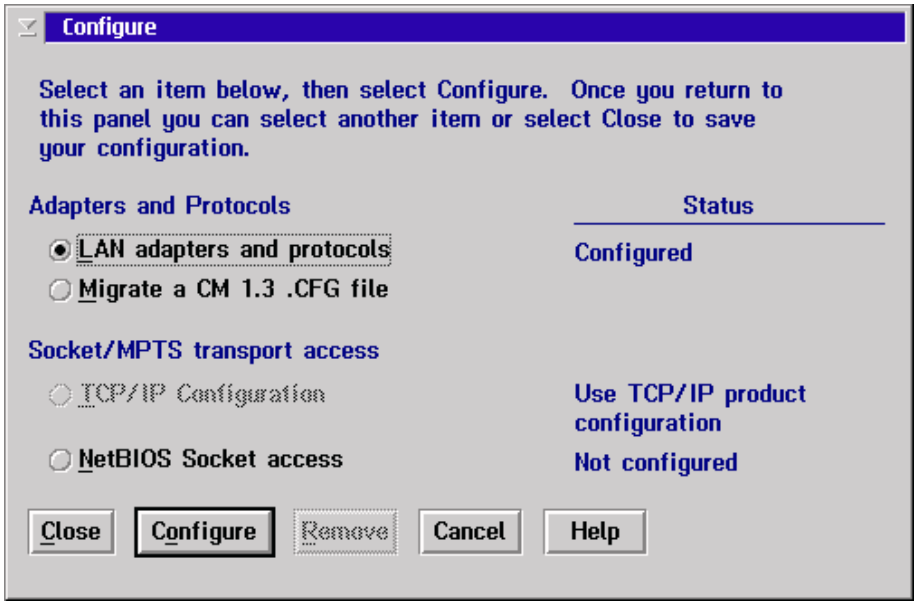

Figure 26. Additional configuration of Adapters and protocol services

Select LAN adapters and protocols and choose **Configure** to bring up the LAPS Configuration window. Use this configuration panel to select the LAN adapter(s) installed on the system and the protocols associated with them. Figure 27 shows an example of what Adapters and protocol services configuration might look like.

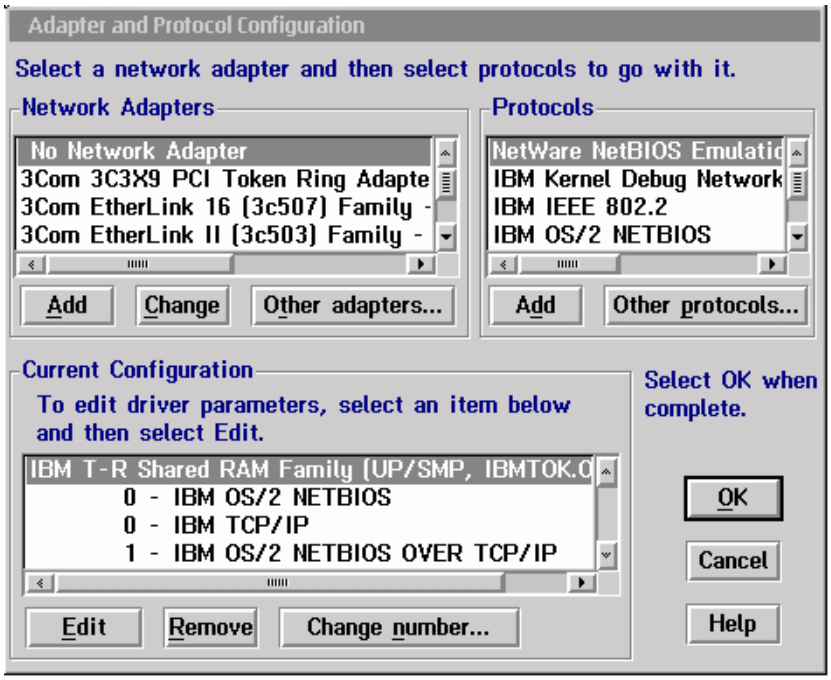

Figure 27. Adapter and protocol configuration

The following table summarizes the use of this configuration menu.

Table 4. Configuration menu items

Configuration Item	Configuration Data
Network Adapters window	
Add	Press this button if you want to add an adapter to your configuration.
Change	Press this button if you want to change a network adapter for the current configuration.
Other adapters...	Press this button if you want to add a new adapter driver that is not supplied with Adapters and protocol services.
Protocols window	
Add	Press this button if you want to add a protocol to your configuration.
Other protocols...	Press this button if you want to add a new protocol that is not supplied with Adapters and protocol services.

Configuration Item	Configuration Data
Current Configuration window	
Edit	Press this button if you want to change the parameters for any item in the configuration list. This will bring up a menu where you can make changes to multiple parameters and then apply all changes at once
Remove	Press this button if you want to remove an item from the configuration list. You can only remove one item at a time. You can only remove an adapter if you have previously removed all protocols that have been associated with the adapter
Change number...	Press this button if you need to change the logical sequence in which a protocol driver will address adapters if this protocol has been associated with more than one adapter. Application programs will use these numbers when they issue calls to the network interface(s).

If you have finished the configuration, press **OK** to save the changes. Press **Close** on the following panel, then press **Exit**. Select to update the CONFIG.SYS file on the OS/2 boot drive so that the configuration changes can be properly applied. You will need to reboot before the changes will become effective.

3.3.1 Configuring Socket/MPTS

Socket/MPTS is configured from the Configuration panel when loading MPTS (see Figure 26 on page 67). Use this panel to select the protocols that you intend to use for Socket access. These selections notify Socket/MPTS to initialize the protocol services required. The selectable protocols are:

TCP/IP Socket access
and
NetBIOS Socket access.

To select TCP/IP Socket access, you must have TCP/IP protocol configured (using the LAN Adapters and protocol services configuration option). To select NetBIOS Socket access, you must have the NetBIOS protocol configured (using the LAN Adapters and protocol services configuration option). You must select at least one protocol for your Socket/MPTS environment. However, you can select more than one protocol.

Note

When you have installed TCP/IP Version 4.21, during the installation of your OS/2 Warp Server for e-business system or at any later time, you must use the TCP/IP Configuration notebook to configure TCP/IP parameters.

3.3.1.1 Configuring TCP/IP Sockets access

Since TCP/IP Version 4.21 and Java 1.1.7 all TCP/IP settings can only be made in the TCP/IP Configuration Notebook.

3.3.1.2 Configure NetBIOS Sockets access

Select NetBIOS Sockets access, then click on **Configure**. The following menu will be shown:

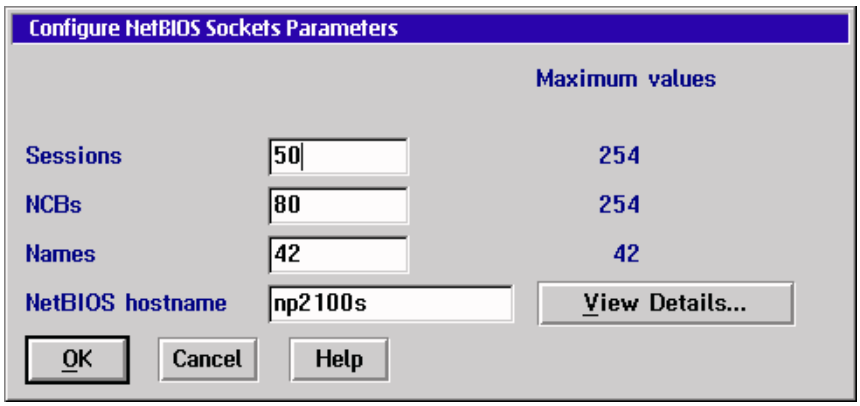

Figure 28. NetBIOS socket access configuration

On this page, you can configure the NetBIOS interface of your OS/2 Warp Server for e-business system for Socket access. The following table

summarizes the configuration parameters of this page and describes their purpose.

Table 5. NetBIOS Interface configuration for socket access

Configuration Item	Configuration Data
Sessions	Specify the number of NetBIOS sessions you want to reserve for Sockets applications. The number of sessions specified here will be taken from the total number of sessions specified in the NETBEUI section of the PROTOCOL.INI file. This means that the amount of sessions available to other NetBIOS applications, such as File and Print Sharing Services, will be reduced by the number specified here.
NCBs	Specify the number of NetBIOS commands (NCBs) you want to reserve for Sockets applications. The number of NCBs specified here will be taken from the total number of NCBs specified in the NETBEUI section in the PROTOCOL.INI file. This means that the number of NCBs available to other NetBIOS applications, such as File and Print Sharing Services, will be reduced by the number specified here.
Names	Specify the number of NetBIOS names you want to reserve for Sockets applications. The number of NCBs specified here will be taken from the total number of names specified in the NETBEUI section in the PROTOCOL.INI file. This means that the number of names available to other NetBIOS applications, such as File and Print Sharing Services, will be reduced by the number specified here
NetBIOS hostname	Enter the hostname that your NetBIOS Sockets application will be using.
View Details...	Click here to see more available hostnames. When the NetBIOS protocol is configured to more than one adapter, Socket/MPTS will use the hostname you have specified for the first interface, and it will add to that name unique identifiers (consecutive numbers) and use them as hostnames for the other interfaces.

When you finished the Sockets Access Configuration, click on Close to return to the Configuration menu.

You can use the `NETSTAT` command to see whether the NetBIOS Sockets interface is initialized and what applications are currently using it.

3.3.1.3 Removing Socket/MPTS configuration

When you want to remove a Socket/MPTS configuration from your OS/2 Warp Server system, you can do it from the Configure panel when loading

MPTS. Select the appropriate Socket access protocol to be removed (TCP/IP or NetBIOS), then click on **Remove**. If you want to remove both protocols, you can do so by removing one after the other.

Note

Removing a Socket access protocol will not remove the protocol driver from the LAN adapter and protocol configuration. It will only update the \MPTN\BIN\MPTCONFG.INI file and remove device drivers from the CONFIG.SYS file.

3.3.2 NETBEUI parameters

The NETBEUI section of the PROTOCOL.INI file has two particular configurable parameters: SIDEBAND and BALANCE. These parameters should not be changed from their default values. When not present in the PROTOCOL.INI file, these two parameters are set to the correct values automatically. The default values (when not present in PROTOCOL.INI) are as follows:

- SIDEBAND = 1
- BALANCE = 2

Setting SIDEBAND to 1 enables a performance enhancement used by File and Print Sharing Services for sending small frames. Setting SIDEBAND to 0 disables this performance enhancement.

BALANCE is used to control how NetBEUI chooses which adapters are used when an `NCB.LISTEN` command is issued on a machine with multiple network adapters. If two network adapters on the same machine are on the same network segment (bridged segment), setting BALANCE to **0** disables load balancing; setting BALANCE to **1** puts it into load balancing mode; and, setting BALANCE to **2** lets NetBEUI decide the appropriate load balancing mode.

3.3.3 Configuring more than four LAN adapters

As discussed below, Adapter and Protocol Services itself is not limited to four adapters as is the NB30 NetBIOS API. Adapter and Protocol Services support includes LAN adapters as well as other NDIS communication adapter drivers, such as asynchronous and parallel port support, WAN, and ISDN. The number of adapter drivers that can actually be used concurrently differs between the NDIS protocol drivers. The TCP/IP, IEEE 802.2, and NetWare Requester Support protocol drivers can support up to 64 adapters. The NetBIOS protocol drivers (IBM NETBEUI, IBM NetBIOS over TCP/IP, and

NetWare NetBIOS Emulation) will only support four adapters, which is the limit of the NB30 NetBIOS API.

Since all NetBIOS drivers that are supplied with OS/2 Warp Server also support the LM10 NetBIOS API, there is no restriction to four adapters when this interface is used by applications. This will increase the number of clients that can be simultaneously connected to OS/2 Warp Server over NetBIOS. It will also help the session load balancing performed by File and Print Sharing Services.

In the following example, we used four DUAL Auto LANStreamer adapter cards that provide two separate token-ring ports each. This gives us a total of eight LAN adapters as seen by NDIS. Figure 29 on page 74 shows how these adapters can be interfaced by NetBIOS applications. NB30 applications will only be able to interface with four adapters, whereas LM10 applications, such as File and Print Sharing Services, will be able to interface with all eight adapters that are bound to the NetBEUI protocol driver.

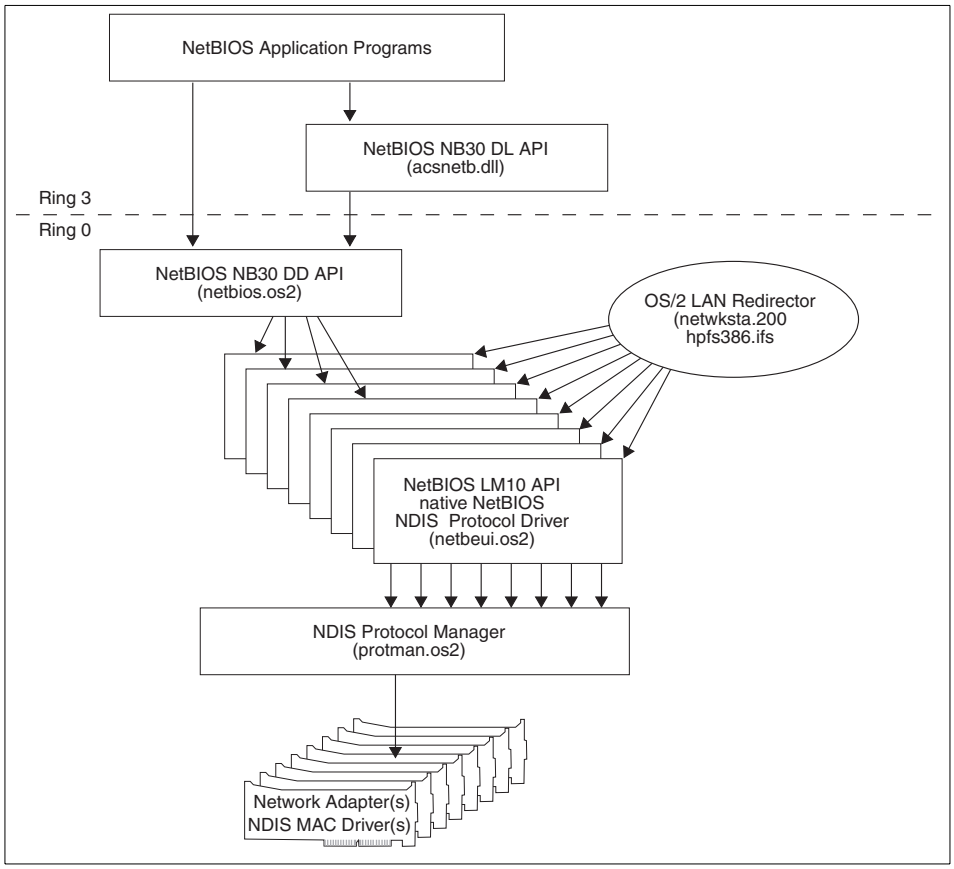

Figure 29. NetBIOS configuration for eight adapters

The following lines are extracted from the CONFIG.SYS file to reflect what drivers must be loaded in order to support the configuration as shown above:

```
...
DEVICE=C:\IBMCOM\MACS\DUALSTRM.OS2 /S:3
DEVICE=C:\IBMCOM\MACS\DUALSTRM.OS2 /S:4
DEVICE=C:\IBMCOM\MACS\DUALSTRM.OS2 /S:6
DEVICE=C:\IBMCOM\MACS\DUALSTRM.OS2 /S:7
DEVICE=C:\IBMCOM\LANMSGDD.OS2 /I:C:\IBMCOM
DEVICE=C:\IBMCOM\PROTMAN.OS2 /I:C:\IBMCOM
...
```

Note

- You have to add the lines for the DUALSTRM.OS2 device driver manually.
- Make sure you add those lines before the PROTMAN.OS2 device driver statement as shown in the example above.
- Include the /S parameter on every statement to specify the slot that the Dual LANStreamer adapter is plugged into.

The following lines are extracted from the PROTOCOL.INI file to reflect what drivers must be loaded to support the configuration as shown above:

```
[PROT_MAN]

   DRIVERNAME = PROTMAN$

[IBMLXCFG]

   netbeui_nif = netbeui.nif
   IBMMPC_nif = IBMMPC.NIF
   IBMMPC_nif2 = ibmmpc.nif
   IBMMPC_nif3 = ibmmpc.nif
   IBMMPC_nif4 = ibmmpc.nif
   IBMMPC_nif5 = ibmmpc.nif
   IBMMPC_nif6 = ibmmpc.nif
   IBMMPC_nif7 = ibmmpc.nif
   IBMMPC_nif8 = ibmmpc.nif

[NETBIOS]

   DriverName = netbios$
   ADAPTER0 = netbeui$,0
   ADAPTER1 = netbeui$,1
   ADAPTER2 = netbeui$,2
   ADAPTER3 = netbeui$,3
;
; only four adapters may be initialized here - that's the NB30 interface
;

[netbeui_nif]

   DriverName = netbeui$
   Bindings = IBMMPC_nif,IBMMPC_nif2,IBMMPC_nif3,IBMMPC_nif4,
   IBMMPC_nif5,IBMMPC_nif6,IBMMPC_nif7,IBMMPC_nif8
;
; the Bindings statement must go on a single line
;
   ETHERAND_TYPE = "I"
   USEADDRREV = "YES"
   OS2TRACEMASK = 0x0
   SESSIONS = 254
   NCBS = 254
   NAMES = 42
   SELECTORS = 15
   USEMAXDATAGRAM = "NO"
   ADAPTRATE = 1000
   WINDOWERRORS = 0
```

```
    MAXDATARCV = 4352
    TI = 30000
    T1 = 1000
    T2 = 200
    MAXIN = 1
    MAXOUT = 1
    NETBIOSTIMEOUT = 2000
    NETBIOSRETRIES = 3
    NAMECACHE = 1000
    RNDOPTION = 1
    PIGGYBACKACKS = 1
    DATAGRAMPACKETS = 10
    PACKETS = 330
    LOOPPACKETS = 8
    PIPELINE = 5
    MAXTRANSMITS = 6
    MINTRANSMITS = 2
    DLCRETRIES = 10
    FCPRIORITY = 5
    NETFLAGS = 0x1000

[IBMMPC_nif]

    DriverName = IBMMPC$
    MaxTransmits = 31
    MaxTxFrameSize = 18000
    MinRcvBuffs = 20
    SizWorkBuf = 2048
    MulticastNum = 16
    EnableTxEofInt = "YES"
    Enet20UTP = "NO"
    EnableHiPriTx = "NO"
    HiPriTxAccess = 5
    HiPriTxThresh = 4
    LLCOnly = "NO"

[IBMMPC_nif2]

    DriverName = IBMMPC$
    MaxTransmits = 31
    MaxTxFrameSize = 18000
    MinRcvBuffs = 20
    SizWorkBuf = 2048
    MulticastNum = 16
    EnableTxEofInt = "YES"
    Enet20UTP = "NO"
    EnableHiPriTx = "NO"
    HiPriTxAccess = 5
    HiPriTxThresh = 4
    LLCOnly = "NO"

[IBMMPC_nif3]

    DriverName = IBMMPC$
    MaxTransmits = 31
    MaxTxFrameSize = 18000
    MinRcvBuffs = 20
    SizWorkBuf = 2048
    MulticastNum = 16
    EnableTxEofInt = "YES"
    Enet20UTP = "NO"
    EnableHiPriTx = "NO"
    HiPriTxAccess = 5
```

```
   HiPriTxThresh = 4
   LLCOnly = "NO"

[IBMMPC_nif4]

   DriverName = IBMMPC$
   MaxTransmits = 31
   MaxTxFrameSize = 18000
   MinRcvBuffs = 20
   SizWorkBuf = 2048
   MulticastNum = 16
   EnableTxEofInt = "YES"
   Enet20UTP = "NO"
   EnableHiPriTx = "NO"
   HiPriTxAccess = 5
   HiPriTxThresh = 4
   LLCOnly = "NO"

[IBMMPC_nif5]

   DriverName = IBMMPC$
   MaxTransmits = 31
   MaxTxFrameSize = 18000
   MinRcvBuffs = 20
   SizWorkBuf = 2048
   MulticastNum = 16
   EnableTxEofInt = "YES"
   Enet20UTP = "NO"
   EnableHiPriTx = "NO"
   HiPriTxAccess = 5
   HiPriTxThresh = 4
   LLCOnly = "NO"

[IBMMPC_nif6]

   DriverName = IBMMPC$
   MaxTransmits = 31
   MaxTxFrameSize = 18000
   MinRcvBuffs = 20
   SizWorkBuf = 2048
   MulticastNum = 16
   EnableTxEofInt = "YES"
   Enet20UTP = "NO"
   EnableHiPriTx = "NO"
   HiPriTxAccess = 5
   HiPriTxThresh = 4
   LLCOnly = "NO"

[IBMMPC_nif7]

   DriverName = IBMMPC$
   MaxTransmits = 31
   MaxTxFrameSize = 18000
   MinRcvBuffs = 20
   SizWorkBuf = 2048
   MulticastNum = 16
   EnableTxEofInt = "YES"
   Enet20UTP = "NO"
   EnableHiPriTx = "NO"
   HiPriTxAccess = 5
   HiPriTxThresh = 4
   LLCOnly = "NO"
```

```
[IBMMPC_nif8]

   DriverName = IBMMPC$
   MaxTransmits = 31
   MaxTxFrameSize = 18000
   MinRcvBuffs = 20
   SizWorkBuf = 2048
   MulticastNum = 16
   EnableTxEofInt = "YES"
   Enet20UTP = "NO"
   EnableHiPriTx = "NO"
   HiPriTxAccess = 5
   HiPriTxThresh = 4
   LLCOnly = "NO"
```

The following lines are extracted from the LANTRAN.LOG file to reflect what messages will be logged when the Adapter and Protocol Services device drivers are initialized and a configuration as shown above is being used:

```
IBM OS/2 LANMSGDD [02/16/99] 5.05 is loaded and operational.
IBM OS/2 NETBEUI 5.50.0
NETBEUI: Using a 32-bit data segment.
Installing NETWKSTA.200 Version 6.0. IBM LAN Redirector ( Feb 16, 1999)

IBM OS/2 NETBIOS 4.0
Adapter 0 has 34 NCBs, 153 sessions, and 28 names available to NETBIOS applications.
Adapter 1 has 34 NCBs, 153 sessions, and 28 names available to NETBIOS applications.
Adapter 2 has 34 NCBs, 153 sessions, and 28 names available to NETBIOS applications.
Adapter 3 has 34 NCBs, 153 sessions, and 28 names available to NETBIOS applications.
NETBIOS 4.0 is loaded and operational.
IBM Streamer Family adapter NDIS device driver Version 4.01.00
Initialization proceeding for section IBMMPC_NIF in PROTOCOL.INI
Initialization proceeding for section IBMMPC_NIF2 in PROTOCOL.INI
Initialization proceeding for section IBMMPC_NIF3 in PROTOCOL.INI
Initialization proceeding for section IBMMPC_NIF4 in PROTOCOL.INI
Initialization proceeding for section IBMMPC_NIF5 in PROTOCOL.INI
Initialization proceeding for section IBMMPC_NIF6 in PROTOCOL.INI
Initialization proceeding for section IBMMPC_NIF7 in PROTOCOL.INI
Initialization proceeding for section IBMMPC_NIF8 in PROTOCOL.INI
IBM LANVDD is loaded and operational.
IBM OS/2 LAN Netbind
Slot 3A:  IBM Streamer Family adapter universal address is 08005a6c072c
Slot 3A:  IBM Streamer Family adapter opened for:  Token Ring, 16 Mbps.
Slot 3B:  IBM Streamer Family adapter universal address is 08005a6c072d
Slot 3B:  IBM Streamer Family adapter opened for:  Token Ring, 16 Mbps.
Slot 4A:  IBM Streamer Family adapter universal address is 08005a6c08a8
Slot 4A:  IBM Streamer Family adapter opened for:  Token Ring, 16 Mbps.
Slot 4B:  IBM Streamer Family adapter universal address is 08005a6c08a9
Slot 4B:  IBM Streamer Family adapter opened for:  Token Ring, 16 Mbps.
Slot 6A:  IBM Streamer Family adapter universal address is 08005a6cd11a
Slot 6A:  IBM Streamer Family adapter opened for:  Token Ring, 16 Mbps.
Slot 6B:  IBM Streamer Family adapter universal address is 08005a6cd11b
Slot 6B:  IBM Streamer Family adapter opened for:  Token Ring, 16 Mbps.
Slot 7A:  IBM Streamer Family adapter universal address is 08005a1e4772
Slot 7A:  IBM Streamer Family adapter opened for:  Token Ring, 16 Mbps.
Slot 7B:  IBM Streamer Family adapter universal address is 08005a1e4773
Slot 7B:  IBM Streamer Family adapter opened for:  Token Ring, 16 Mbps.
```

The following lines are extracted from the IBMLAN.INI file to reflect what statements must be configured to support the configuration as shown above:

```
[networks]

   net1 = NETBEUI$,0,LM10,101,220,14
   net2 = NETBEUI$,1,LM10,101,220,14
   net3 = NETBEUI$,2,LM10,101,220,14
   net4 = NETBEUI$,3,LM10,101,220,14
   net5 = NETBEUI$,4,LM10,101,220,14
   net6 = NETBEUI$,5,LM10,101,220,14
   net7 = NETBEUI$,6,LM10,101,220,14
   net8 = NETBEUI$,7,LM10,101,220,14

   ...

[requester]

   ...
   wrknets = net1,net2,net3,net4,net5,net6,net7,net8
   ...

[server]

   ...
   srvnets = net1,net2,net3,net4,net5,net6,net7,net8
   ...
```

3.3.4 DOS and Windows LAN applications on OS/2

Adapter and Protocol Services also provide virtual device drivers to allow DOS and Windows applications to use the NetBIOS and IEEE 802.2 protocol services. The device drivers are loaded when the following statements are contained in the CONFIG.SYS file:

DEVICE=F:\IBMCOM\PROTOCOL\LANPDD.OS2

DEVICE=F:\IBMCOM\PROTOCOL\LANVDD.OS2

LANPDD.OS2 is the virtual IEEE 802.2 protocol driver, and LANVDD.OS2 is the virtual NetBIOS protocol driver for DOS and WIN-OS2 sessions. Both drivers are required to support any of these interfaces on DOS and WIN-OS2 sessions. These statements will be added by Adapter and Protocol Services automatically when you configure the NetBIOS or IEEE 802.2 protocol for at least one adapter.

To reserve NetBIOS and IEEE 802.2 resources for a DOS or WIN-OS2 session, you have to include the `LTSVCFG` command in the AUTOEXEC.BAT file. This command takes to following parameters:

LTSVCFG

C = number of NetBIOS commands

D = IEEE 802.2 direct station support

N = number of NetBIOS names

N1 = NetBIOS name #1 support

S = number of NetBIOS sessions

/ separator between multiple adapter configurations

The resources specified with the `LTSVCFG` command are taken from the pool of resources that is defined in the PROTOCOL.INI file. The MPTS Configuration Guide on-line book provides configuration and application settings examples for the virtual device drivers.

3.3.5 NetBIOS over TCP/IP on OS/2 Warp Server for e-business

Clients and servers need to know how to find one another in order to share information. The NetBIOS conventions built into DOS and OS/2 clients/servers use 16 byte NetBIOS names, which refer to one another by name. Different applications on the same PC use different names to represent their applications.

NetBIOS names, such as *Corinna's PC*, *Server F.v.S-Marketing* or *LAB-Printer IBM1* can be built into programs or solicited from humans with relative ease. NetBIOS names can be used as unambiguous identifiers even if a station is moved to another location. However, to send one another packets of information, the TCP/IP protocol drivers of the respective PCs must refer to one another by IP address. The problem then exists of having to translate NetBIOS Names into IP addresses in order to effect PC-to-PC communication on an IP network.

To date, this translation has been handled in one of two ways: by use of static tables residing on each client and server, or by the use of (dynamic) broadcast queries (packets sent to every client and server) asking in effect "Where is Corinna's PC?"

The problem with static tables is that they must be continually updated and maintained, an activity far more troublesome than the maintenance of IP addresses alone. Every time any new station is added to the network, all of its applications' names must be added to the static table of each other station that wants to send it data. And with static entries, though the name is always

mappable, there is no telling whether the named application is actually active at the time interaction is desired by another station.

The problem with broadcast queries is that IP networks cannot propagate broadcasts beyond a single (logical) cable segment. Resources located on the other side of a router from the broadcasting station will not receive the query. Every station on the same side of the router will be pestered with queries for which it doesn't know the answer.

A NetBIOS over TCP/IP protocol has been defined by the governing TCP/IP standards body, the Internet Engineering Task Force (IETF), which overcomes each of these problems. The IETF standard describes how NetBIOS stations may interact with a NetBIOS Name Server in order to dynamically register their own application names and to learn the name-to-address mappings of other applications.

Several components of OS/2 Warp Server can use NetBIOS for communications, but they can also use other protocols like TCP/IP or IPX. File and Print Sharing Services remains the only OS/2 Warp Server component that can only use NetBIOS as a programming interface.

The original NetBIOS protocol has some specific characteristics which limit its use in certain wide area network environments:

- The NetBIOS protocol uses a flat name space.
- The NetBIOS protocol relies on the broadcast technique to register/find a name.
- The NetBIOS protocol cannot be routed.

One solution to overcome these limitations can be found in RFCs 1001 and 1002. They describe the standard way to implement the NetBIOS services on top of the TCP and UDP protocols. Adapter and Protocol Services provide a full TCP/IP protocol stack and a TCPBEUI protocol stack, which is a ring 0 implementation of RFC 1001/1002.

Another solution of routing NetBIOS is to use the NetBIOS over IPX protocol driver, which is also supplied with Adapter and Protocol Services.

Note

In order to use NetBIOS over TCP/IP, you do not need to install the TCP/IP Services of OS/2 Warp Server for e-business since the support for this combination of protocols is fully included in Adapters and protocol services. TCP/IP Services of OS/2 Warp Server means TCP/IP applications on top of the TCP/IP protocol, such as FTP, LPR, DHCP and DDNS.

The capability of running NetBIOS applications over routable protocols offers new flexibility when designing OS/2 Warp Server for e-business networks. OS/2 Warp Server for e-business systems, OS/2 Warp Server, Warp Connect Peer workstations, LAN Servers, and LAN Requester workstations can be on remote LAN segments connected by IP routers. This also means that such systems can be introduced into existing TCP/IP networks without introducing an additional network protocol (NetBIOS). There are several defined classes of NetBIOS over TCP/IP implementations specified by RFCs 1001 and 1002.

3.3.5.1 Resolving NetBIOS names to IP addresses

When the client connects to the server, the client must resolve the server's NetBIOS names. There are various ways to resolve NetBIOS names into IP addresses.

The simplest method is to try broadcasting name queries and hope the computer you are trying to communicate with responds in order to do the resolution. Computers running this mode are said to be a Broadcast Node (B-Node). One disadvantage in a large network environment is that broadcast traffic becomes a nuisance or is so dispersed that broadcasts cannot be made to reach all stations.

A better method would be to have a server on the network that could do the NetBIOS name to IP address resolution for us. Such a server is known as a NetBIOS Name Server or NBNS. Computers that are configured to use an NBNS are Peer-to-Peer Note (P-Node) clients. Broadcasting is never used by P-Node clients.

Finally, we could combine the modules to become either Mixed Node (M-Node) or Hybrid Node (H-Node) clients.

An M-Node client will attempt a name query broadcast first, and, if that fails, it will try to use the NBNS.

H-Node clients will attempt a name query broadcast first, and, if that fails, will try a name query broadcast. If an H-Node client detects that an NBNS has

failed, it will continue to poll the NBNS (while using broadcast) so that it knows when to switch back to using the NBNS. H-Node has generally replaced M-Node.

3.3.5.2 Additional information about OS/2 TCPBEUI

OS/2 TCPBEUI is a high performance, ring 0 implementation of NetBIOS over TCP/IP. TCPBEUI provides the LM10 protocol driver interface. It is the same LM10 functionality that is also provided by NetBEUI. Figure 30 shows this interface. TCPBEUI maps NetBIOS API calls into the TCP/IP protocol. NetBIOS over TCP/IP contains enhancements over the B-node standard which improves system performance by decreasing broadcast frames and by expanding communications over routers and bridges. These enhancements are transparent to NetBIOS applications and do not interfere with other B-node implementations that lack similar functions.

RFC 1001/1002 is not an encapsulation technique, but, rather, builds special packets and sends them out via UDP and TCP. For example, once a NetBIOS session has been established, TCPBEUI will use sockets-send commands over a TCP connection to send NetBIOS session data. TCPBEUI builds a four-byte session header that precedes the actual user data. Thus, a NetBIOS Chain Send of 128KB would have an overhead of only four bytes.

TCPBEUI allows peer-to-peer communication over the TCP/IP network with other computers that have compatible services. Figure 30 shows the relationship between the NetBIOS, NetBIOS over TCP/IP, and TCP/IP protocol stacks as implemented in Adapter and Protocol Services.

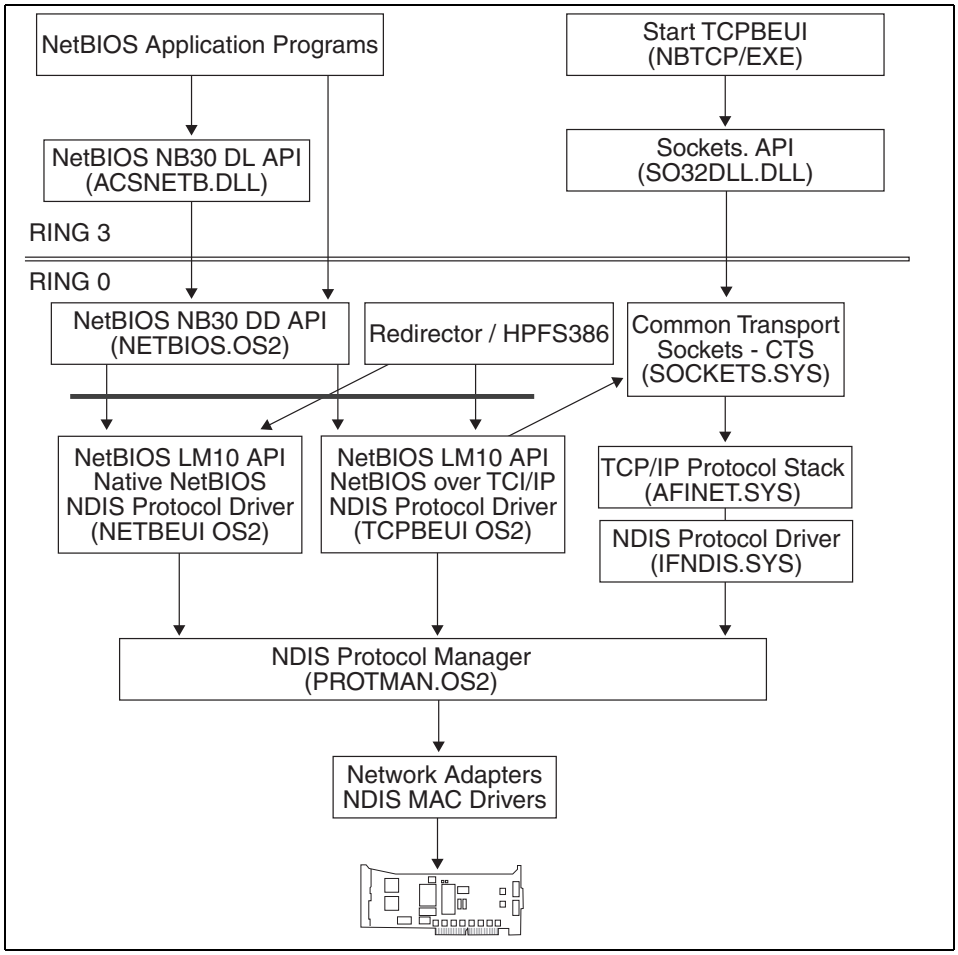

Figure 30. NetBIOS, NetBIOS over TCP/IP and TCP/IP structure

Unlike NETBEUI.OS2, the TCPBEUI.OS2 program does not directly communicate with the NDIS interface. The dotted line in the figure indicates TCPBEUI has a BINDINGS statement in the PROTOCOL.INI file, but a bind process is only required in order to create a control block area.

Figure 30 also illustrates how NetBIOS applications can use both NETBEUI and TCPBEUI protocol stacks. ACSNETB.DLL provides the ring 3 NetBIOS DLL API for application programs. Ring 3 NetBIOS commands are sent to NETBIOS.OS2 for processing. NETBIOS.OS2 provides the ring 0 NetBIOS DLL API for applications and other device drivers to use, and it binds to one or more LM10 (LAN Manager 1.0) transport protocol drivers. The LAN

redirector component of File and Print Sharing Services (NETWKSTA.200) and HPFS386 use the LM10 interface.

Support for NetBIOS over TCP/IP can easily be added to the existing NetBIOS structure since the Warp Server Install program supports up to four LM10 interfaces. It is provided by having NETBIOS.OS2 bind to TCPBEUI.OS2. To enable NETWKSTA.200 to use TCPBEUI, there must be a NETx (where x is 1, 2, 3, 4, for example) statement in the IBMLAN.INI file configured appropriately.

Data transfer to LAN is handled by a MAC device driver, for example, the IBMTOK.OS2 device driver.

3.3.5.3 TCPBEUI coexistence with NetBEUI

Adapters and protocol services provide the capability of configuring NetBIOS applications, especially File and Print Sharing Services, with both NetBEUI and TCPBEUI on the same network interface card. This dual protocol stack configuration will allow local sessions to continue running with NetBEUI performance while also providing wide area network connectivity with NetBIOS over TCP/IP.

Figure 31 shows an example scenario with both TCP/IP and NetBIOS protocols being used and TCP/IP Services installed on a server. In this example, LAN Client A is able to access File and Print Sharing Services resources on the OS/2 Warp Server B on the local LAN segment via NetBIOS, and the OS/2 Warp Server C on the remote LAN segment across the IP network via TCPBEUI. In addition, it is able to use the TCP/IP applications provided by TCP/IP Services to access local and remote TCP/IP hosts via the native TCP/IP protocol.

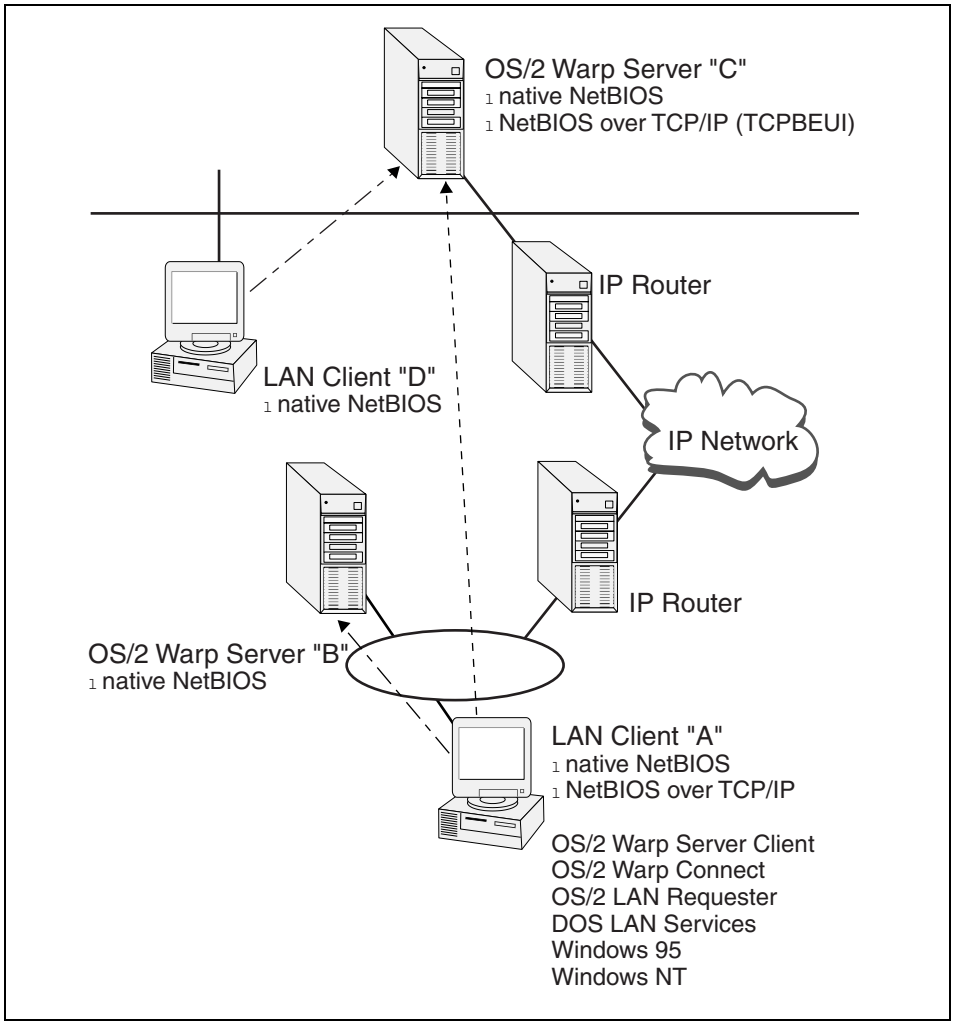

Figure 31. Coexistence TCPBEUI

Note

When configuring Adapter and Protocol Services for both NetBEUI and TCPBEUI, even though a single LAN adapter is present in the workstation, the two protocols need to be configured on different logical adapters. File and Print Sharing Services handle this configuration as if there were two adapters present. Therefore, two NET entries will be made in IBMLAN.INI.

3.4 MPTS - Strong encryption

The strong encryption replacement version of MPTS is part of the IBM OS/2 Warp Server for e-business Security Feature. This consists of two CDs. CD 1 contains the replacement version of Multi-Protocol Transport Services (MPTS). CD 2 contains Lotus Domino Go Webserver 4.6.2.6 and IBM WebSphere (TM) Application Server 1.1.

This version of MPTS contains the SSL and IPSEC libraries used by other applications. The IPSEC libraries provide 56-bit encryption, and the SSL libraries provide 56/128-bit encryption depending on your location. This version of MPTS supports strong encryption for Lotus Domino Go Webserver 4.6.2.6 and IBM WebSphere Application Server 1.1.

To install this version of MPTS, perform the following steps:

1. Insert the OS/2 Warp Server for e-business Security Feature CD 1 into the CD-ROM drive.
2. Open an OS/2 window.
3. From an OS/2 command prompt, type `e:` and press **Enter**. `e` is the drive letter of the CD-ROM drive.
4. Type `cd mpts` and press **Enter**.
5. Type `cd language` and press **Enter**. `language` is the two-character abbreviation of the language version you are installing.
6. Type `install` and press **Enter**.
7. Follow the displayed installation instructions.

This will replace a number of the MPTS component files that are installed onto your hard disk. You will need to reboot to make those changes effective.

3.5 Removing adapters and protocol services

Since Adapters and protocol services are a key feature of OS/2 Warp Server for e-business providing communications support to all other components, it should not be removed. Nevertheless, you can remove MPTS/LAPS by pressing the **Remove** Button in the Configuration Menu.

Chapter 4. Journaled file system and logical volume manager

There has been a rapid explosion in the data storage requirements of today's corporate servers especially when they are used to house information for the Internet. This rise in storage capacity has facilitated the need for improved file system management and recovery. Businesses want more convenient methods to expand with their disk requirements and no longer want to endure the lengthy file system recovery times that are the hallmarks of non-journaled file systems, such as FAT and HPFS.

IBM's OS/2 Warp Server for e-business introduces two key features to help meet these server demands - the Logical Volume Manager (LVM) and the Journaled File System (JFS). JFS is designed for high throughput and reliability as well as providing quick recovery times that are essential factors in improving server availability and performance. LVM helps balance dynamic file requirements by allowing partitions to be expanded dynamically and volumes to span across physical disks.

This chapter begins with an introduction to data storage through a look at the enhanced bundled support of files systems. It will then take an in-depth look into the log-awaited additions of LVM and JFS to OS/2 Warp Server for e-business.

4.1 OS/2 file system support

A file system is essentially a hierarchical structure of directories with each directory containing either more directories (referred to as subdirectories) or files. The main purpose of the file system is to improve management of data by allowing the data to be organized and easily managed.

OS/2 Warp Server for e-business comes bundled with support for the following types of file systems: JFS, FAT, HPFS, HPFS386 (with proof of purchase), NFS and CDFS. This section summarizes the key features of each of the supported file systems and discusses where it would be most appropriate to use them.

4.1.1 File allocation table

One of the simplest methods for storing data on disk is with FAT. FAT is simply an unsorted linear table designed for locating directories and files. It has low memory and disk administrative overheads. Unfortunately, FAT was only designed for small disk systems, namely floppy disks, because it was optimized for performance by reading the FAT content into memory. With the

advent of large hard drives, FAT has become slow because memory cannot hold all the information required for the file transaction in memory. This has meant that the hard drives have had to make several head movements to read a single file, which has had the overall effect of slowing down the system.

The original implementation of FAT had a partition limit of 32MB, which later got increased to 2GB. One of the biggest problems with FAT has been the automatic incremental jumps in cluster size with disk partition sizes (see Table 6). This has lead to large amounts of disk storage wastage for small files. For example, a file which is only 400 bytes long, when stored on a 512MB disk, occupies an 8KB cluster which is a 95 percent waste of space. Another problem stems from the way pieces of data information are written to disk one after another in a chain starting from the outer edge of the disk filling towards the middle then starting at the outer edge again. This quickly leads to disk fragmentation as space forms in the middle of the chain from file deletion. As a file gets modified, parts of it can end up being stored in different areas of the disk resulting in excessive head movement and slow response times. Finally, FAT is restricted to an 8.3 naming convention, which often prevents meaningful names from being given to files.

Table 6. Relationship between partition size and cluster size

FAT Cluster Size	
Partition Size	**Cluster Size**
diskettes	1 sector
< 16MB	2 sectors
16 to 128MB	4 sectors
128 to 256MB	8 sectors
256 to 512MB	16 sectors
512 to 1024MB	32 sectors
1024 to 2048MB	64 sectors

Enhancements have been made to FAT under OS/2 to improve performance, such as large cache, lazy-writing, bad sector bypass and bitmap to track free clusters on disk. OS/2 also allows a FAT partition to have additional Extended Attributes (EAs). Since there is no room for these attributes in the FAT directory, OS/2 creates a separate hidden file, EA DATA. SF, on the disk volume and stores the EA information in this file.

FAT support is, largely, provided for transferring data between different operating systems and is commonly used on floppy disks and on OS/2 clients to allow DOS and Windows applications to run. FAT is also used on OS/2 clients if the system has only a small hard drive (<100MB) and/or does not have much memory (<8MB) since HPFS can end up reducing performance in these situations.

4.1.2 High performance file system

For OS/2, a new file system was created specifically for the multitasking environment and with large disk support in mind. HPFS also followed a new concept in file system implementation called the Installable File System (IFS). An IFS driver, which contains code needed to manage media formats other than DOS to be loaded during the system boot phase.

HPFS has, by design, a better structure that prevents data fragmentation. This is due to the new disk layout whereby the tables that describe the location of files and freespace are positioned at regular intervals throughout the partition from the center of the disk. Also, new file and directory information is written where there is enough freespace. This reduces fragmentation and prevents excessive movement of the disk arm. HPFS maintains a 512byte allocation unit (cluster size) no matter how large the partition becomes and also supports long file names up to 255 characters. Critical for the high performance of HPFS is the caching technology it uses to access the disk. The lazy-write design writes data to a memory cache before writing to disk giving faster overall performance. The drawback is that data is lost if the computer is suddenly turned off preventing the cache from being written to disk.

Data location by the file system is also improved in HPFS through the use of a sorted B-tree (Balanced tree) structure to store the file information. This speeds up the search for files since HPFS can quickly transverse down the correct branch of the tree to locate the file, whereas FAT has to look sequentially through an unsorted linear table. The drawback of the process is that it takes slightly longer to write the file, but this is greatly compensated for by the speed of locating a file.

Note

In the majority of cases, HPFS will be superseded by the JFS with the exception of boot partitions because the JFS is not yet bootable. The boot partition needs to be either HPFS or HPFS386 since it is needed by Java for longname support. Many OS/2 Warp server for e-business configuration notebooks have been written in Java.

4.1.3 HPFS386

With the release of the 80386 (and higher) class of microprocessors, the HPFS design was able to be further enhanced and tailored for the networking environment resulting in a new version of HPFS called HPFS386. Bundled with the Advanced version of the server code, HPFS386 provided extremely fast access to large disk volumes and optimized performance in the server environment.

HPFS386 was developed to tightly integrate with the server code at ring 0 allowing data to be transferred directly from the HPFS386 cache to the network adapter driver. This resulted in accelerated network I/O leading to faster data access by clients. Other features incorporated into HPFS386 were greatly increased cache size, local security, software fault tolerance support for RAID-1 (mirroring and duplexing), directory/user-based disk space limitation, and increased file access related limits. An advantage introduced with HPFS386 was the storage of access control lists (ACLs) within the directory and file structures, thus, making fast user verifications leading to quicker client access to the file. ACL information in HPFS is stored in a separate file called NET.ACC; the overhead of accessing this file slows down verification.

HPFS386 is ideal in print server environments. For file server implementation, the decision between HPFS386 and JFS is dependent upon the server configuration; see Chapter 4.3.8, “Performance considerations” on page 142 for more details. The biggest drawback with HPFS386 comes when there is a system crash of the server with large amounts of disk. Recovery still relies on `CHKDSK` to check and fix disk errors, which can be a very time consuming process.

Note

JFS, in its first iteration, is not intended as a replacement for HPFS386 since it does not support features like Local Security, which HPFS386 provides.

4.1.4 Journaled file system

The JFS provides the same range of supported file system operations for organizing and managing physical files, which was traditionally provided by HPFS. However, in its first release, the JFS is not bootable. The JFS is created in a logical volume and can be expanded across disks to meet disk requirements. It supports both file and partition sizes of up to two attributes and can recover from a system crash within minutes. The JFS has been enhanced for performance scalability on Symmetrical Multiprocessor (SMP)

systems and for TCP/IP transactions making it ideal for Web and Lotus Notes data serving. Another feature of the JFS is its support of Sparse Files making it ideal for large databases. Sparse files allow large database structures to be defined but occupy only the amount of disk space that is consumed by the database, thus, a 2TB file can be created on a 2GB partition if there is less than 2GB of actual data to be stored. The JFS is explored in more detail later in this chapter.

4.1.5 HPFS386 features not available with JFS

HPFS386 has, for some time, been the most significant file system available for OS/2. Many additional features have, over time, been developed based on HPFS386. Some of these features are not yet available for JFS.

4.1.5.1 DASD limits

The current version of JFS shipped with OS/2 Warp Server for e-business does not support DASD limits. Depending on your requirements, there are several possible workarounds; two of them are:

1. Keep the resources that need directory limits on an HPFS386 formatted volume.
2. Use CHKSTOR as a replacement if it is sufficient to send an alert to the Administrator when the limit is exceeded.

4.1.5.2 Fault Tolerance

There is no replacement for the HPFS386 Fault Tolerance feature in JFS. Current server hardware usually comes with RAID adapters that can be used to perform this function as a hardware solution.

If you still need to rely on the software disk mirroring provided by HPFS386, you will need to keep HPFS386.

4.1.6 Comparison of features of FAT, HPFS, and HPFS386 JFS

The following table summaries the difference in features between FAT, HPFS, HPFS386 and JFS.

Table 7. Feature comparison between FAT, HPFS, HPFS386, and JFS.

	FAT	HPFS	HPFS386	JFS
Cache	14Mb	2MB[1]	Unlimited[2]	Unlimited[2]
Max. no of file opens	64000	64000	64000	64000
Max. no of file finds	3072	3072	8192	32768

	FAT	HPFS	HPFS386	JFS
Max. partition size	2GB	64GB	64GB	2TB
ACL Support	NET.ACC	NET.ACC	in Fnode	in Inode
ACL Limit	8192[3]	8192[3]	unlimited	unlimited
Max. no of connections	16384	16384	16384	16384
Max. no of Shares	1500	1500	1500	1500
Max. no of file per directory	512[4] on the root directory	limited by DASD[5] space	limited by DASD[5] space	4 billion
Bad block relocation	No	Yes	Yes	Yes via LVM
Software Fault Tolerance	No	No	Yes	No

[1]With HPFS cache, the 2MB limit is due to limits imposed by the 16-bit IFS architecture.
[2]This would equate to approximately 1 GB, which would be shared between cache and heap.
[3]ACL is limited to 8192 for FAT, CDROMs and HPFS are the information is stored in the last half of NET.ACC.
[4]FAT imposes a limit on the root directory that can be 112 or less (especially if WIN95 long filenames need to be supposed). For Hard disks, it generally has a value of 256 or 512 hardcoded in the boot sector. Although subdirectories are theoretically unlimited, performance can be seriously degraded if the number of files per directory is greater than *number of buffers* (set in config.sys) * *16*.
[5]Direct access storage device.

4.1.7 Network file system

NFS allows files and directories located on remote UNIX file systems to be incorporated into the local file system and accessed as though they were a part of the local server. NFS has been available separately for OS/2, but, with OS/2 Warp server for e-business, it has been integrated into the server. This integration, then, allows the server to share these NFS mounted files and directories just like any other resource to its clients.

4.1.8 CD-ROM file system

This type of file system allows the contents of a CD-ROM to be accessed as if it were part of the local file system.

4.2 Logical volume manager

Even with the rapid advances in DASD (Direct Access Storage Device), modern hard disk capacity still has not managed to satisfy corporate demands. Coupled with continual demands for increasing DASD storage, such as spanning across physical disks, a new solution was needed to simplify the management of DASD within OS/2. An integral part of the solution is the introduction of the Logical Volume Manager (LVM) with OS/2 Warp Server for e-business.

The Logical Volume Manager (LVM) provides an abstraction layer for the file system in that it makes the logical/physical makeup of the volume transparent to both applications and users. The LVM exploits new features of the underlying logical volume management system replacing the fdisk utility, thus, allowing flexible division, allocation, and management of the system disks. Logical Volume Manager(LVM.EXE or LVMGUI.CMD) replaces the Fixed Disk (FDISK) utility from the previous versions of OS/2. It provides all the function of FDISK and it provides additional features, such as disk spanning, dynamic resizing and sticky drive letters.

4.2.1 LVM terminology

There are a number of specialized terms used to describe the various entities that comprise the Logical Volume Manager.

Physical partition

A partition is a portion of the physical disk that functions as though it were physically a separate unit. Each partition can be either a primary or an extended partition. The first physical disk must always have a primary partition. A primary partition is a bootable partition that cannot be subdivided any further, whereas, an extended partition is one that can be subdivided into smaller (logical) partitions.

Logical partition

A logical partition is a subdivision within an extended partition, which is seen as a separate unit.

Logical volume

A logical volume can define a single partition or a collection of partitions. Logical volumes are presented to the user/application as a single unit, thus, the single drive letter. Logical volumes can be of two types: Either Compatibility (to older versions of OS/2) whereby they are formatted FAT or

HPFS, or, they are LVM whereby only the Logical Volume Manager will recognize them. An LVM volume can be formatted either FAT, HPFS, or JFS but cannot be made bootable, and, if formatted, JFS can be dynamically increased.

Figure 32 on page 97 shows the relationship between these component terms.

4.2.2 Overview of LVM

The objective of LVM is to provide a logical layer for the support of Logical Volumes for the OS/2 DASD I/O Subsystem. With the introduction of LVM support, a volume is no longer restricted to a logical partition wholly contained on a physical drive. Disk spanning is available. Also, dynamic partition resizing for JFS volumes is supported. This is best summarized by Figure 32 on page 97; it shows a partition spanned across three disks. Figure 32 also outlines the structure of how information is stored on the disk, the aggregate and the storage of data within it, and the fileset.

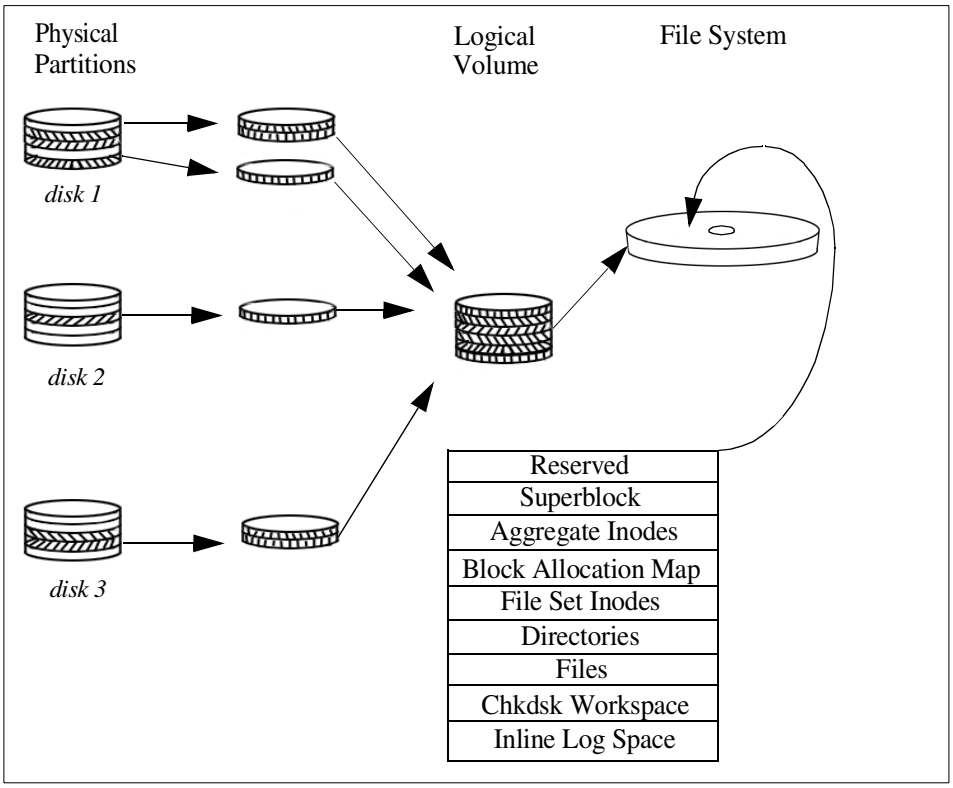

Figure 32. Overview of LVM disk management for a JFS volume.

Note

Only volumes that have not been formatted or are formatted JFS can be expanded. If a volume gets formatted with a file system other than JFS, it can no longer get expanded.

Figure 33 on page 98 illustrates the expansion of a JFS volume using LVM.EXE. In this example, the volume was expanded across two different physical disks.

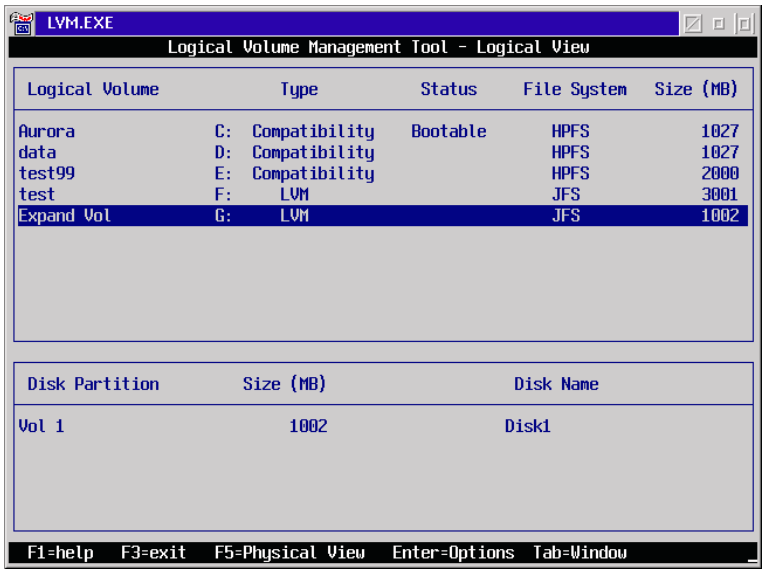

Figure 33. Expanding a volume, initial

Initially, the LVM volume called *Expand Vol* (G:) in Figure 33 on page 98 has only one volume of size 1GB in it on Disk 1 formatted JFS.

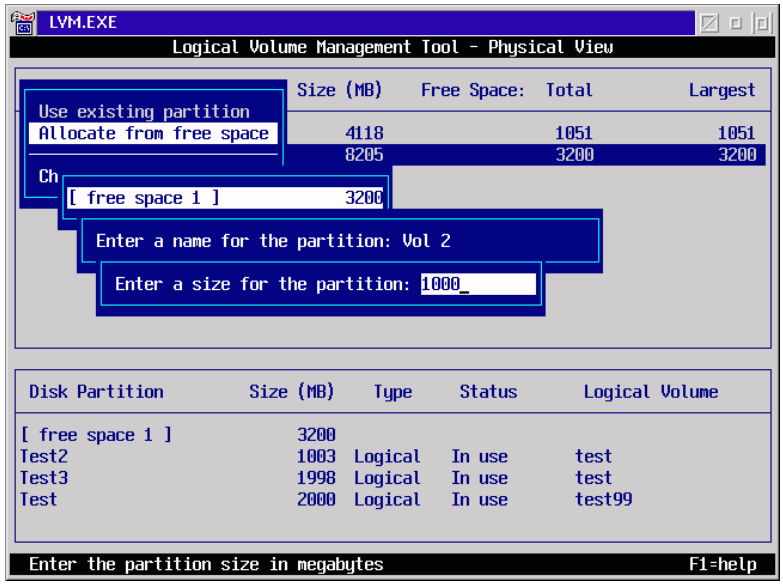

Figure 34. Expanding a volume, adding space

An additional 1GB is added to the volume in Figure 34 on page 98 from the second physical hard disk.

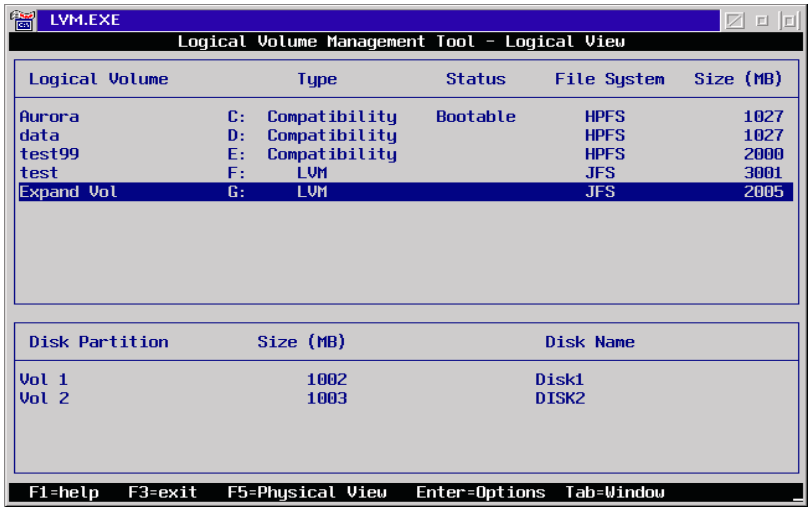

Figure 35. Expanding a volume, completed

In the lower portion of the screen in Figure 35 on page 99, you can now see that the volume Expand Vol now consists of two partitions: Vol 1 (Disk 1) and Vol 2 (Disk 2).

4.2.3 LVM and FDISK

Figure 36 on page 100 is a comparison of LVM and FDISK concepts.

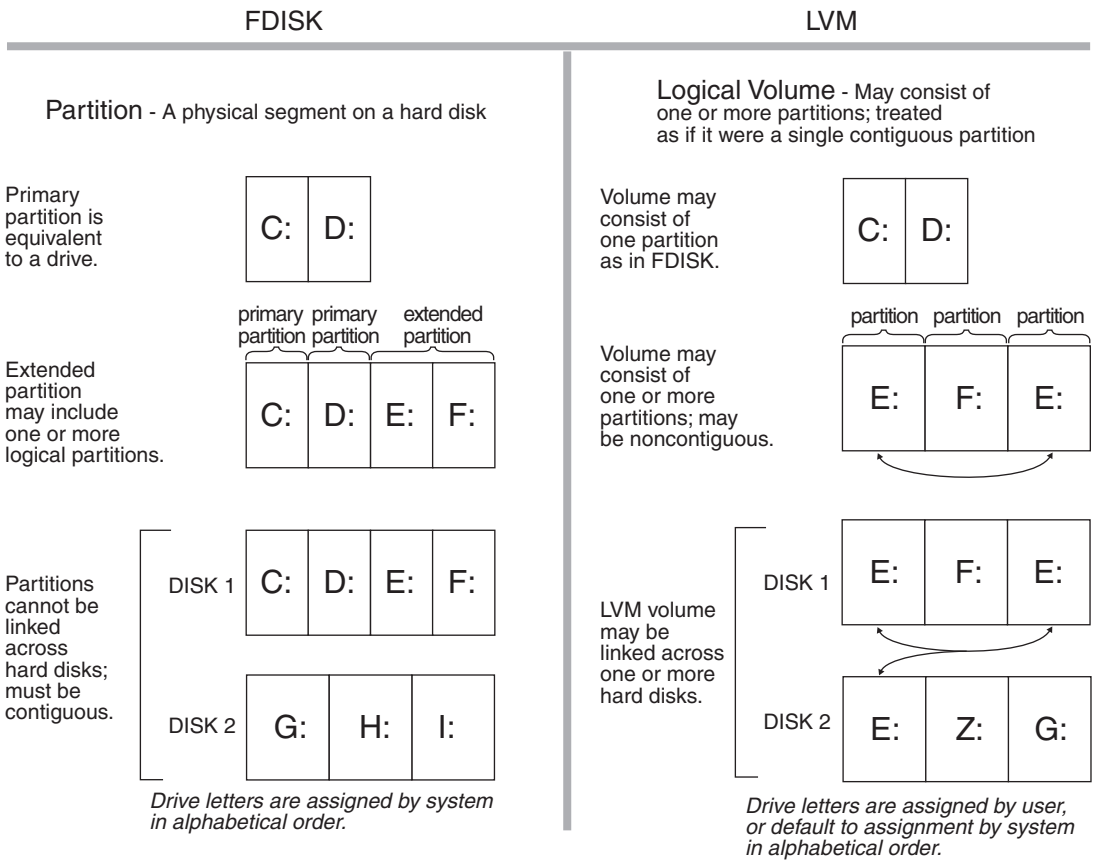

Figure 36. Comparison chart of FDISK and LVM concepts

4.2.4 Key components of the LVM

In order to implement the LVM functionality, some existing OS/2 components have been modified and a number of new components added. This section discusses these components.

4.2.4.1 Installable File System Manager (IFSM)/Kernel

The IFSM no longer communicates with the DASD (Direct Access Storage Device) I/O subsystem in terms of logical/physical location but instead in terms of logical volumes. This is accomplished by introducing an intermediary filter between the disk and the DASD manager (OS2DASD.DMD) called OS2LVM.DMD. OS2LVM.DMD also handles the drive letter assignments reported to OS2DASD, thus, allowing dynamic drive letter allocation. This

gives IFSM the capability to dynamically mount and unmount drives as part of the disk partitioning process and handle removable media with differing numbers of partitions.

The performance and SMP scalability has been improved by converting parts of the existing kernel execution environment (KEE) and IFSM from 16-bit to their 32-bit equivalent (namely KEE32 and IFSM32). This new kernel execution environment is used by JFS, LVM, OS2DASD, and TCP/IP. The 32-bit code provides the following benefits:

- The kernel File System Router layer, JFS, and OS2DASD (I/O path) are 32-bit resulting in improved execution time. The number of segment register loads is minimized.
- The I/O size is not limited to <64 KB in the 32-bit File System Router. Therefore, the router makes a single call to JFS for large application I/O sizes.
- The kernel File System Router will not acquire the spinlock (used to serialize access to a resource in an SMP or UNI environment) prior to calling the 32-bit JFS entry points for DosRead, DosWrite, and DosChgFilePtr. JFS/LVM/OS2DASD will use the new KEE spinlock APIs to provide more granular locking of its data structures resulting in improved SMP scalability. The spinlock will be acquired just prior to calling the DASD ADD and represents only a small percentage of the total I/O path.
- Prior to calling the JFS entry points, the kernel stack selector is switched from a 16-bit to a 32-bit selector. This eliminates the need for JFS/LVM/OS2DASD to perform SS to DS segment register conversions.
- The system services (KEE calls) are 32-bit and do not use a router. They are resolved by the OS/2 loader using KEE.LIB and the fix is applied directly to the calling code. The code path for a call is significantly reduced.

4.2.4.2 Device Manager - OS2LVM.DMD

OS2LVM is a new component. As mentioned above, it sits between the IFSM and OS2DASD providing the logical view of the DASD to the OS/2 file system. OS2LVM also provides Bad Block Relocation (BBR) for JFS; HPFS does its own BBR.

Note

The statement BASEDEV=OS2LVM.DMD must immediately follow the BASEDEV=OS2DASD.DMD statement in config.sys.

4.2.4.3 DASD Device Manager - OS2DASD.DMD

OS2DASD has been modified to become the physical interface to the DASD I/O subsystem. Requests are received from the OS2LVM as physical writes and reads to physical partitions. At initialization, OS2DASD communicates with OS2LVM in order to compile and retain a current logical/physical view of the attached DASD devices.

OS2DASD has been enhanced with 32-bit (called Strat3) entry points for handling calls from OS2LVM for JFS. HPFS and FAT still use the 16-bit Strat1/Strat2 calls. Since these calls are in 32-bit code, they benefit from a 32-bit flat address model and KEE32 kernel calls which result in improved performance.

Note

The 16-bit device driver strategy commands used by FAT & HPFS are referred to as Strat1 and Strat2. Strat1 commands are passed in via Request Packets and block the calling thread until completion while Strat2 (Strat3 is a 32-bit implementation) commands are passed via a Request List and return on the calling thread prior to completion.

This is summarized by Figure 37 on page 103.

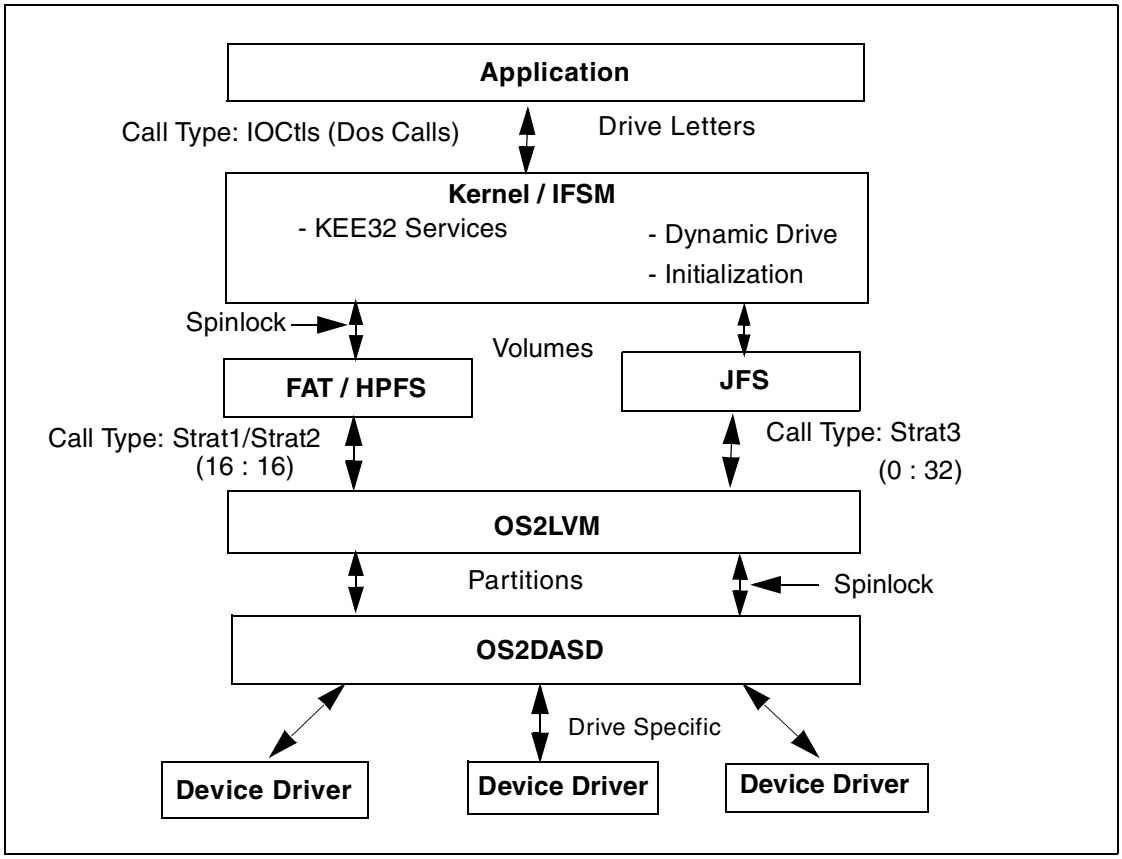

Figure 37. Block diagram summarizing DASD I/O path

4.2.4.4 LVM.EXE - Fixed Disk Utility

The new logical volume management tool, LVM.EXE, is a replacement for FDISK. It provides the capability to manage the system disks, however, it moves the emphasis from working from the physical to the logical view of the disk subsystem. LVM.EXE allows the creation of partitions as well as logical drives now called volumes. The default view is now a logical one for creating volumes; partitions are created during volume creation if necessary. More control of the partitioning can be obtained by changing the view from logical to physical; function key [F5] can be used to alternate between views. The physical view provides the same functionality as FDISK did.

LVM.EXE also provides the capability to exploit new features of the underlying logical volume management system. The main enhancements are the linked volumes and sticky drive letters. A new type of volume called LVM

allows the linking of several partitions from possibly several disks. The LVM volume is presented to the system as a single drive. A drive letter can be assigned to either a new LVM volume or an older compatibility volume. This assignment persists until changed by the user.

LVM is available as command line VIO and GUI (LVMGUI.CMD) and provides the following functionality:

- Create Compatibility Volumes (partitions) that can be seen by Pre-LVM releases of OS/2.
- Create Logical Volumes which span physical disks.
- Expand Logical Volumes (for JFS only).
- Delete Logical Volumes.
- Delete Compatibility Volumes.

Figure 38 on page 104 and Figure 40 on page 105 show the Logical and physical views of LVM respectively. Figure 39 on page 105 and Figure 41 on page 106 show the Logical and physical views of LVMGUI, which is the graphical user interface (GUI) replacement for FDISKPM.

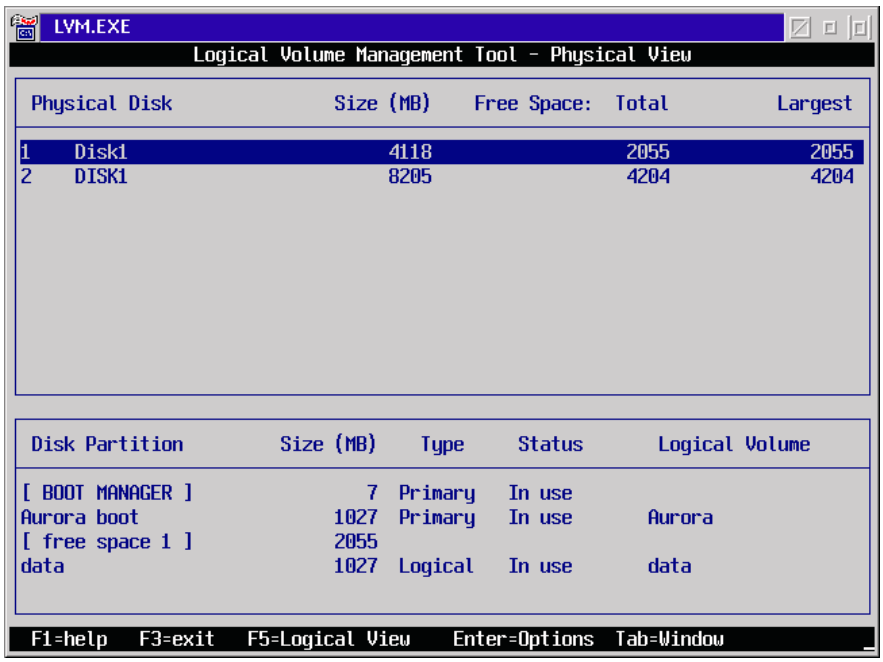

Figure 38. Physical view via LVM.EXE

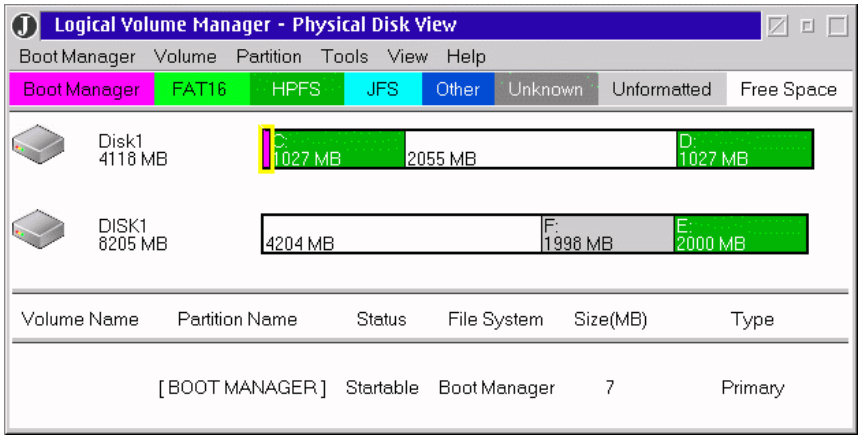

Figure 39. Physical view via LVMGUI.CMD

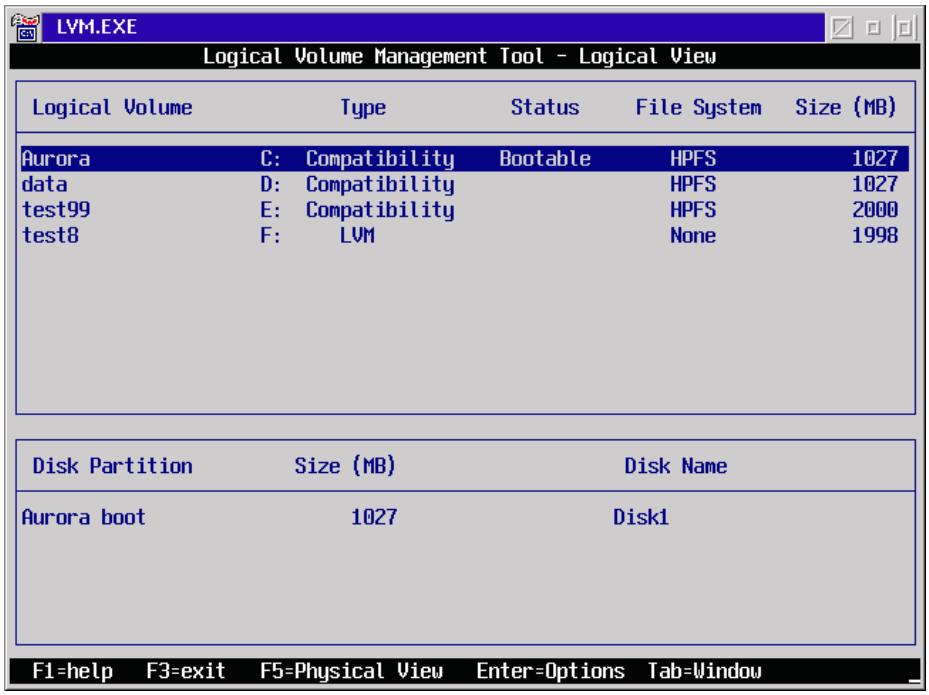

Figure 40. Logical view via LVM.EXE

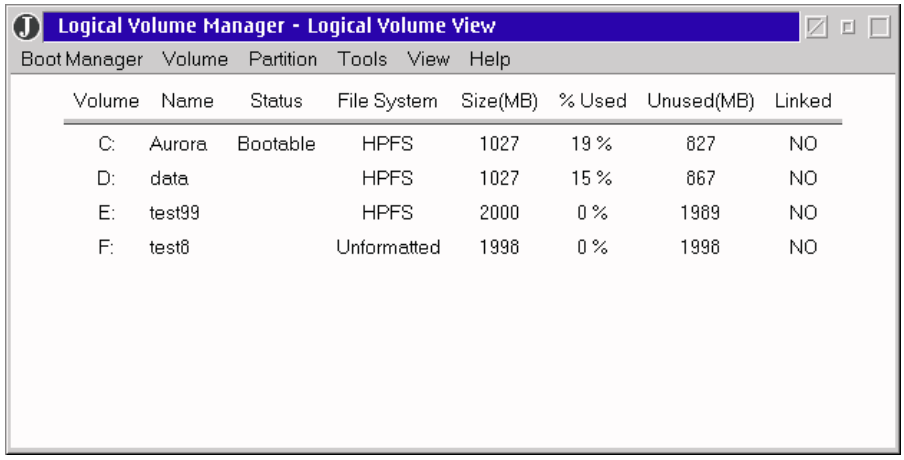

Figure 41. Logical view via LVMGUI.CMD

The syntax associated with the LVM.EXE utility has been redesigned from FDISK to utilize the new functionality being offered by LVM. This functionality is also reflected in the command line interface. The new command line interface syntax is summarized below, and its optional parameters are summarized by Table 8 and Figure 43 through Figure 45.

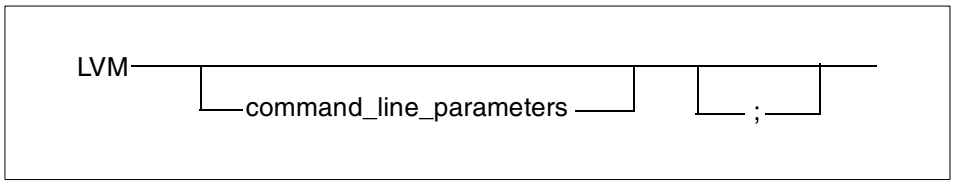

Figure 42. LVM command syntax

Note

You can only have one *command_line_parameter* if the */FILE* option is used.

Table 8. Configuration variables for LVM

Parameters for LVM.exe		
Command line parameters		
Option	Parameters	Description

Parameters for LVM.exe		
/FILE:	<file_name>	Response file with LVM commands.
/QUERY:	<Query_Parameters>	Query configuration.
/CREATE:	<Creation_Parameters>	Create a volume or Partition.
/DELETE:	<Deletion_Parameters>	Delete volumes or Partitions.
/HIDE	<volume_name>	Hide from OS/2.
/BOOTMGR:	physical drive number e.g. 1 [This can only be 1 or 2]	Install the boot manager on the specified physical drive.
/SETNAME:	<Name_Chg_Parameters>	Set volume, partition, or disk name.
/SETSTARTABLE:	<Setstartable_Parameters>	Set the startable flag on primary partitions or bootable volumes.
/NEWMBR	<drive_number>	Write a new Master Boot Record.
/EXPAND	<Expand_Parameters>	Expand a JFS volume.

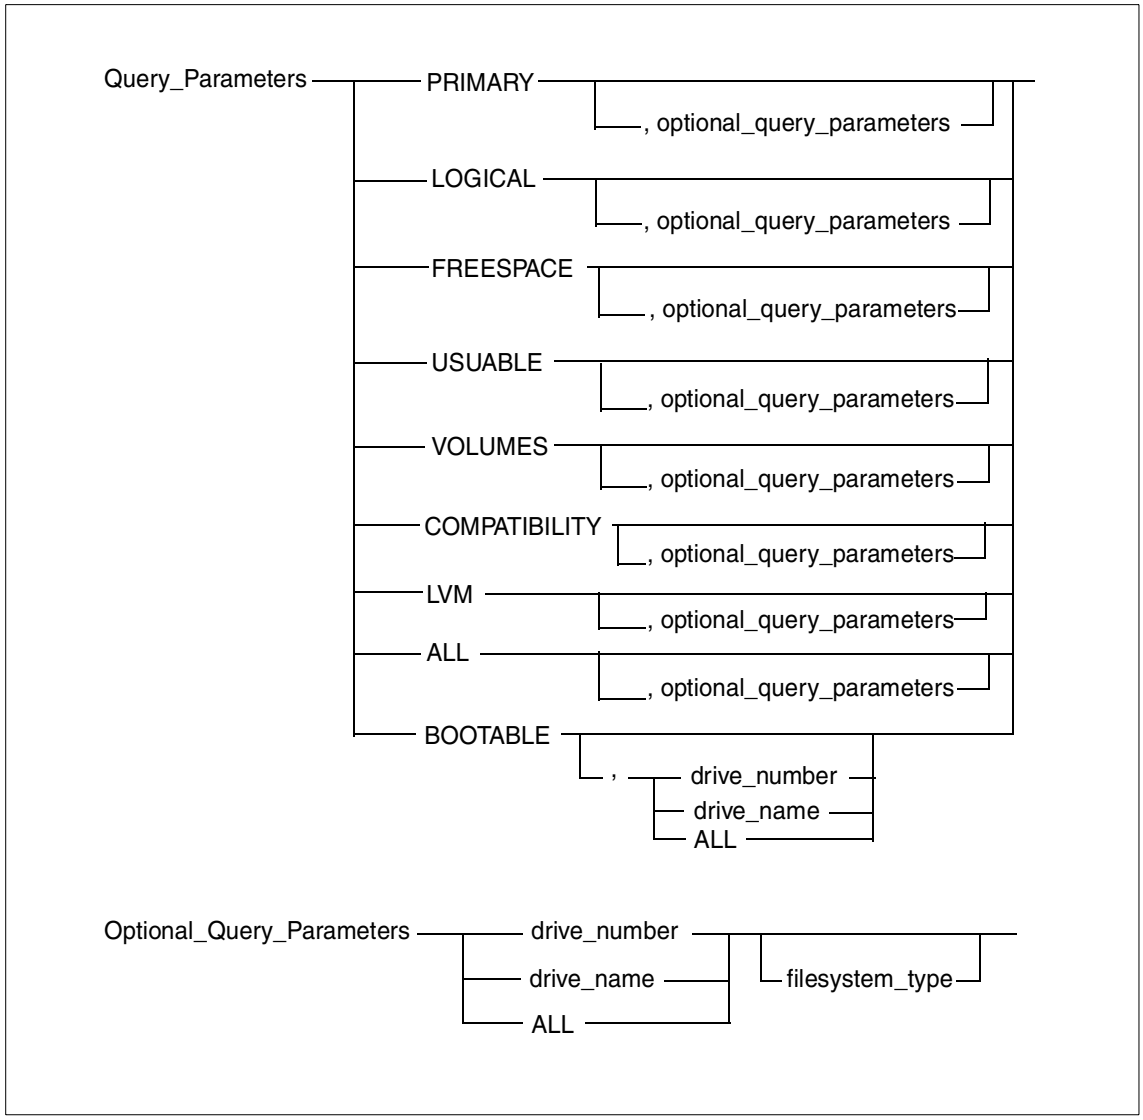

Figure 43. LVM parameter options (1 of 3)

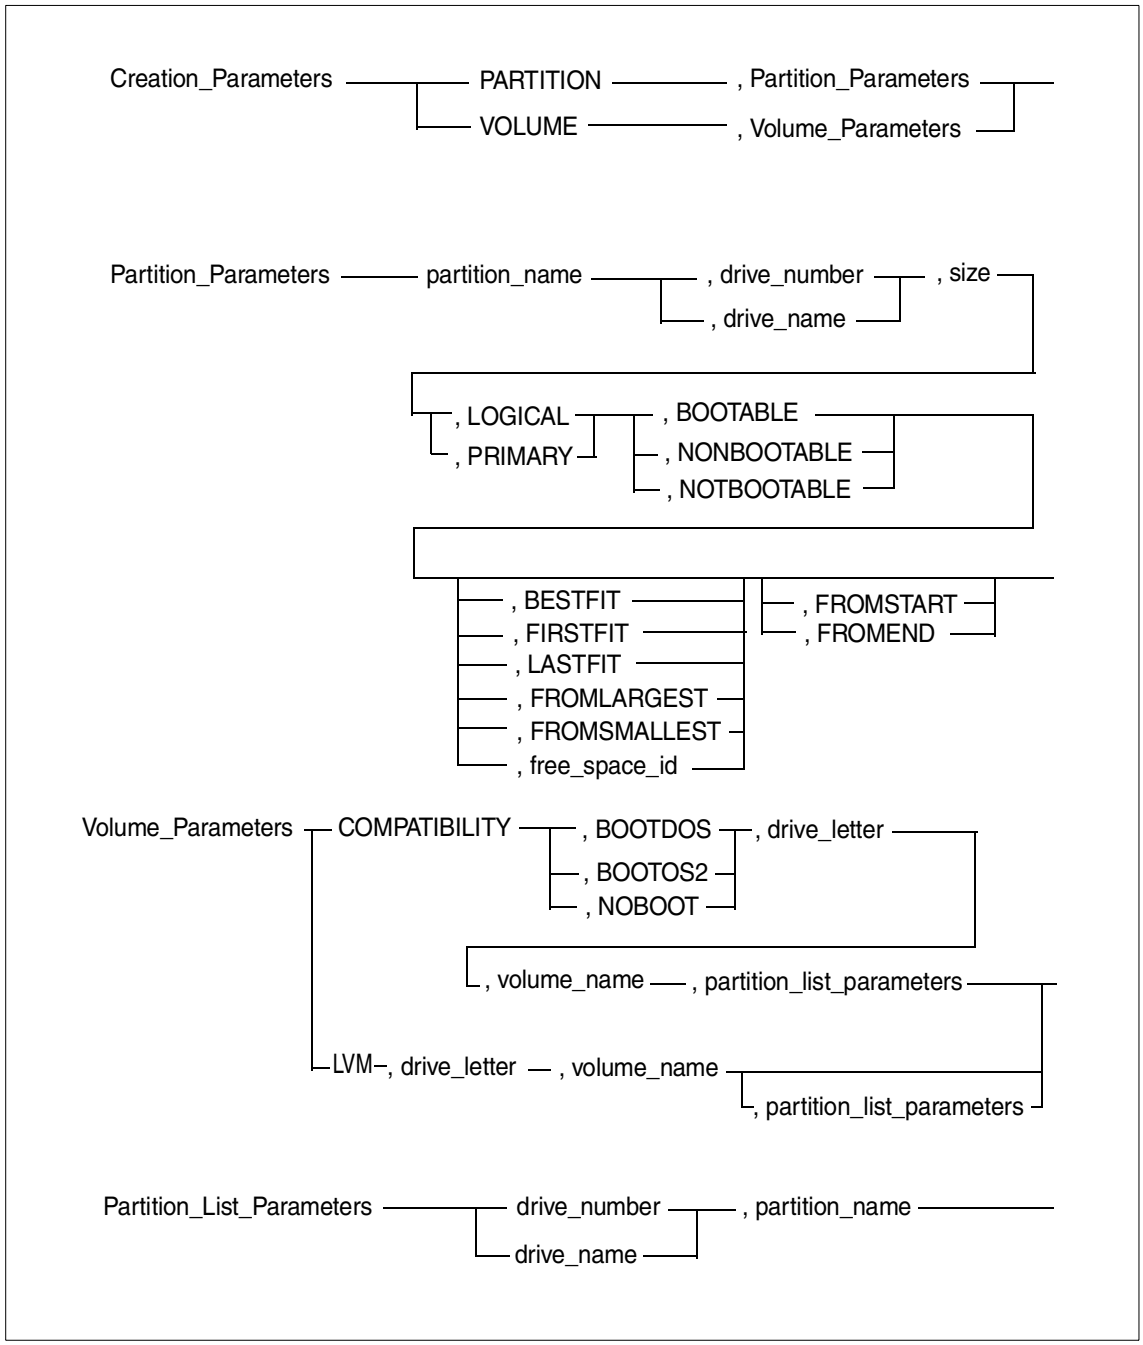

Figure 44. LVM parameter options (2 of 3)

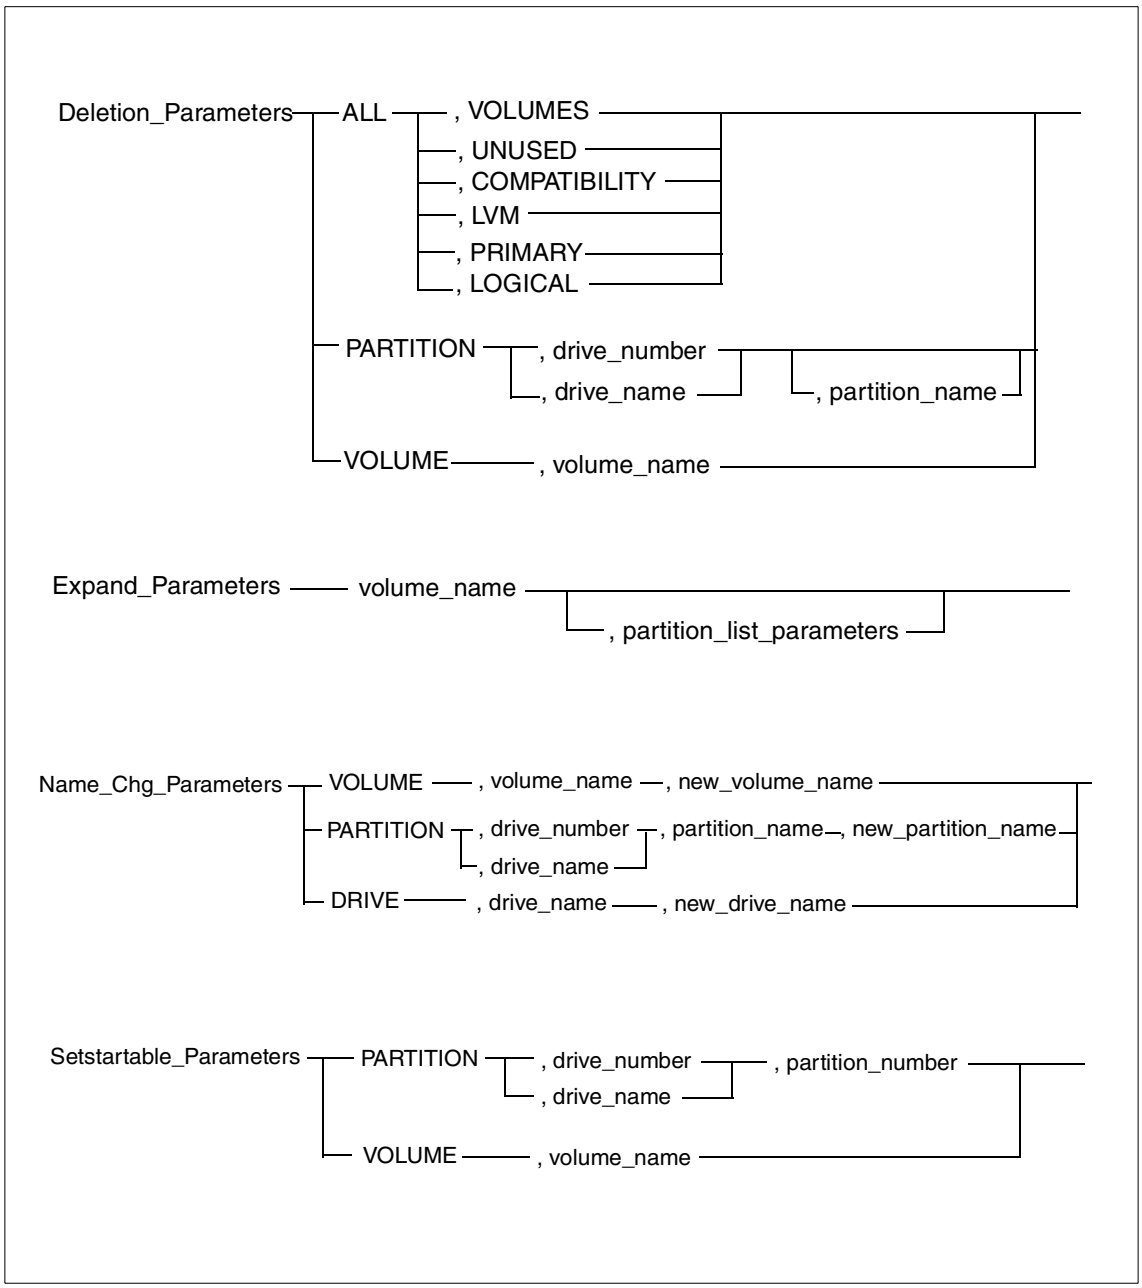

Figure 45. LVM parameter options (3 of 3)

As an example, the LVM syntax will be used to create the partitioning shown in Figure 46 on page 111.

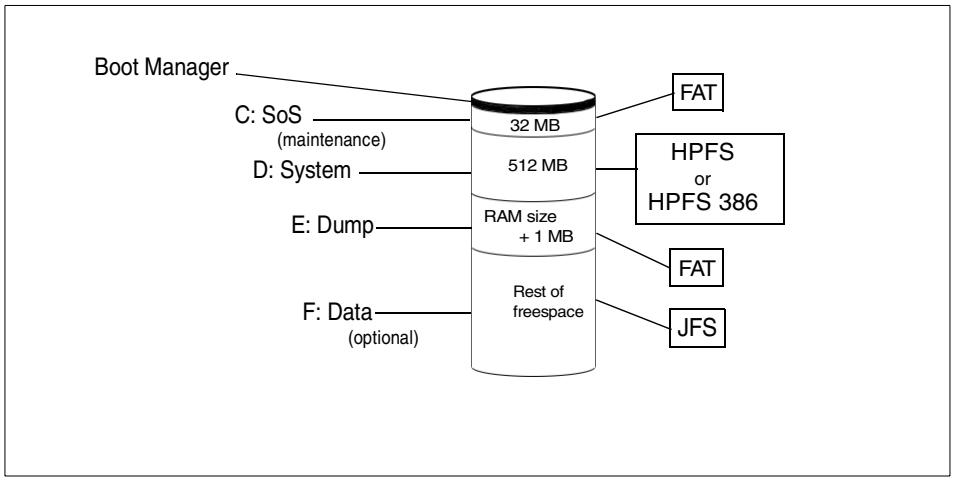

Figure 46. Partitioning used in the CID environment

```
lvm /delete:all,volumes
lvm /delete:all,unused
lvm /delete:all,primary
lvm /delete:all,lvm
lvm /delete:all,logical
lvm /delete:all,compatibility
lvm /bootmgr:1
lvm /create:partition,SoS,1,32,primary,bootable
lvm /create:volume,compatibility,bootos2,c:,SoS,1,SoS
lvm /create:partition,system,1,512,logical,bootable
lvm /create:volume,compatibility,bootos2,d:,system,1,system
lvm /create:partition,dump,1,129,logical,nonbootable,[ FS1 ],fromstart
lvm /create:volume,compatibility,noboot,e:,dump,1,dump
lvm /create:partition,data,1,,logical,nonbootable,[ FS1 ],fromstart
lvm /create:volume,lvm,f:,data,1,data
```

Figure 47. LVM command line commands to partition the disk

The commands listed in Figure 47 are described as follows:

- Lines 1-6 delete any existing partitions.
- Line 7 creates the boot manager partition.

- Line 8 creates the *SoS* partition, and Line 9 creates the volume within this partition.
- Line 10 creates the *system* partition, and Line 11 creates the volume within this partition.
- Line 12 creates the *dump* partition, and Line 13 creates the volume within this partition.
- Line 14 creates the *data* partition, and Line 15 creates the volume within this partition.

4.2.4.5 Bad Block Relocation (BBR) Alarm Utility

The BBR Alarm Utility provides an indication that bad blocks have been encountered and that BBR is happening in some of the accesses to the hard disk, thus, indicating potential hard disk performance degradation and/or potential hard disk failure.

4.2.5 Bad-block relocation

Bad Block Relocation (BBR) allows bad sectors that are discovered after initial burn-in (and hardware recording/relocation) of the drive to be remapped via software. It allows a good sector (replacement sectors are reserved on the disk) to be transparently substituted to an application for the bad sectors discovered during write operations on the disk.

4.2.6 LVM operation

LVM can only work with partitions types it can understand; these are LVM- and LVM-compatible volumes. Hence, the first time OS/2 Warp Server for e-business is installed on a pre-partitioned hard disk, VCU (Volume Conversion Utility) runs. Its function is to detect existing partitions that it understands and convert them into LVM Compatibility volumes. Compatibility volumes are special volumes that are compatible with previous versions of OS/2 or other operating systems. A compatibility volume corresponds to a single physical partition and can be made bootable. Each compatibility volume can exist on only one physical hard drive; they do not support dynamic expansion.

OS2LVM helps handle requests for data blocks to be written (or read) from IFSM by providing a mapping from the new logical sector number (LSN) disk addressing to the actual physical partition sector number (PSN) addressing used by OS2DASD and by the adapter device drivers (ADDs) to access the disks. OS2LVM also provides mapping of contiguous LSN addressing to discontinuous PSN addressing when spanning physical disks for a logical volume.

The interaction between OS2LVM and OS2DASD in detecting DASD can be summarized as follows:

- At boot time, the device drivers locate the devices attached to the system. The storage devices report their findings to OS2DASD.
- OS2DASD examines the storage devices found for partitions.
- OS2DASD creates a table with an entry for each partition it has found.
- OS2DASD passes this table of partitions to the OS2LVM.
- OS2LVM filters the table adding and removing entries as required. It then adds the drive letters found in the LVM data taken from the drives and passes the table to the IFSM.
- IFSM uses the drive letters placed in the table by OS2LVM.

4.2.6.1 LVM GUI

The partitions and volumes can also be managed by using the Logical Volume Manager (LVM) GUI. The LVM GUI icon can be found in the SYSTEM SETUP folder.

The LVM GUI provides two views:

1. The Logical view
2. The physical view

The Logical view

The logical view displays information about the volumes. It displays the list of all volumes that have been created and information regarding each volume. The information includes the drive letter assigned to the volume, the name of the volume whether the volume is bootable or non-bootable, the file system used, the size in MB of the volume, the percentage of the used space and free space on the volume, and the type of volume, such as compatible or LVM.

It allows the user to perform volume management functions. It allows the user to create, delete, rename, expand, and hide volumes. The user can also install or remove Boot Manager.

The **View** menu option allows the user to shift between views.

Physical view

The LVM physical view displays the disk information. The upper half of the screen contains the information about the partitions and the lower half of the screen contains information about the volumes. The partition window displays information about how the disks in the system are partitioned. It displays

information about the partition size and the file system used. The volume window on the lower half of the screen provides information about the volumes and the partitions assigned to them. The LVM physical view allows the user to create and manage partitions.

4.2.6.2 Volume management

A logical volume is a single partition on a hard disk or a collection of partitions spanning multiple hard disks that are presented to a user as a single entity. The user views a volume as a single logical partition even though a volume might be comprised of one or more physical partitions. The user views and accesses this volume using a drive letter. LVM allows the user to create and manage volumes and partitions. It has to be noted that drive letters are assigned to volumes and not partitions.

Note

When you create a partition, the new partition has to be assigned to a volume. Either a new volume has to be created and the partition assigned to the volume or an existing volume can be expanded to include the new partition if the volume is a JFS volume. You cannot use a partition without having it included in a Volume.

Some of the definitions the LVM user needs to be familiar with are explained below:

Volume Name — This is the name assigned to the volume by the user.

Drive letter — A drive letter is assigned by the user to the volume at its creation time. This drive letter is used by the user to view and access information on the volume.

Status — This reflects the volume designation. Four types of volume designation are possible:

- Bootable — This is the volume designation used by the LVM indicating that the volume contains the bootable operating system. When a bootable volume is created, it is automatically added to the Boot Manager menu. For example, if you have two operating systems installed in two different volumes, you can designate each volume as bootable. Bootable volumes exist only when Boot Manager is present; otherwise, the bootable volume is designated as startable.
- Installable — This is a volume designation used by LVM during installation indicating that the OS/2 Warp for

e-business will be installed on that volume. Only one volume can be set as installable. The installable designation is displayed and set only during installation and it implies that the installable volume is also Bootable or Startable depending on whether Boot Manager is installed. If Boot Manager is installed, the volume that is set Installable is also Bootable. If Boot Manager is not installed, the volume that is set as Installable is also startable.

Nonbootable — This is the volume designation used by the LVM indicating that the volume is not set bootable. For example, you might have an operating system installed in one volume and you might save your data in another volume.

Startable — This is a volume designation used by the LVM indicating the volume that is used when the system starts. A startable volume must contain either the Boot Manager or a bootable operating system. If Boot Manager is installed, it is automatically set as a Startable volume and the Boot Manager menu is displayed when the system starts. If Boot Manager is not installed, you must set as startable a volume that contains an operating system. Only one volume can be set as Startable. A Startable volume is the primary partition that is used to start the system by virtue of being the active partition. If any volume other than the Boot Manager partition is set startable, no volume is displayed as bootable, and the Boot Manager will not run. Boot Manager is typically set as the Startable volume but is not displayed in the Logical view.

File system — This indicates the type of file system used on the volume.

Two types of volumes are possible with LVM:

Compatibility Volumes — These volumes can be marked as bootable or startable, are accessible by other operating systems, and can be formatted for 386HPFS, HPFS, and FAT. It cannot be formatted for JFS. The OS/2 base operating system must be installed on a compatibility volume. These volumes can have only one partition.

LVM Volumes	These volumes can be formatted for 386 HPFS, HPFS, FAT, or JFS. A single volume can contain multiple partitions. LVM volumes can be accessed only by OS/2 versions that contain LVM. The LVM volumes cannot be set as bootable.

Logical Volume Manager allows the following operations on a volume:

- Creating a New Volume
- Deleting a Volume
- Expanding a Volume
- Changing the Drive Letter of a Volume
- Hiding the Volume
- Changing the Name of the Volume
- Setting the Type of the Volume

Creating a new volume

A new volume can be created using the LVM GUI:

1. Enter the Logical Volume View. This can be done by selecting the **Logical volume view** entry from the View menu option.
2. Select the **Volume** menu option. Click on the **Create Volume** entry. Select the **Create Bootable** option or the **Create non-bootable** volume. If it is a non-bootable volume, select the type of the volume: **Compatibility** or **LVM** volume.

Note

A single LVM volume can be comprised of multiple partitions, while a single Compatibility volume can only have a single partition. If you are creating a Compatibility volume, you are allowed to select only one partition from the partition list. If you are creating an LVM volume, you are allowed to select multiple partitions.

If you have a CD-ROM drive, it will be assigned a drive letter. When you are creating a new volume, this drive letter is also displayed. If you assign this drive letter to the new volume, you will have to reboot the system when you are finished.

Deleting a volume

- Select the **Logical Volume View** from the View menu option.
- Click on the **Volume** menu option
- Select the **Delete Volume** entry, which displays the list of volume names configured in the system. Select the volume that has to be deleted.

Expanding a volume

You may expand the volume if the volume does not have a file system associated with it (if it is unformatted) or if the volume has JFS associated with it. If the file systems other than JFS are associated with it, you must first delete it and then re-create it.

- Select the **Volume** menu option.
- Select the **Expand Volume** entry. This displays the list of the volumes configured in the system
- Select the volume that is to be expanded
- A window appears that displays the free space and partitions. If a partition is available and you wish to use that to expand the volume, select that partition or select freespace.
- Choose OK to expand the volume.

Changing the drive letter of a volume

- Select the **Volume** menu option from the Logical Volume View
- Select the **Set/change drive letter** entry. This displays the list of the volume names configured in the system. Select the volume that has to be changed. A warning message appears. Read the warning message carefully. Press **OK** to proceed.
- A window is displayed with the list of drive letters. Select the new drive letter and press **OK**.

Changing the name of the volume

Select the Volume menu option

- Select the **Set/Change name on volume** entry. This displays the list of volume names configured in the system. Select the volume that has to be changed.
- Enter the new name of the volume in the window displayed. Choose **OK** to assign the new name to the volume.

Hiding a volume from OS/2

This option is used to make the volume invisible to OS/2, that is, to block access to a volume by the file system. For example, if you have a Windows volume, you might not want to access it from OS/2. Therefore, you need to hide the volume. A volume that is hidden is no longer accessible by OS/2, but all the associated data still exists on the hard disk.

- Select the **Volume** menu option from the Logical Volume view.

- Select the **Hide volume from OS/2** entry, which displays the list of volumes.
- Select the volume that has to be hidden.

Setting a volume startable
Select **Set the volume startable** to set the selected volume startable. A startable volume is a volume that is used to start the system from the hard disk. It is the active partition. Only a compatibility volume on a primary partition can be set startable with LVM. Boot manager, if it is installed, is automatically set as the startable partition. Only one partition can be set as startable; therefore, if Boot Manager is to be active, no volumes can be set as startable. If you set a volume startable with LVM, Boot Manager is disabled. Note that Boot Manager cannot be seen from the LVM logical view because no drive letter is assigned to it, but it can be seen in the LVM physical view.

Setting a volume installable:
A volume can be set installable only during installation of OS/2 Warp Server for e-business. The volume on which the OS/2 Warp Server is to be installed is designated as Installable during the installation of OS/2 Warp Server for e-business.

4.2.6.3 Boot Manager
When you have multiple primary partitions and different operating systems on these partitions, the system has to be told what operating system to boot from at system startup time. Boot Managers are used for this purpose. Boot Manager is a program that normally inserts itself into the very beginning of the boot process by setting up a boot manager partition and making itself the active partition. When the system is started, the Boot Manager runs and analyzes the partitions and presents a menu. This menu lists all the bootable operating system volumes. The user can then select the desired operating system. Control is then transferred to the boot code on that partition and the operating system loaded.

Installing Boot Manager
Select **Install Boot Manager** from the Boot Manager menu to install boot manager on the system. Boot Manager is automatically set as the startable partition. Since only one partition can be Startable, if Boot Manager is installed, no other volumes can be designated as Startable. Boot Manager is installed as a new primary partition at the beginning of the first free space block where a new primary partition is allowed. This partition is the smallest allowed, typically, from 1MB to 10 MB depending on the size of the hard disk.

Removing Boot Manager
Select **Install Boot Manager** from the Boot Manager menu to remove the boot manager.

Note

Once the Boot Manager is removed, one of the volumes must be set as startable; otherwise, it will not be possible to boot from the hard disk.

Adding a volume to the Boot Manager menu
Select **Add the volume to Boot Manager** to add a volume to the Boot Manager menu. Only bootable volumes are displayed.

Removing a volume from the Boot Manager menu
Select **Remove volume from Boot Manager** to remove the volume from the Boot Manager menu.

Setting or Changing the Boot Manager startup values
Boot Manager startup values are used to set the default boot volume, set the timer inactive, change the timeout value, and change the display mode. The following options are available:

- Default Boot Selection: This allows you to choose the default boot volume that boots automatically each time you start the system unless you select another volume from the Boot Manager menu.
- Timer Active: This toggles between YES or NO. Select **YES** when you want the default volume to boot automatically after the timeout period. Select **NO** to specify that the startup menu is to remain displayed until a selection is made.
- Time-out: This value specifies the number of seconds the Boot Manager waits before the default volume is booted automatically. This value is active only when Timer Active is Yes.
- Display Mode: **Normal** mode displays only the names of the volumes on the startup menu. **Advanced** mode gives more information about the volumes on the startup menu. This information includes disk number, drive letter, type of partition, volume size, file system type, and accessibility if hidden.

4.2.6.4 Partition management
The LVM physical view displays the partitions present on each hard disk allowing you to create and manage individual partitions. You can create partitions for other operating systems that do not recognize the LVM volumes.

You can create a partition of a specific size and allocate it from the beginning or end of the free space.

Note

Drive letters are assigned only to Volumes and not partitions. When you create a partition or you want to use a previously created partition, you have to create a volume that can contain this partition. When you create a volume, a drive letter is assigned to it. You cannot create a partition and use it without assigning it to a volume. See the Volume Management section for information about volumes.

Creating a partition

- Select the **Partition** menu option from the physical view.
- The Create partition window is displayed. Select the partition location.
- Select the type of partition: Primary or Logical.
- Select the disk on which the partition has to be created. Specify the size of the partition and the name of the partition.
- Choose OK to create the new partition.

Deleting a partition

Select the **Partition** menu option from the physical view.

- Select the **Delete Partition** entry.
- A window lists all the existing partitions.
- Select the partition that has to be deleted.
- Choose OK to delete the partition.

Note

After volumes have been created with LVM, FDISK should no longer be used to manage partitions.

4.2.7 Logical volume manager benefits

The Logical Volume Manager provides the following enhancements:

- It provides a single interface for configuring both physical and logical volumes.
- Bad Block Relocation (BBR) is provided for JFS; HPFS provides its own BBR. This improves reliability and availability.

- It allows a JFS volume to span physical disks providing volumes up to 2 Terabytes in size.
- It allows partitioning to be done dynamically without having the need to reboot. The LVM has been designed to allow both expansion and reduction of volumes; however, only expansion has been implemented in this release.
- Performance in the I/O path is improved by using 32-bit data and instructions to support the 32-bit I/O path for JFS.
- There are Dynamic/Sticky drive letter assignments regardless of any other system changes.

4.3 Journaled file system

The Journaled File system is an installable file system (IFS) that was built on IBM AIX JFS technology; however, the disk-image layout differs due to the feature set of the file system unique to OS/2.

JFS is intended as a replacement for the High Performance File System (HPFS). One of the shortcomings of HPFS is the amount of time required to recover from a system crash on large hard drives. JFS overcomes this by providing a robust, quickly-restartable, transaction-oriented, log-based, high-performance, 32 bit file system for OS/2. It is primarily tailored for the high throughput and reliability requirements of modern servers especially in the TCP/IP environment.

4.3.1 JFS cache

In most cases, applications requesting data that resides on JFS formatted partitions will receive the data from JFS cache. JFS cache, like HFPS, FAT, and HPFS386 cache is initialized on startup. JFS cache is configured and initialized by the IFS statement in the CONFIG.SYS.

All buffer cache I/O is performed through explicit I/O requests to the underlying DASD device driver using the *strat3* interface shown in Figure 37 on page 103. The buffer cache resides in pinned memory and is shared by all mounted JFS file systems.

By default, the size of the JFS buffer cache will be12.5 percent of the size of real memory. Not all file system I/O will be performed through the buffer cache since the JFS supports the NOCACHE advisory file open mode. When this function has been requested for a file, the JFS will perform I/O directly

between user buffers and the underlying device when processing read and write operations on the file.

To configure the JFS, an IFS= statement must be added to CONFIG.SYS. This statement is described below.

```
IFS= Statement

IFS=pathname [/CACHE:<size in kilobytes>] [/AUTOCHECK:drive[drive...]
/[L:OFF|/L:synctime,maxage,bufferidle]]

Example: IFS=C:\OS2\JFS.IFS /AUTOCHECK:* /CACHE:1024
```

Figure 48. Syntax for JFS initialization

Table 9. JFS initialization parameters

Parameter	Description
pathname	Specifies that pathname of the JFS IFS.
/CACHE:<size in kilobytes>	Specifies the size of the JFS buffer cache in Kilobytes. By default, this size is set to 12.5 percent of real memory.
/AUTOCHECK:drive[drive...]	Specifies a list of JFS file systems, identified by driver letter, to be included in automatic recovery by CHKDSK at the time of JFS initialization. If an asterisk (*) is specified instead of a drive letter list, all JFS file systems will be included in automatic recovery during JFS initialization. If any drive letter is preceded by an arithmetic plus sign (+), CHKDSK will perform a full integrity check on the drive even if journal log replay has restored the drive to an apparently consistent state. If the list of drive letters ends with plus-asterisk (+*), CHKDSK will perform a full integrity check on all JFS file systems not specified in the list even if journal log replay has restored them to an apparently consistent state.
/L:OFF	Specifies the lazy write parameters in seconds. OFF forces all to be synchronous.

Parameter	Description
/L:<synctime,maxage,bufferidle>	Synchtime is the interval at which the sync thread runs, default is 32. Maxage is the longest time that a modified file is kept in cache, default is synchtime*4. Bufferidle is the time indicating a "recent" change. Changes newer than this value are not written unless the last write was older than maxage, default is MIn(1, synchtime/8).

4.3.2 JFS structure

The Journaled File System is implemented through a set of operating system commands that allow the creation, management, and deletion of files. Special 32-bit extensions have been added to optimize JFS performance. A JFS is created inside a logical volume and is organized as shown in Figure 32 on page 97. To enable JFS support, you first need to create a volume using LVM, which is a non-bootable LVM volume. Within this, you can install JFS support.

In JFS, data storage space is organized in three levels: files (and directories), filesets, and aggregates. The file is a unit of user data. Files are organized in a hierarchical directory structure. The directories themselves are located in filesets. The filesets reside on aggregates. A fileset is a collection of directories and files managed as a single unit. The aggregate is the unit of disk storage that contains one or more filesets; initially, only one fileset per aggregrate is implemented in the initial release of JFS. Filesets can vary in size but cannot be larger than the housing aggregrate. The concept of an aggregrate has been introduced to support the Distributed Computing Environment Distributed File System (DCE DFS).

When a new aggregate is created, it gets registered with the LVM as a logical volume.

4.3.3 JFS disk layout

JFS divides the Logical Volume (also known as a disk partition) into a number of fixed-size units or logical blocks. An aggregate has an array of disk blocks with specific formats that include a superblock and an allocation map. The superblock identifies the Logical Volume as a JFS aggregate while the allocation map describes the allocation state of each data block within the aggregate. The format also includes the initial fileset and the control structures necessary to describe it. The fileset is a mountable entity containing files and directories. Files and directories are represented

persistently by inodes; each inode describes the attributes of the file or directory and serves as the starting point for finding the file or directory's data on the disk. JFS also uses inodes to represent other file system objects, such as the map which describes the allocation state and location on disk of each inode in the fileset. Extent-based addressing structures rooted in the inode are used for mapping file data to disk.

In addition to files, user data may exist in the form of extended attributes that allow applications to attach named data to the file or directory. Extended attributes are stored separately from file data or directory entries and are maintained through the associated inode.

Together, the aggregate superblock and disk allocation map, fileset descriptor and inode map, inodes, and addressing structures represent the JFS control structures or meta-data.

JFS logs are maintained in each aggregate and used to record information about operations on meta-data. JFS uses techniques originally developed for databases to log information about operations on the file system meta-data as atomic transactions. In the event of a system failure, a file system is restored to a consistent state by replaying the log and applying log records for the appropriate transactions. However, JFS only logs operations on the meta-data; so, replaying the log only restores the consistency of the structural relationships and resource allocation states within the file system. It does not log data or recover this data to a consistent state. Consequently, some data may be lost if stale after recovery; thus, users with a critical need for data consistency should use synchronous I/O. JFS does not log the extended attributes; instead, it replaces the old version with the new version as an atomic operation.

4.3.4 JFS system structure

This section describes the major aspects of the JFS on-disk layout. It describes the extent-based file geometry, directory formats, the formats of block allocation maps, inodes, and so on.

JFS uses extent-based addressing structures and aggregate block allocation policies to produce compact, efficient, and scalable structures for mapping logical offsets within files to physical address on disk. An extent is a sequence of continuous blocks allocated to a file as a unit and is described by a triple consisting of <logical offset, length, physical address>. The addressing structure is a B^+-tree populated with descriptors (the triples above), rooted in the inode, and keyed by the logical offset within the file.

Thus, a file can be viewed as being allocated in a sequence of extents where each extent defines a contiguous variable-length sequence of aggregate blocks allocated as a unit. An extent can range in size from 1 to 2^{24}-1 aggregate blocks. Therefore, the size of the maximum extent depends on the aggregate block size. With a 512 byte aggregate block size (the smallest allowable), the maximum extent is 512*(2^{24}-1)bytes (almost 8GB). With a 4096 byte aggregate block size, the maximum extent is 5096*(2^{24}-1)bytes (almost 64GB). These limits only apply to a single extent; it does not have any effect on limiting the overall file sizes. With smaller size extents, there is less internal fragmentation for aggregates with large numbers of small-size files.

In general, the allocation policy for JFS tries to maximize contiguous allocation by allocating a minimum number of extents, with each extent as large and contiguous as possible. This results in larger I/O transfer, which results in improved performance.

Fragmentation can still occur, and, for this reason, a defragmentation utility, defragfs， is provided to reduce fragmentation in the internal data structure, which occurs from the dynamic allocation/deallocation of variable size extents. The allocation/deallocation results in disconnected variable size extents all over the aggregate. The defragmentation utility coalesces small free extents into single larger extents.

JFS divides the Logical Volume into a number of fixed size units or logical blocks. Each logical block comprises a fixed number of aggregate blocks, which can be 512, 1024, 2048, or 4096 bytes. The aggregate block size is defined at creation time via extra parameters to the `format` command, which defines the smallest unit of space allocation supported on the aggregate.

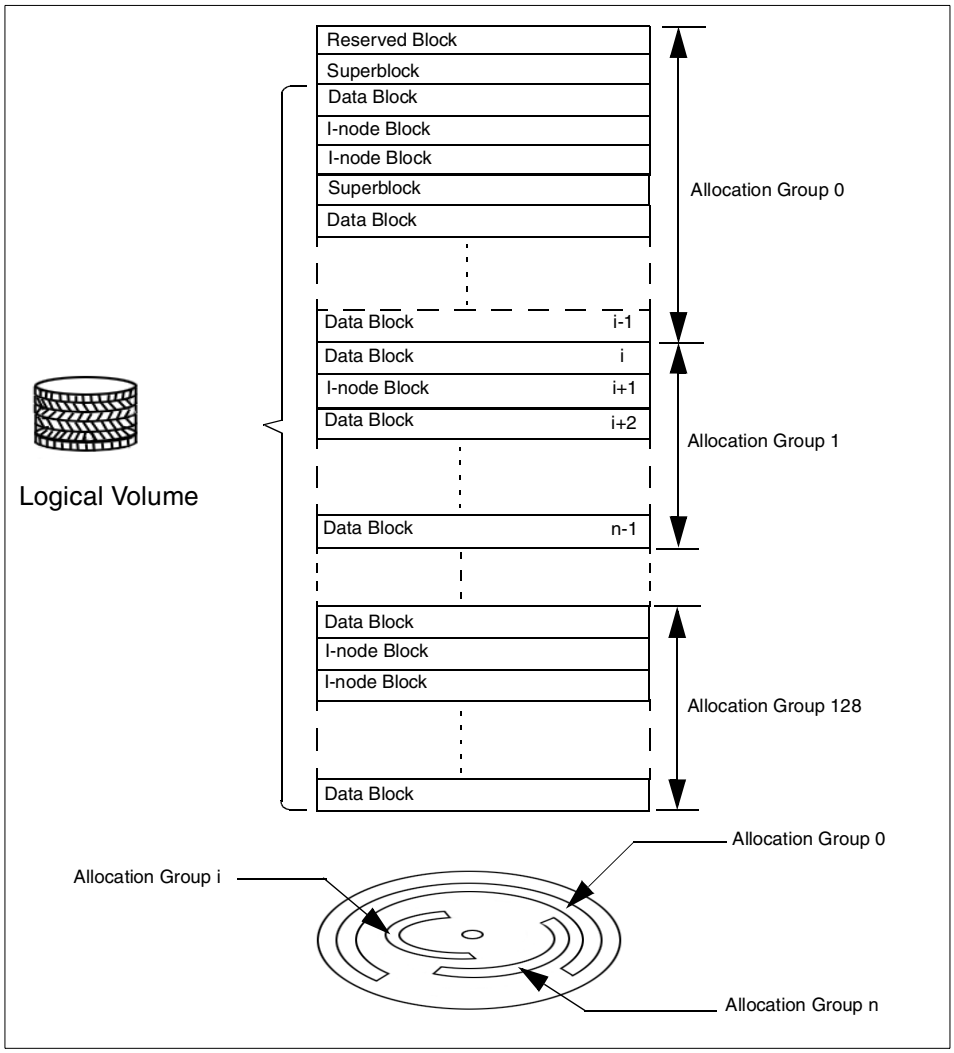

Figure 49. JFS physical organization in a logical volume.

The logical blocks in the file system are comprised of:

Reserved Block 0

This logical block is reserved and largely unused. Its origins is from the AIX implementation of JFS; if it is at the start of the disk, 1K is used for the bootstrap program. LVM data is also stored in this block. The LVM data resides just before the bootstrap area.

Superblock

The first and fourth logical blocks are reserved for the superblock, with the fourth block being a backup copy of the first. The superblock contains information, such as the aggregrate size, the file system name, version number, and file system state.

Allocation groups

The rest of the logical blocks are divided into a number of allocation groups. Allocation groups are comprised of data blocks and i-nodes, which reference these data blocks. Allocation groups allow resource policies to be applied, which results in good I/O performance. The purpose of these allocation policies is as follows:

- To improve locality of reference

  Files created within a directory will be maintained in an allocation group with that directory. The policy tries to allocate contiguous logical blocks, clustering together disk blocks and disk i-nodes of related data to achieve good locality for the disk. This results in less disk head movement, thus, quick access to the data.

- To ease file system extension

  Extending the file system is easier because a new allocation group of i-nodes and data blocks can be added maintaining the relationship between i-nodes and file system size simply. Without allocation groups, the file system would either have to be reorganized to increase the number of i-nodes, or, the extension could only increase the number of data blocks available, thereby, conceivably, limiting the number of files and directories in the file system.

In order to minimize the number of updates required when an aggregate is expanded (or shrunk), the allocation groups have been limited to a maximum number of groups, 128. Additionally a minimum of 8192 aggregate blocks has been imposed in the size of the allocation group.

I-nodes

Basically an i-node is a pointer to a file. An i-node contains information on the file, such as the type of file, the size in bytes, permissions, owner, access permissions for the file, the number of blocks allocated to the file, creation date, last modification date, last accessed date, and pointers to the blocks that actually contain the file.

Inodes are allocated dynamically by allocating inode extents that are simply a contiguous chunk of inodes on the disk. By definition, a JFS inode extent contains 32 inodes. With a 512 byte inode size, an inode extent is, therefore, 16 KB in size on the disk.

Information within an i-node is divided into two parts. The first part contains information, such as permissions and owner for the directories or file. The second part contains an array of pointers to the actual disk addresses of the logical blocks that make up the file or directory.

The extent is represented on the aggregrate disk structure by an extent allocation descriptor (xad) structure. The xad structure is as follows:

```
struct xad {
        unit8 xad_flag;
        unit16 xad_reserved;
        unit40 xad_offset;
        unit24 xad_length;
        unit40 xad_address;
   };
```

where:

`xad_flag` is an 8-bit field containing miscellaneous flags. These flags can indicate copy-on-write for DFS if the extent is allocated but not recorded for databases, information for compression, help for journaling, and so on.

`xad_reserved` is a 16-bit field reserved for future use. It is always zero.

`xad_offset` is a 40-bit field containing the logical offset of the first block in the extent.The logical offset is represented in units of the aggregate block size, that is, to get a byte offset, `xad_offset` must be a multiplied by the aggregate block size.

`xad_length` is a 24-bit field containing the length of the extent. The length is represented in units of aggregate block size.

`xad_address` is a 40-bit field containing the address of the extent. The address is represented in units of the aggregate block size.

Files that can fit within the array storage area, such as most links, are actually stored in the i-node itself, thus, saving space.

The following four examples illustrate the use of the extent descriptors in file representation in the aggregate. In the examples below, the aggregate block size is 1KB.

- A 1041377 byte file allocated contiguously:

  This file requires 1017 1KB aggregate blocks (with 31 bytes in the last aggregate block lost to internal fragmentation). Only one xad structure is required to describe this contiguous file:

xad_flag	<not discussed here>	
xad_offset	0	/* the beginning of the file */
xad_length	1017	/* 1017 1KB aggregate blocks */
xad_address	nnnnn	/* aggregate block # */

Note

This same xad structure could be used to represent any contiguous file of size 1040385 (1016*1024+1) to 1041408 (1017*1024), because extent descriptors only represent sizes down to aggregate block size granularity.

- A 1041377 byte file allocated in three pieces:

  Assuming the above file is split into three separate extents on the disk: One 495 aggregate blocks long, one 22 blocks long, and one 500 blocks long. It requires three xad structures to represent this file - one per physical extent.

  xad # 0

xad_flag	<not discussed here>	
xad_offset	0	/* the beginning of the file */
xad_length	495	/* 495 1KB aggregate blocks */
xad_address	nnnnn	/* aggregate block # */

xad # 1

xad_flag	<not discussed here>	
xad_offset	495	/* bytes starting at #506880 */
xad_length	22	/* 22 1KB aggregate blocks */
xad_address	mmmmm	/* aggregate block # */

xad # 2

xad_flag	<not discussed here>	
xad_offset	517	/* bytes starting at #529408 */
xad_length	500	/* 500 1KB aggregate blocks */
xad_address	jjjjj	/* aggregate block # */

In this case, xad # 0 describes the first 495 physical aggregate blocks of the file. The xad_offset field contains zero because this xad describes the bytes starting at logical offset zero. Next, xad # 1 describes the next 22 physical aggregate blocks of the file. The xad_offset field contains 495 because this xad describes the bytes starting at logical offset 506880 (495*1024); the previous bytes being described by xad 0. The final xad describes the last 500 blocks of the file. The xad_offset field here is 517. Notice, that for files which are not sparse, the xad_offset field of a given xad is equal to the sum of the lengths of the previous xad structures (517=495+22) for this example.

- 1041377 byte sparse file:

  For this example the file has two bytes of data starting at logical byte offset zero, and three more bytes starting at logical byte offset 1,041,374, and has all zero (sparse) in between. The file is 1041377 bytes long. The volume has been formatted for sparse files by using the */s* option when formatting.

  When formatted sparse JFS does not allocate physical disk space to hold byte ranges of a file that has never been written to. Therefore it will take two xad structures to represent this file, one for the first two bytes and one for the last three bytes of data.

xad # 0

xad_flag	<not discussed here>	
xad_offset	0	/* the beginning of the file */
xad_length	1	/* 1 1KB aggregate blocks */
xad_address	nnnnn	/* aggregate block # */

xad # 1

xad_flag	<not discussed here>	
xad_offset	1016	/* bytes starting at # 1040384*/
xad_length	1	/* 1 1KB aggregate blocks */
xad_address	nnnnn	/* aggregate block # */

In this case, the first extent (xad 0) contains two bytes of data and 1022 bytes of zero. The last extent (xad 1) contains 990 bytes of zero followed by 3 bytes of data. The remaining 31 bytes in the 1KB extent are not part of the file.

Note that, in this case, the xad_offset fields are necessary because they are the only way to know that xad 1 represents a sequence of bytes that are at an unexpected logical offset within the file. That is, the offset for xad 1 does not equal the offset of xad 0 + length. This is how a sparse file is represented.

- 16GB file allocated contiguously:

  The length field in an xad structure is only 24 bits long; therefore, it can hold a value up to 2^{24}-1. With an aggregate block size of 1KB, the longest extent a single xad can represent is $(2^{24}-1)*2^{10}$ = 1KB less than 16 GB. By implication, this is also the largest extent a single xad structure can represent. Thus, if a file is large enough, it will require multiple xad structures to represent it, even if the file is contiguous on the disk. Thus, the continuously allocated 16GB file starting at aggregate block number 12345 and going for 16777216 1KB aggregate blocks would be represented as follows:

xad # 0

xad_flag	<not discussed here>	
xad_offset	0	/* the beginning of the file */
xad_length	1677715	/* 16777215 1KB aggregate blocks */
xad_address	12345	/* aggregate block # */

xad # 1

xad_flag	<not discussed here>	
xad_offset	16777215	/* bytes starting at 16777215K */
xad_length	1	/* 1 1KB aggregate blocks */
xad_address	16789560	/* aggregate block # */

In this case, whether or not the file is contiguous on the disk, it takes at least two xad structures to represent it due to the length limitation of individual extents.

Generally, if a file can be represented by eight or less extents (xad structures), it will be stored within the inode itself as illustrated in Figure 50 on page 133 as direct referencing; In this case, up to eight would be direct references. If the number of extents required to represent the data is between 9 and 2032 extents, indirect referencing would be used. In this case, each of the xad in the i-node will point to 254 leaf node (the header of the leaf node is 32 bytes long, an xad is 16 bytes long, thus, the number of xad entries in a 4K page = (4K-32)/16=254). The points, in turn, point to data blocks. If the number of extents is > 2032 and <= 516128 allocated extents, double indirect referencing is used. Here, the xad in the i-node point to leaf nodes that, in turn, point to xad internal nodes that point to the data blocks.

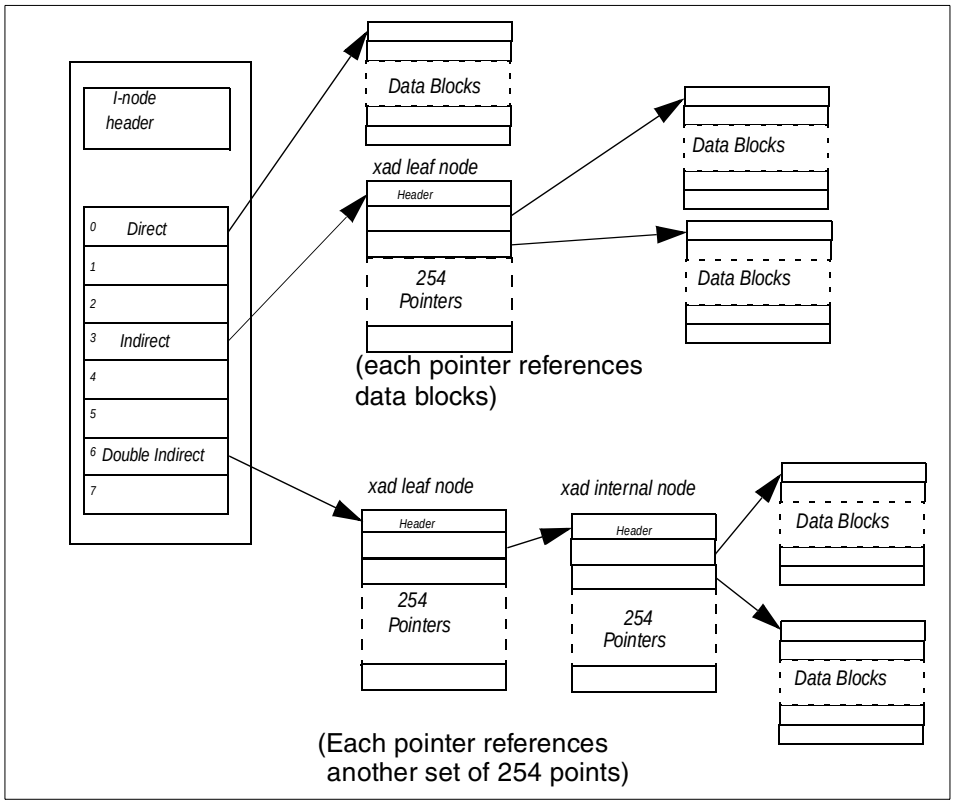

Figure 50. Anatomy of an i-node

4.3.4.1 In-depth look at the aggregate structure

An aggregate has:

- A 32K reserved data area at the front of it.
- Primary and Secondary Aggregate Superblocks. These superblocks contain aggregate-wide information, such as the size of the aggregate, the size of the allocation groups, the aggregate block size, and so on. The secondary aggregate superblock is a direct copy of the primary aggregate superblock used if the primary becomes corrupted. These superblocks are at fixed locations on the disk.
- An Aggregate Inode Table. This table contains an array of inodes describing the aggregate-wide control structures. This table is critical for finding any file system information.

- A Secondary InodeTable. This is a replica of the Aggregate Inode Table. The actual data for the inodes is not repeated, just the addressing structures used to find the inodes themselves.
- An Aggregate Inode Allocation Map. This map describes the Aggregate Inode Table and contains allocation state information on the aggregate inodes as well as their on-disk location. Since the inode allocation is dynamic, there is not any relationship between an inode number and the disk address of the inode, thus, the map is required.
- A Secondary Aggregate Inode Allocation Map. This map is a duplication of the Aggregate Inode Allocation Map.
- *A* Block Allocation Map. This map describes the control structures for the allocation and deallocation of aggregate disk blocks within an aggregate. It tracks the allocated and freed disk blocks for the entire aggregate.
- A Fileset Inode Table. This table contains an array of inodes that describe the fileset-wide control structures.
- A Fileset Inode Allocation Map.This map describes the fileset Inode table. It contains allocation state information on the fileset inodes as well as their on-disk location. There is also a secondary version of this map that points to the same data. When the fileset is initially created, the first inode extent is allocated, and additional inode extents are allocated and deallocated dynamically as required.
- File and Directory logical blocks.
- fsck (chkdsk) Working Area. This area provides space for check disk to be able to track aggregate block allocations. This space is used when dealing with very large aggregate block allocations since there may not be sufficient memory to do the tracking. There is a one-to-one relationship of bit to aggregate block. The space is described by the superblock and is always found near the end of the aggregate just before the In-line log.
- In-line Log. This space is used for logging the changes to the meta-data of the aggregate. This space is described by the superblock and is found at the end of the aggregate.

Figure 51 on page 135 summarizes this structure and details the role of some of the inodes.

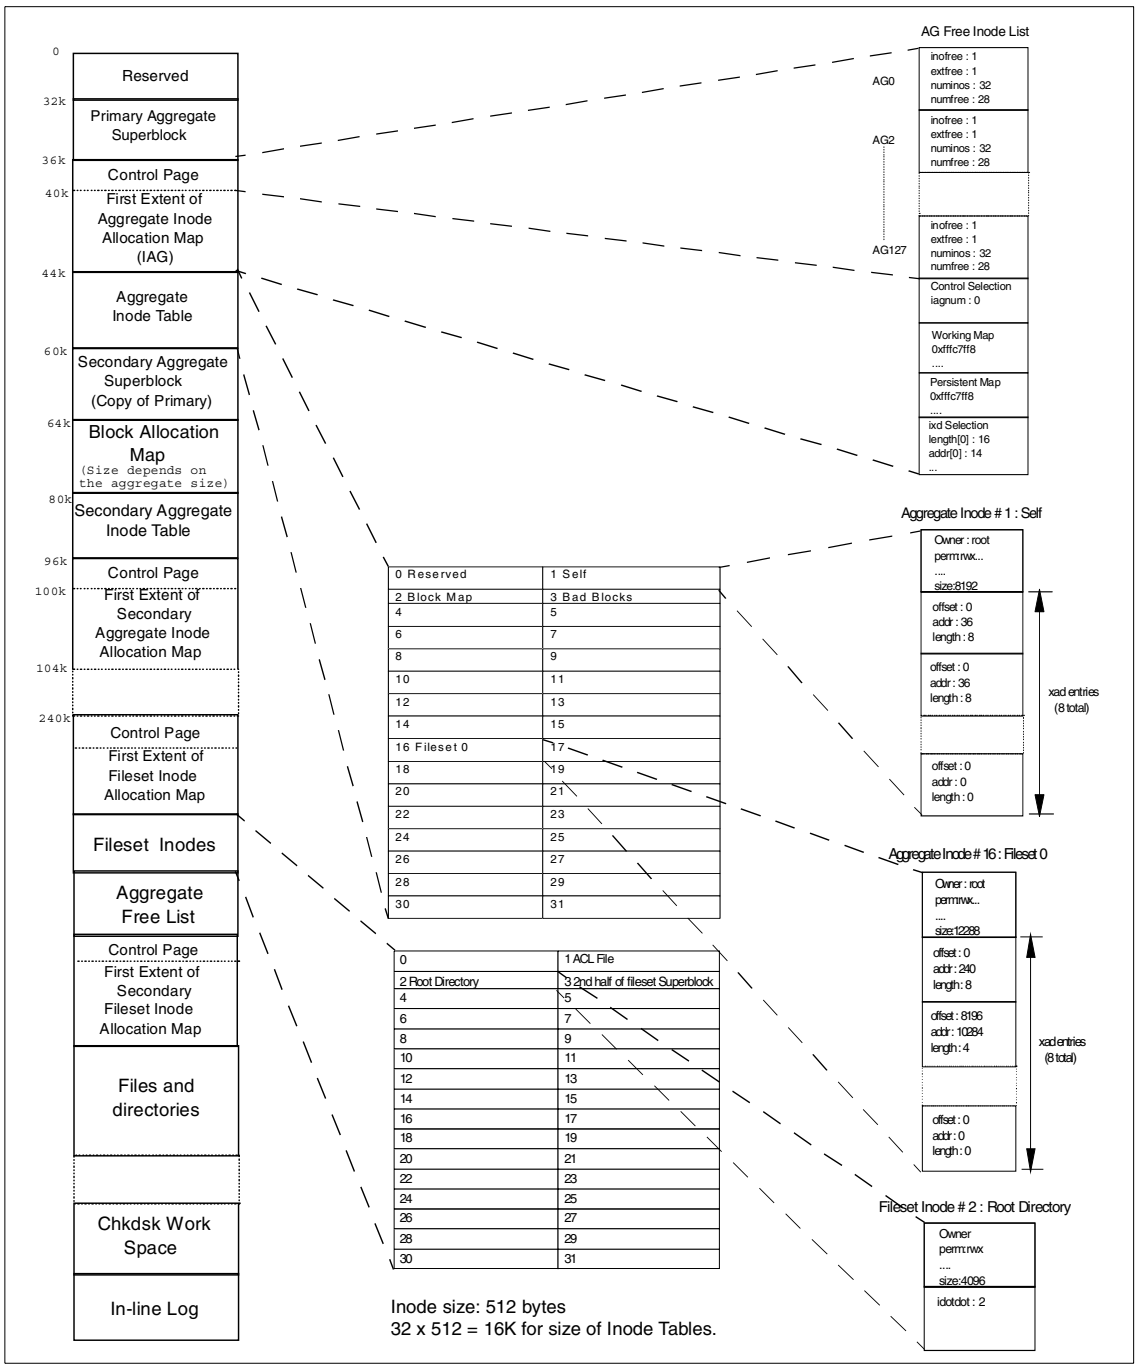

Figure 51. Aggregate details of a JFS partition.

4.3.5 JFS utilities

The JFS component is made up of two pieces. The first has already been covered; it is the Installable File System (IFS) that provides read/write file access and persistent storage on local machines while supporting OS/2 file system semantics (that is, functional behaviors and error returns) and coexisting with other IFS implementations including FAT and HPFS. Details of the IFS framework can be found in section 4.2.4, “Key components of the LVM” on page 100. The second piece is a set of utilities for creating and maintaining JFS at the lowest (media) level.

The JFS support utilities have been integrated into the IFS framework to extend the `format` and `chkdsk` commands. Two new utilities, defragfs and extendfs, have also been introduced to assist in the maintenance of logical volumes. Defragfs defragments the file system's free space. Extendfs increases the size of the file system after it has been expanded. Both of these new utilities are stand-alone 32-bit commands that interact with the OS/2 kernel and underlying device driver. Figure 52 illustrates the integration of the JFS support utilities into the IFS framework.

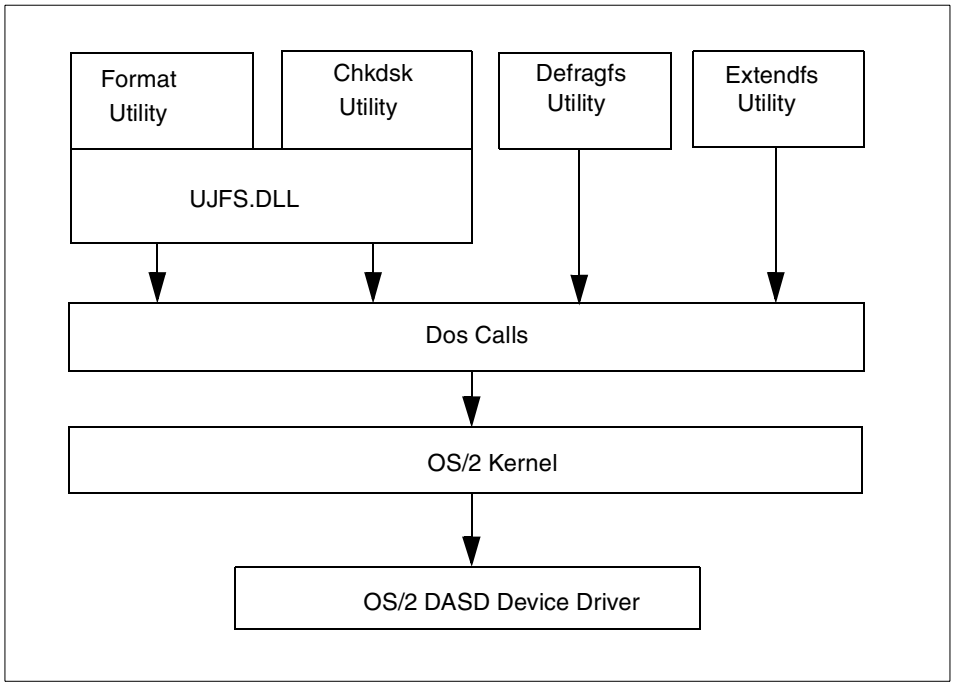

Figure 52. File system utility framework.

4.3.5.1 JFS-specific FORMAT Utility Options

The `format` command has a number of additional parameters added to support JFS.

The JFS-specific syntax for FORMAT is as follows:

FORMAT drive [/FS:JFS] [/V:<volume label>] [/BS:<size_in_bytes>]

[/LS:<size_in_megs>] [/S] [/L]

Table 10. Parameters for JFS format

Parameter	Description
drive	Specifies the drive to run FORMAT on.
/FS:JFS	Specifies that the drive should be formatted as a JFS file system.
/V:<volume label>	Specifies the volume label. The volume label can be up to 11 characters long. To include blanks in a volume label, enclose the label in quotation marks. If you do not specify a volume label, it will be set to null.
/BS:<size_in_bytes>	Specifies block size to use for the file system. If not specified, FORMAT will set the file system block size to 4096 bytes. The block size specified must be one of 512, 1024, 2048, or 4096 and cannot be smaller than the device's sector size.
/LS:<size_in_megs>	Specifies the size of the journaling log to create in megabytes. If not specified, FORMAT will create a log of 0.4% of the file system size.
/S	Specifies that the file system created should support sparse rather than dense files. If not specified, dense file will be supported.
/L	Specifies that FORMAT should verify the disk media, that is, long format.

4.3.5.2 JFS-specific CHKDSK Utility Options

The `chkdsk` command has a number of additional parameters to support JFS.

The JFS-specific syntax for CHKDSK is as follows:

```
CHKDSK [/C] [/F[:{0|1|2|3}] [/V] [/B] drive
```

The following table lists the CHKDSK parameters:

Table 11. Parameters for CHKDSK

Parameters	Description
/C	Perform verify/repair functions only if drive is in an inconsistent state.
/F:0	Analyze only. No write access given.
/F:1 or /F:2 or /F:3 or /F	Repair all problems detected. Initiate /B processing (see below) after all other analysis and repair processing is completed and the volume is mounted.
/V	Verbose messaging.
/B	LVM Bad Block List Processing. This is the only chkdsk processing which both writes to the drive and is performed while the file system is active. In effect, this processing makes bad blocks permanently unavailable so that they can be removed from the LVM's bad block list, which is fixed in length.
drive	Specifies which drive to run CHKDSK on.

Note

If the F option is not specified then /F:0 is performed.

Consistent with HPFS, the V option has no actual effect on messaging.

For chkdsk during boot processing (also known as autocheck), the /C and /F options should be used. This will enable log replay processing; if that fails, chkdsk will repair the file system.

In addition to the options described above, the JFS chkdsk utility may be invoked with the undocumented debug (/d or /D) option specified. This option causes the chkdsk utility to process the file system as usual, but, in addition to being written to the JFS chkdsk service log, all log messages are issued to stdout in English.

4.3.5.3 JFS DEFRAGFS utility syntax

DEFRAGFS defragments a file system's free space. This is a description of the DEFRAGFS utility options.

The syntax for DEFRAGFS is as follows:

DEFRAGFS {[/Q]} drive

Table 12. Parameters for DEFRAGFS

Parameters	Description
/Q	Specifies query only. It will give the current file system status and perform no defragmentation activities.
drive	Specifies the drive on which to run DEFRAGFS.

4.3.5.4 JFS EXTENDFS Utility syntax

EXTENDFS extends the length of the file system to use the entire partition on which it residues. Extending the length of the file system is automatically taken care of. Should a reboot be necessary, this command is automatically placed in the CONFIG.SYS and automatically runs checking all partitions. Should you remove it from the CONFIG.SYS, you may need to run it manually.

This is a description of the EXTENDFS utility options.

The syntax for EXTENDFS is as follows:

EXTENDFS [/LS:<size_in_megs>] drive

The following table lists the parameters for EXTENDFS:

Table 13. Parameters for EXTENDFS

Parameters	Description
/LS:<sixe_in_megs>	Specifies the (total) size of the journaling log to create in megabytes. If not specified, EXTENDFS will create a log of 0.4 percent of the file system size.
drive	Specifies the drive on which to run EXTENDFS.

4.3.5.5 JFS-specific configuration options

To configure JFS, an IFS= statement must be added to CONFIG.SYS. This section describes the syntax of this statement.

The JFS configuration is described in Chapter 4.3.1, “JFS cache” on page 121.

4.3.6 A closer look at chkdsk

The chkdsk utility has been enhanced for JFS volumes. JFS features allow chkdsk design to balance processing time and workspace size requirements. If the workspace was too small, chkdsk would not be able to process the file system at all. If chkdsk took to long to process (like the old chkdsk did), it would be useless in processing large file systems. The following features were selected to achieve a balance between time and size:

- A portion of the chkdsk workspace is within the file system itself dramatically reducing chkdsk's consumption of dynamic storage.
- chkdsk optimizes processing for the correct file system at the expense of processing time for error cases.
- This implementation of chkdsk uses 12 bytes per i-node and completes its processing in one sequential pass in which all i-nodes are read and one sequential pass where only directory i-nodes are read.

4.3.6.1 chkdsk processing flow

Preliminary: Analyze the parameters specified at invocation. Open the aggregate. Determine whether chkdsk has write access. Verify/Correct the aggregate superblock.

Phase 0: Replay the journal log.

Phase 1: During Phase 1, processing consists of validation, recording information in the workspace, and recording blocks allocated. Initialize the chkdsk workspace. Process aggregate i-nodes. Process fileset super i-nodes. If any multiple-allocated blocks have been identified at this point, chkdsk exits. Process all fileset i-nodes. Count child links.

Phase 2: Count parent links.

Phase 3: This phase locates the first reference to any multiple-allocated blocks identified as Phase 1. For each directory i-node in use, verify that it has only one parent and that the stored parent i-node number matches the i-node number of the parent chkdsk observed in phase 1.

Phase 4: For each i-node record in chkdsk's workspace i-node map which chkdsk has flagged for some kind of repair (excepting a directory with illegal hard links), flag any implied repairs.

Phase 5: Detect problems related to i-node connectedness. Identify each unconnected in-use i-node and flag it for reconnection. Identify each directory i-node with multiple parents and flag it for repair.

Phase 6: Perform all approved i-node corrections.

Phase 7: Rebuild/Verify i-node allocation maps.

Phase 8: Rebuild/Verify disk allocation maps.

Phase 9: Reformat the journal log if journal log replay had an error.

Final: Refresh redundant copies of the superblock and Aggregate I-node Table.

4.3.7 System limits

The Majority of the JFS size limits are imposed by the OS/2 kernel and system interfaces rather than the JFS itself. These limitations, both actual (OS/2 imposed) and potential (the JFS limit), are detailed below:

Buffer Cache The size of the buffer cache is configurable and is limited by the amount of virtual memory that can be allocated to the kernel. This limit is approximately 1GB. The JFS cache is designed to support a cache of up to 2 GB.

Note

The default JFS cache size is 12.5 percent of the system's physical memory size. To change the JFS default cache size from its default value, edit the JFS.IFS statement found in CONFIG.SYS to read:

DEVICE=C:\OS2\JFS.IFS /CACHE:(XX) /AUTOCHECK:D

where XX is the size in kilobytes and D: is the JFS volume.

File System Size The maximum file system size imposed by the width of the device driver I/O request interface is 2 TB. However, the JFS is designed to support files of 512 TB (with a block size of 512 bytes) to 4Petabytes (with a block size of 4KB).

File Size The file size limit has been increased to 2 TB; this is imposed through the width of the file offsets and sizes beginning at the DOS APIs. The JFS file size is governed by the underlining file system.

Number of File System Objects The maximum number of files and directories within a single JFS system is well over 4 billion. This limit is imposed by the JFS and is due to the size of the unique ID (inode number) used by the JFS to identify files and directories within a file system, since the JFS represents inodes as a standard, unsigned, 32-bit integer.

Extended Attributes The extended attributes for any single file or directory is limited to 64 KB. This limit is imposed by the width of extended attribute sizes beginning at the DOS APIs. However, JFS, internally, has a size equal to the amount of storage that can be represented through a single extent. The size of an extent varies from 8 GB (with a block size of 512 bytes) to 64 GB (with a block size of 4 KB).

4.3.8 Performance considerations

As already mentioned, the JFS code has been highly optimized for the 32-bit environment. Much extra 32-bit code has been introduced into the OS/2 base code (such as KEE32 and IFS32; see section 4.2.4 on page 100) so that JFS gains performance from the flat 32-bit memory address space. The 32-bit implementation has also allowed JFS to scale better in SMP environments.

Additional APIs have been added to the instruction set to allow TCP/IP applications, such as Notes or Web servers, to have gains in performance. This results from the direct access some TCP/IP instructions have to data stored in the JFS cache. Preliminary tests have shown that JFS serving local applications (thus acting as an application server) outperforms HPFS386, almost doubling the performance. When used to serve Notes to a large number of users, it was shown that the response time of JFS was quicker than even Windows NT NTFS.

Unfortunately, JFS performance as a file and print server is not as easy to compare. First, when any performance comparisons between HPFS386 and JFS are made, networking consideration need to be removed. The network bottleneck will impede server performance long before the file system will, thus, a fast network is recommended. In fast networks, it is the hardware configuration of the server that is going to govern the choice of file systems. If the server is only uniprocessor, JFS can have up to a 40 percent degradation over HPFS386. Figure 53 on page 144 is provided to help better understand the mechanics of OS/2 Warp Server for e-business as a file server.

The reason for the difference in file serving throughput between JFS and HPFS386 is increased path length for JFS due to the additional ring transitions it must do. These extra ring transitions require extra cpu cycles.

Examples of paths traced through the system are:

- Local access to HPFS file in HPFS386 cache: f-g
- Local access to JFS file in JFS cache: f-h (f-h is faster than f-g)
- Local access for HPFS file which is not in the HPFS386 cache: f-g-j
- Local access for JFS file which is not in the JFS cache: f-h-k (this is faster than f-g-j)
- Local access to FAT file: f-i
- Network access to HPFS file in HPFS386 cache: a-b-c
- Network access to HPFS file not in HPFS386 cache: a-b-c-j
- Network access to JFS file in JFS cache (NetBIOS): a-b-d-e-h
- Network access to JFS file not in JFS cache (NetBIOS): a-b-d-e-h-k
- Network access to JFS file in JFS cache (TCP/IP): m-n
- Network access to FAT file: a-b-d-e-i

JFS scales better in an SMP environment, thus, at 4-way throughput between HPFS386 and JFS, it has been shown to be similar, with HPFS386 only slightly better, but, in 5-way or better, JFS is expected to excel. If the server CPU is underutilized (less than 50 percent free most of the time), JFS performance could be equivalent to HPFS386 in 4-way or less scenarios because the CPU will have enough time to process the ring transitions.

Another point that may need to be considered is if the hardware configuration provides fault tolerance. If not, then HPFS386 is still the better option since it provides RAID-1 (both mirroring and duplexing) fault tolerance. This will need to be considered in small server installations.

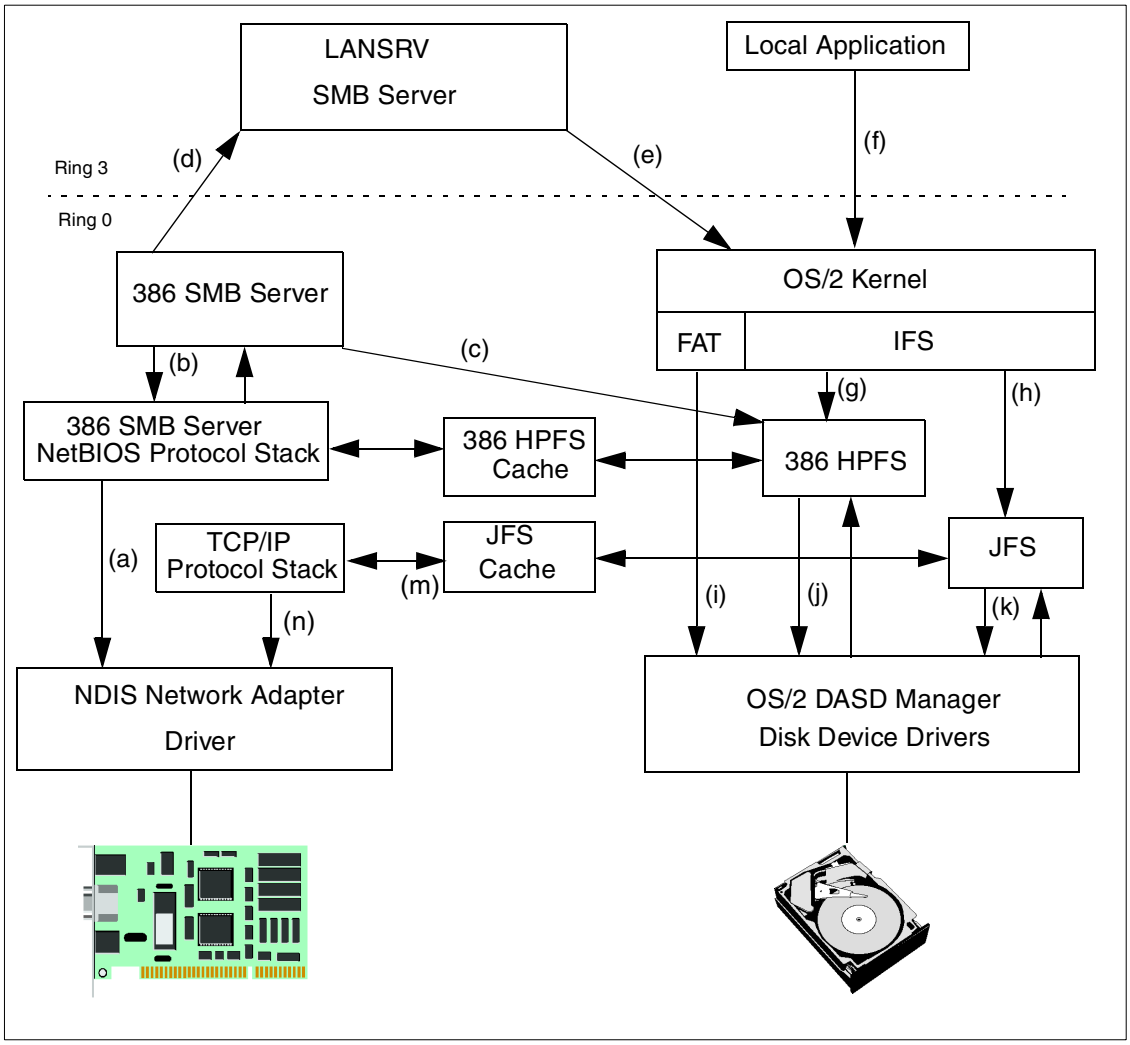

Figure 53. SMB system architecture

Chapter 5. File and print services

This chapter describes the File and Print Sharing service of OS/2 Warp Server for e-business. The File and Print Sharing Service is a Local Area Network application. It allows you to share hardware and software resources that are located on the server. Once you have connected to a network resource, you may use that resource in the same way as you use a local resource.

The File and Print Service, originally named LAN Server, was a core component of OS/2 Warp Server as well as OS/2 Warp Server for e-business. There are two comprehensive guides to this service: *Inside OS/2 LAN Server 4.0*, SG24-4428, and *Inside OS/2 Warp Server: Exploring the Core Components*, SG24-4602. Both of these books contain in-depth descriptions of the File and Print Service.

The File and Print service component within OS/2 Warp Server for e-business contains a number of service and functional rollups. Functional enhancements to the File and Print service, available for download from software choice, have been integrated into OS/2 Warp Server for e-business. They are:

- IBM Neighborhood Browser Enabler 1.0 for OS/2 Warp Server and OS/2 Warp Server SMP
- IBM Networks Client 4.1 for Windows 95 (now includes support for Windows 98)
- IBM Networks Coordinated Client 4.2.3 for Windows NT
- IBM Networks Primary Logon Client 4.2.3 for Windows NT

A number of new features have also been included with this release. These include:

- A number of capacity enhancements
- Support for Multiple Server Names
- New Command switch (net use /perm)
- Windows NT Server Integration

This chapter will attempt to provide a basic overview of the service and focus on the rollups and the new features that have been added. The Windows NT Server Integration function is described completely in Chapter 6, “Integrating Windows NT Servers” on page 181.

5.1 Review

The File and Print service component consists of a number of subcomponents. A high level architectural view of the service with all the subservices is shown in Figure 54 on page 148.

Each of these services are programs that are part of LAN Server. The available programs are identified in the\IBMLAN\IBMLAN.INI file. To control which services are started at startup, as well as the configuration of each of the services, edit the IBMLAN.INI file on your workstation. Temporary changes to network services can be made through the OS/2 Warp Server Administration or the command line (instead of editing the IBMLAN.INI file). Using the NET command, services can be started, stopped, and some paused and continued.

Some of the network services are briefly described in the following list. (For detailed descriptions of the services, view the on-line guides.)

Alerter — This service notifies selected user IDs when problems occur. It also notifies the Generic Alerter service when certain LAN problems are detected or anticipated. This service cannot be paused.

DCDB Replicator — This service copies the domain control database from a primary domain controller to one or more backup domain controllers.

Generic Alerter — This service enables the server to build and send Systems Network Architecture (SNA) alerts. The Alerter service notifies the Generic Alerter service when certain LAN problems occur. This service cannot be paused.

LSserver — This service provides DOS LAN Services support and logical server functions. The logical server supports remote requests from clients for activities such as spooling, querying users, logon, and logoff. The default value for LSserver is to start when OS/2 Warp Server or LAN Server is installed. The user should never change the default. If you stop LSserver, the Server service automatically stops. Therefore, this service should never be stopped and cannot be paused.

Messenger — This service supports the reception of messages at a client or server. This service cannot be paused.

NetLogon — This service copies the master user and group definitions file located on the domain controller to network servers.

This service is available only on servers. In order for the NetLogon service to replicate user and group definitions across servers in a domain, digit 3 of the srvheuristics parameter (on domain controllers) and digit 8 of the wrkheuristics parameter (on additional servers) must be left at their default values. If you change the defaults for these parameters, user and group data may not be replicated from the domain controller resulting in unknown user IDs or group IDs on the additional servers.

Netrun — This service handles requests for running programs remotely on a server.

Network Neighborhood Browser Enabler — This service allows OS/2 Warp Server to function as a master browser for Windows Clients. The master browser function provides Windows 95 and Windows NT clients the ability to view the domain's LAN Server machines and their resources via the Network Neighborhood object.

Remote IPL — This service allows the Server service to support remote initial program load (remote IPL) of workstations. This service corresponds to the Remoteboot section of the IBMLAN.INI file.

Replicator — This service copies files from a master location on a server to one or more servers or clients requiring a copy of the data. This service cannot be paused.

Requester — This service redirects requests for files, printers, and serial devices from one workstation to another workstation.

Server — This service receives and responds to network requests for files, printers, and serial devices. The service checks the requests against its database of user IDs and access permissions.

Timesource — This service designates a server as a source of a reliable time and date with which other workstations on the network can synchronize. The Timesource service does not keep time. It allows other workstations on the network to identify a server with a reliable clock. The default value for the Timesource service is to start on domain controllers when OS/2 Warp Server or LAN Server is installed. This service cannot be paused.

UPS — This service provides protection against power failure. If power is interrupted, the UPS service keeps the server running until the service can shut down the server safely or until an administrator stops the server.

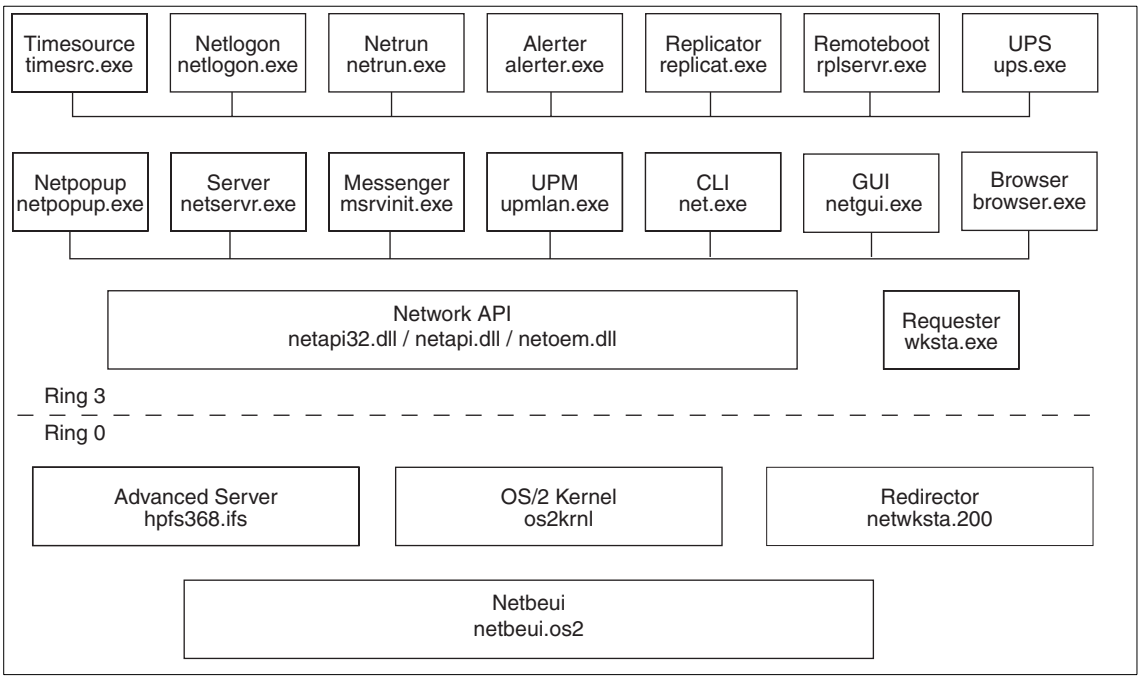

Figure 54. File and print service, architecture

The file and print services run at the application privilege level. OS/2 services are used to satisfy network file I/O requests session setup and resource sharing. All Ring 3 requests include requests for files stored on a FAT, JFS or HPFS partition, print requests, and serial device requests.

Network file I/O requests from client workstations and responses are sent using server message blocks (SMBs). The Server processes SMBs using internal network buffers. The IBMLAN.INI file parameters that define the size and number of network buffers on the server are sizreqbuf and numreqbuf**.** The IBMLAN.INI file parameters that define the size and number of network buffers on the requester are sizworkbuf and numworkbuf.

An SMB received from the network is copied into the adapter receive buffers by the network adapter. The NetBIOS device driver using a global descriptor table (GDT) selector copies the data from the adapter's receive buffers into

an available server network buffer. The NetBIOS device driver can acknowledge the message or piggyback the acknowledgment on a subsequent network message. The PROTOCOL.INI file contains the configuration information for the NetBIOS device driver.

The SMB is passed through the redirector to the server. The redirector is a requester component that directs file system request traffic between the server, the file system, and the network. Three types of SMB protocols exist that can be used for transferring data between a requester and a server:

- Core SMB protocol
- Read and Write (RAW) SMB protocol
- Multiplexed SMB protocol

The Server is designed to optimize the movement of file I/O from the server to the requester. It supports JFS, FAT, HPFS, and HPFS386. For more information on these, refer to Chapter 4.1.6, “Comparison of features of FAT, HPFS, and HPFS386 JFS” on page 93.

The following figures show the overall component structure of the server and an OS/2 requester.

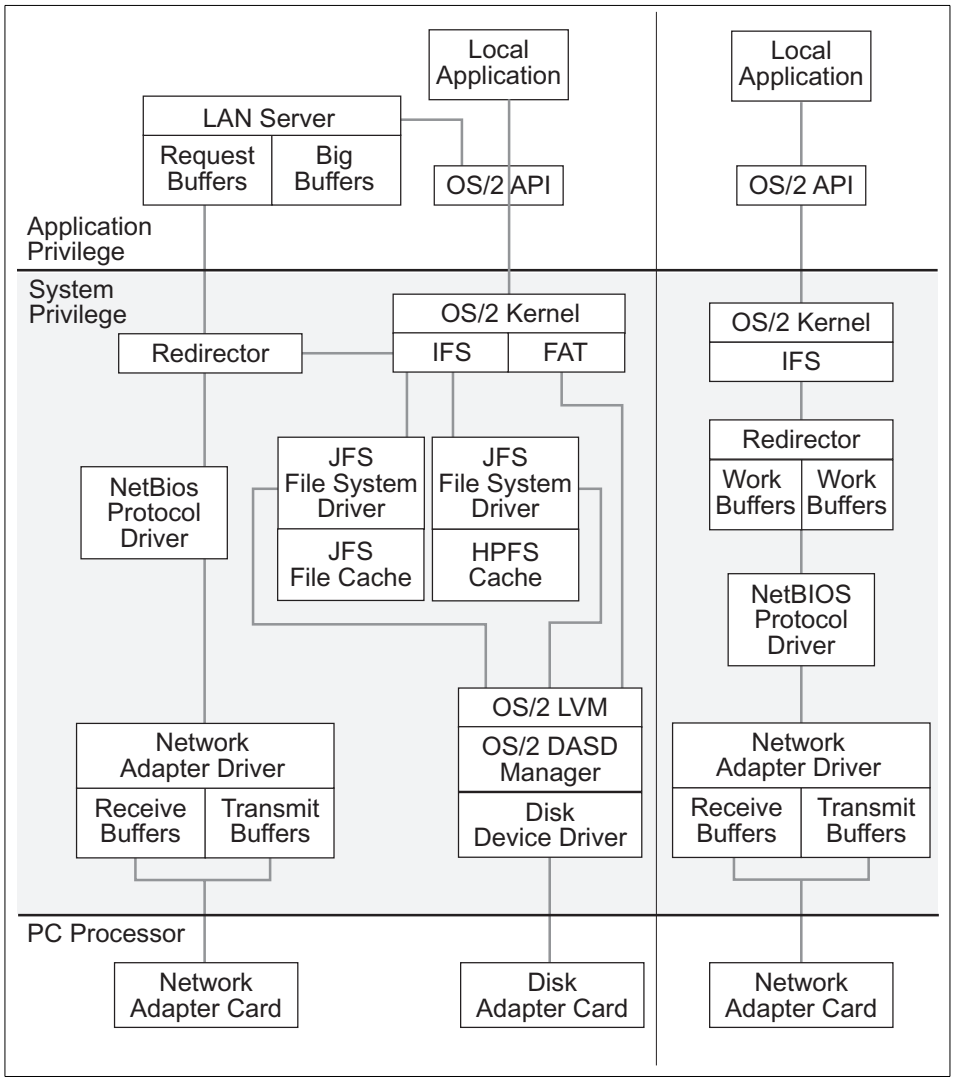

Figure 55. OS/2 client and server architecture

5.2 Installation

The File and Print service can be installed as part of the OS/2 Warp Server for e-business integrated installation program. The service can also be installed in attended or unattended mode once the base operating system has been installed.

As described above, the File and Print Services component includes a number of supporting services such as Netlogon, Messenger, Alerter, and so on. Some of these services are mandatory, and others are optional. The installation procedure provides the user with an opportunity to select/deselect the services.

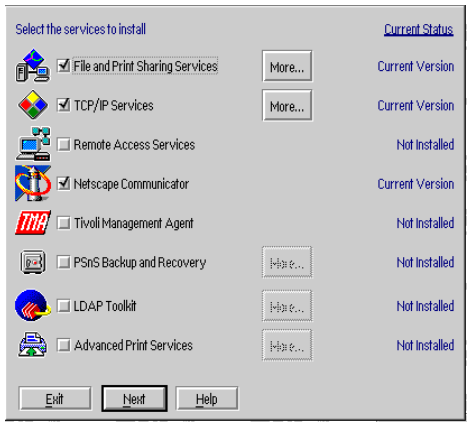

Figure 56. Selective install for networking

During an attended installation, or when using the Selective Install for Network services, the screen shown in Figure 56 on page 151 is displayed. This lists all the components that can be selected for installation. One of them is the File and Print Services component.

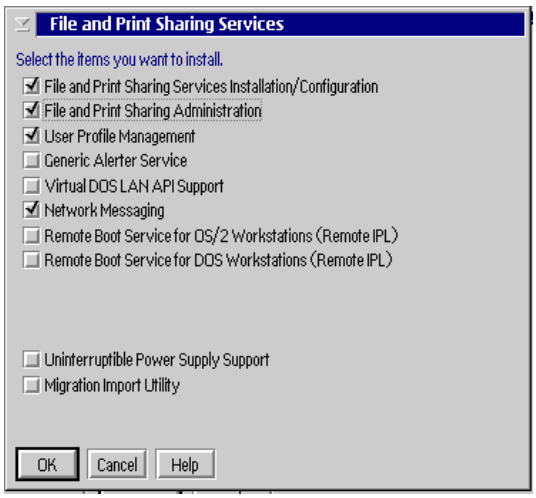

Figure 57. File and pint services features

Figure 57 on page 151 shows the services that can be selected for installation.

5.3 File and print services configuration

During the configuration process, you are given the opportunity to configure the network features, the server name, the domain name, the server role, as well as choose a user ID and password.

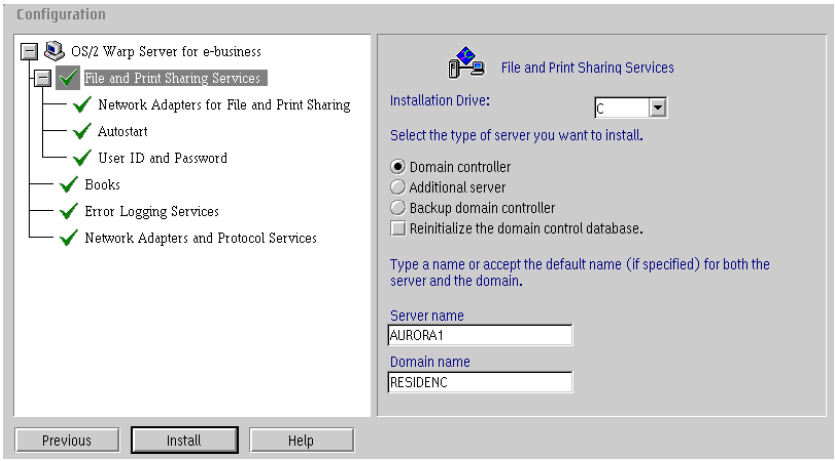

Figure 58. File and print services, configuration.

Selecting the File and Print Sharing Services brings up the File and Print Services Configuration Window on the right half of the configuration screen.

The user can select the installation drive where these services can be installed. The Default drive is the drive where the operating system is installed.

The server role has to be specified. A server in a domain can take on different roles.

Domain Controller: There should be only one Domain Controller in the domain. This is crucial for the domain because it maintains the user database and does the validation of users and servers.

Backup Server: This server has a read only copy of the user database, which is constantly replicated from the Domain Controller. The Backup server also handles logon validation.

Additional server: Additional server cannot perform logon validation. It uses and shares the data on the network

Depending on the type of server that has to be installed, select the role. Domain Controller is the default selected server role.

The name of the server and the domain into which this server has to be configured have to be entered.

Note

The Server name must be unique in the Network. If the server role is Primary Domain Controller, the Domain name must also be unique.

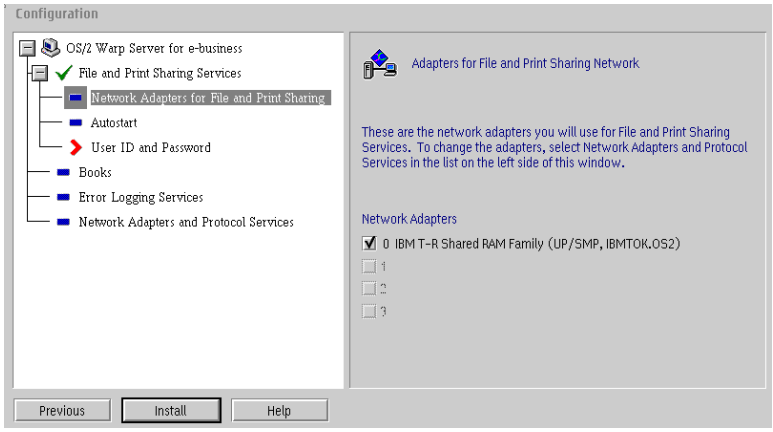

Figure 59. File and print services, selecting a network adapter.

The Austostart section gives the user the choice of selecting the services that need to be started at server startup time. Services that are not started at server startup can be started manually from the command line.

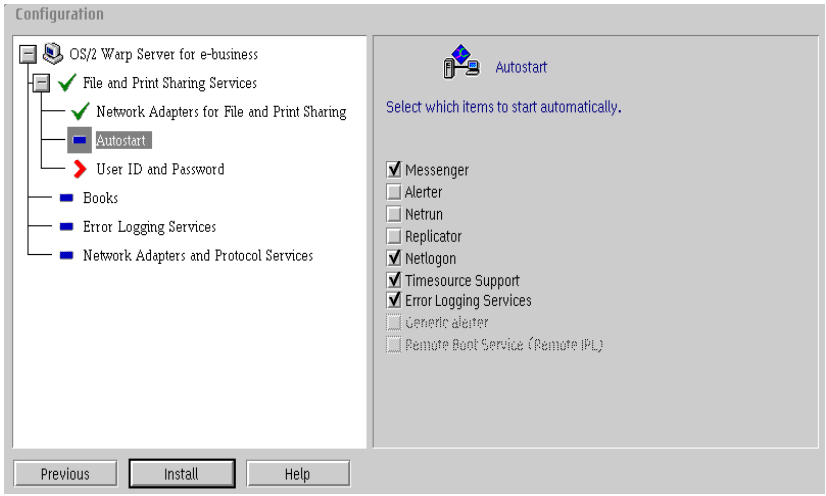

Figure 60. Autostart configuration

After the Installation and configuration is complete, the following icon is added to the desktop:

Figure 61. File and print services icon

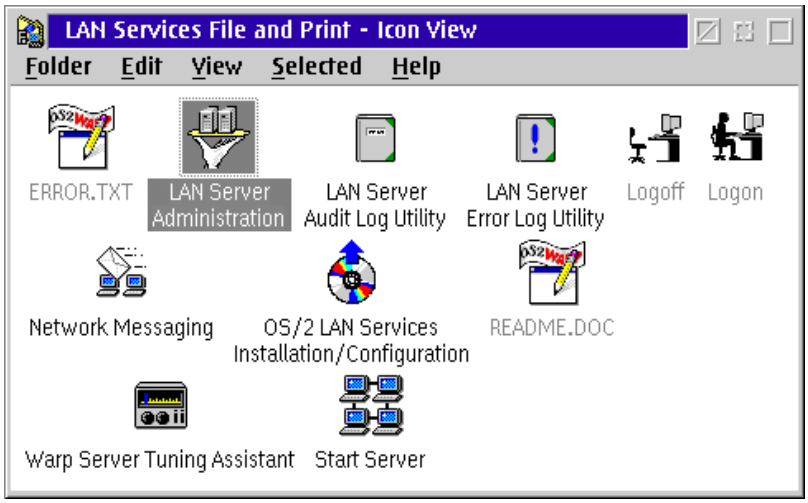

Figure 62. LAN services, file and print

The File and Print Services folder shown in Figure 62 on page 155 provides access to the following functions:

- LAN Server Administration allows you to administer the network. It is easy to use. Managing users, groups, or shares is as easy as drag and drop.
- LAN Server Audit Log Utility for auditing.
- LAN Server Error Log Utility for error logging.
- Logon allows you to perform domain logon.
- Logoff allows you to log off the domain.
- Network Messaging for sending and receiving network messages.
- OS/2 LAN Services Installation/Configuration: Use this function to install, reconfigure, or remove the File and Print Sharing services. It also provides you with the option to create response files for CID installation.
- Warp Server Tuning Assistant: This application assists with performance and capacity tuning.
- Start Server: Select this object to start the File and Print Sharing Services on the server.

5.4 File and print services administration

The File and Print service is administered via an easy to use graphical user interface. This interface was introduced with OS/2 LAN Server 4.0 and is object-oriented allowing the user or system administrator to configure and use the network via the manipulation of visual objects.

The paradigm used has a consistent look and feel with the Workplace Shell (WPS).

For a detailed description of how to perform each of the administrative tasks, refer to the on-line manual or to the redbook *Inside OS/2 LAN Server 4.0*, SG24-4428. Some of the tasks that a File and Print administrator may be expected to perform are listed below:

- Create and manage users, assign home directory, and logon applications.
- Manage resources that are to be shared on the network, known as aliases.
- Define permissions to allow restricted access to users of the shared resources.
- Define and provide access to applications that are stored on server workstations but may be executed from requesters.
- Manage server workstations on the network.
- Define Users and Groups

Each of the above tasks can be performed by using the GUI as shown in the figure below. To get to the GUI, double click on the domain icon (castle) that displays the tools for the OS/2 Warp Server for e-business domain administration.

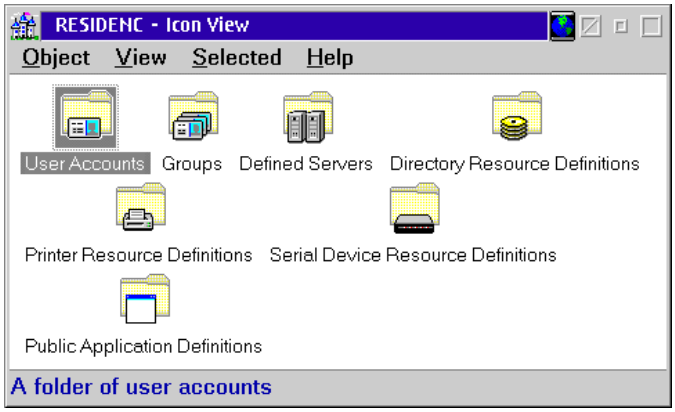

Figure 63. Domain administration

The figure below shows the contents of the User Accounts folder. To create a new user ID, drag the **UserID** template into an open space in the folder and fill out the details in the notebook.

.

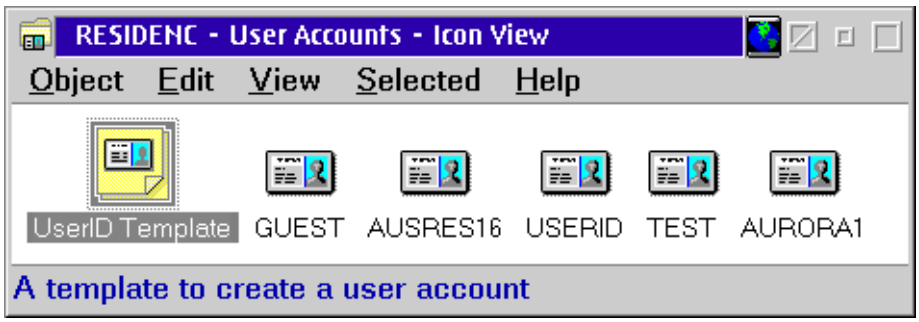

Figure 64. User's window

Other tasks are performed in a similar way.

The Command Line

Each of the tasks performed using the GUI can also be performed from the command line. REXX command files can be created for batch processing as well as automating processes.

For example,

Use Net Start <service name > to start a specific service

Use Net Stop <service name > to stop a specific service.

5.5 386 HPFS

As described in the introduction, 386 HPFS is available as a separate feature for OS/2 Warp Server for e-business. You must install OS/2 Warp Server for e-business before you install the 386 HPFS file system. Since Fault Tolerance and Local Security are dependent, they can only be installed with or after 386 HPFS has been installed.

386 HPFS is a server-optimized variant of HPFS. It provides improved access to large disk volumes, and it optimizes performance in a server environment where many files are opened simultaneously by clients. If 386 HPFS is installed, all HPFS volumes are managed by 386 HPFS. If you have already installed 386 HPFS on your system, the volume is formatted with 386 HPFS if you select the file system type of HPFS.

Fault Tolerance
Fault Tolerance lets your server handle disk hardware problems without significantly interrupting system performance or losing system data. Fault Tolerance provides drive mirroring and duplexing as well as error logging, alerting, and monitoring of disk activity. Drive mirroring and drive duplexing provide duplication of data stored on disk, thereby, improving data integrity.

If you install Fault Tolerance Support, you must run the FTSETUP utility after completing the 386 HPFS installation to enable and configure Fault Tolerance. For more information about installing and configuring Fault Tolerance, see the product documentation *Network Administrator Tasks*.

Local Security
Local Security allows you to restrict access by local users (users working at the server itself) to files on the server and on HPFS volumes.

5.5.1 Attended installation of 386 HPFS with fault tolerance

1. Close all applications on the server.
2. Open an OS/2 window.
3. Stop the server by typing NET STOP SERVER.
4. Insert the 386 HPFS Upgrade for OS/2 Warp Server for e-business CD into the CD-ROM drive. This is a two part installation so do not remove the 386 HPFS Upgrade for OS/2 Warp Server for e-business CD from the CD-ROM drive until the installation is complete.
5. Type `e:` and press **Enter**, where `e:` is the CD-ROM drive letter.

6. Type `cd\xx\install` and press **Enter**, where `xx` is the two-character abbreviation of the language version you are installing.
7. Type `install` and press **Enter**. The OS/2 Warp Server for e-business Setup and Installation window is displayed.
8. Select **386 HPFS**, and then click **More**. You will get the panel displayed in Figure 65 on page 159.

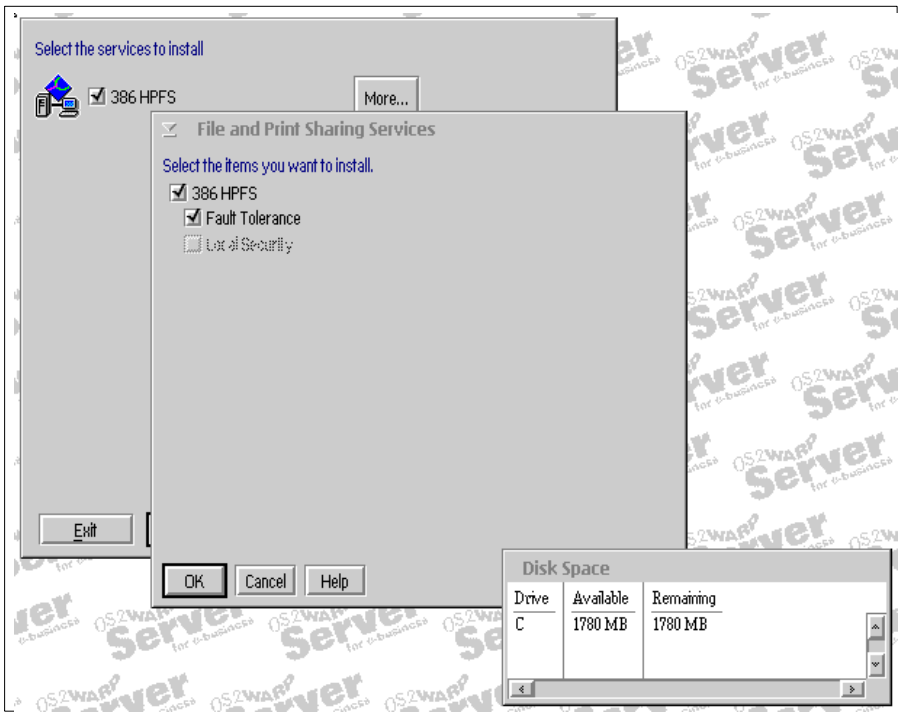

Figure 65. 386 HPFS installation panel.

9. If you want to install support for the Fault Tolerance component now, select it.
10. You cannot select the Local Security component until 386 HPFS is installed.
11. Click **Next** to continue the installation.
12. The Configuration window is displayed allowing you to configure the Cache, Lazy Write, and Heap settings.

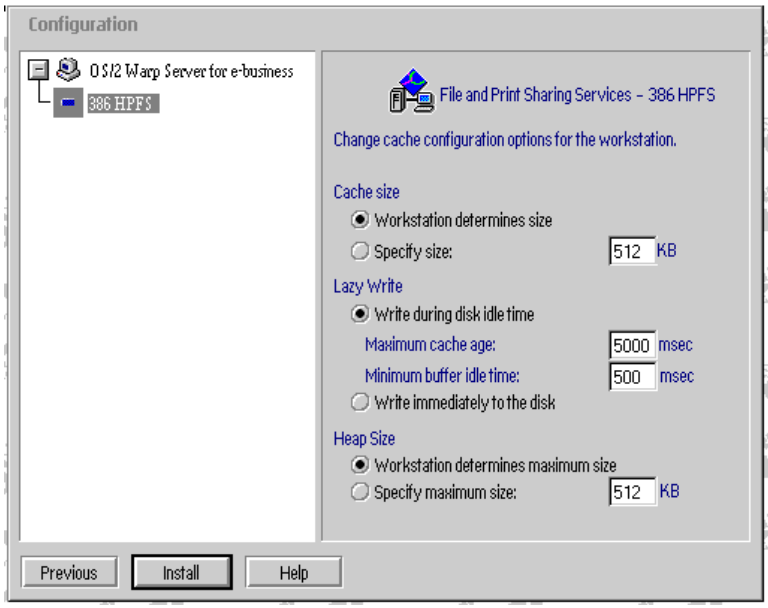

Figure 66. 386 HPFS configuration panel.

13. Click **Install**, and then click **OK** to complete the installation of 386 HPFS. The server is restarted twice during the 386 HPFS installation process.

5.5.2 Installing local security

If you want to use the Local Security component, you must install it after the 386 HPFS Upgrade is installed. You will go through the same procedure as described in Chapter 5.5.1, “Attended installation of 386 HPFS with fault tolerance” on page 158.

5.5.3 Log files created during installation

The following log files are created in the IBMINST\LOGS\HPFS386directory during installation of 386 HPFS:

- FS386ERR.LOG
- FS386HST.LOG

5.5.4 Installing 386 HPFS using CID

CID gives you the ability to install the 386 HPFS Upgrade, Fault Tolerance Support, and the Local Security component.

Previously, these components were installed as part of the IBM LAN Server installation (LANINSTR) on OS/2 Warp Server Advanced and previous IBM LAN server versions. It must now be installed separately from the IBM LAN server component during OS/2 Warp Server for e-business installation.

Note

These components must be installed *after* the IBM LAN server compone installation. If these components are installed before LAN server is installed, the system might become inoperable.

You will need to use the Feature Installer and the CLIFI utility to install these components. You need to provide CLIFI two response files as input parameters: fs386cid.rsp and fs386.rsp. The first file, fs386cid.rsp, might need modifications. Do not modify the fs386.rsp file.

To install 386 HPFS from a CID Code Server, perform the following steps:

1. Xcopy the files from the files from the HPFS386 directory to the CID\SERVER\HPFS386 directory on the code server. If you are upgrading a system that already has 386 HPFS installed on it, proceed directly to step 3. Otherwise, continue to step 2.
2. You must modify the CID response file fs386cid.rsp located in the HPFS386 directory. The keyword values must be updated to appropriate values for which component to install (386 HPFS, Local Security, or Fault Tolerance). Local Security may be installed only after 386 HPFS is installed and running. For the install keywords, the value must be 1 to install the component and 0 to not install the component.
3. Change fs386cid.rsp to meet your requirements. For Local Security support, 386 HPFS must first be installed, and the installation needs to be done in two phases.
4. The following steps provide an example of modifying the fs386cid.rsp file to install 386 HPFS, Fault Tolerance Support, and Local Security.
   1. Update the following information to fs386cid.rsp:

```
        Install386HPFS.Selection=1
InstallFaultTolerance.Selection=1
InstallLocalSecurity.Selection=0
WkStaDeterminesCacheSize.Selection=1
WkStaDeterminesHeapSize.Selection=1
ConfigLazyWrite.Selection=1
...
HPFS386_Top.InstallDrive=<C:>
...
HPFS386_Top.isIntegratedInstall=NO
...
```

where <C:> is the drive where 386 HPFS will be installed.

2. Copy the changed fs386cid.rsp file to another file name, and modify this new file for Local Security support. (If you do not need Local Security support, you can skip this step.)
3. Make the following changes to the file:

```
Install386HPFS.Selection=0
InstallFaultTolerance.Selection=0
InstallLocalSecurity.Selection=1
WkStaDeterminesCacheSize.Selection=0
WkStaDeterminesHeapSize.Selection=0
ConfigLazyWrite.Selection=0
...
HPFS386_Top.InstallDrive=<C:>
...
```

   where <C:> is the drive where 386 HPFS (with Local Security support) will be installed.

5. Run CLIFI to install 386 HPFS and Fault Tolerance Support using the fs386cid.rsp file modified in step 1. If you want to install Local Security support, reboot your system and run CLIFI again using the response file created in step 2.
6. Sample CLIFI invocation:

```
clifi /a:c
/r:<x:>\cid\server\hpfs386\fs386.rsp
/l1:<c:>\os2\install\fs386err.log
/l2:<c:>\os2\install\fs386his.log
/s:<x:>\cid\server\hpfs386
/b:<c:>
/r2:<x>\cid\server\hpfs386\fs386cid.rsp
```

   where <x:> is the server drive and <c:> is the boot drive. If you want to reinstall 386 HPFS, you must first uninstall it.

Thin386 changes

If you are using the Thin386 utility, note the addition of a new switch called /386Path. This switch should be used to point to the directory where the 386 HPFS files are installed (Example: C:\IBM386FS).

New response file keyword

The following keyword specifies the location of LAN Server:

HPFS386_TOP.LanDrv=E:

where E: is the drive where LAN Server is installed.

The default value for this keyword is the drive where the operating system is installed. Use this keyword if LAN Server is installed on a drive other than the operating system.

Determining the cache size
386 HPFS uses a default cache size if one is not specified. To view the cache size, type the following at the command line: `CACHE386 /O`. The cache size is controlled by settings in the file \IBM386FS\HPFS386.INI.

AIC78U2.ADD device driver
If your system configuration requires the AIC78U2.ADD device driver, and, if you install 386 HPFS, download the latest version of the driver from the Adaptec Web page:

`http://www.adaptec.com`

5.5.5 Existing ACLs

Any existing Access Control Lists (ACLs) that reference HPFS drives before the installation of 386 HPFS Upgrade are removed and saved in<LAN Drive>\IBMLAN\ACCOUNTS\<Drive>.acl.

One of these files is created for every drive with ACLs that is formatted for HPFS. After installation, you may restore the ACLs using the `PREPACL` command with syntax similar to the following:

```
PREPACL /R /B:<LAN Drive>\IBMLAN\ACCOUNTS\<Drive>.acl /L1:<error log path>
/L2:<history log path>
```

5.6 Functional rollups

As described previously, a number of components have been made available via software choice. These components have now been integrated into OS/2 Warp Server for e-business. This section will describe these components and how they may be enabled

5.6.1 IBM Neighborhood Browser Enabler 1.0

The ability to share resources is one of the biggest advantages of networked computer systems. To provide a way to determine what resources are available for Microsoft Windows Clients, the IBM Neighborhood Browser Enabler was developed. The IBM Neighborhood Browser Enabler helps you to maintain a centralized list of available resources and servers in your domain. The IBM Neighborhood Browser Enabler eliminates the need for every Microsoft Windows Client to maintain a list of shared network

resources. This lowers the network traffic required to build and maintain the list and frees the CPU time each Microsoft Windows Client will need to create a network resource list. To enable an OS/2 Warp Server to function as a master browser for Microsoft Windows Client, the IBM Neighborhood Browser Enabler for OS/2 Warp Server was developed. We refer to it as the IBM Neighborhood Browser Enabler. The Master Network Browser functionality allows Microsoft Windows NT and Microsoft Windows 95 clients to view the domain's OS/2 Warp Server machines and their resources. This is done by the Network Neighborhood object.

The installation, configuration, and use of the IBM Neighborhood Browser Enabler is covered in-depth in the redbook *OS/2 Warp Server Functional Enhancements: Part 1*, SG24-2008.

5.6.1.1 Enabling the Neighborhood Browser Enabler

The IBM Neighborhood Browser Enabler is automatically installed but not automatically enabled. The Neighborhood Browser Enabler should not be autostarted on every server on which it is installed. At most, it should be autostarted and run on one out of every 20 servers in a domain.

To start the Neighborhood Browser Enabler manually, perform the following steps:

1. Open an OS/2 Window.
2. Type `net start browser` and press **Enter**.

A message stating that the BROWSER service was started successfully is displayed.

To autostart Neighborhood Browser Enabler, perform the following steps:

1. Open the x:\IBMLAN\IBMLAN.INI file (where x: is the drive where file and print services is installed) into a text editor.
2. Type `BROWSER` at the end of the SRVSERVICES = line.
3. Save the IBMLAN.INI file and exit the text editor.
4. The next time the system is restarted, it will autostart the Neighborhood Browser Enabler.

5.6.2 IBM Network Client 4.1 for Windows 95

The installation, configuration, and use of the Windows 95 client is described in depth in the redbook *Network Clients for OS/2 Warp Server: OS/2 Warp 4, DOS/Windows, Windows95/NT, and Apple Macintosh*, SG24-2009.

5.6.3 IBM Network Client for Windows NT

There are two OS/2 Warp Server clients available for the Windows NT Workstation:

- IBM Networks Coordinated Logon Client for Windows NT 4.0
- IBM Networks Primary Logon Client for Windows NT 4.0

The differences between the two clients are summarized in Table 14 on page 165.

Table 14. IBM Network Clients for Windows NT – Summary

Primary Logon Client	Coordinated Logon Client
Network Environment	
OS/2 Warp Server domains	Windows NT Workstation clients, Windows NT Server and OS/2 Warp Server domains.
Network Protocols	
NetBEUI, NetBT (1)	NetBEUI, NetBT[1]
Logon Validation	
OS/2 Warp Server domain	Windows NT Workstation clients or Windows NT Server, OS/2 Warp Server domain.[2]

[1]NetBT is Microsoft's abbreviation for NetBIOS over TCP/IP (TCPBEUI).
[2]There is no need for a Windows NT Server to validate a logon request. The logon request can be validated by the user's Windows NT Workstation and should always do so in a OS/2 Warp Server network.

5.6.3.1 IBM Networks Coordinated Logon Client for Windows NT 4.0

The IBM Networks Coordinated Logon Client is intended for Windows NT workstations in a mixed LAN environment consisting of Windows NT and OS/2 Warp Server domains. A separate logon validation on the OS/2 Warp Server domain is performed in addition to the normal Windows NT logon validation.

5.6.3.2 IBM Networks Primary Logon Client for Windows NT 4.0

The IBM Networks Primary Logon Client has the same GUI functions as the IBM Networks Coordinated Logon Client. The main difference is in how it handles the logon process. This client is recommended for Windows NT workstations operating in a LAN environment consisting of OS/2 Warp Server

domains only. If there are also Windows NT domains, the IBM Networks Coordinated Logon Client should be considered.

Note

The IBM Networks Primary Logon Client requires Version 4.0 of Windows NT. Version 3.x is not supported.

5.6.3.3 For more information

The installation, configuration, and use of the Windows NT client is described in depth in the redbook *Network Clients for OS/2 Warp Server: OS/2 Warp 4, DOS/Windows, Windows95/NT, and Apple Macintosh*, SG24-2009.

5.7 Capacity enhancements

With the introduction of JFS in OS/2 Warp Server for e-business, it became necessary to make improvements upon some of the existing practical limitations regarding the usage of the resources. In the previous implementation of OS/2 Warp Server, there are limitations on the number of files that can be opened simultaneously, file/directory searches, and the maximum number of available shares. These limitations have been improved upon in the new OS/2 Warp Server for e-business.

File and Print for JFS has been implemented in OS/2 Warp Server for e-business as a Ring 3 server. This makes it subject to the limits of HPFS or FAT. Since this new file system has several distinct advantages, the aim of the capacity enhancements is to bring the Ring 3 Server limits up to the current HPFS386 or Ring 0 server limits.

5.7.1 Overview

In order to understand the implications of the new enhancements, a brief description of the requester and server communications is described below.

Requesters and servers establish a *virtual circuit* which is a reliable communication channel between two machine names where received data is acknowledged. Requesters and servers use *sessions* to send and receive data. A session is a connection that has a link (virtual circuit) and a user-name associated with it. One session will link a user to all of the resources being used on one server. In other words, one session will be used for all Net Use resources for a user to a server.

If the user uses resources on another server, a different Virtual Circuit /session is established between the workstation and server. When a user logs

on to the network, the user is establishing an association with the workstation being used to the domain. A session is not necessarily established. However, if a user logs on to a server and uses a resource, a session exists.

The enhancements that are described below are tunable parameters in the IBMLAN.INI file. When the File and Print services code starts, usually, by issuing a `Net Start service` command, it reads the x:\IBMLAN\IBMLAN.INI file and initializes internal structures accordingly. Most of the internal structures set up at initialization time cannot be changed without stopping and starting the LAN code or, at times, rebooting the system.

The parameters below existed in previous versions. The maximum values of these parameters have been changed; default values are assigned to them during installation.

5.7.2 Maximum number of connections

This parameter specifies the maximum number of connections that requesters can have to the server. This is the number of `NET USE` commands the server can handle.

For example, a user issuing five `NET USE` commands needs five connections. Five users who each issue one `NET USE` command need five connections. Increase this parameter value if many users access the server. This parameter value must be greater than or equal to the maxusers parameter.

The previous version of the Warp Server has a limit on the number of connections a client can make to a server. Ring 0 server, which is the HPFS386 server, allows 4096 connections to HPFS resources and Ring 3 server allows 2048 connections to non-HPFS resources. If a customer is using JFS and not HPFS386, the customer will be limited to 2048 connections allowed by Ring 3 server. With a maximum of 1000 users, this would leave two connections per client putting a severe restriction on the number of connections.

To alleviate this restriction, OS/2 Warp Server for e-business supports an increased number of connections. Up to a maximum of 16384 connections are supported. So, in a server environment with 1000 users, the new limit will allow 16 connections per user irrespective of whether HPFS386 is running or not.

The previous limit was artificially imposed by the server. By changing the way in which IDs for connections are built and used, the number of connections could be increased.

Table 15. Maxconnections, details.

Default	Minimum	Maximum
300	1	16384

5.7.3 Maximum open files

The MAXOPENs parameter sets the maximum number of files, pipes and devices the server can have open at any time. An open is counted whenever a user access a resource, such as a file. For example, the value of this parameter must be greater than or equal to 100 for a user opening 100 files. If 100 users access the same file, it counts as a 100 opens even though the file was open before.

Note

The maximum number of open files is 8192. However, the maximum number of unique open files is 1279. The first opening of a file counts against the maximum of 1279. Additional openings of the same file count against the maximum of 8000.

In the previous version of the Warp Server, the server can only allow a maximum of 8192 files/pipes/devices open at a time. The introduction of the JFS in the OS/2 Warp Server e-business encourages the user to have huge storage space. This invariably increases the usage of resources, which results in the increased number of access to the resources. A maximum of 8192 opens may sometimes not be sufficient.

The OS/2 Warp Server for e-business has improved upon the maximum open files/pipes/devices to 64-1K.The server can now support 65535 open files. It means that the server can have 64-1K resources open concurrently.

The MAXOPEN parameter in the IBMLAN.INI file is the variable that is used to set a value that reflects the maximum open files at a time. This is a tunable parameter that is assigned a default value when the server is installed.

Table 16. Maxopens, details.

Default	Minimum	Maximum
256	1	65536

5.7.4 Maximum search value

The parameter in the server section of the IBMLAN.INI that represents the maximum directory/file searches the server can do simultaneously is maxsearches. The searches are executed when a user does a wildcard search of a directory, for example, dir C:\myfile.*

In the previous implementation of OS/2 Warp Server, the maximum number of directory/file searches that the client can request from the Ring 3 server is limited to 1927. The limit for directory/file searches in the Ring 0 Server is 8192.

But, the OS/2 Warp Server for e-business allows the clients to make a directory/file search request up to a maximum of 16384 searches.

If the server files are heavily used the value of the parameter has to be increased.

Table 17. maxsearches, details.

Default	Minimum	Maximum
150	1	16384

The maxsearches parameter is linked to the srvheuristic setting. OS/2 Warp Server for e-business also allocates memory for the maxsearches in a slightly different way.

The previous implementation of OS/2 Warp Server allocated memory for server search structures up front during server initialization based on the value assigned to the MAXSEARCHES parameter in the IBMLAN.INI. For a default MAXSEARCHES=700 value, the server allocates 23.2 K, and, for the maximum value, MAXSEARCHES=1927, the server allocates 64K of non-heap memory. If the srvheuristic 7 (in the IBMLAN.INI) is set to a default value of 1, when MAXSEARCHES= number of searches used up, the memory is allocated dynamically up to a maximum of 1927. If srvheuristic is set to 0 and when MAXSEARCHES= is used up, no more memory is allocated for the server search structures and, thus, no more searches are allowed.

OS/2 Warp Server for e-business allocates memory for the server structures dynamically. Instead of allocating and maintaining a fixed number of server structures, now they are allocated from the server heap on a need basis. It maintains a pointer table which points to the server structures. Whenever a server structure is allocated from the server heap, the pointer table is updated with the new pointer to the server search structure. Instead of having

a memory segment where a fixed number of server search structures were allocated at server initialization time, the OS/2 Warp Server for e-business has a table of pointers of fixed size allocated during server initialization, and the entries point to server search structures allocated dynamically from the server heap. For a default value, MAXSEARCHES=700, the server allocates 2.7K of memory for the pointer table and 1008 bytes for the initial allocation of server structures. For a maximum value, MAXSEARCHES=16384, the server would allocate 64K for the pointer table and 576K from the server heap for the server search structures. For a minimum value, MAXSEARCHES=1, the memory allocated would be 4 bytes for the pointer table and 36 bytes for one server search structure

If the srvheuristic bit 7 in the IBMLAN.INI is set to 1 (its default value), when the MAXSEARCHES= is used up, the pointer table is grown to allow more searches up to a maximum of 16384. When the srvheuristic bit 7 is set to 0 and the MAXSEARCHES= is used up, no more searches are allowed because the pointer table cannot be expanded to include more pointers to server search structures.

5.7.4.1 New parameter: KEEPDOSSEARCH

In the previous implementation of the Warp Server, the timeout for the inactive DOS searches is 10 minutes, and this value is not configurable. Because of this high timeout value, the server may end up allocating more searches than are necessary. In order for the inactive searches to timeout earlier, OS/2 Warp Server for e-business has a new configurable IBMLAN.INI parameter: KEEPDOSSEARCH. The new parameter, KEEPDOSSEARCH=, would allow the timeout value to be specified for the inactive DOS searches. The default value is 600 seconds(10 minutes).

In some server environments with DOS clients, it might be better to configure a lower value that will allow the timeout of inactive searches sooner and allow the server to reuse them rather than having to allocate additional server structures from the heap. Lowering the values will also help in closing some of the outstanding OS/2 search handles. The same benefit of allowing inactive core searches to be timed out sooner, if necessary, will apply to the Ring 0 Server.

Note

The srvheuristics parameter sets a variety of server fine-tuning options. Each digit of the srvheuristics parameter has an independent meaning. Exceptions noted, each digit of the parameter is a binary digit: 0 means off, 1 means on.

Digit position 012345678901234567890

Default value 11110141111311001331 1

For the meaning of each digit, refer to the OS/2 LAN Administrators book. Bit 7 is for MAXSEARCHES

5.7.5 Maximum shares

This is the IBMLAN.INI configurable parameter that specifies the maximum number of resources the server can share with the network. If n resources on the server are being shared, this parameter value should at least be n. For example, if one user is using five resources on the server, the value of this parameter must be at least 5; if the same resource is shared by 5 users, then, the value can be 1.

Note

The number of shared resources displayed by the `NET CONFIG SRV` command will be different than the number specified with the maxshares parameter. This is because the number of shared resources displayed by the NET CONFIG SRV command also includes default system shares (ibmlan$, admin$, and so on), and one share for each partition on the server (a$, b$, and so on).

The OS/2 Warp Server limit on the number of shares that can be made available is a maximum of 1000 shares. In certain server environments with 1000 users and with each user having a home directory that uses a share, all 1000 shares will be used up.

OS/2 Warp Server for e-business has expanded the limit from 1000 to 1500.

MAXSHARES. The server can now share a maximum of 1500 resources.

Table 18. maxshares, details.

Default Value	Minimum Value	Maximum Value
192	2	1500

The MAXSHARES value can also be set from the command line.

Type `NET START SERVER /MAXSHARES:xxx`

where `xxx` is a value less than or equal to 1500.

5.7.6 Summary

Table 19 on page 172 summarizes the new capacity enhancements that have been introduced with OS/2 Warp Server for e-business.

Table 19. Summary of new capacity enhancements.

Parameter	Default Value	Minimum Value	Maximum Value
Maxconnections	300	1	16384
Maxopens	256	1	65536
Maxsearches	150	1	16384
Keepdossearch	600	0	
Maxshares	192	2	1500

5.8 Multiple server names

The multiple server names feature has been developed mainly for high-availability solutions. A current IBM Business Partner solution, Vinca Co-Standby Server solution, provides safeguards for companies that cannot tolerate down time. Vinca uses a server-mirroring technique which defines clustered resources on two nearly identical systems creating high availability for both servers. In this configuration, one of the servers is designated to be the Standby Server and the other the Primary Server. When the primary server goes down, the Standby Server changes its role to be the Primary Server until the Primary Server becomes available. Data is automatically mirrored between the machines via a high-speed dedicated link.

An improvement on the current model is the bi-directional failover. This failover solution is also known as an *Active-Active* solution. In an Active-Active configuration, both servers are fully functional. When a server in this configuration fails, a standby server takes on the additional responsibility of handling the failed server's functions as well as handling its own functions. In this case, both servers are fully functional

Two fully functional servers stand in for each other when the server fails, that is, if server A fails, the standby server B takes over the functions of A and handles the functions of B.

To implement an Active-Active solution in an OS/2 Warp Server domain, the server needs to be able to respond to multiple server names, that is, respond to its own name and also the failed server name.

OS/2 Warp Server for e-business now provides the support for multiple server names. Each server is able to respond to multiple NETBIOS names and shares associated with those names. The support for multiple server names has been included to enable the implementation of Active-Active solutions. The implementation of an Active-Active solution involves:

- Failure detection
- Access to hardware drives of the failing system
- Configuration of a server to handle multiple names
- Migration of the shares to the new server
- Management of access controls on the new server

OS/2 Warp Server for e-business does not support any of the above-mentioned functionality. It is the responsibility of products that provide failover support to provide this functionality. OS/2 Warp Server for e-business implements only that part of the functionality where the server responds to multiple NETBIOS names when it has been configured to do so.

The multiple NETBIOS names function can also be used for other purposes, such as manual failover support, merger, and migration of servers which is discussed later in this chapter.

5.8.1 Configuration

There are multiple ways in which a server can be configured to have multiple names.They can be added either during server installation time or dynamically when the server is running.

The additional server names can be added during server initialization time by:

1. Setting the othsrvnames= parameter in the [server] section of IBMLAN.INI

   othsvrnames = srv1, srv2, srv3

   The othsrvnames parameter can take up to a maximum of seven server names.

2. Using the NET START SERVER /OTHSRVNAMES option.

   NET START SERVER /OTHRSRVNAMES:SRV1,SRV2,SRV3

The othsrvnames parameter can take up to a maximum of seven server names.

The names can be added/deleted dynamically when the server is running by:

1. Using the `NET CONFIG SERVER` command with the option /OTHSRVNAMES

   NET CONFIG SERVER */OTHSRVNAMES:SRV1,SRV2,SRV3*

   The othsrvnames parameter can take up to a maximum of seven server names.

2. Using the new API NetServerNameAdd ( ) / NetServerNameDel ( )

   Refer to the *API REFERENCE MANUAL* for the syntax and usage of these APIs.

The validity of OTHSRVNAMES parameter is based on the following:

- A maximum of seven secondary server names are allowed.
- There can be no name repeated in the list.
- Each name must be a valid computername.
- No name in the list can be the same as the primary computername.

5.8.2 Scenarios

Let us discuss the different scenarios under which multiple server names can be employed for different purposes

5.8.2.1 Simple duplicate server fails

Add the shares from the failed server as share names on the duplicate server using either the API or the `Net Share` command

Add the failed server's name to the remaining server. If either the *NetServerNameAddAPI* or *Net Config Server /OTHSRVNAMES:* command is used to add the server name, the name will be added and the server will not have to be restarted. If the name is only added to the OTHSRVNAMES parameter in IBMLAN.INI, the name will not be automatically added; the user will have to stop and restart the server in order for the name to be added.

5.8.2.2 Simple duplicate server comes on line

Remove the failed server's share name from the remaining server using either the API or the `Net Share` command.

Remove the failed server's shares from the remaining server. If either the `NetServerNameDel` or `Net Config Server /OTHSRVNAMES:` command is used to remove the server name, the name will be removed and the server will not

have to be restarted. If the name is only removed from the OTHSRVNAMES parameter in IBMLAN.INI, the name will not be automatically removed; the user will have to stop and restart the server in order for the name to be removed.

Start the server on the recovered system.

5.8.2.3 Migration of servers to a single system

A customer might find that high-powered hardware is now within their budget and might want to combine the functionality of several existing low-powered servers onto one high-powered machine. Merging the domain with all of its server names onto one server is a significant reconfiguration task and may disrupt the users as the namespace is modified. The customer could, instead, use the multiple server names feature to have one machine handle the requests for several previously-existing servers. This would require less configuration of the domain name space and would have a lesser impact on the customers.

5.8.2.4 Failure of a Domain Controller

When a Domain Controller (DC) or a Backup Domain Controller (BDC) fails, the remaining server cannot assume the functions provided by NETLOGON on the failed server, namely, logon validation. The remaining server can only assume the functions of File and Print sharing of the failing server name. It cannot perform the logon validation functions for the server name that it assumes. The remaining server can provide logon validation on its own server name if it is configured to do so.

Normal LAN Server domain controller failure functions will still be available, that is, BDCs will cover for failing DC. For example, server A is the DC and server B is the BDC. When server A fails, the network name A is added to server B. Server B now takes over the File and Print services of server A. But, it will not respond to any of the logon validation requests that come in for A. Calls for logon validation will appear as though DC failed. Server B, which is the BDC, will perform its function as a BDC and handles the logon validation. Adding the DC name A to the BDC server B does not have any impact on the functions of the BDC. Server B can have its role changed over as DC. This is part of the normal LAN server administration.

5.8.2.5 Failure of the REPLICATOR exporter

A server that is running the replicator service and providing files to clients is known as an exporter. When an exporter fails, the remaining server can provide the exporter function for the failed server, but only under certain conditions:

1. The remaining server must have access to the failed server's drives. Without access to the drives, the remaining server cannot export the files.
2. The remaining server must not already be an exporter. Exporters can only export one tree. If the remaining server is already an exporter, it cannot take on the responsibility of exporting an additional tree.

The user has to update the IBMLAN.INI file to put in the correct information for the replicator service and then start the replicator service.

5.8.2.6 Failure of a NETRUN server

A remaining server can cover for a failed server that was providing the NETRUN service provided that the runpath= parameter in its IBMLAN.INI can be set to be the same as or equivalent to the runpath= parameter of the failed server. The directories in the runpath= must either be copies of the directories in the runpath= on the failed server, or they can be directories on drives shared between the servers.

The user has to update the IBMLAN.INI to put in the correct runpath for the netrun service and, then, start the netrun service.

5.8.2.7 Failure of a Time Source server

The remaining server can cover for a failed server that was providing the time source service. There are no configuration issues. The remaining server merely has to have the Time Source service running.

Note

While a system is handling multiple server names, it will appear on the network as two servers that are sharing the same resource. This may have unforeseen consequences for other system management software. There are no requirements for any IBM system management software to detect or handle this situation in a special manner.

5.8.2.8 Server roles

When the server is configured for multiple names, the secondary names do not have any roles. If one of the servers is a Primary Domain Controller (PDC) and the other server is a Backup Domain Controller (BDC) and if the PDC fails, the BDC takes over but will not handle the role of the PDC. The role of the BDC will still be as BACKUP.

5.8.2.9 Miscellaneous

- Some services, such as the requestor service (wksta.exe) and logical server service (lsserver.exe), have been made *multiple server name*

aware. Other services, such as NETLOGON.EXE, DCDBREPL, ALERTER, REMOTEBOOT, and MESSENGER are not multiple server name aware and see only the primary name.

- If two servers are exporting directories and files via the replicator service, when combined, only one tree will be able to be exported under the primary name.
- Under a light load, a server that takes on additional names will perform as well as independent servers with separate names. It is expected that the system that responds for multiple server names will not perform as well under a heavy load as multiple independent servers would. Simply having multiple names does not affect performance.
- If a server is stopped and started, any server name that was previously added dynamically through the API or NET START SERVER /OTHSRVNAMES or NET CONFIG SERVER /OTHSRVNAMES will be lost. Either the names need to be in the OTHSRVNAMES= in the IBMLAN.INI or re-added through a command line option. If a name is dynamically deleted before a server is stopped and restarted, the name can be re-added through the NET CONFIG SERVER /OTHSRVNAMES or NET START SERVER /OTHSRVNAMES, or NetServerNameAdd() API if names are not kept in the OTHSRVNAMES= parameter in the IBMLAN.INI. If the name dynamically deleted has been originally added through the OTHSRVNAMES= in the IBMLAN.INI and the intention was to delete it because the name was no longer required, the name should be removed from the OTHSRVNAMES= in the IBMLAN.INI; otherwise, the name will be re-added on the next server start.
- The server initialization module will also delete any server names created if initialization fails so as not to leave any extra names on the card. The APIs NetServerNameAdd() and NetServerNameDel() rollback the changes if the addition or deletion of the names does not succeed.

5.9 Vinca StandbyServer

Vinca Corporation currently has a product, StandbyServer, for OS/2 Warp, which is a premier high-availability server-mirroring system for servers running Warp, Warp Server, Warp Server Advanced, and Workspace on demand. With StandbyServer, you can protect your production server by mirroring data in real-time between the primary server and a second, up-to-date, standby machine. When the primary server fails, the standby machine automatically assumes the role of the primary server including the failed server's identity, functions, and data set.

IBM and Vinca are working together to enable StandbyServer for OS/2 Warp on an OS/2 Warp Server for e-business platform. For further details, contact Vinca corporation at the following Web site:

`www.vinca.com.`

5.10 New NET USE switch

The `NET USE` command lists the network resources in use and connects or disconnects your workstation to or from shared resources on the network.

Normally, a NET USE is in effect until you log off or until you stop it. To make the connection permanent until reboot, you can now use the /PERM switch.

To connect to resources on a server on another domain where you are defined with a different user ID, use both /PERM and /USER options.

For example, if a user FRED wants to connect to a resource called projects and make it a permanent connection, he would perform the following steps:

1. Logon to DOM1 with the command:

   ```
   logon FRED /p:ab12cd /d:DOM1
   ```

2. If he wanted this to be a permanent connection, he would use the command

   ```
   NET USE x: projects /perm
   ```

   Note the following:

   - If you issue a net use to check the status of your connections, the permanent connections will not display. In Fred's case above, if he issued a net use, he would get the message:

     There are no entries in the list.

   - The connection can be deleted by issuing the command:

     ```
     net use x: /d
     ```

   - If the connection is not deleted, it will only disconnect when the workstation is rebooted.

A password can be used on the `NET USE` command to override the password that was registered at logon time. This feature is useful for connecting to remote peer workstations or servers where the password is different than the password of the local peer workstation or LAN Server domain. Refer to Table

20 on page 179 to determine the correct NET USE password syntax when accessing external resources.

Table 20. External Resources on an OS/2 Warp Server domain

Password on Logon Domain	Password on External Domain	Passwords are the same	Password Syntax on first net use
Yes	Yes	Yes	(Blank)
Yes	Yes	No	External Password
Yes	No	N/A	Double Quotes ("")
No	Yes	N/A	External Password
No	No	N/A	(Blank)

In the above table, the cases where you would not provide a password are trivial. You would follow the standard NET USE syntax.

In cases where you need to supply a password, the syntax is the same. For example, if a user FRED is defined in domain DOM1 with password - ab12cd and would like to use a resource tool in domain DOM2 where he is defined with password - fg34hi, he would do the following:

1. Logon to DOM1 with the command

   ```
   logon FRED /p:ab12cd /d:DOM1
   ```

2. Connect to the resource on the external domain by providing his password:

   NET USE x: tools fg34hi /d:DOM2

3. If he wanted this to be a permanent connection, he would use the command

   ```
   NET USE x: tools fg34hi /d:DOM2 /perm
   ```

If the NET SHARE password is the same as the password for the user ID, no password is required on the `NET USE` command.

Chapter 6. Integrating Windows NT Servers

With the growing popularity of Windows NT, there have been a number of attempts to install and integrate Microsoft NT Servers into existing OS/2 Warp Server environments. In most cases, this has been a difficult road. IBM OS/2 Warp Server for e-business can bring a measure of order to this chaos.

Many organizations have complex computing environments. They may have many different kinds of servers and clients that were introduced over time. Their Information Technology (IT) strategies were based on the most current technologies available at the time. Unfortunately, technologies evolve so rapidly that most strategies were obsolete before they were ever implemented. Most medium to large organizations also have exceptions to their strategic IT policies. This adds yet another level of complexity. If all this sounds familiar or, perhaps, even, describes your organization's current state of computing, you should read this chapter.

This chapter will describe the IBM Networks User Account Manager which helps simplify the task of OS/2 Warp Server and Windows NT Server integration.

Note

OS/2 Warp Server for e-business is also referred to by the terms Warp Server and Warp unless otherwise stated. Microsoft NT Server 4.0 is also referenced by the terms NT and NT Server.

6.1 Overview

OS/2 Warp Server and Windows NT, despite their common ancestry, did not integrate well. Integration efforts were limited to a measure of dual maintenance. Users who required access to resources on both Windows NT and OS/2 Warp Server were required to be defined on both systems with identical passwords. This is a nightmare to manage for both the administator and the user who needed to keep his passwords synchronized.

The IBM Networks User Account Manager, which is shipped with OS/2 Warp Server for e-business, alleviates some of these problems. This product is installed on Windows NT Server 4.0 and its key objective is to allow single point management of user and group accounts for both Warp Server and NT Server 4.0.

With this feature, User and Group accounts created on the Warp Server domain controller are propagated to the NT Server. This feature also allows for the creation of an alias to reference a shared directory or printer on an NT Server. From the end user perspective, resources are accessed seamlessly from either Warp Server or NT Servers in the network.

6.1.1 Key features

This feature allows for integration of two components between the two server types, that is, user management and resource management.

The key functions of the user management feature are that Administrators will be able to create, update, and delete user IDs and passwords as well as to create, update, and delete Groups. These user and group changes will be propagated to the designated NT servers.

Alias support has been extended to include shared resources on NT servers that are defined in the domain. An alias that is created from either the NET command line interface or the Warp Server Administrative GUI will be allowed to reference a shared directory or printer on an NT Server.

NT does not support the sharing of serial devices; an attempt to share a serial device will result in an error return code.

6.1.2 Benefits

You can view this new domain environment from different perspectives. The users will now have a unified resource environment from which applications, files, and printers will be seamlessly available. The Administrators will now have a way to do single-point accounts management within their integrated domain(s).

The user perspective

The user can now easily access resources from all servers including NT Servers. Some of the advantages are:

- Seamless access to Warp Server Domain resources, which may include resources on NT Servers. These resources could be:
  - Home Directories
  - Logon Assignments
  - OS/2 Applications
  - Microsoft Application
  - Win 32 application through CITRIX Based Products

- Warp or NT Server File and Printer Resources

Figure 67 shows the OS/2 Warp Server for e-business Domain with an integrated NT Server.

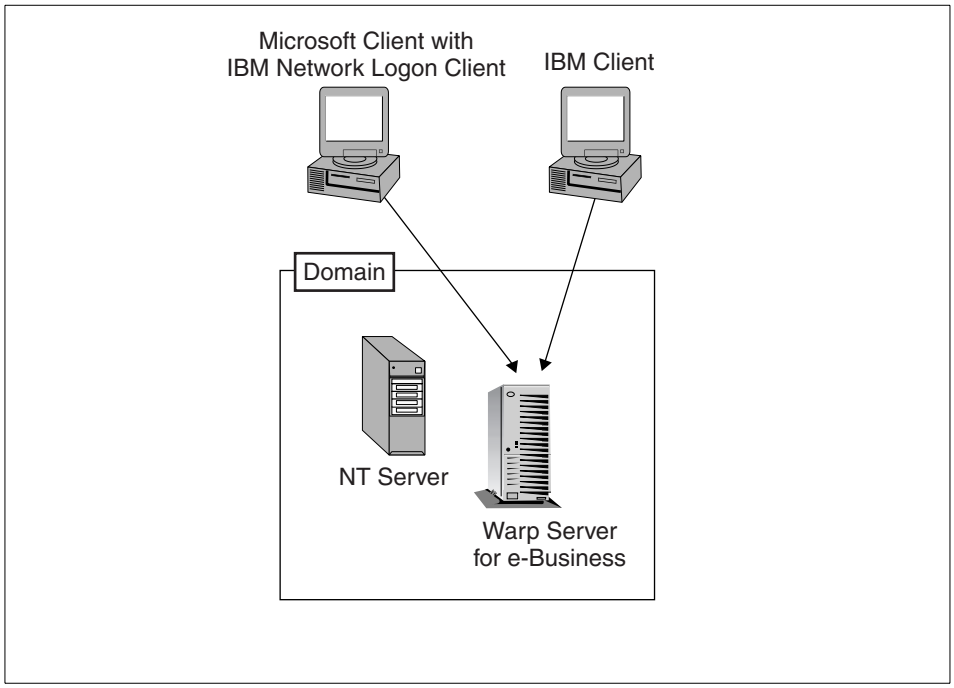

Figure 67. Seamless access to the integrated server domain

The client workstations can be any client supported by OS/2 Warp Server for e-business. The client does single-point logon to access all their resources seamlessly. The resources could be files, printers, or applications residing on Warp Servers or NT Servers.

6.1.2.1 The administrator perspective

The administrator can now easily administer the accounts management database and resources in a mixed server environment. Some advantages are:

- Single point administration of user and group accounts via the Warp Server User Accounts Management
- Administration of Warp Server resources via GUI, command line, and APIs
- Administration of NT resources via NT GUI, command line, and APIs

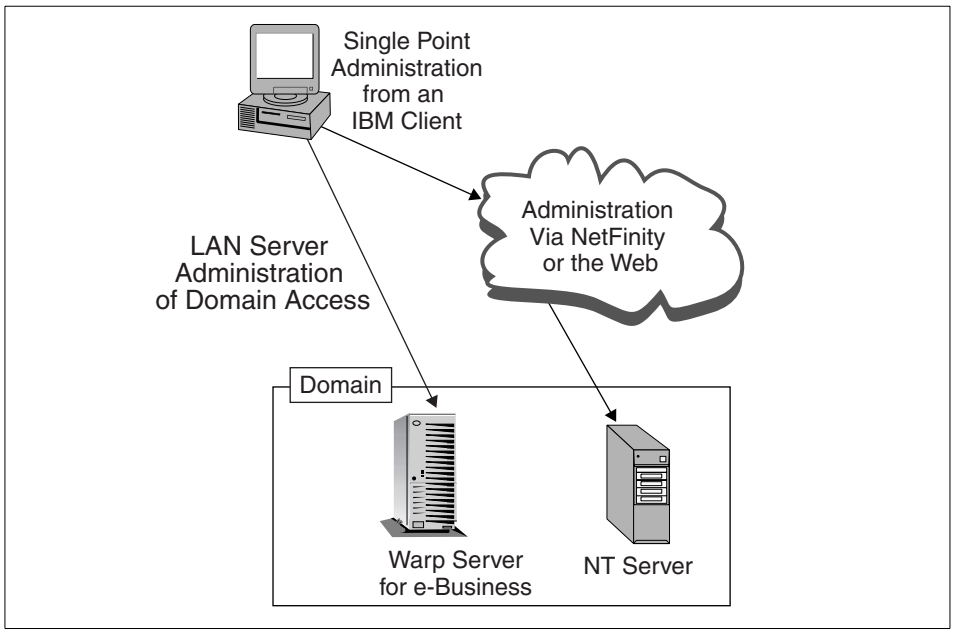

Figure 68. Single point administration

Figure 68 shows the single-point administration model. There are a variety of administration tools available to ensure single-point administration. These tools and utilities are discussed in Appendix A, “More on Windows NT administration” on page 407 of this redbook.

- NT Servers in Warp Server Domains provide:
  - Windows applications for Microsoft clients
  - Access to Win32 applications through CITRIX for Non-Microsoft clients
  - NT file and printer resources for all clients
- Integrated Server Management
  - Domain User and Group administration
  - Domain password administration
  - Domain resource administration including:
    - Alias support for NT Server resources
    - Logon assignments of NT resources

- Application definitions, which may now include Windows-based applications and resources that reside on NT Servers

6.1.3 Limitations

No provisions have been made to support the following NT Server functions in the OS/2 Warp Server for e-business Domains:

- Security Access Control Management of NT Resources
- Start/Stop of NT Services
- Viewing Event Logs
- Remote IPL
- Replication

Note: Most of these can be managed by existing or additional tools and utilities that are discussed in Appendix A.

Functional enhancements to the Warp Server product provide for better client services and single-point administration. What follows explains what these enhancements will provide for the client while giving more effective domain administration of NT Server resources.

6.2 Architecture

The architecture of the IBM User Accounts Manager is described in Figure 69 on page 186. This architecture is for the IBM User Account Management. The support for aliases on Windows NT Server did not require any application service on Windows NT Server.

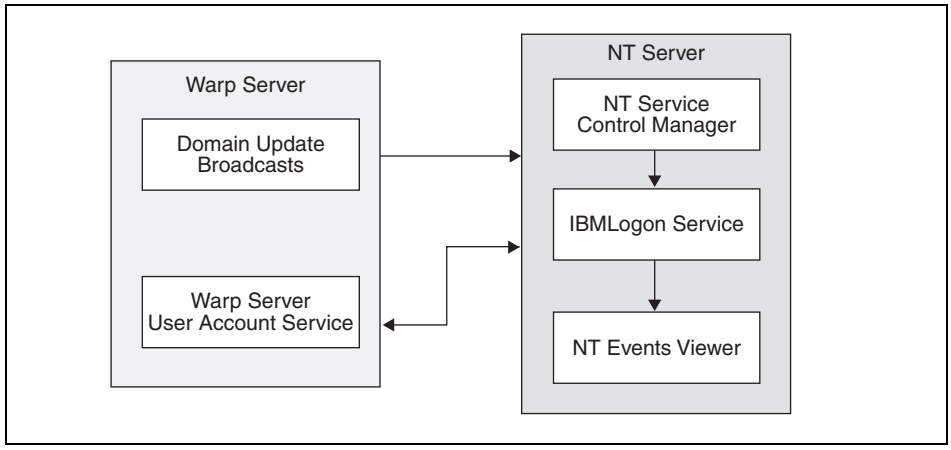

Figure 69. The IBM user accounts manager for Windows NT Server

6.2.1 User accounts

The IBM User Accounts Manager is designed as a Windows NT service that has the following attributes:

- The IBM Logon Service runs as a Windows NT Service and is registered with the Service Control Manager.
- The service Logs is also registered with the NT Event Viewer to which it can log errors.
- It responds to Start and Stop commands sent by the Control Manager.
- The service must run with Administrator Privilege
- The service responds to the OS/2 Warp Server Domain Controller for UPM update queries:
  - Request a Delta UPM update
  - Request a Full Synchronization
- The service is able to Create, Delete, and Update Windows NT user and group accounts
- The service also has debug tracing capabilities to a trace log file.

6.2.2 Windows NT Server alias support

The alias creation function did not require any additional code; however, existing code had to be modified to cater to the Windows NT Server limitations.

Code that allows the creation of an alias to a Windows NT server from the OS/2 Warp Server command line has been added to NET.EXE. Code that disables unsupported administrative functions via the Warp Server administrative GUI has been added to the NETGUI.EXE.

The NETGUI does the following:

- It checks whether the server is a Windows NT Server and sets a flag if it is.
- It then disables or modifies the menu items if it is a Windows NT Server
  - No Serial Devices
  - No Statistics
- Windows NT Servers Open with the setting view by default. Details and Tree views have been removed.
- The GUI only allows *Share at Server Start-up* mode for alias creation

6.3 Concepts

This topic will discuss the concepts around Windows NT Server and OS/2 Warp Server Management architectures. The intention is to provide the reader with some background information with regard to the latter.

Windows NT Server and OS/2 Warp Server share a lot in common with regard to the basic file and print architecture. However, there are some significant differences in how environments are set as well as how security is administered.

6.3.1 Microsoft domain models

Windows NT and OS/2 both share the concept of a domain. The implementation has some significant differences. As with Warp Server, with Windows NT Server, a domain is a logical grouping of network servers and other computers that share common security and user account information. Within this grouping, administrators can create one user ID for each user. Users then log on to the domain and are able to seamlessly access any resources within the administrative unit.

What makes this work in Windows NT Domains is the directory database or what is sometimes referred to as the Security Accounts Management (SAM) database. The Warp Server counterpart to the SAM is a combination of the NET.ACC and the Domain Controller Data Base.

Theoretically, one domain on Windows NT Server can accommodate up to 26,000 users and approximately 250 groups. OS/2 Warp Server can handle

up to 128 servers and 254 groups per domain. Neither of these limitations may be practical depending on your environment. OS/2 Warp Server only supports the single domain model.

Microsoft has architected, within NT, a number of domain models. The four domain models are:

- Single Domain
- Master Domain
- Multiple Master Domain
- Complete Trust Domain

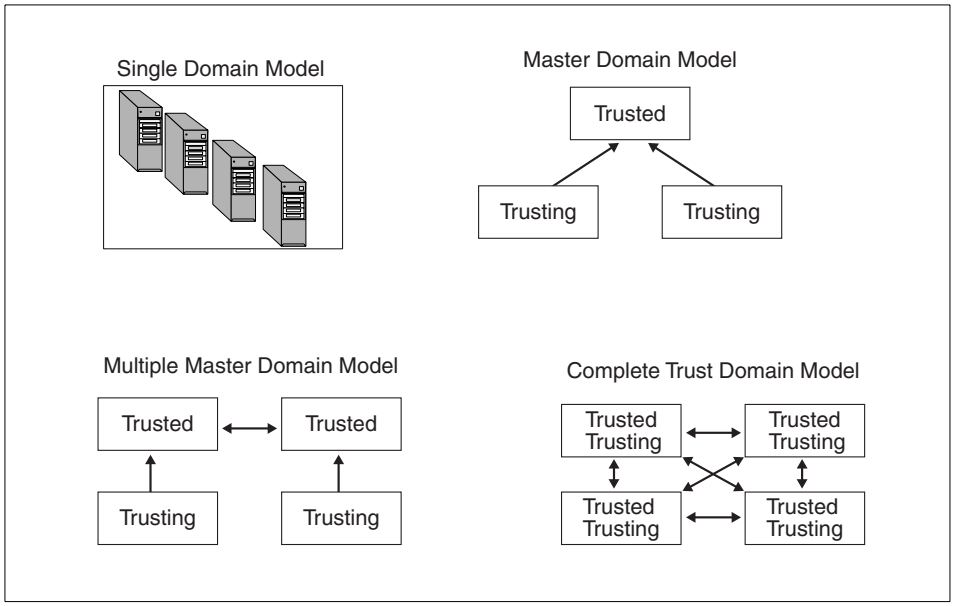

Figure 70. Microsoft NT domain models

It is not the intention of this redbook to discuss these models in any detail. Each of these models is intended to support a different size and scope of the client server environment. Each of these models also has its own strengths and weaknesses. Each succeeding model in the list above is more complex than its preceding model. The Complete Trust domain is a complex structure and is very difficult to manage. However, these models were designed by Microsoft to support domains that number from only a few clients up to thousands of clients.

Note

The redbook *OS/2 Warp Server, Windows NT, and NetWare: A Network Operating System Study*, SG24-4786, contains details about Windows NT Domain Models.

Our discussion will be limited to NT Domains in general and not to any specific domain model. When you integrate NT Servers, you must share information or synchronize the information in these databases.

6.3.2 Microsoft workgroup model

Microsoft defines a workgroup as an organizational unit of computers that do not belong to a domain. In a workgroup, each computer tracks its own user and group account information and, in contrast to domain controllers, does not share this information with other workgroup computers.

Users of computers defined within a Workgroup can log on only if they have been defined on the computer that they are logging on to. The Workgroup does not share user or group information with other members of the same workgroup.

Computers running Windows NT Workstation, Windows NT Server, Windows for Workgroups, or Windows 9x can be configured to participate in either a domain or a workgroup. When setting up one of these computers for networking, specify a computer name and a workgroup name. If the workgroup name matches a domain name, the computer name appears in the browse list for that domain and can browse computers running Windows NT Server and Windows NT Workstation whether participating in a domain or a workgroup. To determine whether the computer participates in a domain or a workgroup, during setup, specify that the computer logs on to either a Windows NT Server domain or a workgroup.

6.3.3 User authentication

Warp Server Domains have their accounts database in the NET.ACC file. A portion of this file contains the accounts management database. This database must be replicated to all Servers in an OS/2 Warp Server Domain. In addition to this, the Domain Controller Database, which is a collection of files, must be replicated to all Backup Domain Controllers in a Warp Server Domain. This is shown in the diagram below.

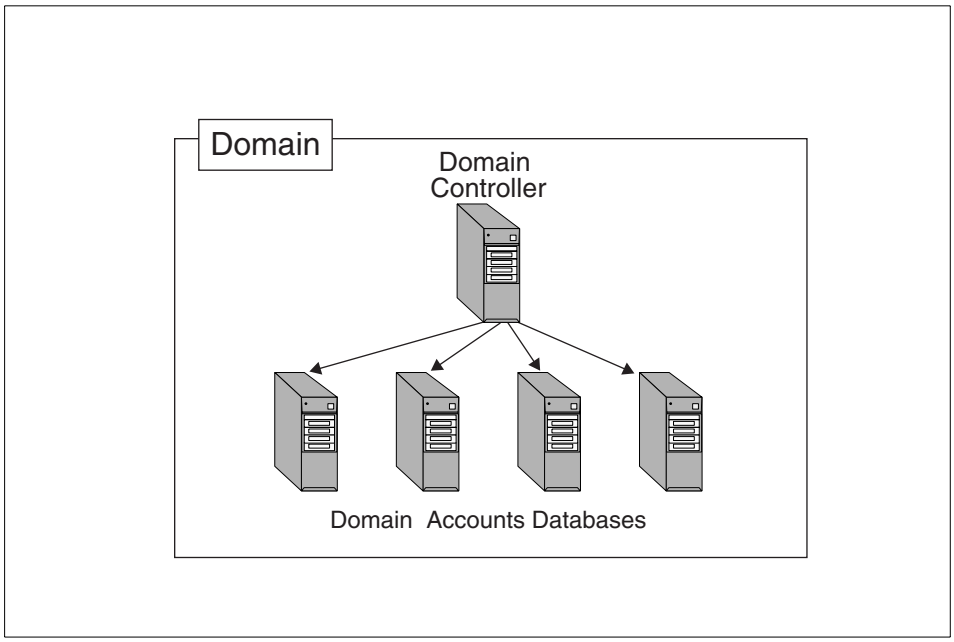

Figure 71. Warp Server accounts database synchronization

Note

As most Warp Server Administrators know, the NET.ACC files on the Additional Servers of a domain, are *NOT* identical. They contain unique Access Control Profiles (ACPs) in addition to the accounts database.

Additional Servers in a Warp Server Domain must have a copy of the domain's accounts management database. This database is a portion of the master NET.ACC file, which resides on the Primary Domain Controller in the Warp Server Domain.

With Windows NT, Member Servers do not need to hold the entire accounts database of the NT Domain. NT Servers that join an NT Domain set up a secure channel to the accounts database located in the SAM of the Domain Controllers. The following figure illustrates this.

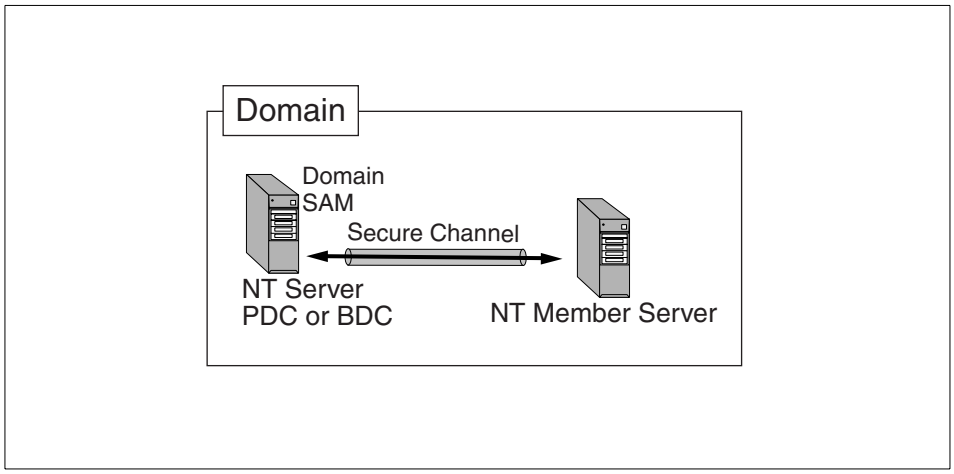

Figure 72. NT Member server's access to the domain SAM

The SAM (accounts database) is shared among all Domain Controllers that belong to the same domain. The accounts database is accessed by the Member Servers to authenticate domain users and groups that need access to their resources. This is like a pass-through authentication through the implicit trust that exists with the domain controller. True pass-through authentication occurs in multi-domain NT environments via explicit trusts, which is beyond the scope of this discussion.

Because of the way that NT Domains handle authentication through secure channels, there is less overhead in keeping the accounts database on all servers of the domain. This would lead us to believe that NT Domains could grow much larger than Warp Server Domains. Factors other than authentication methods usually limit domain sizes. A bigger model than the single domain model is needed to handle thousands of clients.

6.3.4 The IBM Networks User Account Manager

In the diagram above, you can think of any one or all of the Additional Servers as NT Servers (with IBM Networks User Account Manager). They will get a copy of the accounts database from the Warp Server Domain Controller. The NT Server does not have a NET.ACC file; so, it simply integrates the account information into its SAM. Therefore, the NT Server is *NOT* a Member Server of an NT Domain, but, rather, it is an NT Server that happens to have its accounts synchronized with Warp Server Domain accounts. This makes it an *Additional Server* in the Warp Server Domain.

The Windows NT Server should be configured to participate in a workgroup. During configuration, you are required to specify a computer name and a workgroup name. To use the IBM Networks User Account Manager, you should put the domain name of the OS/2 Warp Server for e-business domain as the workgroup name. This enables the IBM Networks User Account Manager to receive updates from the OS/2 Warp Server domain and propagate these through to the Windows NT Server.

The installation process is described in more detail in section 6.5, “Installation of IBM Networks User Accounts Manager for NT” on page 201.

6.4 Access control

What follows is a description of the manner in which resources are secured on Windows NT Server and OS/2 Warp Server. Although the accounts database is automatically sychronized using the User Accounts Manager, access control must be separately administered.

In order to properly integrate and administer file resource servers, you must have a good understanding of how the security access works in both kinds of servers. The access security discussed here is primarily for NTFS, HPFS386, and JFS volumes used within the network environment.

Note

In the following discussion, HPFS386 and JFS should be considered as having the same access control characteristics. Where appropriate, the differences are mentioned.

6.4.1 An overview of access control on NT

On Windows NT Servers, file resources have two levels of access control. They are called Share Permissions and Permissions on directories and files. Volumes that are formatted NTFS may make use of both levels of access control. Volumes that are formatted for FAT can only use Share permissions.

6.4.2 Share permissions

Share permissions determine who can use shared directories over the network and in what manner. To set share permissions for either NTFS or FAT, use the Sharing tab in the directory property sheet on the shared directory. When you share a directory, you can grant each group and user one of four types of permissions for the share and all of its subdirectories and

files: Full Control, Change, Read, or No Access. Each of these are explained in the table below.

Table 21. Sharing a directory

Attribute	Description
Full Control	User can read and change files, add new ones, change permissions for the directory and its files, and take ownership of the directory and its files.
Change	User can read and add files and change the contents of current files.
Read	User can read the contents of files in this directory and run applications in the directory.
No Access	User cannot access the directory in any way even if the user is a member of a group that has been granted access to the directory.

For FAT volumes, share permissions provide the only way to limit access to network files. You can specify one set of share permissions on a shared directory that applies to users for all files and subdirectories of the shared directory. So a user that has Full control of a parent directory will have full control of all the child directories unless another share is defined at the child level limiting the security.

To secure shared directories keep the following in mind:

- The default permissions set on a newly created share are Full Control for Everyone.
- Permissions set through a shared directory are effective only when the directory is reached over the network.
- Permissions set through a shared directory apply to all files and subdirectories in the shared directory.
- Permissions set through a shared directory in an NTFS volume operate in addition to NTFS permissions set on the directory itself. (This is described later).

6.4.3 Permissions on files and directories

On NTFS volumes, you can set permissions on directories and files. These permissions apply to users accessing the files at the server as well as users accessing the shared directory over the network. Windows NT Server offers a set of standard permissions for NTFS directories and files. The standard permissions are combinations of specific types of access, which are called individual permissions. The individual permissions and their abbreviations are Read (R), Write (W), Execute (X), Delete (D), Change Permissions (P), and Take Ownership (O). Each of these is explained in the table below.

Table 22. File and directory permissions on Windows NT Server

Permission	Meaning
Directory	
No Access (none) (none)	User cannot access the directory in any way, even if the user is a member of a group that has been granted access to the directory.
List (RX) (Not Specified)	User can list only the files and subdirectories in this directory and change to a subdirectory of this directory. User cannot access new files created in this directory.
Read (RX) (RX)	User can read the contents of files in this directory and run applications in the directory.
Add (WX) (Not Specified)	User can add files to the directory but cannot view the contents of the directory.
Add & Read (RWX) (RX)	User can add files to the directory and read current files but cannot change files.
Change (RWXD) (RWXD)	User can read and add files and change the contents of current files.
Full Control (All) (All)	User can read and change files, add new ones, change permissions for the directory and its files, and take ownership of the directory and its files.
File	
No Access (none)(none)	User cannot access the file in any way even if the user is a member of a group that has been granted access to the file.
Read (RX)	User can read the contents of the file and run it if it is an application.
Change (RWXD)	User can read, modify, and delete the file.

Permission	Meaning
Full Control (All)	User can read, modify, delete, set permissions for, and take ownership of the file.

In the first column of the first table (for directory permissions), the first set of parentheses following the standard permission indicates the individual permissions for the directory itself. The second set of parentheses indicates the individual permissions that apply for new files subsequently created in the directory.

When you set a standard permission, the abbreviations for the individual permissions appear beside the standard permission. For example, when you set the standard permission Read on a file, the abbreviation RX appears beside it.

In addition to setting standard permissions, you can set special access permissions. Special access permissions allow you to define a custom set of individual permissions for directories and files.

To work with NTFS security effectively:

- Users can use a directory or file only if they have been granted permission to do so or if they belong to a group that has permission to do so.
- Permissions are cumulative, but the No Access permission overrides all others. For example, if the coworkers group has Change permission for a file, and the finance group has only Read permission, and John is a member of both groups, John will be granted Change permission. However, if the finance groups permission for the file is changed to No Access, John will be unable to use the file despite his membership in the coworkers group.
- When you create files and subdirectories in a directory, they inherit permissions from the directory. For example, if you add a file to a directory that allows the coworkers group Change permission and the finance group Read permission, those same permissions apply to the file.
- The user who creates a file or directory is the owner of that file or directory. The owner can always control access to the file or directory by changing the permissions set on it. Users who are members of the Administrators group can always take ownership of a file or directory.
- File permissions always override directory permissions.
- The easiest way to administer security is by setting permissions for groups rather than individual users. Typically, a user needs access to

many files. If the user is a member of a group that has access to the files, you can end the users access by removing the user from the group rather than changing the permissions on each of the files. Setting permissions for an individual user does not override the access granted to the user through groups to which the user belongs.

- Every file and directory on an NTFS volume has an owner. The owner controls how permissions are set on the file or directory and can grant permissions to others.
- When a file or directory is created, the person creating the file or directory automatically becomes its owner. It is expected that administrators will create most files on network servers, such as when they install applications on the server. Therefore, most files on a server will be owned by administrators except for data files created by users and files in users' home directories.

Ownership can be transferred in the following two ways:

- The current owner can grant the Take Ownership permission to other users allowing those users to take ownership at any time.
- An administrator can take ownership of any file on the computer. For example, if an employee leaves the company suddenly, the administrator can take control of the employee files.

Note

Although an administrator can take ownership, the administrator cannot transfer ownership to others. This restriction keeps the administrator accountable.

Keeping this in mind, let us point out some differences in how NT Servers permit IDs (users or groups) to access resources. NT Servers allows users to access resources based on their membership in groups and, sometimes, based on the user's ID. This works the same in Warp Servers, except that permissions granted to individual user IDs override group permissions. This is not so on NT Servers. This is one big difference in the way the access control mechanism works on NT Servers versus Warp Servers.

Note

In this context(not the card game), the permissions associated with a user ID will be used in place of any other permissions that may happen to exist for that user ID. Including any permissions that may have been granted to groups of which the user is a member.

Further explanation is in order. If you permit individual users and the groups to which they belong access to an NT file resource, the access permissions accumulate. If, for instance, you give a user called user1 read access to a resource called data1. You then give a group called group1 (of which user1 is a member) write access to the same resource (data1); then, user1 would have read and write access to data1. On a Warp Server, if the same permissions were applied for user1, user1 would only have read access to data1. In short, the user permissions override the group permissions.

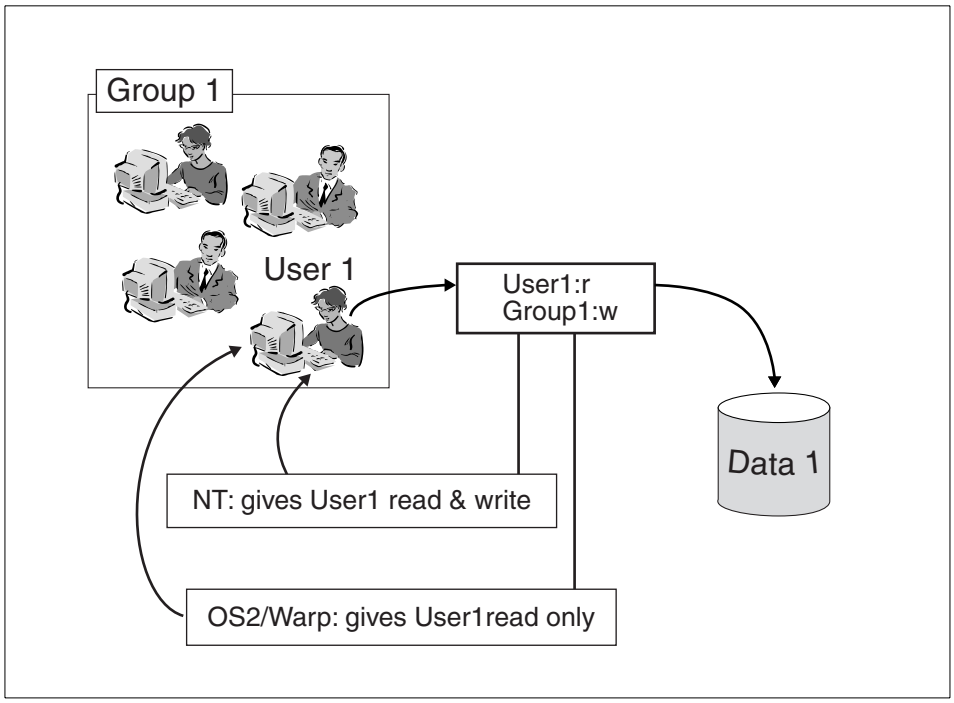

Figure 73. NTFS access control example

Perhaps this is why, in NT, you hardly ever grant access permissions based on individual user IDs. It really does not buy you much. With OS/2 Warp Server, you can grant access based on an individual user ID, thus, tailoring a user's access. It should also be mentioned that resource management should be done through the use of local groups in single domain environments. Global groups are used primarily to allow domain-wide access to resources on a server. They are also used in trust relationships to allow users defined in trusted domains access to resources in trusting domains. Since we are talking here about stand-alone NT servers being integrated into Warp Server Domains, the groups used for resource management should be *local* groups.

Another big difference is that administrators do *not* always have full access permission to all resources on the NT Server. However, an administrator can take ownership of a resource and can, therefore, change the permissions of the resource. This resource lockout for administrators does not exist on Warp Servers.

Note

Administrators do not have unrestricted access. As an NT Administrator, if you have ever tried to access or move a users home directory, you will understand this.

One last comment about file access on NT Servers is that *no access* always overrides all other access permissions. Whether the access permission is applied at the share level or the NTFS security level. No access means no access.

The following are some strategies for using NTFS file permissions:

- Grant permissions to groups not individual users.
- Create local groups and assign permissions to them rather than assigning permissions directly to global groups.
- When you create and share a file or directory on a server, grant Full Control to the Administrator's local group. This ensures that all administrators of that domain can change permissions for and, otherwise, administer the file or directory in the future.

In general, all of the above is true with respect to file resources on your servers. However, there are different file systems in NT and Warp. NTFS, HPFS386, and JFS are the file systems that should be used on your file servers. Both NT and Warp also support FAT partitions, but you cannot apply and maintain access control permissions to the same degree on FAT volumes; so, this redbook assumes NTFS, HPFS386, or JFS volumes. All of these file systems support Access Control Lists (ACLs) as an integral part of their file systems as opposed to external linkages to their ACLs. A little more information about ACLs and their use on NT Servers is in order.

6.4.3.1 Access Control Profile

An Access Control Profile (ACP) is created for each network resource and contains permissions governing the level of access each user and group will have to that resource. The ACPs contain access permissions assigned to a resource. The ACPs of HPFS386 resources are stored in the file system and the ACPs of non-HPFS resources are stored in NET.ACC.

The following permissions can be used while creating an ACP.

- Attributes (A) - Attributes of the files in the corresponding directory can be changed. This option is valid for drives, directories, and files.
- Create (C) - Files and subdirectories can be created. This option is valid for drives, directories, files, printers, and serial devices.
- Delete (D) - Files and subdirectories in the corresponding directory can be deleted. This option is valid for drives, subdirectories, and files.
- Permissions - Permission for the specified resource can be changed. This option is valid for drives, directories, files, printers, and serial devices.
- Read (R) - Files in the corresponding directory can be read but not changed.This option is valid for drives, directories, files, and serial devices.
- Write (W) - Files in the corresponding directory can be written but not read or executed. This option is valid for drives, directories, files, and serial devices.
- Execute (X) - Programs contained in the corresponding directory can be run. This option is valid for drives, directories, and files.

6.4.4 ACLs, SIDs, and security related issues

Those who have worked with either IBM or Microsoft Servers know that permissions are controlled through Access Control Lists (ACLs). These are simply lists of accounts that have access to a resource. An ACL is composed of entries called Access Control Entries or ACEs. Both NT and Warp Servers use ACLs, however, the ACLs for NT resources allow access via Security Identifiers or SIDs. So, what is an SID?

If you have been working with NT, you have probably discovered that NT Servers use SIDs as the keys for their resources. Every ID that is created in an NT Server gets an SID assigned to it. These are the keys for unlocking the resources on the server and the access to the resource is based on the SID *not* the account name.

Note

Instead of using the account name directly, HPFS386 uses a number, which is formed from a hashed value of the account name and a serial number, however, JFS uses the account name.

This brings us to some potential problems. Assuming all account management is done on the Warp Server Domain, what happens if you

re-create an account (deleting the account, then creating it again with the same name)? If you delete an account in the Warp Server Domain, the NT Server will reflect the deletion when the accounts database synchronizes with the master database. That will eliminate the account on the NT Server and all of the ACEs that may exist for the account. ACEs are the list entries in the ACL. Now, you re-create an account with the same name. The new account gets created on the NT Server upon synchronization. All the ACEs that were associated with the original account no longer exist. If this was a group account that was controlling major resources on that server, you would probably have a number of unhappy users.

The ACEs are always deleted (or obsoleted) when you delete an account. When you create a new account with the same ID, you must create new ACL entries that re-associate the account ID with the resource. The new ACEs will contain new SIDs for the account and all should be well again. Even if an ACE does not get deleted when an account is deleted, the old SID will not allow the new account to access the resource.

If you think about the way the IBM Networks User Accounts Manager works, it synchronizes the accounts database on the NT servers with the master database on the Warp Server Domain. This happens on a periodic basis by the netlogon service. You may be painfully aware of account synchronization problems that have always plagued domain environments. This dates back to the first domain implementations; so, it only stands to reason that synchronization will, probably, be a problem here as well. The key is to be prepared

Note

If you have saved the ACLs for NTFS file resources, access control can be restored in the event of account corruptions, which may require account re-creation and ACL replacement.

You should have some method of logging and tracking your accounts and their associated resource requirements. You will need tools that allow you to restore accounts and tools that allow you to restore ACLs throughout your Domain(s). Do these tools exist? The answer is yes; there are, probably, too many of them.

On the Warp Server side, the standard ACL utilities that come with the product are BACKACC and RESTACC. These Warp Server utilities (commands) are explained in the Warp Server online command reference documentation; so, we will not discuss them any further here. In addition to

BACKACC and RESTACC, there is a set of tools called LAN Server Maintenance Tools (LSMT), which allow you to save and restore accounts and permissions. LSMT is Employee Written Software (EWS) available from IBM Support sites on the Internet.

For NT Servers, you also have a number of tools/commands (for example, CACLS and PERMS) that will assist you with ACLs. However, these standard NT Server tools are, sometimes, not adequate. You need to go beyond the standard set of tools that come with NT and its resource kit. The standard tools and utilities plus a number of additional tools are referenced in Appendix A. Whichever set of tools you choose, at a minimum, they should include a tool that will back up and restore the ACLs on your NT Servers. One such tool is Super CACLS available from the Internet at the following Web site:

`http://www.trustedsystems.com`

6.4.5 Summary of domain security management

This topic covered file security on domain servers. With an understanding of file security as implemented on both Warp Server and NT Servers, you can better decide which tools can help you administer these resources. You can also plan for backup and recovery of your security keys, which are the ACLs of your server's file resources.

6.5 Installation of IBM Networks User Accounts Manager for NT

The following is a walk through of the installation. For a successful installation, you will need to perform the following tasks:

1. Go through the prerequisite checklist.
2. Install the network service on the NT Server.
3. Add the NT Server to the Warp Server Domain.
4. Check the server to make sure it is functioning correctly within the domain.

The next few sections will discuss the above steps in greater detail.

6.5.1 Prerequisites tasks

Before beginning the installation of this service, the targeted NT Server must have the following prerequisites:

- Be at Version 4.0 (No Service Pak requirements; however, SP4 was installed on our NT Servers.)

This can be checked with the VER command as follows:

- Click on the **Start** button and select the **Command Prompt**.
- Type VER and then press **Enter**. The product version should display as Windows NT Version 4.0.

- It must be configured as an additional or stand-alone server.

  This can be checked with the NET ACCOUNTS command as follows:

  - Type NET ACCOUNTS and then press **Enter**. The Computer role parameter should be configured as SERVER. Exit the Command Prompt.

- It must have a functioning network adapter.

  This is best checked with the Event Viewer by clicking on **Start -> Programs -> Administrative Tools (Common) -> Event Viewer**. Any errors associated with the network adapter would be logged here.

- It must have a supported protocol.

  This can be checked through the Network Properties.

  - Click the **Start** button again, point to Settings, then click **Control Panel**.
  - Double-click **Network**.
  - On the Services and Protocols pages, you can view the components that are currently installed. Make sure you have a supported protocol NetBEUI and/or TCP/IP protocol. You should note that the same protocol needs to installed on the OS/2 Warp Server Domain Controller that is going to communicate with this server, that is, NETBIOS or NETBIOS over TCP/IP (TCPBEUI).
  - As an additional test, if you are using TCPBEUI, ping the OS/2 Warp Server machine from the Windows NT Server.

6.5.2 Installation on the NT server

The following is an example of a typical install and configuration of the User Accounts Manager on an NT Stand-alone Server.

1. Log on to the Windows NT system with a user ID that has administrator privileges.
2. Click the **Start** button, point to Settings, and then click **Control Panel**. The Control Panel window will be displayed.
3. Double-click **Network**. The Network notebook will be displayed.
4. Click on **Services**

5. Click on **Add**. The Select Network Service window will be displayed as follows:

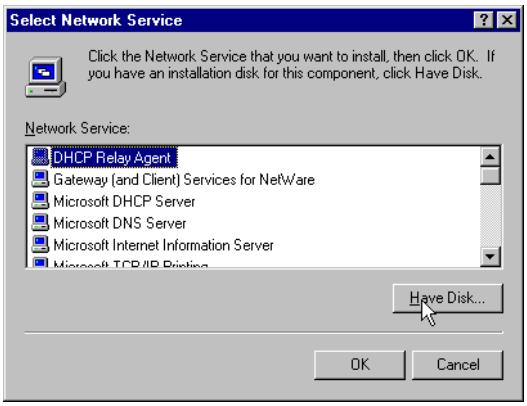

Figure 74. Windows NT installation, select network services

6. Click on **Have Disk**. The Insert Disk window will be displayed.

Note

To make the diskette, simply copy the files from the \CID\WINDOWS\MANAGENT subdirectory of the CDROM onto a formatted diskette.

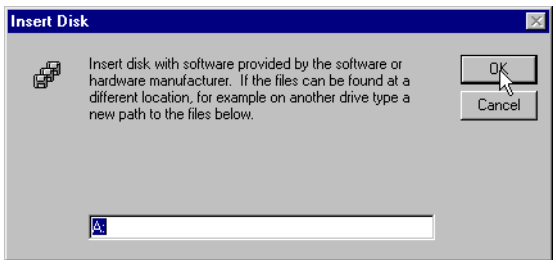

Figure 75. Windows NT installation, insert disk

7. If you have copied the User Accounts Manager installation files to a diskette, simply insert the diskette and click on **OK**. You can also type in the directory path to the installation files or to the UNC network path, which contains these files. Once this is done, the files will be copied, and you will see the following:

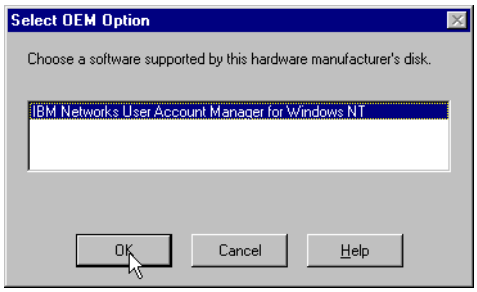

Figure 76. Windows NT installation: Select OEM option

8. Clicking on **OK** will bring you to the User Account Manager Properties notebook as follows:

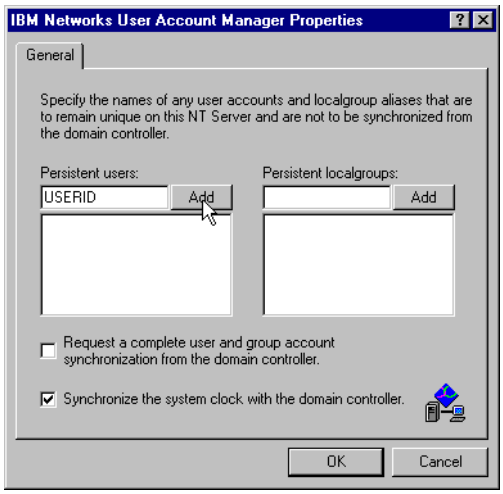

Figure 77. IBM Networks user account manager properties

9. Add the names of persistent user accounts and persistent localgroups by typing them into the appropriate fields and clicking **Add** for each entry.

Persistent Accounts?

Persistent users and localgroups are accounts that will be retained in the Security Access Management database of the NT Server and will not be updated or deleted by the Warp Server's Netlogon service. Therefore, if you want to keep, for instance, an NT administrator account on this server, you should make that account a persistent user.

10. After you finish entering the persistent accounts and groups, you will receive the following setup messages:

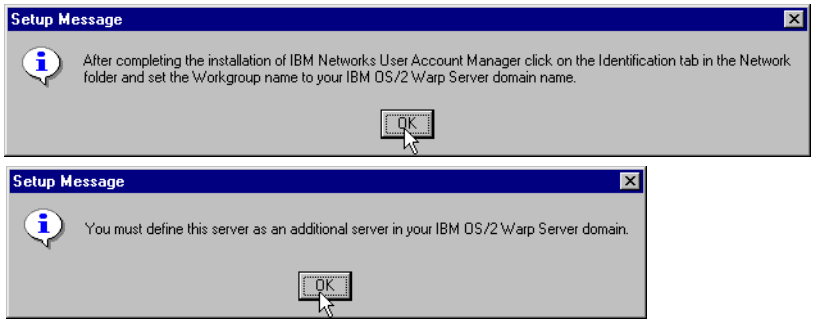

Figure 78. Windows NT installation, setup messages

11. As indicated, you must complete the server identification on the NT Server and also in the Warp Server Domain, which we will do after we complete the installation on this server. Simply select **OK** twice.

12. The next step is to identify the NT server as a server participating in a *workgroup*, the name of which is the domain name of the Warp Server Domain. Click on **Change** in the identification window.

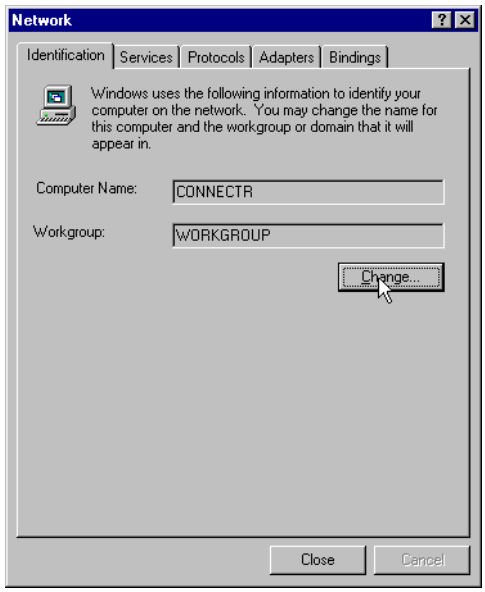

Figure 79. Network identification

13.Complete the workgroup identification as follows:

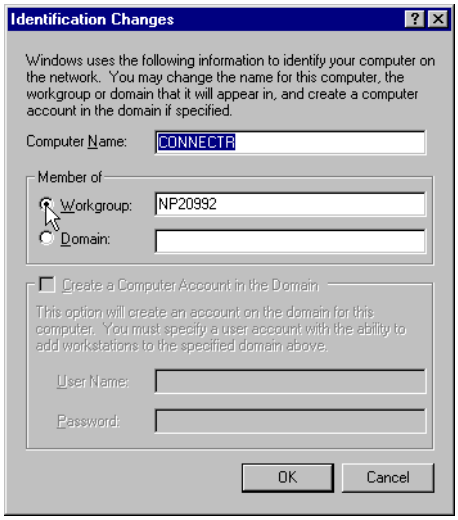

Figure 80. Identification changes

14.After completing this identification, click on **OK** and back out until you are prompted to restart your computer. Doing so completes the required steps on your NT Server. Next, you must complete the server definition on the Warp Server domain.

6.5.3 Defining the NT server on the Warp Server domain

The following procedure will complete the integration of the NT Server by defining it on the Warp Server domain as an additional server.

There are two methods that can be used to add the Additional Server to the Warp Server Domain. They are:

- LAN Server Administration Graphical User Interface(GUI)

  or

- Command Line Interface (CLI) Administration

Note

If you are familiar with NET commands, you should refer to Part 6.5.3.2, “Alternate method” on page 209

Both methods require you to logon to the Warp Server Domain as an Administrator. The GUI method follows.

6.5.3.1 The GUI method

From a workstation, log on to the Warp Server Domain with a user ID that is part of the Admins group in that domain. The workstation itself must be capable of doing administration in the domain. This could be a Warp V 4.0 client, Windows or a DOS Client with IBM LAN Client capability, or, perhaps, a Warp Server workstation. In any case, invoke the LAN Server Administration and you should see a window similar to the following:

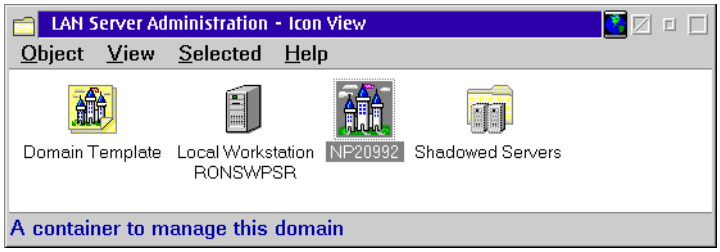

Figure 81. LAN server administration

Note

If you do not see your domain, there could be a number of reasons:

- Your user ID does not have administrative rights.
- You may have local logon status only.

  From a command prompt, enter the command logoff /L to check logon status.
- You are not logged onto the desired domain.
- You may need to create another domain from a domain template, and you must also have administrator rights in that domain.

The quickest and easiest solution, if you do not see the desired domain, is to log off and log on in the desired domain. The domain should then appear.

15. Now, you should select the Domain, and, provided you have administrator access, you will be able to select an object called Defined Servers as shown in the next figure:

Figure 82. Domain view

16. You can now use a template from the Defined Servers Window to create a server definition for the NT Server you are adding to this domain. Remember, this is a template, and new objects are created from it. That is, drag and drop.

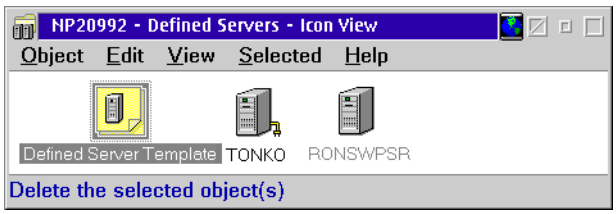

Figure 83. Defined servers

17. The final step is to fill in the ComputerName of the NT Server. This is actually the NetBios name of the NT Server workstation. Complete the entries as illustrated below.

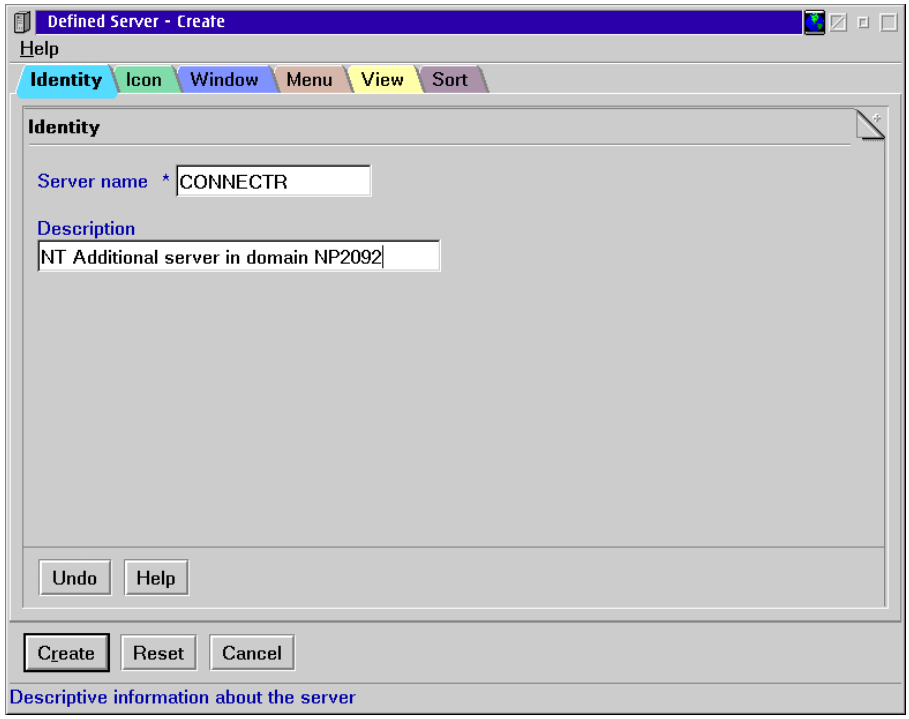

Figure 84. Defined server - create

18. You will now see your NT Server (in our case CONNECTR) in the Defined Servers Group of the Warp Server Domain.

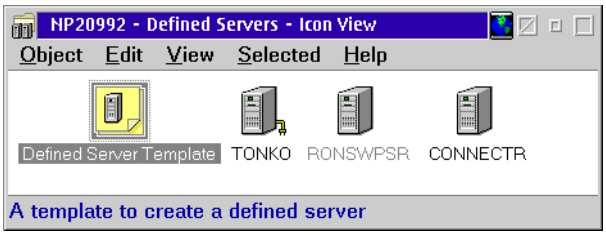

Figure 85. Defined servers

6.5.3.2 Alternate method

There is an alternate way to add Additional Servers to Warp Server Domains. It is much faster and may be necessary if your workstation does not support the LAN Server Administration GUI. For example, you may be sitting at an NT workstation that has the IBM Primary Logon Client network services installed.

You cannot run NETGUI (the OS/2 graphical admin interface) on an NT workstation. In this case, you would need to do something like the following.

Here is an example of accomplishing the same results as above from within a command prompt window on an NT Workstation. It is assumed that you have logged onto the Warp Server Domain using an Administrator ID.

```
C:\>NET ADMIN \\RONSWPSR /C
Type EXIT or press Ctrl+Z to exit.

[\\RONSWPSR] NET USER CONNECTR /ADD /PASSWORDREQ:NO
The command completed successfully.
[\\RONSWPSR] NET GROUP SERVERS CONNECTR /ADD
The command completed successfully.
[\\RONSWPSR] EXIT

C:\>
```

In our example, the server RONSWPSR is the domain controller of the Warp Server Domain. The Domain Controller contains the master accounts database (NET.ACC file).

6.5.4 Post installation

Now that you have completed the installation of IBM Networks User Accounts Manager for NT, you will observe some changes to your NT Server's accounts. The NT Server's accounts will have been synchronized with the Warp Server Domain's accounts. The result will be a combination of both the NT and Warp Server's accounts. The user accounts that remain on the NT Server after synchronization are:

- All Warp Server Domain user/machine accounts
- Administrator
- Guest
- IBMLogon
- All persistent user accounts

The group accounts that remain are also a combination of the Warp Server Domain and the NT Server accounts. They will include:

- All Warp Server Domain group accounts including the default accounts of GROUPID, LOCAL, and SERVERS
- Administrators

- Backup Operators
- Guests
- Power Users
- Replicator
- Users
- All persistent localgroup accounts

6.5.5 Problem determination

6.5.5.1 The Event Viewer

The NT event log is always a good place to start. Access the Event Viewer by clicking on **Start -> Programs -> Administrative Tools -> Event Viewer**.

> **Note**
>
> The IBM Networks User Accounts Manager Service is referred to as the IBMLogon Service from within the Event Viewer.

Figure 86 shows you a sample of the typical entries showing a normal start-up of the IBMLogon service. The event log contains information shown in the following figure:

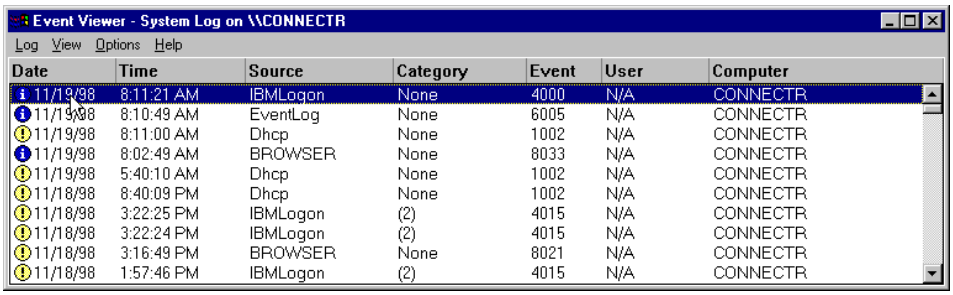

Figure 86. Event Viewer - system log

Details of event log indicate a successful start.

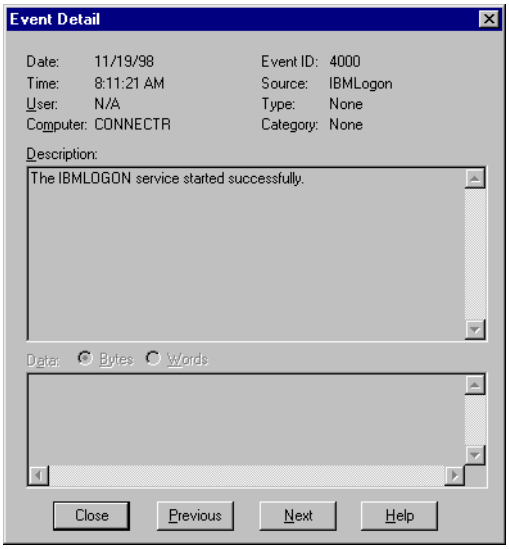

Figure 87. Event Detail - IBMLOGON service started successfully

After the service has started successfully, you will see further entries in the event log that indicate synchronization with the Warp Server Domain. These entries will indicate whether the account is being deleted or left unchanged. Samples of these Event Details are shown in Figure 88.

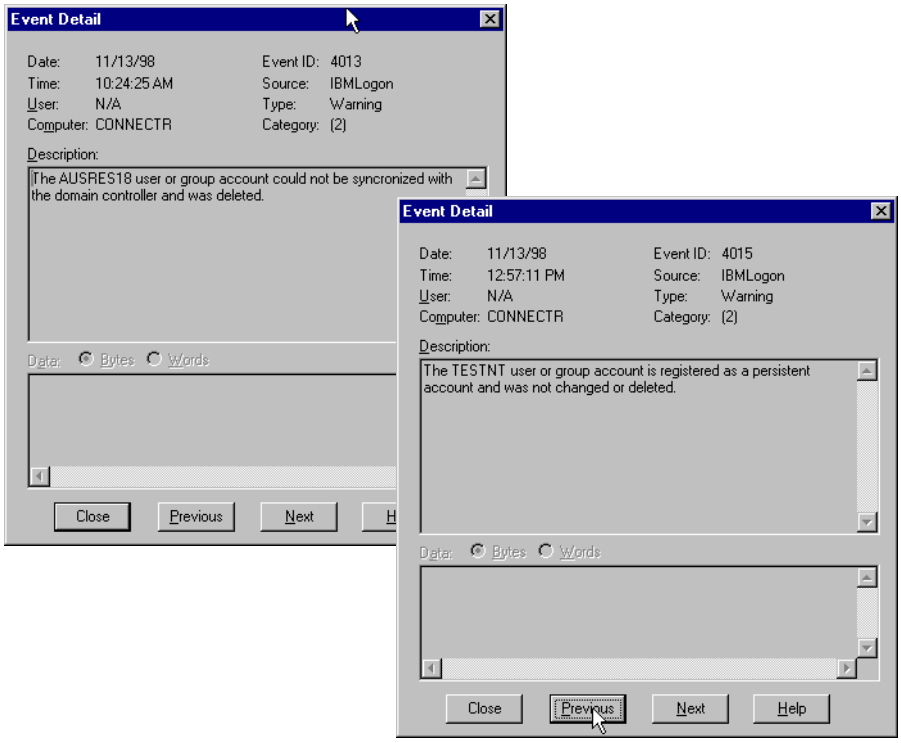

Figure 88. Event Detail - account synchronization

Again, the event log should be the first place you look if you suspect problems. If the IBMLogon Service did not start successfully, you will see entries to indicate this.

After you have installed the IBM Networks User Accounts Manager for NT service on the NT Server and added that server to the domain, synchronization should occur. If this does not happen, it will most likely be the result of a failure to authenticate the server in the domain. The event log events shown in Figure 89 would be indicative of this.

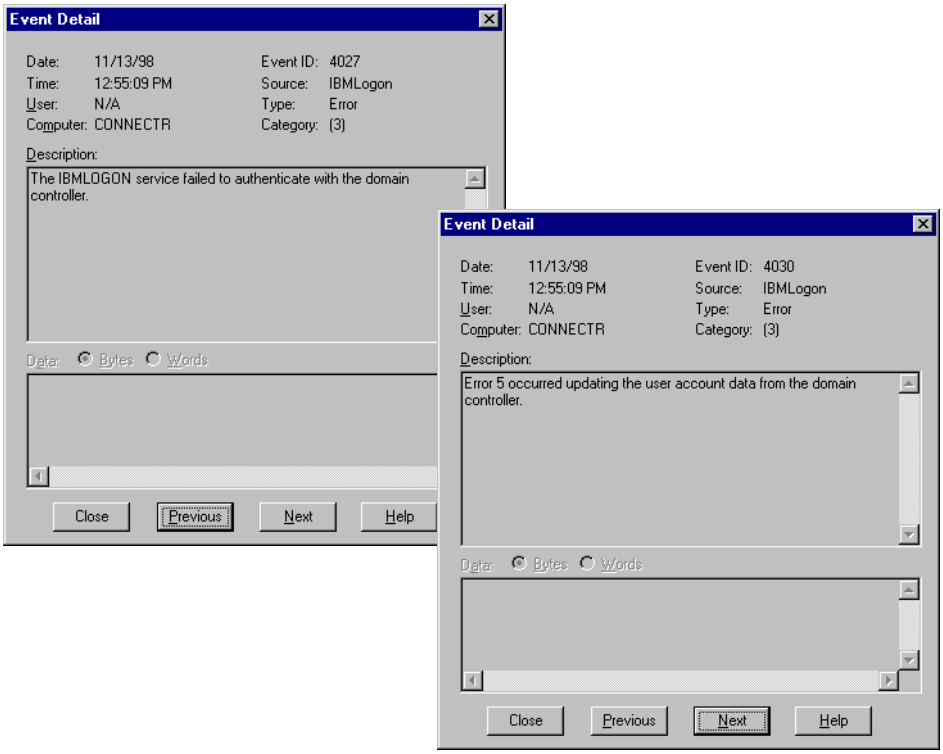

Figure 89. Event Detail - synchronization failure

If you need to analyze the account synchronization further, the Network Monitor utility may be helpful.

6.5.5.2 Network monitor

The name of this utility understates its functionality. This utility, in addition to monitoring network performance, will also trace, capture, and analyze network traffic.

If your NT Server will not join or synchronize with the Warp Server Domain, this utility could help you. A simple version of Network Monitor comes with the NT Server V 4.0. A more comprehensive version comes with the NT Systems Management Server. The simple version is only capable of capturing network traffic to or from itself (the NT server on which you installed network monitor); However, this could be helpful in determining why an NT Server isn't joining the domain.

The Network Monitor utility can capture and analyze protocols. Virtually, any protocol that NT handles can be interpreted. For example, LLC, NetBios, IP,

and SMP to mention just a few. The first thing that an NT Server does when it joins the Warp Server Domain is to synchronize its database with the Warp Server Domain's accounts database. If this is not happening, you can use this tool to analyze the SMB frames. If you have any experience with sniffers, you will find this a somewhat familiar and useful tool.

Some examples of using this tool are contained in the redbook *OS/2 Warp Server, Windows NT, and NetWare: A Network Operating System Study*, SG24-4786 (Included on the CD with this redbook).

6.5.6 Summary of installation

The installation is a two-part process. You must first install the IBM Networks User Accounts Manager service on the NT Server; then, you must add the NT Server to the domain. After the installation is completed, the NT Server synchronizes its accounts database with the domain's accounts. The event log will indicate the status of this. If the NT Server has problems joining the domain, the event log should indicate where the problem exists. Further problem determination methods were also discussed. The Network Monitor can also assist in problem determination.

6.6 NT file and print servers

File and Print servers are the backbone servers of most client server environments. These servers are, usually, based on Warp, NT, or NetWare. Both IBM and NT have ways of integrating NetWare Resource Servers by using gateway services and/or migration utilities. This effectively leaves just Warp and NT resource servers.

Resources can be managed from the Warp Server GUI; however, you cannot manage the access control for NT resources without using a method outside of the Warp Server GUI. When you create an alias for an NT resource, it creates the alias and its associated share, but it *does not* create the Access Control Profile for the resource. This means you must use some method of creating the Share Access and/or the NTFS permissions. The methods of managing these Access Permissions are discussed in Appendix A.

The following will discuss the setup and management of a file server and a print server along with the possible problems associated with each.

6.6.1 Defining an alias on an NT file server

In the following example, we will bring the resources of an NT 4.0 file server into the domain. Specifically, we will:

1. Determine the resources to be made available on the NT Server
2. Define an Alias for each NT resource
3. Manage access permissions of the NT Server resources
4. Assign the resource as a network assignment
5. Test the results with various clients

6.6.1.1 Determine which NT Server resources to share

Perhaps you have a set of document templates that you wish to make public; so, let us go through the steps listed above to create and implement a documents database on an NT Server in our domain.

6.6.1.2 Define an alias for each NT resource

We have already decided the data will be document templates, so the next step is to create the Alias for the resource. This is done on the Warp Server Domain from a Warp Client (you could do it from a Warp Server, but it is a somewhat questionable practice because of the amount of resource required to run the administration GUI). Start the **LAN Server Administration GUI -> Domain Object -> Directory Resource Definition -> Create a directory from a Directory Template**. Then, finish the process by filling in the Directory Create notebook. Our example below shows only the final step (the notebook).

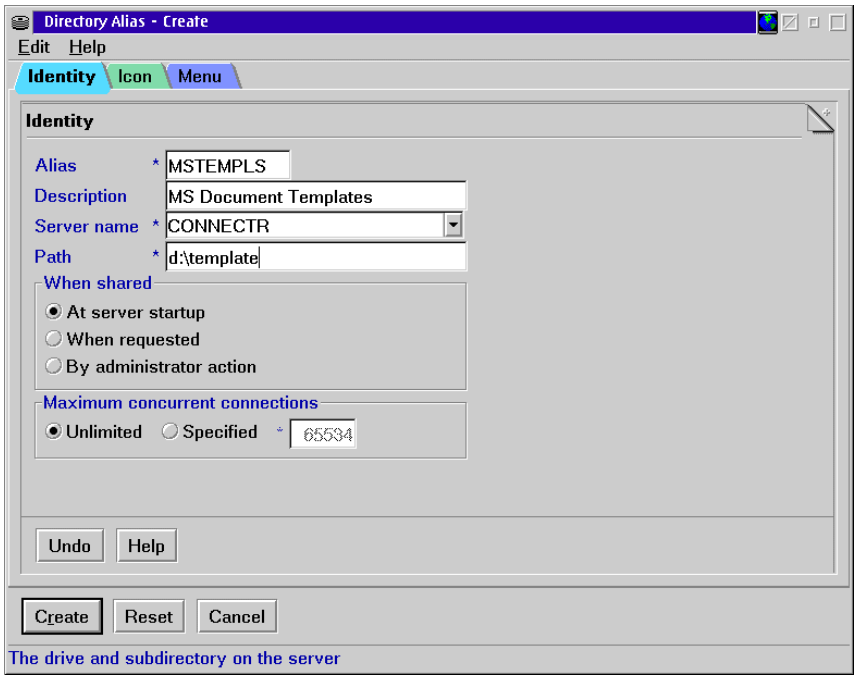

Figure 90. Directory alias for NT resource

The Alias should now exist in the Warp Server Domain, and the share on the NT Server (CONNECTR) should have been created. Next, the access controls for the share must be modified on the NT Server CONNECTR.

6.6.1.3 Manage access permissions of the NT Server resources

If you have the Netfinity Client installed on the NT Server, you can use the Netfinity Manager to access the NT Server and modify the share and security access permissions for the share (MSTEMPLS in our example).

Note

If you do not have Netfinity, you can do this kind of administration with the tools mentioned in Appendix A, such as The Microsoft Web Administration Tool.

What follows is an example of doing this for our MSTEMPLS share.

The Netfinity Manager must be started. The process varies depending on the client you are using (NT, Warp, and so on). Without getting to granular, the following example shows the process of selecting a remote workstation (the

NT Server) on which we will be modifying the access controls for the MSTEMPLS resource.

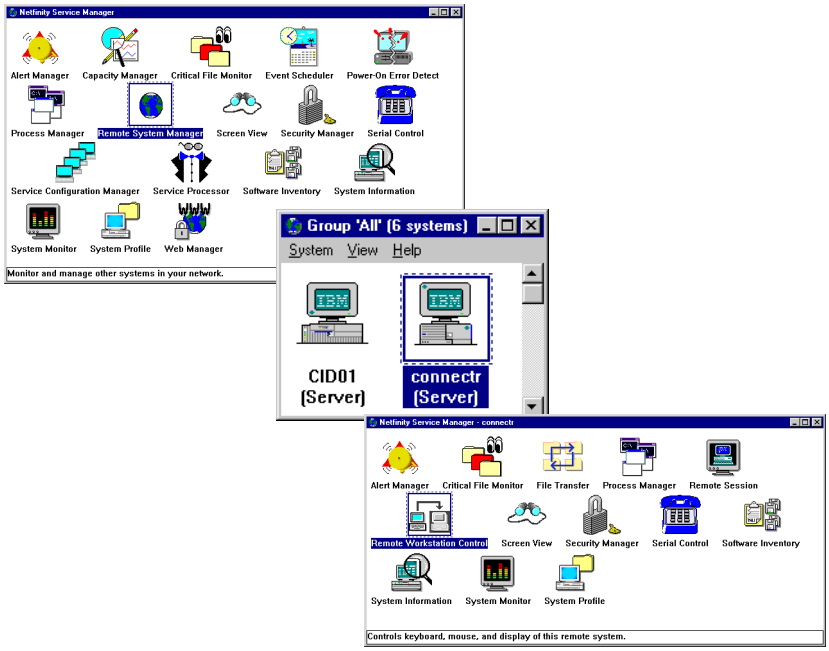

Figure 91. Manage the access of the NT resource

After you have selected **Remote Workstation Control**, you can follow normal NT Administration techniques to change the access permissions of the resource. In our example, we set the access permissions of the MSTEMPLSs resource to allow the MSUSERS group read access.

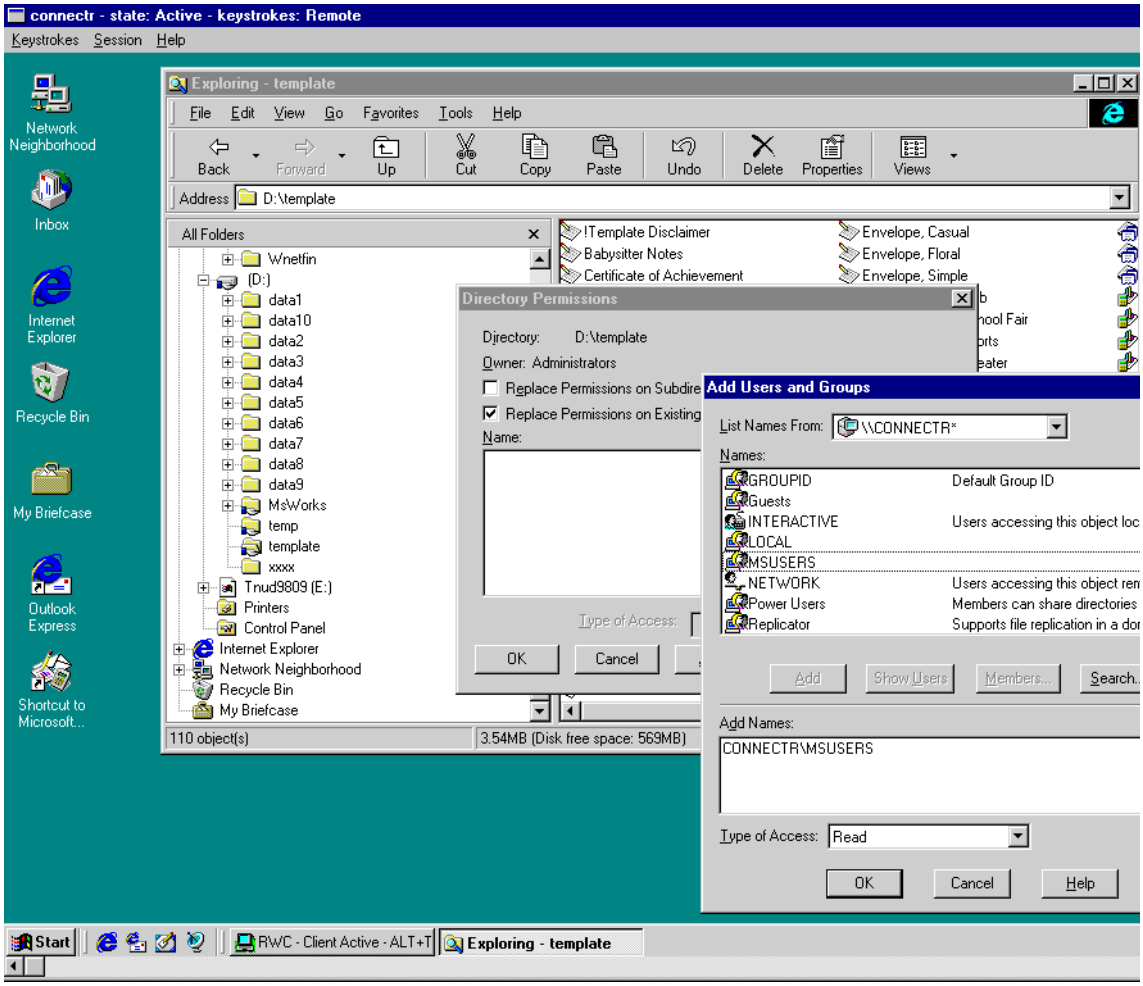

Figure 92. Setting the access permissions

With this completed, the Netfinity Service Manager can be closed, and the resource is ready for use by the clients.

Note

Do not expect stellar performance while controlling the Netfinity Client. The amount of graphics and network bandwidth dictate the performance.

6.6.1.4 Assign the resource as a network assignment

This is normal Warp Server Administration. The resource could be a logon assignment for the group of MSUSERS or it could be used as a network assignment associated with an application like Microsoft Works.

Before you actually use the resource with an application, it would be good practice to try accessing the resource from a client typical of the clients that require the resource. Just doing a NET USE of the resource would be a simple test. If the client can read the files in the directory, the resource can be tested with an application.

6.6.1.5 Test the results with various clients

Testing the results now becomes a matter of setting up an application that utilizes the resource as a network assignment or perhaps a logon assignment for some clients.

6.6.1.6 Example of NT resource assignment and usage

In our example, we will use the resource as a logon assignment, which we will assign through the MSUSERS group. This group contains only Microsoft clients in our domain. We will then logon to some of these clients and verify that we have access to the resource and can use it with Microsoft applications.

The screen below contains our setup by using NET Commands:

- NET ADMIN to the AURORA1 server, which is the Domain Controller
- NET GROUP of the MSUSERS group to check the members
- NET USER to assign the MSTEMPLS resource to T:
- NET USER command to check the user ID: MSUSER1

```
H:\>net admin \\aurora1 /c
Type EXIT or press Ctrl+Z to exit.

[\\aurora1] net group msusers
Group ID     MSUSERS
Comment

Members
-------------------------------------------------------
MSUSER1                  MSUSER10                 MSUSER2
MSUSER3                  MSUSER4                  MSUSER5
MSUSER6                  MSUSER7                  MSUSER8
MSUSER9
The command completed successfully.
[\\aurora1] net user msuser1 /assign t:mstempls
The command completed successfully.
[\\aurora1] net user msuser1
User ID                              MSUSER1
Full Name
Comment
User's comment
Parameters
Country code                    000 (System Default)
Privilege level                 USER
Operator privileges             None
Account active                  Yes
Account expires                 Never

Password last set               12-10-98 04:08pm
Password expires                Never
Password changeable             12-10-98 04:08pm
Password required               Yes
User may change password        Yes

Requesters allowed              All
Maximum disk space              Unlimited
Preferred logon server          Any
Logon script
Home directory
Last logon                      12-10-98 04:13pm

Logon hours allowed             All

Group memberships               *MSUSERS
                                *USERS

Logon assignments for MSUSER1:
   MSTEMPLS    T:
Applications assigned to MSUSER1:
   MSWORKS
The command completed successfully.
[\\aurora1]
```

Note

You can use the NETGUI in place of Net Commands; however, some familiarity with the command line interface is quite helpful for quick checks and small changes.

All that remains is to logon to the domain from a Microsoft Client and see if the network drive is properly assigned and if the access permissions are *read only*.

You may have noticed that MSWORKS had been previously assigned as a public application for this user. You could invoke the MSWORKS application and see if the templates could be used by the application.

6.6.1.7 Review of NT File Resource Integration

As outlined at the beginning of this topic, the five steps used to create and test the NT file resources are:

1. Determine the resources to be made available on the NT Server.
2. Define an Alias for each NT resource.
3. Manage access permissions of the NT Server resources.
4. Assign the resource as a network assignment.
5. Test the results with various clients.

6.6.2 Integrating NT print servers

The integration process is very similar for Windows NT print servers; of course, the real difference is the type of resource. The resource definitions and share permissions for printer resources are similar to those for file resources. The big difference is how the clients will utilize the resource and, for printers, that usually comes down to whether the proper *printer driver* is being used with the printer. In fact, most problems with network printing are *print driver* related.

Printer drivers are not binary-compatible across hardware-processor platforms and OS platforms. The Windows NT print server deals with this by providing the proper print driver for the client or by expecting a certain format of printer data stream from the client.

Microsoft Windows NT and Windows 95/98 Client computers can connect to a network printer served by a Windows NT print server. When these clients initiate a print request, the required printer driver is downloaded from the Windows NT print server if it is not already on the client's hard disk. All other

client computers must *create* the printer. That is, they must install the printer driver directly on their hard disks, specify a port, name the printer, and so on.

Note

Windows 95 clients do not obtain printer drivers on Windows NT version 4.0 print servers in the same way that Windows NT clients use the printer drivers. Windows NT clients download the printer driver from the server if a newer version has been installed on the server. However, Windows 95 clients use a technology called Point and Print to download the printer driver and some printer settings to the client only when the client runs the Windows 95 Add Printer Wizard.

Microsoft clients can also create a remote printer served by a Windows NT print server. Each method has advantages and disadvantages. Connecting to a remote printer is easier and faster than creating one. If the Windows NT Client has connected to a printer, the print job does not spool on the client machine; so, no spool options are available. Windows 95 clients always spool locally and again remotely. The *connected* client also cannot queue print jobs locally. Creating a printer gives the user more control, but that control is not always needed.

Microsoft NT print server interacts very well with Microsoft clients, and it does provide a reasonable amount of support of other client types, such as Netware and TCP/IP LPR. However, it does not support all clients equally well.

Similar driver rules exist for OS/2 clients that use Warp Server printer resources. However, the two worlds of network printers do not mix extremely well because of printer driver dissimilarities between OS/2 and Microsoft clients. The best way to implement Warp Server network printing is though Advanced Print Services, which is a very good enterprise print solution that can use legacy printers as if they were host printers or host printers instead of legacy network printers.

If you are supporting both OS/2 and Microsoft clients, you now have the option of using and managing your network printers within a domain while avoiding some of the classic network printer problems, such as printer driver mismatching. You could elect to continue using NT print servers for your Microsoft clients while using Warp Servers for your OS/2 clients. You could also do mix and match or migrate from one to the other.

Chapter 7. TCP/IP Version 4.21 new functions

Most of the functions discussed in this chapter have already been introduced in TCP/IP 4.1 or 4.2. There have been a number of improvements to the TCP/IP stack - most of them geared toward performance, efficiency, and reliability. Most of these will be discussed during the course of this chapter.

Because this is not a chapter that discusses TCP/IP in OS/2 from the ground up, knowledge is assumed of one of the following products:

- TCP/IP 2.0
- TCP/IP 3.0 as part of OS/2 Warp Connect 3.0
- TCP/IP 3.1 as part of OS/2 Warp Server 4.0 and OS/2 Warp Server 4.0 Advanced
- TCP/IP 3.5 as part of OS/2 Warp Server 4.0 SMP
- TCP/IP 4.0 as part of OS/2 Warp 4.0
- TCP/IP 4.1 as available from IBM's Software Choice Catalog
- TCP/IP 4.2 as part of WorkSpace On-Demand 2.0

For other publications relating to TCP/IP on OS/2, see Appendix D, "Related publications" on page 433.

7.1 TCP/IP 4.21

TCP/IP 4.21 is a 32-bit, SMP-capable, stack implementation. The stack is more accurately SMP-enhanced in that a number of modules specifically exploit SMP hardware. It features an enhanced 32-bit Ring 0 library that allows Ring 0 applications (applications running at the same privilege level as the operating system) to make calls directly into the library for entry into the protocol stack. Performance is further enhanced by the preallocation and reuse of connection resources. These are specifically aimed at Web Server performance. The TCP/IP stack structure is shown in Figure 93 on page 226.

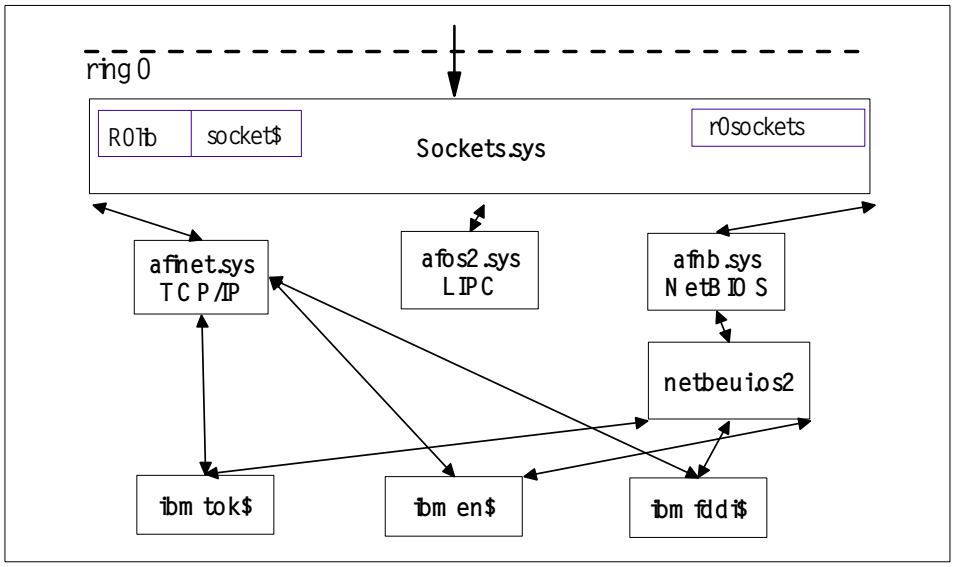

Figure 93. TCP/IP stack structure.

It is 4.4BSD Lite-2 compliant and preserves existing 16-bit and 32-bit compliance. It provides backward binary support. INETCFG has also been enhanced. The DLL interface for the various application types is shown in Figure 94 on page 227.

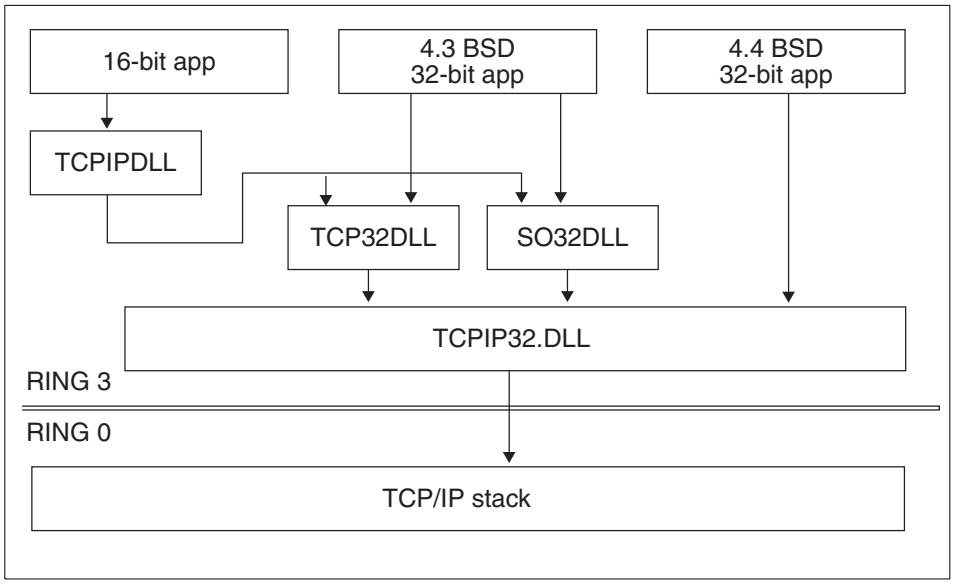

Figure 94. TCP/IP, DLL interface to stack

Security is also included to prevent SynAttack which use SynCookies.

7.1.1 Changes in TCP/IP 4.21

The following table lists the main differences between TCP/IP versions.

Table 23. Differences between TCP/IP versions

Feature	TCP/IP version							
	2.0	3.0	3.1	3.5	4.0	4.1	4.2	4.21
32-bit BSD 4.4 lite based stack						X	X	X
Send_File() API								X
Accept_and_Receive() API								X
SMP Enablement				X		X	X	X
Feature Install based Installation						X	X	X
Java-based configuration program						X	X	X
Remote configuration support						X	X	X
NFS Client/Server	X						X	X

Feature	TCP/IP version							
	2.0	3.0	3.1	3.5	4.0	4.1	4.2	4.21
PMX Client/Server	X							
DHCP/DDNS Clients			X	X	X	X	X	X
DHCP/DDNS Servers			X	X		X	X	X
Virtual Private Networking (VPN)						X	X	X
Multi-threaded FTPD and TFTPD							X	X
Streaming LPD and LPRPORTD							X	X
Time Server (TIMED)							X	X
BootP Server (BOOTPD)	X	X	X	X	X	X		
DHCP / BOOTP Relay Agent							X	X
TCP/IP Development Toolkit	X	X*			X*	X*		X*
TN3270 and TN5250		X	X	X				
Personal Communications 4.1 lite					X	X	X	
WebExplorer		X	X	X	X			
NewsReader/2		X	X	X	X			
Gopher		X	X	X	X			
UltiMail Lite		X	X	X	X			

* These Toolkits are available separately.

The following related products are also included with OS/2 Warp Server for e-business. They are:

- Netscape Communicator 4.04

  Required to install TCP/IP 4.21 from the Client CD-ROM included with OS/2 Warp Server for e-business

- Java 1.1.7

  Required for the new Java based TCP/IP Configuration Notebook

Netscape Communicator can be used as a replacement for WebExplorer, NewsReader/2, Gopher, and UltiMail Lite.

7.2 TCP/IP 4.21 installation

This section will document the installation of TCP/IP 4.21 using either *Selective Install for Networking* to install it on the server or the new *Feature Install* installation for installing it on a client.

7.2.1 Selective install for networking

During installation of OS/2 Warp Server for e-business, the screen shown in Figure 95 will be presented. Alternatively, you can return to it by starting *Selective Install for Networking* from the *Install/Remove* folder in *System Setup*.

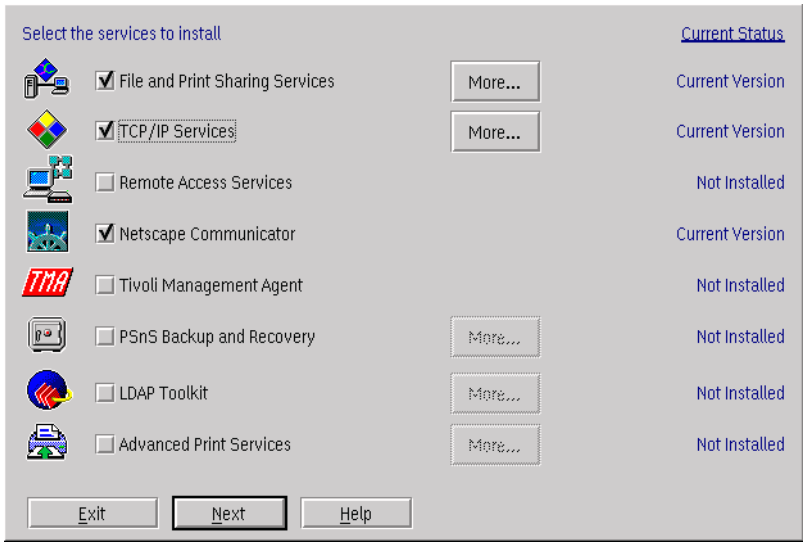

Figure 95. Selective install for networking: Selection screen

Make sure **TCP/IP Services** is selected, and, if the **More** button next to it is selected, the dialog box shown in Figure 96 will be presented.

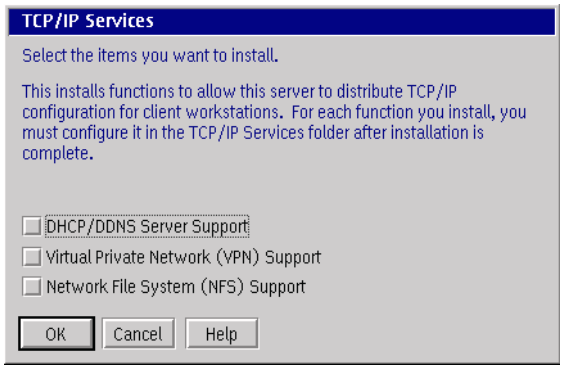

Figure 96. Selective install for networking: TCP/IP 4.21 options

The available options are:

Table 24. Selective install for networking: TCP/IP services

Option	Description
DHCP/DDNS Server Support	This installs the DDNS and DHCP servers for automatically configuring clients.
Virtual Private Network (VPN) Support	This installs the VPN support for building a secure private network over an insecure link like the Internet.
Network File System (NFS) Support	This installs support to mount remote NFS filesystems and to create NFS shares.
OK	This keeps changes made and returns.
Cancel	This discards changes made and returns.
Help	This is the present online help for the dialog box.

On the Selection screen, pressing **Next** will show the configuration screen shown in Figure 97.

Further configuration can be done after installation using the TCP/IP Configuration Notebook as discussed in Section 7.3, “TCP/IP configuration notebook” on page 236.

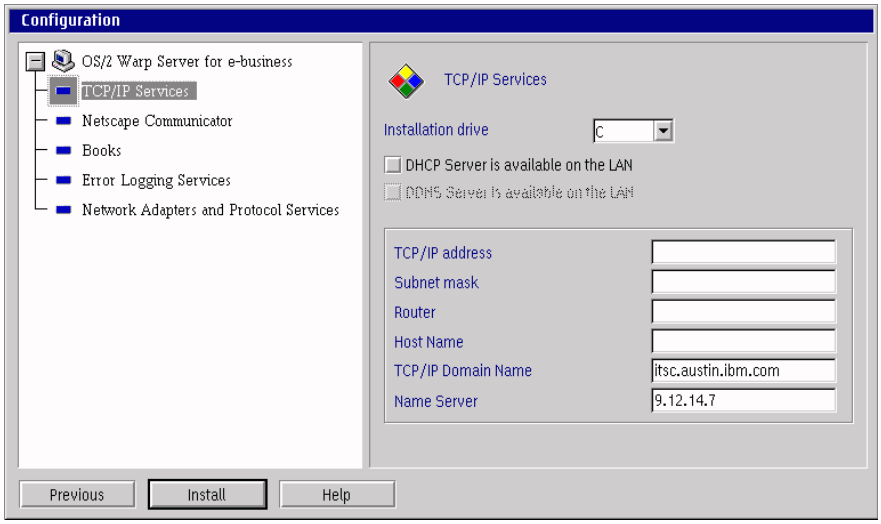

Figure 97. Selective install for networking: TCP/IP configuration

When finished entering the TCP/IP configuration and selecting the right network adapter, press **Install**.

Table 25. Selective install for networking: TCP/IP configuration

Option	Description
Installation drive	Local drive \TCPIP should be installed to.
DHCP Server available on the LAN	Indicate that configuration should be retrieved from a DHCP server on the LAN
DDNS Server available on the LAN	register hostname with a DDNS server
TCP/IP Address	TCP/IP address, like 10.1.2.3
Subnet mask	Subnet Mask, like 255.255.255.0
Router	The IP address for the default router
Host Name	The TCP/IP hostname of the local machine
TCP/IP Domain Name	The TCP/IP domain name (without hostname)
Name Server	The IP address of the nameserver

7.2.2 Installation from the client CD-ROM

Installation of TCP/IP 4.21 is supported on the following Operating Systems:

- OS/2 Warp 4
- OS/2 Warp Server 4.0
- OS/2 Warp Server 4.0 Advanced
- OS/2 Warp Server 4.0 SMP

The minimum requirements for the computer are a Pentium processor with at least 32MB of memory (16MB per processor with SMP) and sufficient diskspace to install all of the components required.

The installation from the client CD-ROM supplied in the OS/2 Warp Server for e-business package first requires that the following components be installed in this order:

- Netscape Communicator 4.04

  Located on the client CD-ROM in \OS2\NS\EN as OS2EN40.EXE.

  Execute the self-extracting file to unpack and start INSTALL.EXE. This requires about 10MB of diskspace to install.

- Feature Installer 1.23 (FI)

  Located on the client CD-ROM in \OS2\FI\EN as FIRUNPKG.ZIP.

  UNZIP this file and execute FISETUP.EXE to install. This requires about 1MB of diskspace to install.

- Java 1.1.7

  Presuming the Runtime with UNICODE font is required, this is located on the client CD-ROM in \OS2\JAVA as JAVAINUF.EXE.

  Execute the self-extracting file to unpack and start INSTALL.EXE. This requires about 24MB of diskspace to install.

- MPTS 5.50

  Located on the client CD-ROM in \OS2\MPTN\EN as MPTS550X.ZIP.

  UNZIP this file and execute INSTALL.CMD to install.

The above assumes the English language version is to be installed; if it is not, replace the EN part in the directory names with DE for German, NL for Dutch, FR for French, and so on.

After that, TCP/IP 4.21 install can be started by unpacking TCPAPPS.ZIP from \OS2\TCPIP\EN (presuming English language version) and starting

INSTALL.CMD. This causes Netscape to be started for the installation process. After the initial welcome screen, the choice for a Guided Path or Advanced Path will be given.

Note

If Guided Path is selected, the opportunity to configure certain TCP/IP settings will not be presented. Instead, TCP/IP needs to be configured using the TCP/IP Configuration Notebook.

If Advanced Path is selected, the screen shown in Figure 98 is presented where selections can be made for the components to be installed.

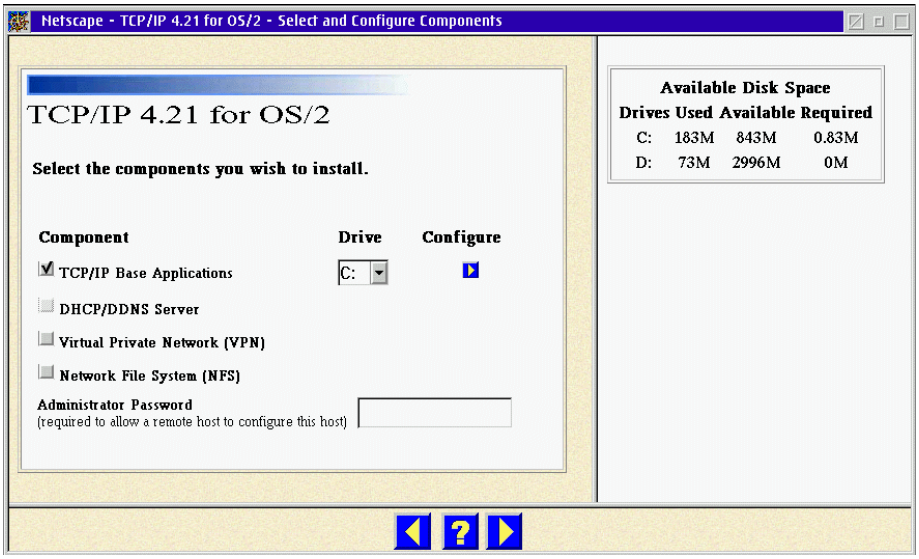

Figure 98. TCP/IP 4.21 advanced path selection screen

Table 26. TCP/IP installation: Select and configure components

Option	Description
TCP/IP Base Applications	Base applications, such as FTP, TELNET, and the configuration notebook
DHCP/DDNS Server	DHCP and DDNS servers for clients to get there TCP/IP configuration from over the LAN
Virtual Private Network (VPN)	SecureIP support to create secure network over an insecure network like the Internet

Option	Description
Network File System (NFS)	NFS client and server support to mount UNIX filesystems and create NFS shares.
Administrator Password	Optional administrator password for remote TCP/IP configuration of the machine.

Press the **Configure** button; you will have the option to enter some basic TCP/IP configuration.

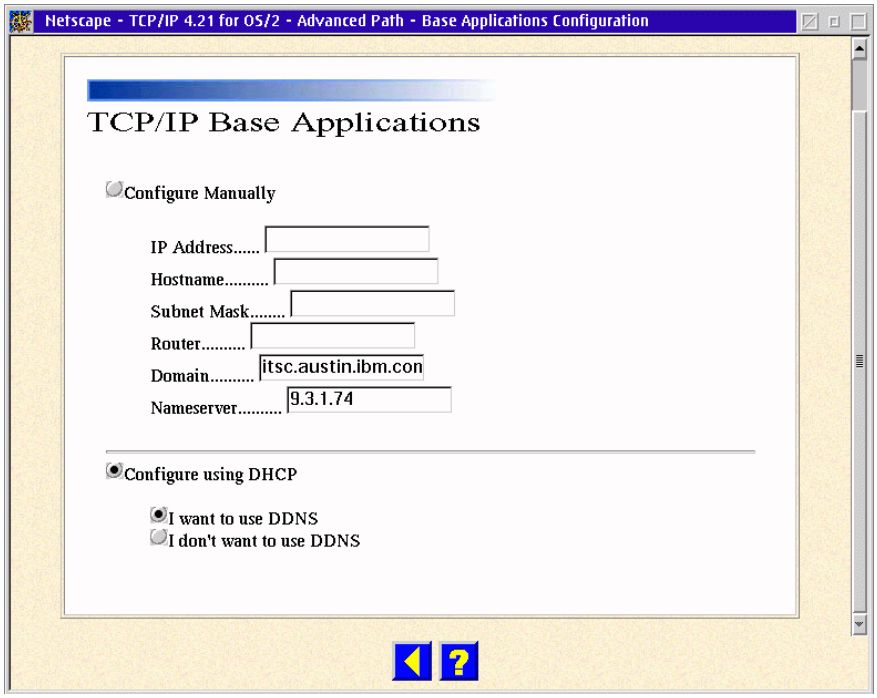

Figure 99. TCP/IP configuration during advanced installation

Table 27. Configuration during advanced installation

Option	Description
Configure Manually	Configure TCP/IP manually
IP Address	The IP address, like 10.1.2.3
Subnet Mask	The subnet mask, like 255.255.255.0
Hostname	The hostname
Router	The IP address of the default router
Domain	The TCP/IP domain name (without the hostname)
Nameserver	The IP address of the nameserver
Configure using DHCP	Get TCP/IP configuration from a DHCP server on the LAN
I want to use DDNS	Do not register the hostname with a DDNS server
I do not want to use DDNS	Register the hostname with a DDNS server

Further configuration can be done after installation using the TCP/IP Configuration Notebook as discussed in Section 7.3, “TCP/IP configuration notebook” on page 236.

Press the previous screen arrow to go back once you have finished entering the details.

Note

There are two new files. The afinet.sys and socket.sys files are used on uniprocessor machines. SMP servers will use afinetk.sys and socketsk.sys. These files require that the server use the SMP kernel.

All the files are copied to the hard disk.

7.3 TCP/IP configuration notebook

For configuration, there is a Java 1.1-based configuration utility that was first introduced with TCP/IP 4.1 that allows both local and remote configuration.

7.3.1 Local configuration

For local configuration, start *TCP/IP Configuration (Local)* from the TCP/IP folder or **TCPCFG2** from an OS/2 Command line.

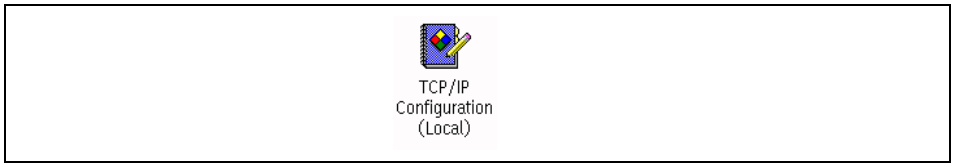

Figure 100. TCP/IP configuration (local)

7.3.2 Remote configuration

To allow remote configuration, first, the administrator password needs to be defined.

This can either be done through the TCP/IP Configuration Notebook on the **Security** tab and then under the **Admin PW** sub-tab, or by using *Create TCP/IP Administrator Password* from the TCP/IP folder.

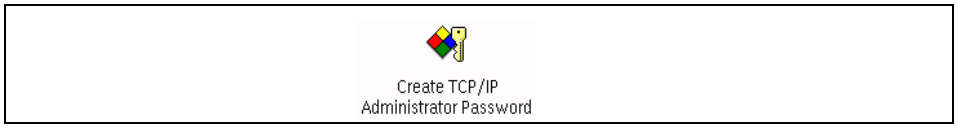

Figure 101. Create TCP/IP administrator password

Note

The Admin PW subtab on the Security tab will not be present when doing remote configuration. This can only be configured on the local machine.

Both of these functions will not prompt for an existing password before allowing the password to be changed.

When a password has been defined, the *Allow Remote Configuration* process needs to be started (TCPCFG2D from an OS/2 command line).

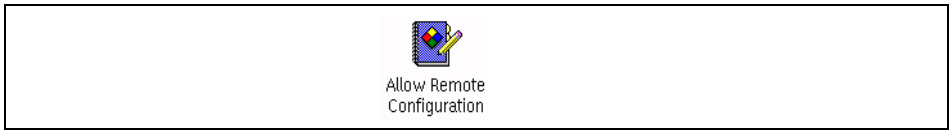

Figure 102. Allow remote configuration

Note

As long as the Allow Remote Configuration process is running, neither TCP/IP Configuration (Local) nor TCPCFG2 can be used.

A workaround without terminating the Allow Remote Configuration process is to use Configure Remote System and enter the machines hostname or localhost.

After this, Configure Remote System or TCPCFG2R from an OS/2 Command line can be used to configure the system from another machine.

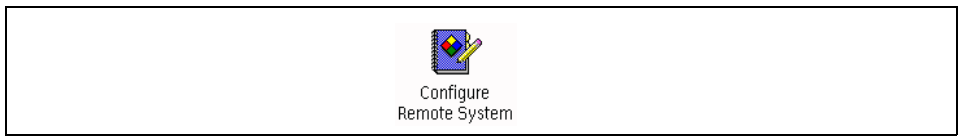

Figure 103. Configure remote system

Then, the dialog box shown in Figure 104 will be presented. It asks for the host to configure and the password. After successful verification, the configuration notebook will be started with the configuration of the remote machine.

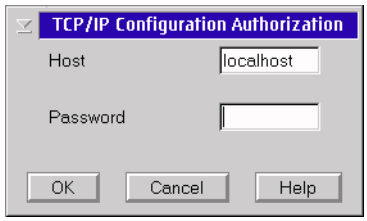

Figure 104. TCP/IP configuration authorization

7.3.3 Network tab

Here, the network adapters that have TCP/IP enabled can be configured. The configuration for this tab will be saved in \MPTN\BIN\SETUP.CMD.

When changes are made to this tab and the configuration notebook is closed, you will be prompted to reboot the system. However, just running MPTSTART.CMD from an OS/2 command line will accomplish the same thing.

7.3.3.1 Basic sub-tab

The tab shown in Figure 105 allows the selection of a LAN interface on the left side and the configuration of its corresponding settings on the right.

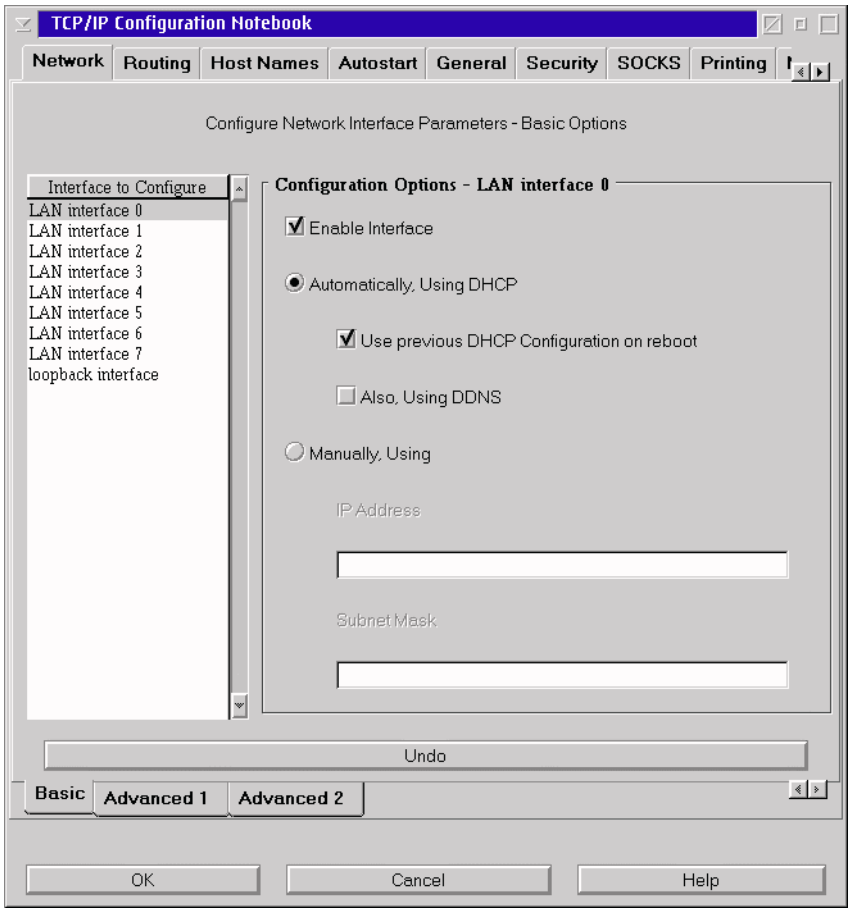

Figure 105. TCP/IP configuration notebook: Network, basic

Unless multiple network adapters need to have TCP/IP, only LAN Interface 0 needs to be configured. This corresponds to the network adapter in MPTS that has a 0 in front of the IBM TCP/IP stack.

The options available are:

Table 28. TCP/IP 4.21 settings: Network, basic

Option	Description
Interface to Configure	List box with IP interfaces that can be configured. The numbering (0 - 8) corresponds to the numbering of the TCP/IP interfaces in MPTS.
Enable Interface	This enables the selected interface.
Automatically, Using DHCP	Request configuration for the selected interface from a DHCP server
Use previous DHCP configuration on reboot	If the DHCP client has a valid lease, it will contact the DHCP server requesting a renewal. If this fails because the DHCP server is down, it will use the previous configuration.
Also, Using DDNS	Use a Dynamic DNS server in combination with the DHCP server. This will cause the configured hostname to be registered with the IP address assigned by the DHCP server at the DDNS server.
Manually, Using	Specify the IP address and subnet mask for the selected interface manually
IP Address	The IP address for the selected interface like 10.1.2.3
Subnet Mask	The subnet mask for the selected interface like 255.255.255.0
Undo	Undo any changes since this tab was last accessed.

7.3.3.2 Advanced 1 sub-tab

The values here are for the LAN Interface selected on the **Basic** sub-tab as indicated by the LAN Interface 0 message near the top of the page.

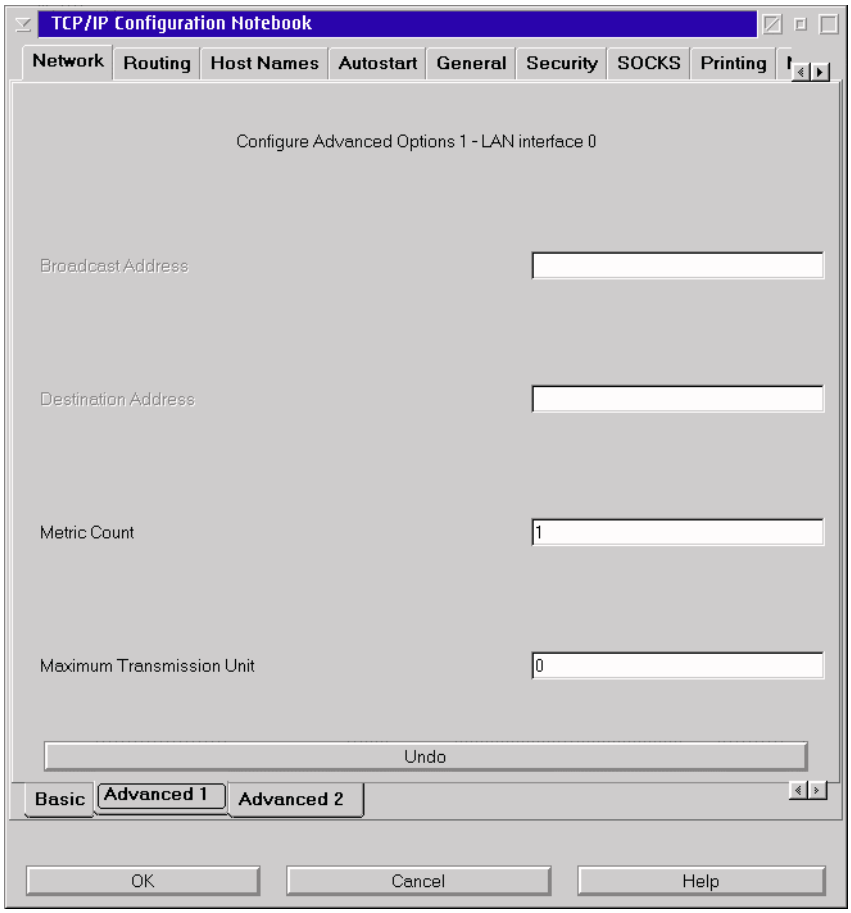

Figure 106. TCP/IP configuration notebook: Network, advanced 1

The available options are listed in the following table:

Table 29. TCP/IP 4.21 settings: Network, advanced 1

Option	Description
Broadcast Address	Broadcast address for the network; by default, this will be calculated by TCP/IP. Not available when DHCP is selected for the currently selected interface.
Destination Address	Destination address when using a point-to-point link. Not available when DHCP is selected for the currently selected interface.
Metric Count	The number of hops; zero indicates a direct connection
Maximum Transmission Unit	The MTU size; default is 1500 bytes. Maximum is 17800. Enlarging this value on some types of LANs like Token Ring can increase performance but can also cause problems because, when the packet needs to be forwarded to another type of network like Ethernet, which has a maximum MTU of 1500, the Router will have to fragment the packets into smaller ones or may even discard the packet.

7.3.3.3 Advanced 2 sub-tab

The values here are for the LAN Interface on the selected **Basic** sub-tab as indicated by the LAN Interface 0 message near the top of the page.

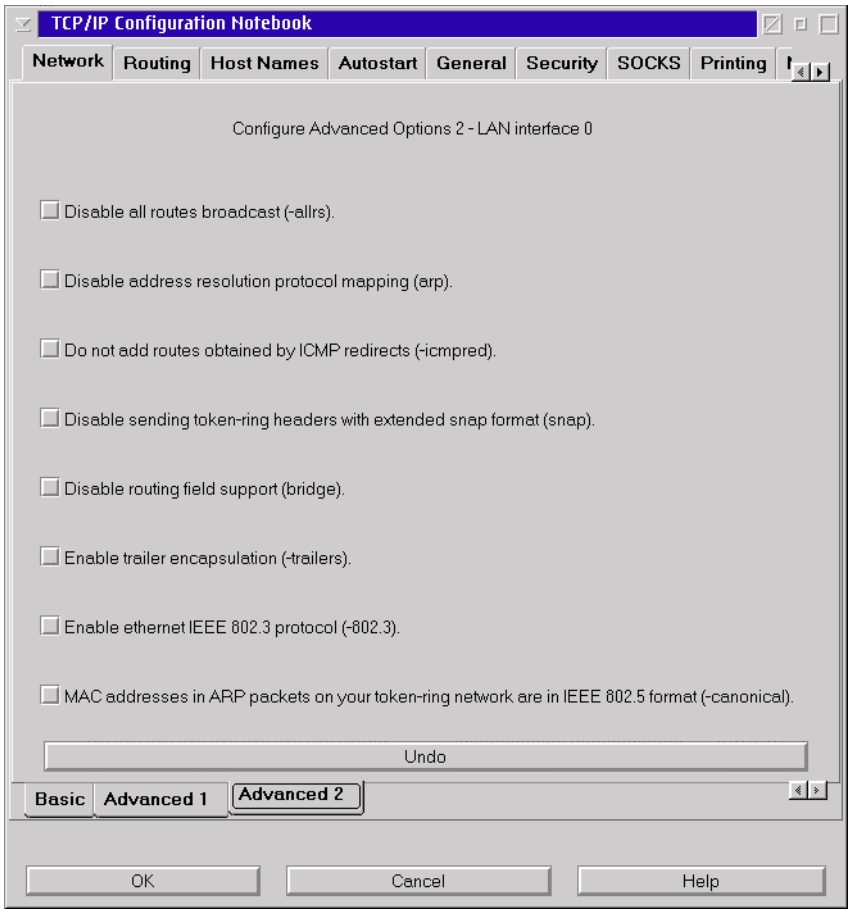

Figure 107. TCP/IP configuration notebook: Network, advanced 2

The available options are listed in the following table:

Table 30. TCP/IP 4.21 settings: Network, advanced 2

Option	Description
Disable all routes broadcast (-allrs)	Used to disable broadcasts over all routes
Disable address resolution protocol mapping (-arp)	Used to disable mapping between IP addresses and Token Ring or Ethernet MAC addresses.
Do not add routes obtained by ICMP redirects (-icmpred)	This is to disable adding routes to the routing table obtained via ICMP redirects.
Disable sending token-ring headers with extended snap format (-snap)	SNAP is part of the IEEE standard for Token Ring and needs to be enabled to communicate computers that use the extended SNAP format such as AIX.
Disable routing field support (-bridge)	to disable routing field support used with bridges
Enable trailer encapsulation (-trailers)	Enables trailer encapsulation and reduces the number of memory-to-memory copy operations a receiver must perform. When used on a network that supports ARP, the system will request that the other host also use trailer encapsulation.
Enable ethernet IEEE 802.3 protocol (-802.3)	Checking this box will use the IEEE 802.3 Ethernet protocol instead of the default DIX 2 protocol.
MAC addresses in ARP packets on your token-ring network are in IEEE 802.5 format (-canonical)	To indicate that the MAC address in ARP packets on this Token Ring network are in canonical IEEE 802.5 format instead of the default non-canonical.

7.3.4 Routing tab

This page allows the definition of routes to reach a host. These settings will be saved in \MPTN\BIN\SETUP.CMD.

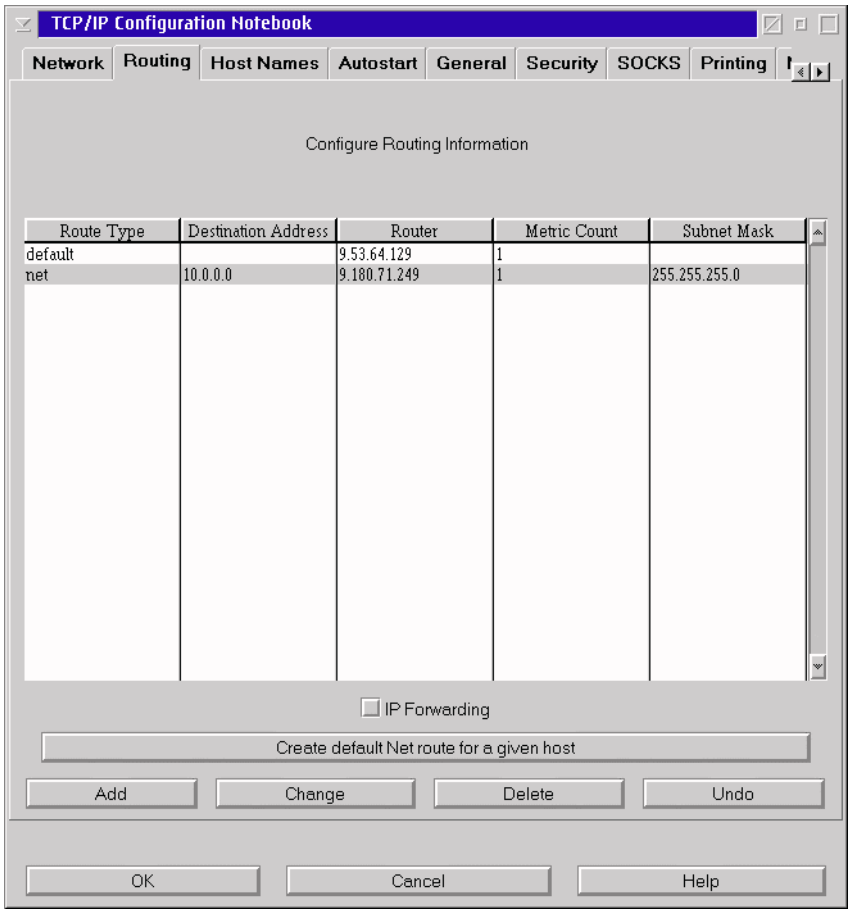

Figure 108. TCP/IP configuration notebook: Routing

Table 31. TCP/IP configuration notebook: Routing

Option	Description
IP Forwarding	To enable IP forwarding between interfaces, this will make it work as a basic IP router.
Create default Net route for a given host	If using a simultaneous LAN and dial-up connection using the Internet Connection Kit, select this option to have TCP/IP calculate the route needed to access specific hosts.

Option	Description
Add	To add a routing entry above or below the currently selected entry
Change	To change the currently selected routing entry
Delete	To delete the currently selected routing entry
Undo	Undo any changes since this tab was last accessed

Pressing **Add** will present the dialog shown in Figure 109. This will allow the creation of a net, host or default routing entry.

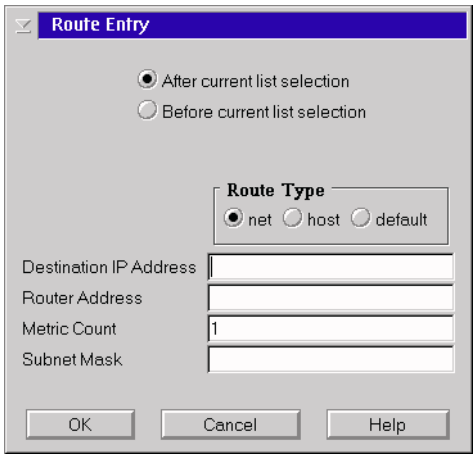

Figure 109. TCP/IP configuration notebook: Routing, route entry

Table 32. TCP/IP configuration notebook: Add routing entry

Option	Description
After current list selection	To add the new entry after the current selected one
Before current list selection	To add the new entry before the current selected one
Route Type	Selection for net, host, or default route type
Destination IP Address	For net and host routes, the IP address of the destination
Router Address	The IP address of the TCP/IP router
Metric Count	The hopcount
Subnet Mask	Only for net entries, the subnetmask of the network

7.3.5 Host names tab

This tab allows the configuration of everything that has to do with host name resolution.

7.3.5.1 Name resolution sub-tab

Here, the machine's hostname, domain name, name servers, and LAN domain searchlist can be defined. The hostname is saved in the CONFIG.SYS in the SET HOSTNAME= variable, the rest are saved in \MPTN\ETC\RESOLV2

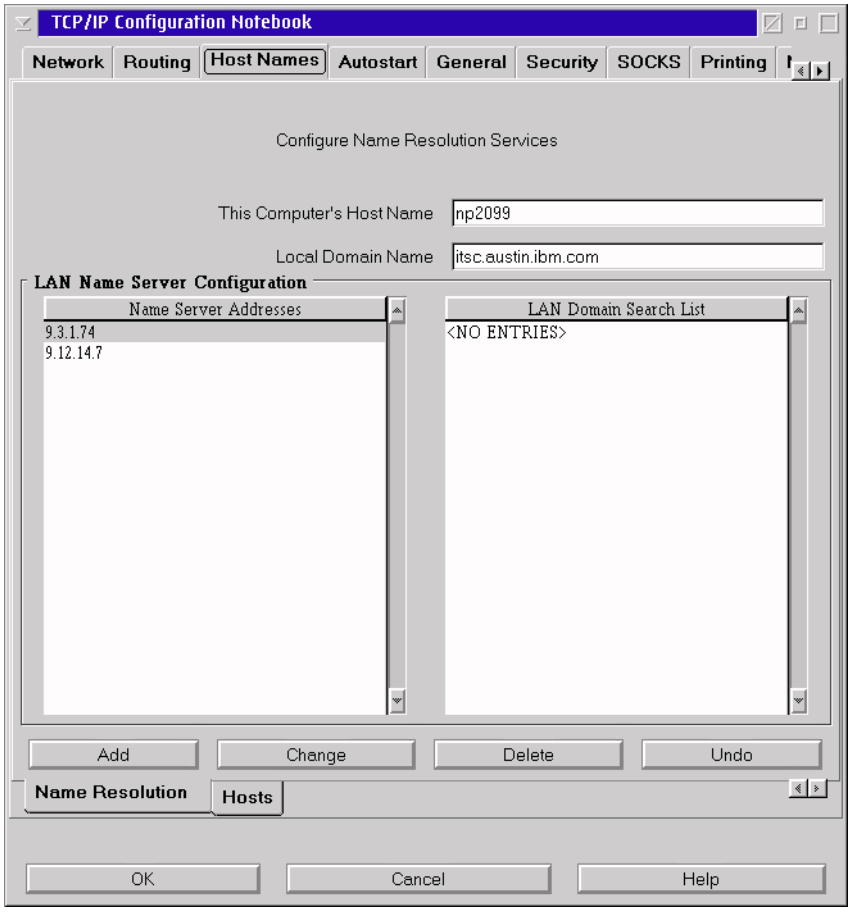

Figure 110. TCP/IP configuration notebook: Host names, name resolution

Options available are:

Table 33. TCP/IP configuration notebook: Host names, name resolution

Option	Description
This Computer's Host Name	The Host name for this PC, this does not include the domain name.
Local Domain Name	The TCP/IP domain the PC is in
Name Server Address	List box with nameservers listed in order of preference.
LAN Domain Search List	List box with domain names to be searched when an application requests a hostname without a domain. listed in order of preference.
Add	To add a entry to the nameserver address list or the LAN domain searchlist depending on the list that is selected.
Change	To change the selected entry
Delete	To delete the selected entry
Undo	Undo any changes since this tab was last accessed.

7.3.5.2 Hosts sub-tab

Here, static hostnames can be defined. This is saved in \MPTN\ETC\HOSTS. Changes are effective immediately (no reboot required).

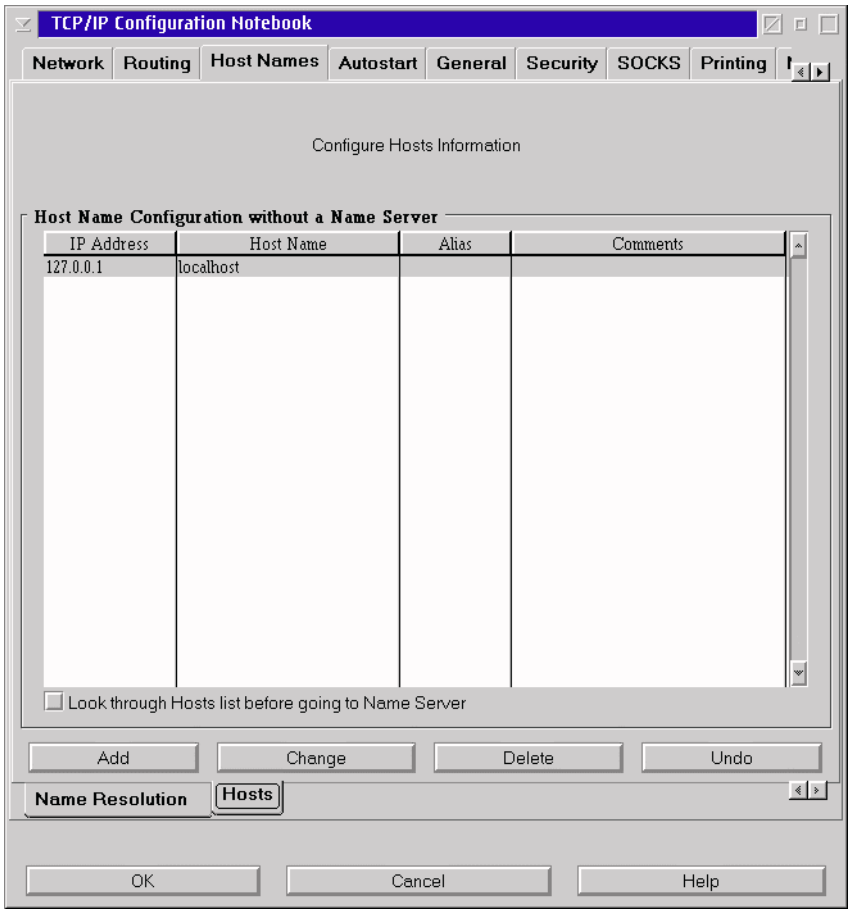

Figure 111. TCP/IP configuration notebook: Host names, hosts

The available options are listed in the following table:

Table 34. TCP/IP configuration notebook: Host names, hosts

Option	Description
Host Name Configuration without a Name Server	List of hosts with IP addresses for which the nameserver does not have entries.
Add	Add a new host entry
Change	Change the current selected host entry
Delete	Delete the current selected host entry
Undo	Undo any changes since this tab was last accessed.
Look through Hosts list before going to Name Server	Default is to first check the nameserver and then look for a local HOSTS file; if selected, the HOSTS file will be checked before going to the nameserver.

There are a couple of situations where defining a list of hostnames with IP addresses on the local machine might be useful.

- A certain machine on the network does not have a hostname entry associated with it in the DNS.
- For giving a machine an easy-to-remember nickname
- For accessing a machine in another domain with just the hostname instead of the fully qualified name (without having to use the LAN domain search list).
- A DNS server is not available on the LAN.

The disadvantage of this is that if an IP address changes all machines that have references to it using a local HOSTS file will have to change their configuration. This only needs to be done once for the DNS.

7.3.6 Autostart tab

Here, TCP/IP services (Daemons) can be autostarted.

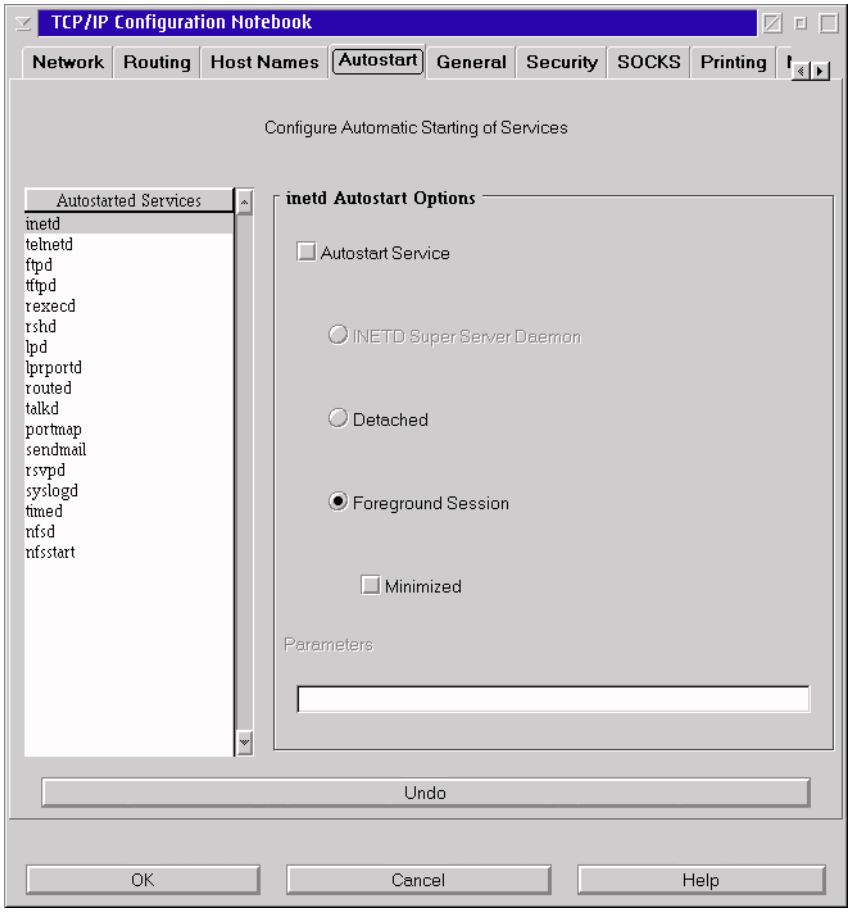

Figure 112. TCP/IP configuration notebook: Autostart, inetd

Table 35. TCP/IP configuration notebook: Autostart, inetd

Option	Description
Autostart Service	To enable autostarting of the selected service on the left.
INETD Super Server Daemon	Start the service using INETD; this requires that INETD be autostarted.

Option		Description
Detached		Start the service detached; it will not show up in the tasklist and cannot be terminated unless a special process monitor is used that can kill processes.
Foreground Session		Start the session in the foreground. This will cause a OS/2 command session to be started with the service.
Minimized		This will minimize the foreground session automatically.
Parameters		Parameters to be passed to the service
Autostarted Services	inetd	The inetd super server. This is a process that monitors for incoming connections and starts the appropriate service when needed, thus, saving resources.
	telnetd	Accepts incoming telnet sessions
	ftpd	Accepts incoming ftp sessions
	tftpd	Handles incoming tftp requests
	rexecd	Handles incoming rexec requests
	rshd	Handles incoming rsh requests
	lpd	Handles incoming lp print jobs
	lprportd	Gives \PIPE\LPDx ports needed for printing to an LPD server using WPS printer objects.
	routed	Router service using RIP
	talkd	Accepts incoming talk sessions
	portmap	Port Mapping service for RPC
	sendmail	Handles incoming and outgoing SMTP sessions
	rsvpd	Handles incoming Quality of Service reservations
	syslogd	Server for the LOG service.
	timed	Time server
	nfsd	Handles incoming NFS sessions
	nfsstart	Starts the client control program, NFSCTL.EXE, and mounts NFS shares defined in FSTAB.INI. More info can be found in Section 7.7.2.2, “NFSSTART service” on page 297.

7.3.7 General tab

This allows the definition of a username for RSH and LPR sessions and the selection of the timezone the machine is in.

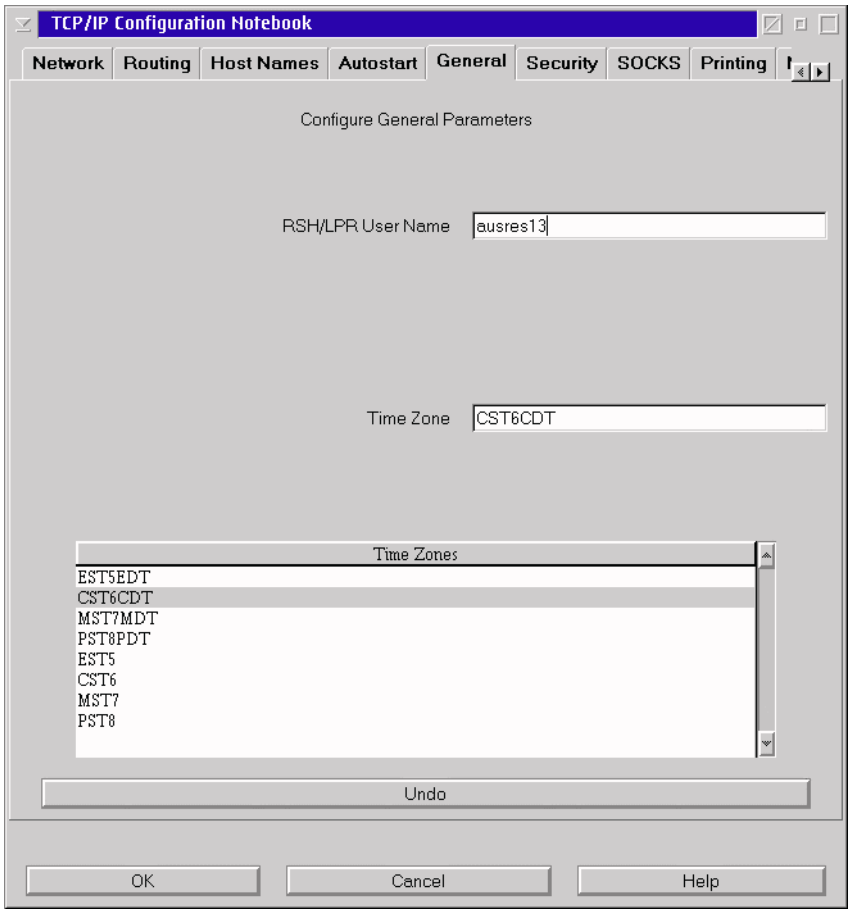

Figure 113. TCP/IP configuration notebook: General

The available options are listed in the following table:

Table 36. TCP/IP configuration notebook: General settings

Option	Description
RSH/LPR User Name	The user name used when making an RSH connection to a remote host and when creating printjobs using the line printer.
Time Zone	The time zone, select a US timezone from the list box, or enter your own values as described below.
Undo	Undo any changes since this tab was last accessed.

The selection of the right timezone can be important for some applications or services. An example of this is NFS.

For help on writing a timezone entry, select the on-line help. What follows are a couple of sample timezones that you might want to use.

Table 37. Sample timezones.

Timezone	Location
GMT0BST	UK and Portugal
MET-1DST	Western Europe (BENELUX, France, Germany, Italy, Spain, and so on)
WET-2WET	Finland
NZS-12NZD	New Zealand

7.3.8 Security tab

This tab is for configuring access rights.

7.3.8.1 User access sub-tab

Here, users can be defined and access granted to FTP, TELNET, REXEC and NFS.

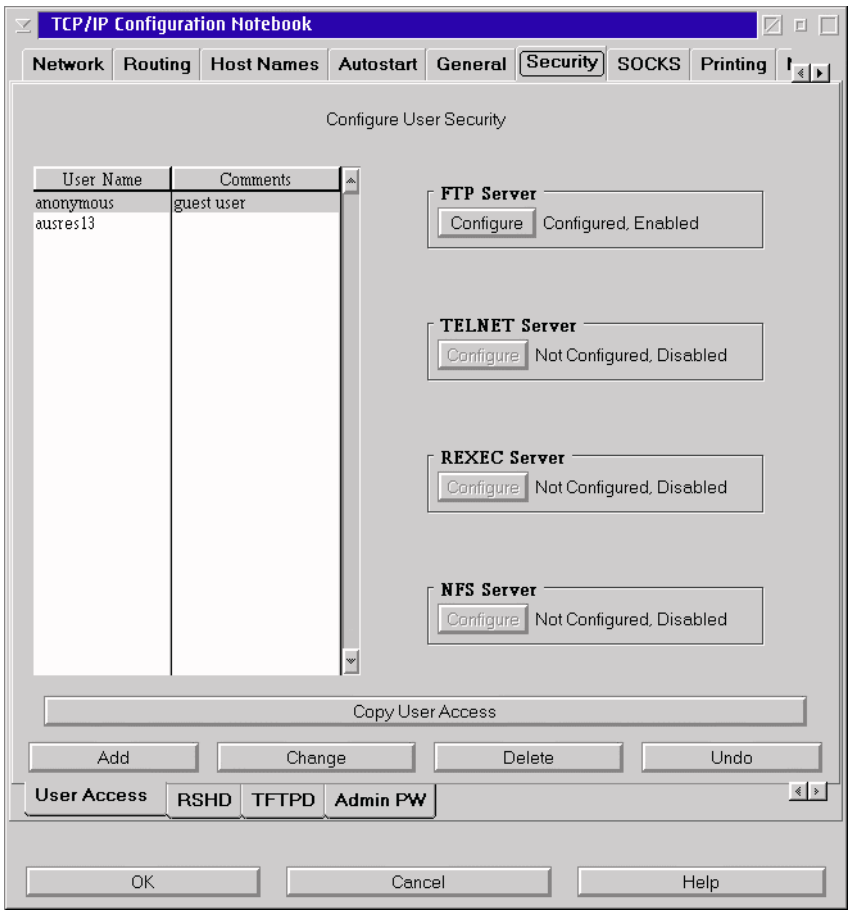

Figure 114. TCP/IP configuration notebook: Security, user access

Note

The *anonymous* user name is a special one that can only be given access to FTP. All other options will be grayed out.

The options available are:

Table 38. TCP/IP configuration notebook: Security, user access

Option	Description
FTP Server	Enable and/or configure FTP server for the user selected in the listbox.
TELNET Server	Enable and/or configure TELNET server for the user selected in the listbox.
REXEC Server	Enable and/or configure REXEC server for the user selected in the listbox.
NFS Server	Enable and/or configure NFS server for the user selected in the listbox.
Copy User Access	Copy the access rights of the selected user and specify new user details.
Add	Add a new user ID.
Change	Change the selected users details.
Delete	Delete the selected user.
Undo	Undo any changes since this tab was last accessed.

Add user

This is the dialog box for adding a new user.

Figure 115. TCP/IP security, add user

The available options are listed in the following table:

Table 39. TCP/IP security, add user

Option	Description
User Name	A short name or ID for the user.
Comment	A comment, for instance the full name of the user.
Password	The users password, this can only be modified from the this configuration notebook. The user cannot change it himself.
Verify Password	Verify the password entered.
Home DIR	The home directory for the user. This is the directory he/she ends up in when they establishes an FTP or TELNET session to the machine.
OK	Keep changes.
Cancel	Discard changes.
Help	Display on-line help.

Configure FTP access

This allows the configuration of FTP access for the selected user.

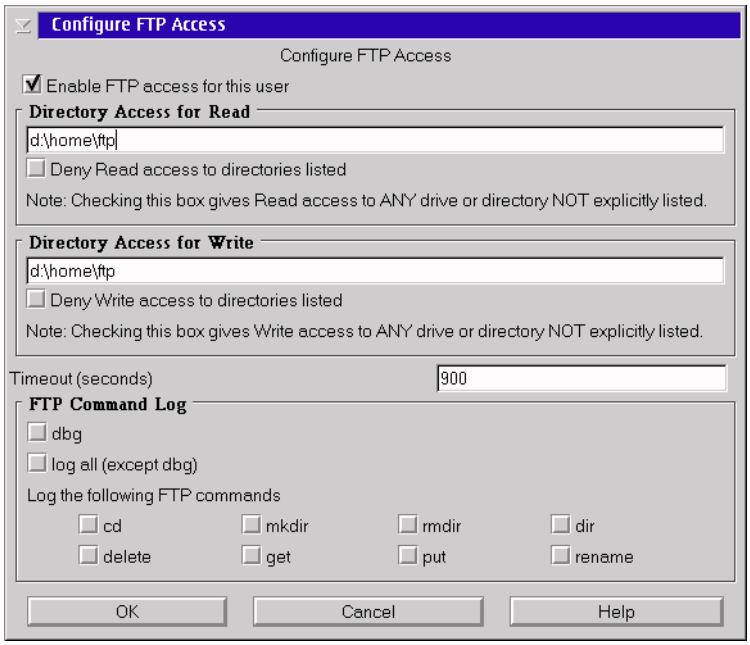

Figure 116. Configure FTP access

For clients to be able to access the FTPD server, first the FTPD server needs to be started. It can be autostarted on bootup of the machine on the Autostart tab, or it can be started manually by starting FTPD from an OS/2 command line.

The available options are listed in the following table:

Table 40. Configure FTP access

Option	Description
Enable FTP access for this user	Enables the user to be able to FTP to this machine disabled by default.
Directory Access for Read	Specify directories the user can read separated by spaces.
Deny Read access to directories listed	Gives the user read access to everything except the directories specified for read.
Directory Access for Write	Specify directories the user can write to separated by spaces.

Option	Description
Deny Write access to directories listed	Gives the user read access to everything except the directories specified for write.
Timeout (seconds)	Specifies, in seconds, how long the before the connection gets closed on no activity.
dbg	full debugging mode, everything is logged. Requires that SYSLOGD is running.
log all (except dbg)	Logs all FTP commands; requires that SYSLOGD is running.
Log the following FTP commands	Log only the FTP commands specified; requires that SYSLOGD is running.
OK	Keep changed settings and return to previous screen.
Cancel	Discard changed settings and return to previous screen.
Help	Present on-line help for the dialog box.

See Section 7.9, “FTPD improvements” on page 306 for details on new features.

Configure TELNET access

This allows configuration of TELNET access for the selected user (requires that TELNETD be started).

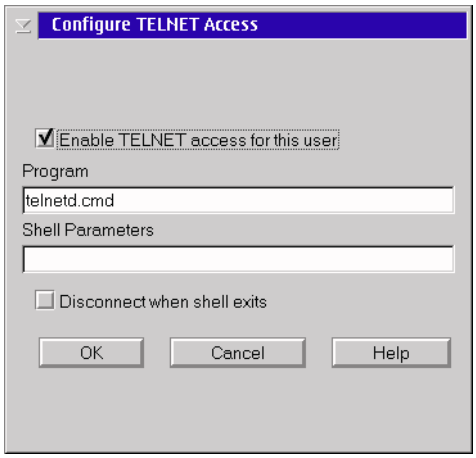

Figure 117. Configure TELNET access

The following options are available:

Table 41. Configure TELNET access

Option	Description
Enable TELNET access for this user	Enable the selected user to TELNET into the machine, disabled by default.
program	Specify the shell to run for the user. Normally, telnetd.cmd, but can be another OS/2 VIO program.
shell parameters	Specify parameters to pass to the program.
Disconnect when shell exits	Terminates the users session when the shell ends. Use this when a user needs to be given access to a certain OS/2 VIO program, but is not allowed to terminate it and get to a command prompt. Do not use with telnetd.cmd since, then, the user will not even get a command prompt.
OK	Keep changed settings and return to previous screen.
Cancel	Discard changed settings and return to previous screen.
Help	Present on-line help for the dialog box.

Configure REXEC access

When enabled, this allows the user to execute REXEC commands.

Figure 118. Configure REXEC access

The following options are available:

Table 42. Configure REXEC access

Option	Description
Enable REXEC access for this user	Give the selected user access to REXEC, disabled by default.
OK	Keep changed settings and return to previous screen.
Cancel	Discard changed settings and return to previous screen.
Help	Present on-line help for the dialog box.

For clients to be able to execute RECEC commands, the RECECD server first needs to be started. It can be autostarted on bootup of the machine on the Autostart tab or started manually by starting RECECD from an OS/2 command line.

Configure NFS access

This enables mapping between a UID/GID and a user ID and password for PCNFSD-enabled NFS clients.This is saved to the TCPNBK.LST file located in \MPTN\ETC.

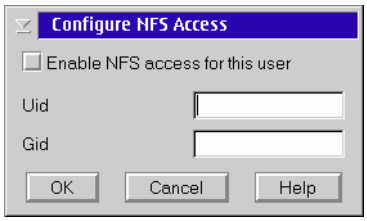

Figure 119. Configure NFS access

The available options are listed in the following table:

Table 43. Configure NFS access

Option	Description
Enable NFS access for this user	Give the selected user access to NFS shares, disabled by default.
Uid	Enter the UNIX UserID (UID) for the user.
Gid	Enter the UNIX GroupID (GID) for the user.
OK	Keep changed settings and return to previous screen.
Cancel	Discard changed settings and return to previous screen.
Help	Present on-line help for the dialog box.

For clients to be able to access the NFS server, the NFSD server first needs to be started. It can be autostarted on bootup of the machine on the Autostart tab, or started manually by starting NFSD from an OS/2 command line.

See Section 7.7, “Network file system” on page 294 for more information on NFS.

7.3.8.2 RSHD sub-tab

Here, hosts or specific users at those hosts can be given the right to execute RSH (Remote Shell) commands on this machine.

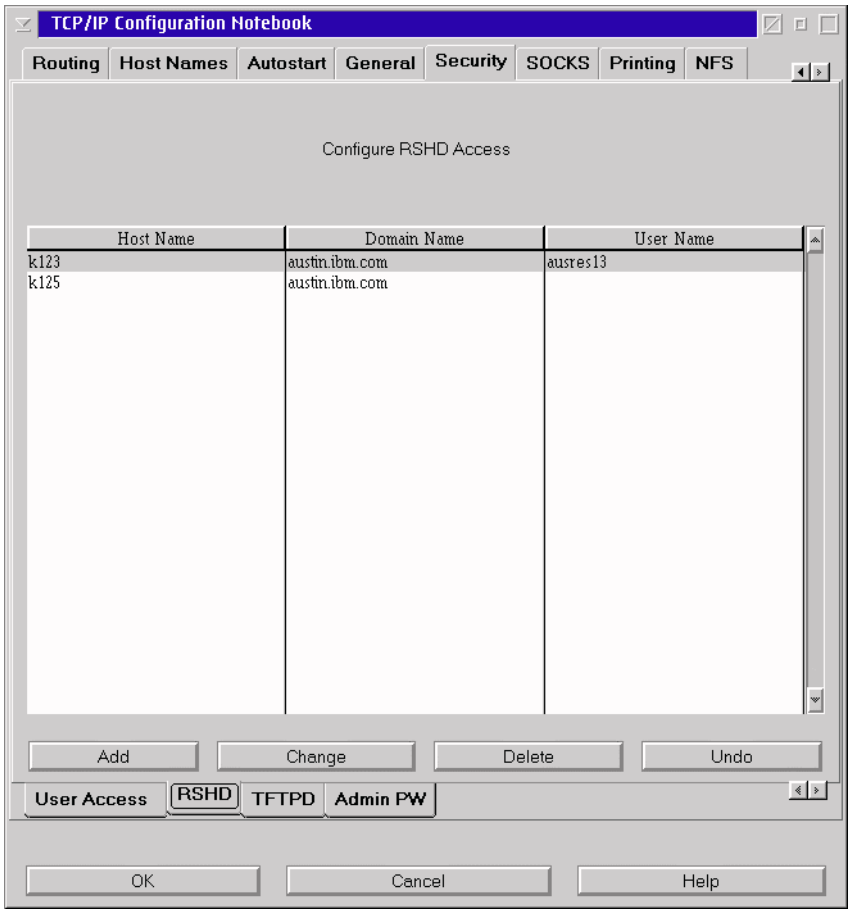

Figure 120. TCP/IP configuration notebook: Security, RSHD

For clients to be able to execute RSH commands, the RSHD server first needs to be started. It can be autostarted on bootup of the machine on the Autostart tab, or started manually by starting RSHD from an OS/2 command line.

The available options are listed in the following table:

Table 44. TCP/IP configuration notebook: Security, RSHD

Option	Description
Add	Add a new entry in the list
Change	Change the selected entry
Delete	Delete the selected entry
Undo	Undo any changes since this tab was last accessed.

Note

When a User Name is not configured for a given host, every user at that host will be able to execute RSH commands on your machine.

7.3.8.3 TFTPD sub-tab

Here, directories can be defined with read/write or read-only access that the TFTPD server will give access to.

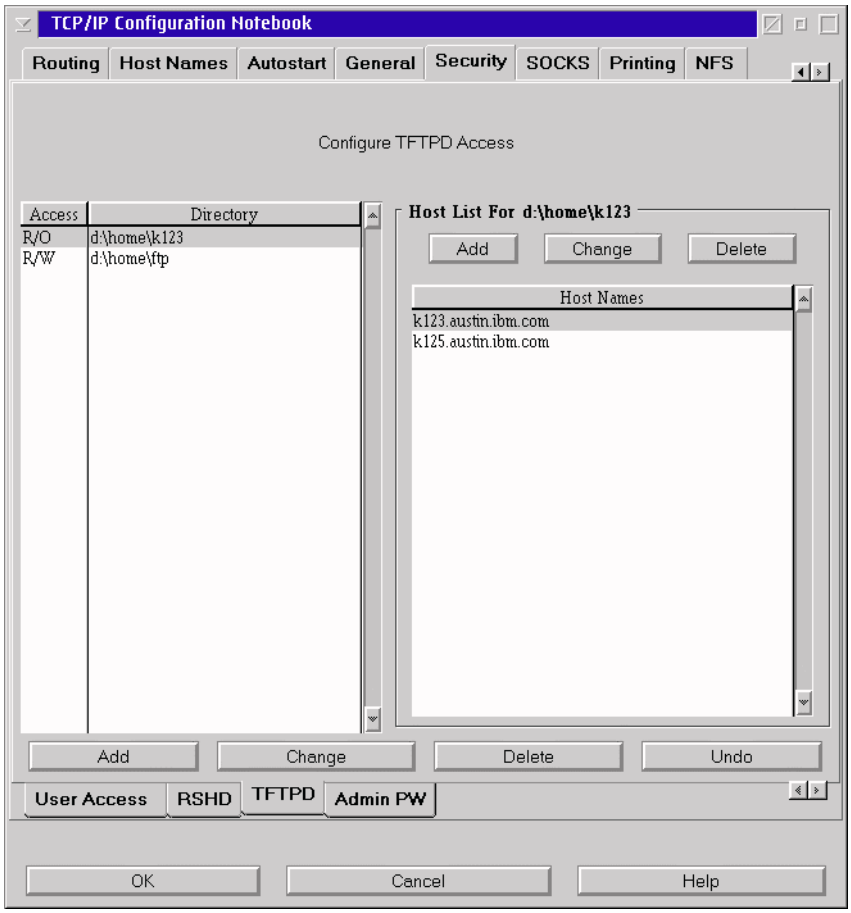

Figure 121. TCP/IP configuration notebook: Security, TFTPD

For clients to be able to access the TFTPD server, the TFTPD server first needs to be started. It can be autostarted on bootup of the machine on the Autostart tab, or started manually by starting TFTPD from an OS/2 command line.

Options available are:

Table 45. TCP/IP configuration notebook: Security, TFTPD

Option	Description
Add	Add a directory that can be accessed via TFTP.
Change	Change the selected directory entry.
Delete	Delete the selected directory entry.
Undo	Undo any changes since this tab was last accessed.
Host List for - Add	Add a hostname entry for the directory selected.
Host List for - Change	Change the selected hostname entry.
Host List for - Delete	Delete the selected hostname entry.

Note

It is not possible to define a directory with R/O or R/W access and not configure at least one host that can access it. In other words, public access to TFTPD is not possible.

7.3.8.4 Admin PW sub-tab

This allows the administrator password to be defined for allowing remote configuration.

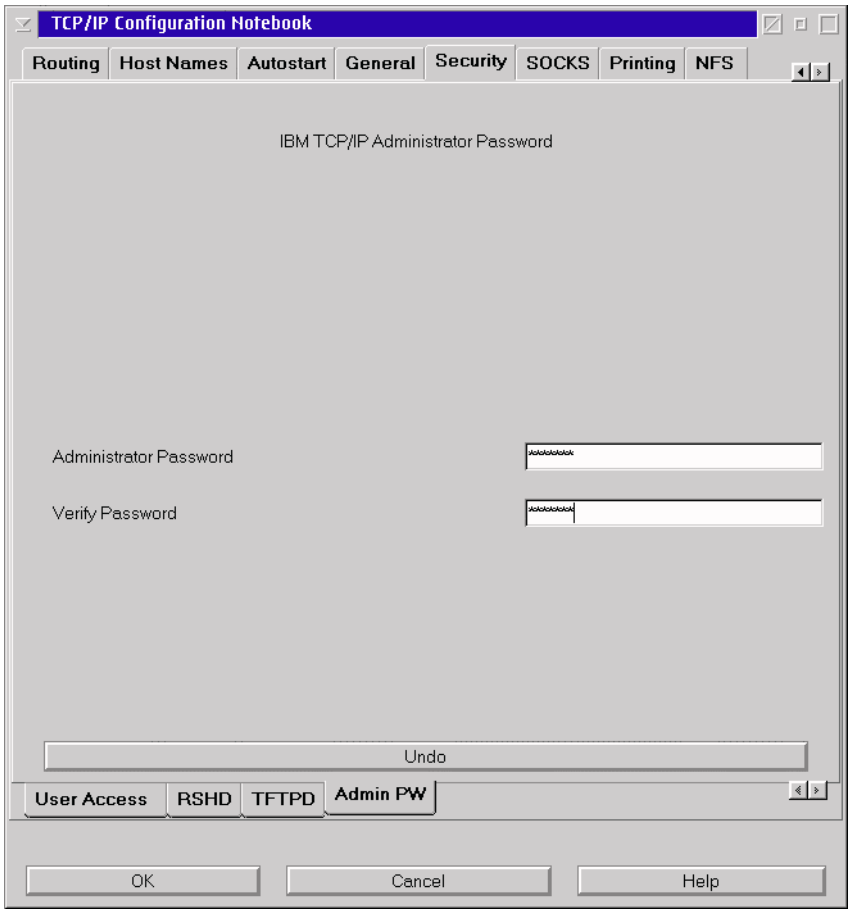

Figure 122. TCP/IP configuration notebook: Security, Admin PW

For more information on remote configuration, see Section 7.3.2, “Remote configuration” on page 236.

Note

This page will only be present when configuring TCP/IP on the local machine. When using the remote configuration function, it will be missing as a security precaution.

The available options are listed in the following table:

Table 46. TCP/IP configuration notebook: Security, Admin PW

Option	Description
Administrator Password	A new administrator password for remote configuration.
Verify Password	Re-type password to verify.
Undo	Undo any changes since this tab was last accessed.

7.3.9 SOCKS tab

The SOCKS tab is for use in private networks where a SOCKS server is available to access the internet. When properly configured, this will allow any application on the machine that uses TCP or UDP packets to access machines outside the firewall (not the other way around).

7.3.9.1 Defaults sub-tab

Examples of applications that can make use of this are FTP and TELNET. PING and TRACERTE are examples that will not work because they use ICMP packets. When used with Netscape, the PROXY or SOCKS settings in Netscape should not be enabled.

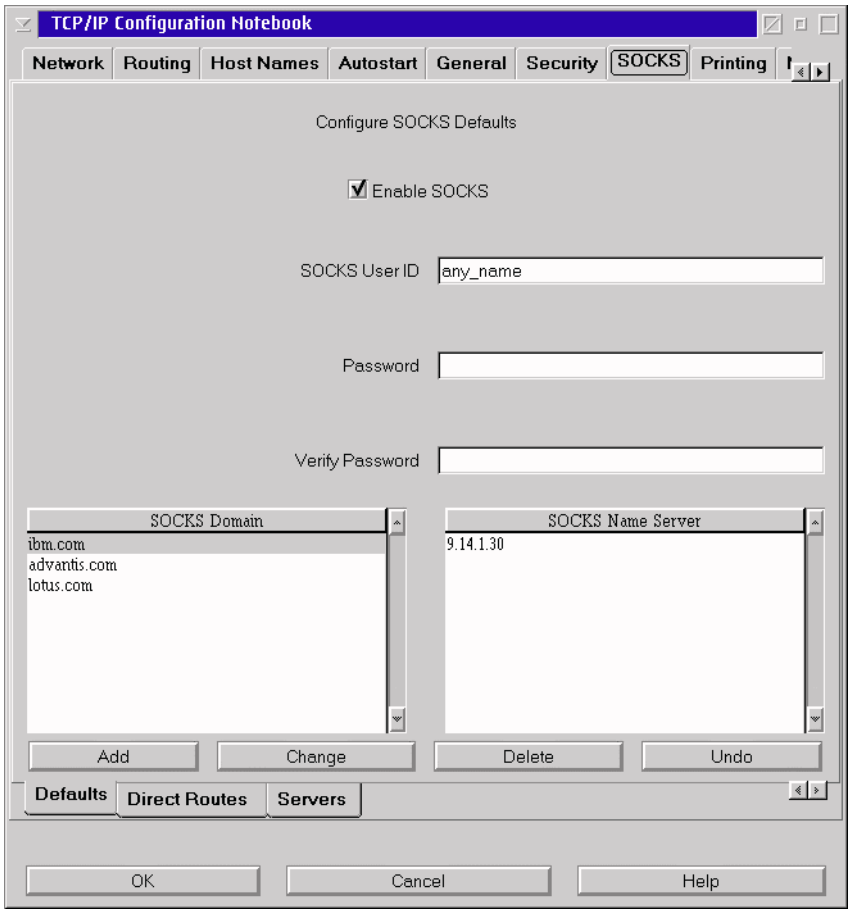

Figure 123. TCP/IP configuration notebook: SOCKS, defaults

The configuration is saved in SOCKS.ENV and SOCKS.CFG in the ETC directory. Once SOCKS is configured on one machine, the files can be just be copied onto other machines and optionally have the userid/password modified if needed.

This SOCKS client implementation is compatible with SOCKS version 5.

The available options are listed in the following table:

Table 47. TCP/IP configuration notebook: SOCKS, defaults

Option	Description
Enable Socks	Enable the SOCKSified TCP/IP stack, disabled by default.
SOCKS User ID	A user ID that is registered at the SOCKS server; this is an optional field. Requirement depends on the SOCKS server.
Password	The password for the user ID, this is an optional field length: 1-254 characters, case sensitive.
Verify Password	Verify the password entered.
SOCKS Domain	List of domains that do not require a SOCKS server to access them.
SOCKS Name Server	SOCKS name servers that will be used to resolve hostnames outside your private network.
Add	Add an entry to the domain or nameserver list depending on the list selected.
Change	Change the selected domain or nameserver entry.
Delete	Delete the selected domain or nameserver entry.
Undo	Undo any changes since this tab was last accessed.

7.3.9.2 Direct Routes sub-tab

This tab allows the definition of direct routes, in other words, routes to hosts in your private network that do not require SOCKS.

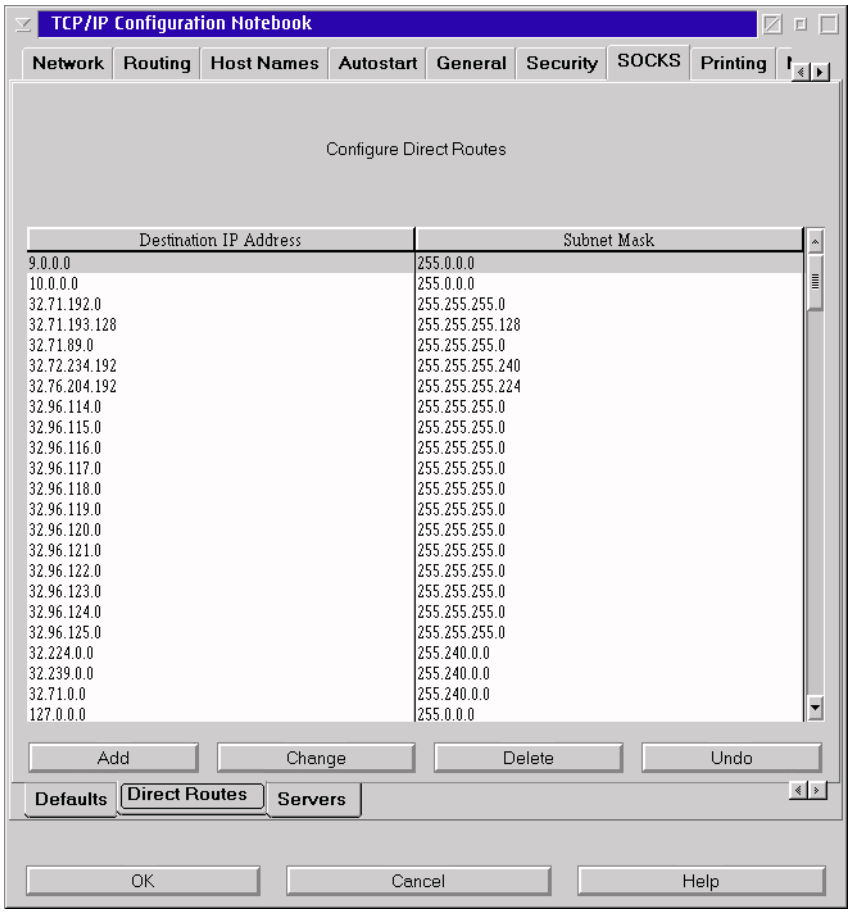

Figure 124. TCP/IP configuration notebook: SOCKS, direct routes

Table 48. TCP/IP configuration notebook: SOCKS, direct routes

Option	Description
Add	Add a direct route entry.
Change	Change the selected entry.
Delete	Delete the selected entry.
Undo	Undo any changes since this tab was last accessed.

Figure 125 shows the dialogbox to enter a new direct route entry. Routes will automatically be sorted in the listbox depending on IP address.

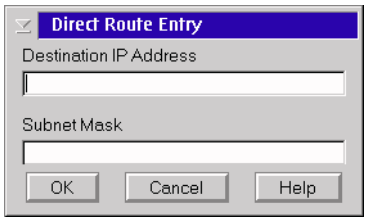

Figure 125. TCP/IP configuration notebook: SOCKS, direct routes, add entry

Table 49. TCP/IP configuration notebook: SOCKS, direct routes, add entry

Option	Description
Destination IP Address	The IP address of the destination host or network.
Subnet Mask	The subnet mask of the destination host or network.
OK	Keep the settings and close the dialog box.
Cancel	Discard the settings and close the dialog box.
Help	Provide on-line help for the dialog box.

7.3.9.3 Servers sub-tab

This page allows the definition of the actual SOCKS servers to be used. Several entries can be made.

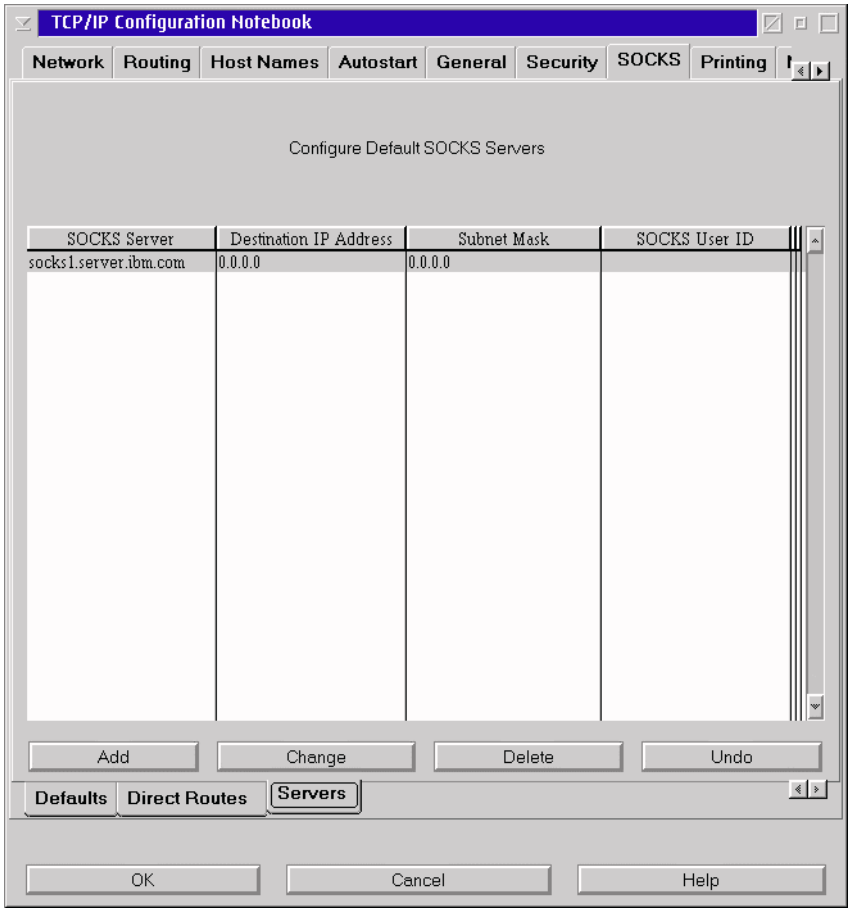

Figure 126. TCP/IP configuration notebook: SOCKS, servers

Table 50. TCP/IP configuration notebook: SOCKS, servers

Option	Description
Add	Add an entry.
Change	Change the selected entry.
Delete	Delete the selected entry.
Undo	Undo any changes since this tab was last accessed.

When adding a new entry, the default SOCKS user ID and password on the **Defaults** tab can be overruled by defining a SOCKS user ID and password for the new entry.

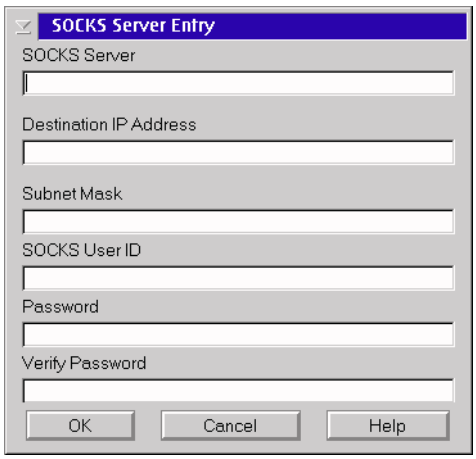

Figure 127. TCP/IP configuration notebook: SOCKS, servers, add entry

7.3.10 Printing tab

This page allows the definition of a default LPD printserver and queue for use with LPR, LPQ, LPRM and LPRMON commands, as well as the number of LPD ports that should be serviced by LPRMORTD.

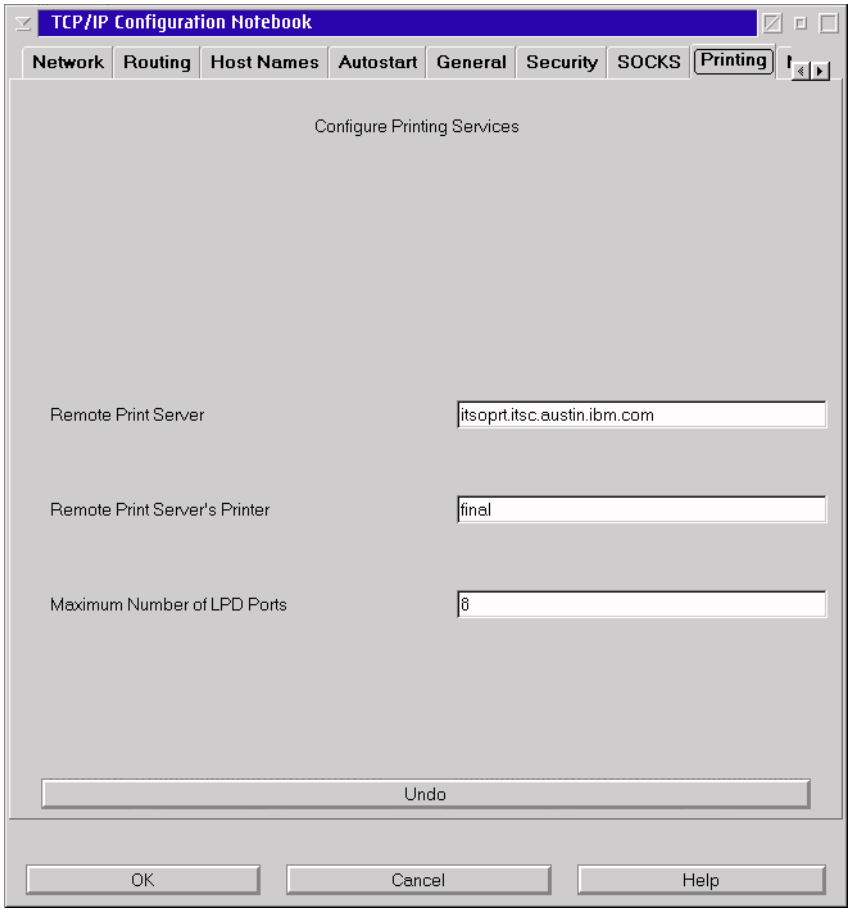

Figure 128. TCP/IP configuration notebook: Printing

Table 51. TCP/IP configuration notebook: Printing

Option	Description
Remote Print Server	The IP address or hostname of the default LPD server.
Remote Print Server's Printer	The default queue on the default LPD server.
Maximum Number of LPD Ports	The number of LPD ports LPRPORTD should service.
Undo	Undo any changes since this tab was last accessed.

Note

When using DHCP for one of the interfaces, the default Remote Print Server and default queue specified here will be overwritten.

7.3.11 NFS tab

This tab allows the definition of NFS shares that are to be available from this machine. This will be saved in the EXPORTS file in \MPTN\ETC. Changes will not take effect until the NFSD server is restarted.

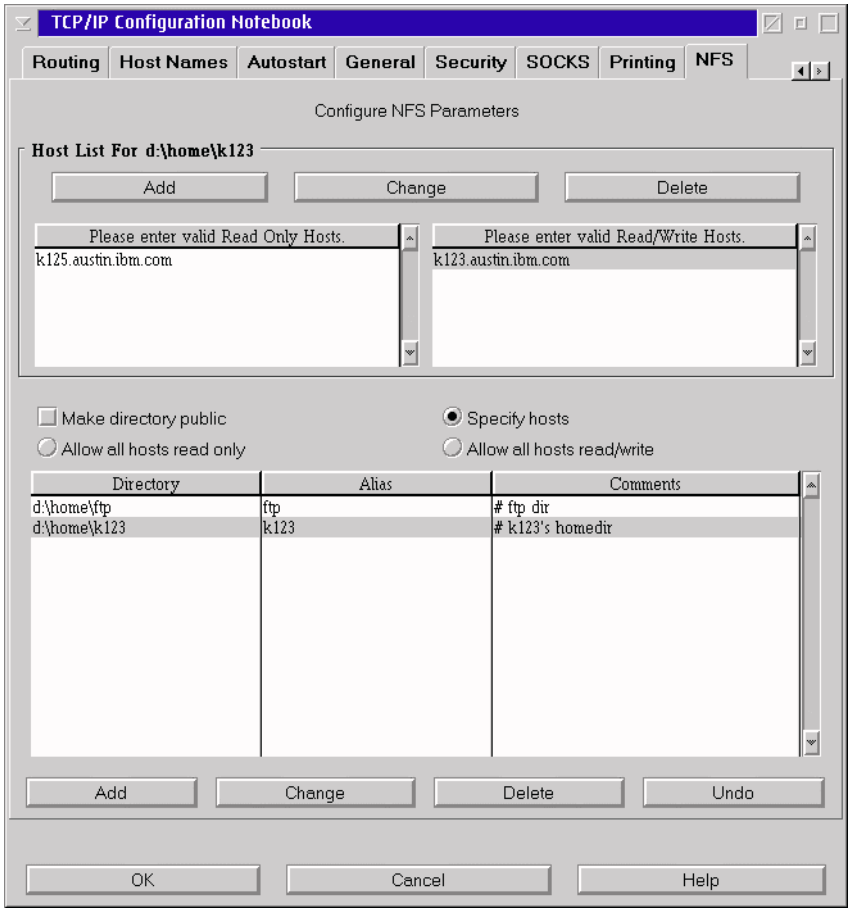

Figure 129. TCP/IP configuration notebook: NFS

More information about NFS can be found in Section 7.7, “Network file system” on page 294.

The available options are listed in the following table:

Table 52. TCP/IP configuration notebook: NFS

Option	Description
Host List for - Add	Add a hostname entry for the selected directory.
Host List for - Change	Make changes to the selected hostname entry.
Host List for - Delete	Delete the selected hostname entry
Make directory public	This is for WebNFS and gives public access to the selected directory. Only one directory can be made public this way. If multiple directories are configured public for WebNFS and the NFSD service is started, it will only make the first entry public and disregard the rest.
Specify hosts	Indicates that you only want the specified hosts to have access.
Allow all hosts read only	Allow all hosts read-only access to the selected directory.
Allow all hosts read/write	Allow all hosts Read and Write access to the selected directory.
Add	Add a directory entry to be shared.
Change	Change the selected directory entry.
Delete	Delete the selected directory entry.
Undo	Undo any changes since this tab was last accessed.

When the **Add** option is selected for defining a new directory, the dialog box shown in Figure 130 is presented.

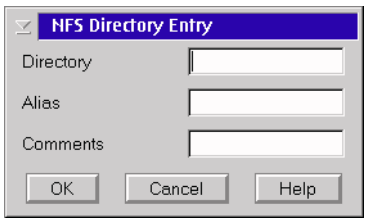

Figure 130. TCP/IP configuration notebook: NFS, add directory

The options available are listed in Table 53

Table 53. TCP/IP configuration notebook: NFS, add directory

Option	Description
Directory	The directory to be made available via NFS.
Alias	An optional Alias for the directory to make it easier to mount.
Comments	An optional comment needs to be written with, firs, a hashmark (#), then a space, and then the comment.
OK	Keep settings and close dialog box.
Cancel	Discard settings and close dialog box.
Help	Present on-line help for the dialog box.

When the Add button is pressed to add a host to be given Read Only or Read/Write access (depending on whether the Read Only or Read/Write listbox is selected) to the selected directory is pressed, Figure 131 is presented.

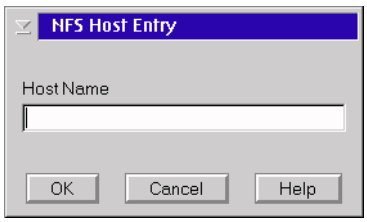

Figure 131. TCP/IP configuration notebook: NFS, add host

The options available are listed in Table 54

Table 54. TCP/IP configuration notebook: NFS, add host

Option	Description
Host Name	A hostname or IP address to be given Read Only or Read/Write access (depending on the listbox selected).
OK	Keep settings and close the dialog box.
Cancel	Discard settings and close the dialog box.
Help	Present on-line help for the dialog box.

7.4 DHCP/DDNS servers

The DHCP and DDNS server functions provided by TCP/IP 4.2 and 4.21 are essentially the same as that of TCP/IP 4.1.

This has been described in great detail in the redbook *Beyond DHCP - Work Your TCP/IP Internetwork with Dynamic IP*, SG24-5280.

7.5 Virtual private networks

Initially, companies were using the internet to promote their company's image, products, and services by providing World Wide Web access to corporate Web sites. Today, however, the Internet potential is limitless, and the focus has shifted to e-business using the global reach of the Internet for easy access to key business applications and data that reside in traditional I/T systems. Companies can now securely and cost-effectively extend the reach of their applications and data across the world through the implementation of secure virtual private network (VPN) solutions.

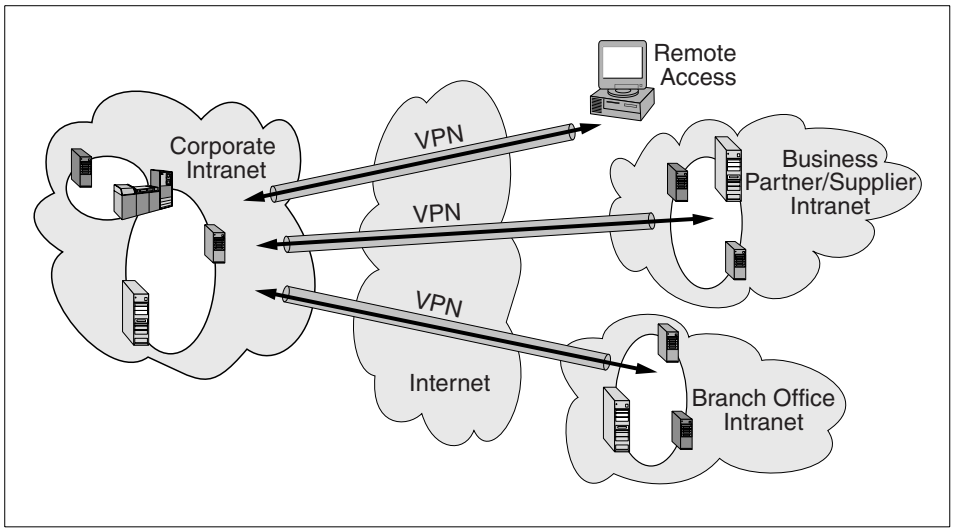

Figure 132. Virtual private networks, implementation possibilities.

A virtual private network (VPN) is an extension of an enterprise's private intranet across a public network, such as the Internet creating a secure private connection, essentially, through a private tunnel. VPNs securely convey information across the Internet connecting remote users, branch

offices, and business partners into an extended corporate network as shown in Figure 132 on page 279.

The technology to implement these virtual private networks, however, is only now becoming standardized. Some networking vendors today are offering non-standards-based VPN solutions that make it difficult for a company to incorporate all its employees and/or business partners/suppliers into an extended corporate network. However, VPN solutions based on Internet Engineering Task Force (IETF) standards will provide support for the full range of VPN scenarios with more interoperability and expansion capabilities.

Vendor's VPN offering can be categorized in a number of ways. In our opinion, the most important differentiator is the protocol layer on which the VPN is realized. In this context, there are the following approaches to VPN implementation:

- Network layer based (IPSec-based)
- Data link layer based (layer 2-based)

There are other methods that operate on upper layers and complement a VPN solution, such as SOCKS, Secure Sockets Layer (SSL), or Secure Multipurpose Internet Mail Extension (S-MIME). Some vendors' solutions use only the upper layer protocols to construct a VPN, usually a combination of SOCKS V5 and SSL. Within the layered communications protocol stack model, the network layer (IP in the case of the TCP/IP stack) is the lowest layer that can provide end-to-end security. Network layer security protocols provide blanket protection for all upper layer application data carried in the payload of an IP a datagram without requiring a user to modify the applications.

The solutions are based on the IP Security Architecture (IPSec) open framework defined by the IPSec Working Group of the IETF. IPSec is called a framework because it provides a stable, long-lasting base for providing network layer security. It can accomodate today's cryptographic algorithms and can also accomodate newer, more powerful algorithms as they become available. IPv6 implementations are required to support IPSec, and IPv4 implementations are strongly recommended to do so. In addition to providing the base security functions for the internet, IPSec furnishes flexible building blocks from which robust, secure, virtual private networks can be constructed.

The principle IPSec protocols are:

- IP Authentication Header (AH) provides data origin authentication, data integrity, and relay protection.

- IP Encapsulating Security Payload (ESP) provides data confidentiality, data origin authentication, data integrity, and relay protection.
- Internet Security Association and Key Management Protocol (ISAKMP) provides a method for automatically setting up security associations and managing their cryptographic keys.

This Chapter will give an example of how to set up VPN between two OS/2 clients running TCP/IP 4.1 or higher and give an example of how to turn the provided VPN function into a mini-firewall. The functional differences between the VPN provided with TCP/IP 4.1 and 4.21 are also briefly discussed.

For more comprehensive information on Virtual Private Networks and how to connect to an AIX implementation, see the redbook *A Comprehensive Guide to Virtual Private Networks, Volume 1: IBM Firewall, Server and Client Solutions*, SG24-5201

7.5.1 Changes in VPN versions

The functional difference between the VPN function provided with TCP/IP 4.1 and that provided with 4.2 and 4.21 is minimal. The changes are as follows.

SSL support has been updated from v3.1 to 3.2

Security in the base package of TCP/IP 4.1, WorkSpace On-Demand 2.0, and OS/2 Warp Server for e-business will be limited to 40 bits because of U.S. export regulations. A separate CD will be available in certain countries that will contain a version of the stack (part of MPTS) that supplies 56 bits

Note

For some countries a separate Security CD-ROM will be available to add 56 bits encryption capabilities to the IPsec stack that is part of MPTS.

This CD-ROM might also contain, in some cases, the 128bit secure version of Netscape Communicator version 4.04 for OS/2.

7.5.2 Configuring IPSec clients

The IPSec code in the OS/2 TCP/IP V4.21 IPSec Client provides the functionality of being a dynamic tunnel client to an eNetwork Firewall for AIX as well as the functionality of being configured for manual IPSec tunnels. We provide the necessary information to configure such manual tunnels, which will then allow the following IPSec combinations:

- Tunnel between OS/2 TCP/IP V4.21 IPSec Client and eNetwork Firewall for AIX
- Tunnel between OS/2 TCP/IP V4.21 IPSec Client and OS/390 Server
- Tunnel between OS/2 TCP/IP V4.21 IPSec Client and AIX V4.3
- Tunnel between two OS/2 TCP/IP V4.21 IPSec Clients
- Tunnel between OS/2 TCP/IP V4.21 IPSec Client and eNetwork Communication Suite

Finally, the packet filtering capability of the OS/2 TCP/IP V4.21 IPSec Client can be used to turn OS/2 into a mini firewall. First, the manual setup is described, then, the possible combinations, followed by the packet filtering capabilities of OS/2 TCP/IP V4.21

Note

Do not use the Secure Remote Client Configuration program because it can only be used to connect to an eNetwork Firewall for AIX via SSL and obtain a dynamic tunnel configuration.

7.5.3 Configuring IPSec filters and tunnels

IPSec on OS/2 TCP/IP V4.21 consists of four device drivers and associated configuration utilities. Two of the device drivers are required, and the other two are implementations of encryption and authentication algorithms (CDMF and MD5) and are, therefore, optional. The two required device drivers for IPSec are a filter device driver, FWIP.SYS, and a device driver that provides the framework for IPSec, IPSEC.SYS. To manually configure IPSec tunnels on OS/2, three components must be configured:

1. Filters to filter IP packets through a tunnel
2. Tunnel policy to define tunnel endpoint IP addresses and whether to apply encryption and/or authentication and in what order, and tunnel modes
3. Tunnel context containing keys, tunnel ID, and algorithms to be applied.

The format of the filters and tunnel definition files is very similar to that on the eNetwork Firewall for AIX and OS/390 Server.

Note

If you have a system running eNetwork Firewall for AIX, configure the OS/2 system as a remote IPSec client to the Firewall (see "Defining the IPSec Connection to the Firewall" in *A Comprehensive Guide to Virtual Private Networks, Volume 1: IBM Firewall, Server and Client Solutions*, SG24-5201) and activate an IPSec connection (see "Activating and Deactivating the Dynamic Tunnel", in the same manual). This will transfer tunnel policy and tunnel context file to the OS/2 system and create the necessary IPSec filters file. Save those files before you close the tunnel (whereupon, they will be deleted) and use them as a base to configure your manual IPSec connections.

7.5.3.1 Setting up IP filtering

The filter device driver, FWIP.SYS when active, is called by the IP layer for each IP packet coming into the workstation and going out of the workstation. The filter device driver will determine, based on configuration rules, whether an IP packet is allowed to continue. The packet is compared to the rules starting with the first and then through each subsequent rule until a match is found. When an exact match is found, the action (permit or deny) is performed. If no rules match, the packet is denied by default.

Filter rules are described fully in *Firewall for AIX, Reference Guide*, SC31-8418. What is described here is what is necessary to filter IP packets through an IPSec tunnel.

The only tool we have today to build filter rules on OS/2 is an editor. Once the filters file is built, the executable CFGFILT.EXE can be called to process the file and deliver the configuration to the filter device driver. CFGFILT.EXE takes the following parameters:

-u -i -f[file] -c -m[#] -p -d[{start|stop}] <no parameter>

Table 55. Configure filter, parameter description

parameter	action
-u	Updates the filter rules in the device driver using the contents of \MPTN\ETC\SECURITY\FWFILTRS.CNF

parameter	action
-i	To initialize the filter device driver; must be used with -u
-f[file]	To check a set of rules. Default is \MPTN\ETC\FWFILTRS.CNF
-c	To deactivate active filter rules and go to default rules of deny everything
-m[#]	Maximum concurrent Real Audio connections
-p	Real Audio port
-d[{star|stop}]t	Start or stop filter logging; default -d with no option is to start
<no parameter>	Default action is to dump active rules (no options)

The filter rules to support a local host (1.2.3.4) to remote host (4.3.2.1) VPN (tunnel 10) are the following (lines have been indented to fit the page):

```
# This is a comment as are all lines that start with #
# These are the rules for the gateway with non-secure address
001.002.003.004

 # ip packets that are ipsec esp protocol from/to a known tunnel end on
        the non-secure interface
permit 1.2.3.4 255.255.255.255 4.3.2.1 255.255.255.255 esp any 0 any 0
          non-secure local outbound
permit 4.3.2.1 255.255.255.255 1.2.3.4 255.255.255.255 esp any 0 any 0
          non-secure local inbound

# ip packets that are ipsec ah protocol from/to a known tunnel end on the
          non-secure interface
permit 1.2.3.4 255.255.255.255 4.3.2.1 255.255.255.255 ah any 0 any 0
          non-secure local outbound
permit 4.3.2.1 255.255.255.255 1.2.3.4 255.255.255.255 ah any 0 any 0
          non-secure local inbound

# ip packets that are from/to non-secure interface to/from known tunnel
          end through tunnel 10
permit 1.2.3.4 255.255.255.255 4.3.2.1 255.255.255.255 all any 0 any 0
          non-secure local both t=10
permit 4.3.2.1 255.255.255.255 1.2..3.4 255.255.255.255 all any 0 any 0
```

```
            non-secure local both t=10

   # default final rule is deny everything
```

Optional values to add to the end of a rules lines are:

l=<log control>

f=<fragmentation control>

t=<tunnel id>

Note

The format of filter rules on OS/2 resembles that of the eNetwork Firewall for AIX. You can take a look at the /etc/security/fwfilters.cnf on a firewall to get a better idea of how to configure filter rules on OS/2.

The rules for the other end of the tunnel would be the same with reversal of source and destination addresses and masks.

All interfaces (IP addresses) on the system are assumed to be non-secure. To define a secure interface, see Section 7.6, “Creating a mini-firewall” on page 293.

To define and activate filters on a tunnel endpoint, perform the following steps:

1. Edit the file MPTN\ETC\FWFILTRS.CNF adding the above lines with correct IP addresses and masks.
2. On a command line, enter `cfgfilt -i -u -d` to activate filters, load the filter rules, and activate logging.

There is a switch available in the IP layer to control filters. To have the IP layer invoke filters when FWIP.SYS is loaded, enter the following on the command line: `inetcfg -s firewall 1`

To prevent the IP layer from invoking filters, enter the following on the command line: `inetcfg -s firewall 0`

To check the current setting of this switch, enter the following on the command line: `inetcfg -g firewall`

CFGFILT, when loading the filter rules issues inetcfg -s firewall 1.

7.5.3.2 Setting up IPSec tunnel policies

The tunnel policy is implemented in FWIP.SYS along with filters and consists of settings that control how IPSEC.SYS processes outbound IP packets that are filtered through a tunnel and indicates the level of protection required to FWIP.SYS for inbound IP packets that are filtered through a tunnel.

The policy file is created with an editor and has the following format (line has been indented to fit the page):

```
<Source IP Address><Target IP Address><Context ID><Encrypt/Auth>
    <Encrypt Mode/Auth Mode> <MAC first>
```

Table 56. IP Sec policy file, parameter description

entry field value	explanation
Source IP address	The IP address, either in quad format or alias name, of the source host associated with the policy entry (this host IP address)
Target IP address	The IP address, either in quad format or alias name, of the target host associated with the policy entry (other end of the tunnel host IP address
Context ID	The tunnel context ID that identifies the tunnel ID in the policy file
Encrypt	A y or Y in this field indicates encrypted data. The other option available is n or N
Auth	A y or Y in this field indicates authenticated data. The other option available is n or N
Encrypt Mode	A y or Y in this field indicates transport mode and a n or N indicates tunnel mode for encrypted packet
Auth Mode	A y or Y in this field indicates transport mode and a n or N indicates tunnel mode for authenticated packet
MAC first	A y or Y in this field indicates authentication is done before encryption. A n or N indicates encryption is done first.

For example: 1.2.3.4 4.3.2.1 10 y/y n/n y

The above example would be the policy for tunnel endpoint 1.2.3.4. It indicates the tunnel endpoints as 1.2.3.4 and 4.3.2.1 and the tunnel ID as 10. To perform encryption and authentication, use tunnel mode (IP packet in an IP packet) and perform authentication before encryption, that is, authenticate clear text.

To enter the tunnel policy into the FWIP.SYS device driver, use the utility FWINSERT.EXE. The only parameter FWINSERT.EXE takes is the file name of the file containing the policy.

To create and load a tunnel policy do the following:

1. Edit the file MPTN\ETC\SECURITY\POLICY, adding the desired policy statement for each tunnel supported, on a new line.
2. On a command line, enter `fwinsert mptn\etc\security\policy`

Each time FWINSERT is called, it completely replaces the existing policy entries with the contents of the given file. If a tunnel policy indicates that transport mode is to be used, MTU path discovery must be turned off. To do this, enter the following on the command line: `inetcfg -s mtudiscovery 0`

To turn mtudiscovery back on, enter the following on the command line:
`inetcfg -s mtudiscovery 0`

7.5.3.3 Setting up the IPSec tunnel context cache

Tunnel context cache is implemented in IPSEC.SYS along with the framework for IPSec. A tunnel context entry defines all the parameters needed by IPSEC.SYS to encrypt, decrypt, and authenticate an IPSec IP packet.

To create a tunnel context entry, a file must be edited to contain the following fields, each on a new line. Any line beginning with a # is treated as a comment.

Line 01: This is the IP address, in dotted format or alias name, of the other end of the tunnel.

Line 02: This is the IP address, in dotted format or alias name, of this host.

Line 03: This is the tunnel ID.

Line 04: This is this host's Security Parameter Index for ESP (Encryption).

Line 05: This is this host's Security Parameter Index for AH (Authentication).

Line 06: This is the other end of the tunnel host's Security Parameter Index for ESP (Encryption).

Line 07: This is the other end of the tunnel host's Security Parameter Index for AH (Authentication).

Line 08: This is this host's encryption algorithm. It is supported in the Internet distributed version of TCP/IP CDMF.

Line 09: This host's encryption key length must be 8.

Line 10: This host's encryption key (hexadecimal integer) must be hex characters so that it will be 16 characters in length.

Line 11: The other end of the tunnel host's encryption algorithm is supported in Internet distribution version of TCP/IP CDMF.

Line 12: The other end of the tunnel host's encryption key length must be 8.

Line 13: The other end of the tunnel host's encryption key must be hex characters so it will be 16 characters in length.

Line 14: This is this host's Authentication (mac) algorithm. Supported: KEYED_MD5

Line 15: This host's Authentication (mac) key length must be 16.

Line 16: This host's Authentication (mac) key must be hex characters so it will be 32 characters in length.

Line 17: This is the other end of the tunnel host's Authentication (mac) algorithm. Supported: KEYED_MD5

Line 18: The other end of the tunnel host's Authentication (mac) key length must be 16.

Line 19: The other end of the tunnel host's Authentication (mac) key must be in hex characters so it will be 32 characters in length.

Line 20: This is the start (Time). If it is zero, use current time Represented as the number of seconds since Jan 1, 1970. This value is used with the value on line 21 to determine the tunnel end time. The tunnel context for this record will be active as soon as hand_k is run.

Line 21: This is the end (Time). If the start time is zero, it is key life time in seconds. Represented as the number of seconds since Jan 1, 1970. The recommended key lifetime is 8 hours (28800 seconds) or less.

Line 22: Reserved. Must be 0.0.0.0

For example:

```
# Tunnel Context for tunnel 10
4.3.2.1                                 # 1. The other tunnel endpoint's IP
1.2.3.4                                 # 2. This host's IP
10                                      # 3. Tunnel ID
400                                     # 4. This host's Sec. Par. Index (Encr)
400                                     # 5. This host's Sec. Par. Index (Auth)
668829                                  # 6. Other host's Sec. Par. Index
                                            (Encr)
668829                                  # 7. Other host's Sec. Par. Index
                                            (Auth)
CDMF                                    # 8. This host's Encryption Algorithm
8                                       # 9. This host's Encryption Key Length
0x000079140006a0ac                      #10. This host's Encryption Key
CDMF                                    #11. Other host's Encryption Algorithm
8                                       #12. Other host's Encryption Key
                                             Length
0x1234da442443212c                      #13. Other host's Encryption Key
KEYED_MD5                               #14. This host's Authen. (MAC) Alg.
16                                      #15. This host's MAC length
0x618a751f4411818d41ce388ed7bfc9fd # 16. This host's MAC key
KEYED_MD5                               #17. Other host's MAC Alg.
16                                      #18.Other host's MAC length
0xa8e3377a51605476683fb8d269c21023 # 19. Other host's MAC Key
0                                       #20. Start - Current Time
28800                                   #21. ENd - 8 hours from now
0.0.0.0                                 #22. Reserved - must be 0.0.0.0
```

Note

Care must be taken when creating the contents of these files. Any fields not having the correct length or specifications will cause the HAND_K, CFGFILT and FWINSERT commands to fail.

The format of tunnel context files on OS/2 resembles that of the eNetwork Firewall for AIX and OS/390 Server, but line 22 must be added manually. Modify an AIX V4.3 tunnel export file according to the format shown in the example above.

There are a few utilities to manage the contents of the tunnel context cache.

The HAND_K.EXE command is used to enter entries into the tunnel context cache. It takes as input a file containing one or more tunnel context entries. This interface was meant to be a programmed interface to a GUI and does not produce error messages that are particularly helpful. Successful processing of the file will result in no messages. Any other messages should be considered as error, and, if a message is unclear, it is often caused by a duplicate entry that is already in the cache.

To add an entry to the tunnel cache, perform the following steps:

1. With an editor, create a file mptn\etc\security\fwmctx.man containing the fields described above.
2. On the command line, enter the following:

```
hand_k mptn\etc\security\fwmctx.man
```

The tunnel context cache entry for the host that is the other end of the tunnel will have the same content with the following lines exchanging values:

line 1 and line 2
line 4 and line 6
line 5 and line 7
line 8 and line 11
line 9 and line 12 (Should be the same. There is no need to change.)
line 10 and line 13
line 14 and line 17 (Should be the same)
line 15 and line 18 (Should be the same)
line 16 and line 19

Note

When the other end of a tunnel is either an eNetwork Firewall for AIX, an OS/390 Server, or an AIX V4.3, define a manual tunnel on the partner system and use the export file for OS/2 configuration. Therefore, you do not have to worry about the other end of the tunnel because that has already been taken care of.

The `FWD_K.EXE` command is used to delete all the tunnel context entries that match the given source and destination IP addresses from the cache.

To delete all entries in the tunnel cache associated with a set of addresses, perform the following:

On the command line, enter `fwk_k <Source IP Address> <Destination IP Address>`. Both IP addresses *must* be in dotted format, for example: fwd_k 1.2.3.4 4.3.2.1.

The `ADMIN.EXE` command is used to delete or read one tunnel context entry from the cache. ADMIN takes a file as input with the following format to read a tunnel context entry. (The line has been indented to fit the page).

comments

<Source IP Address> <Destination IP Address> <Context ID> <SAID>
<USE_MY_SAID> <USE_ESP_SAID>

Table 57. Admin command, parameter description

Field Contents	Explanation
Source IP Address	IP address of this host. This is a mandatory field.
Destination IP Address	IP address of the host at the other end of the tunnel. This is a mandatory field.
Context ID	Tunnel ID. This is an optional field
SAID	Security association ID/security parameter index. This is an optional field.
USE_MY_SAID	The character y or Y in this field tells the program to use this host's SAID instead of the remote host's SAID. This is an optional field

Field Contents	Explanation
USE_ESP_SAID	The character y or Y in this field tells the program to use security parameter index for encryption rather than security parameter for index for authentication. This is an optional field

For example, the file should contain entries like the following:

1.2.3.4 4.3.2.1 10 400 y y

If the file contains a line of the above format, the default action of reading the tunnel context is done. If a particular tunnel context is to be deleted, the format of the line is the following. (The line has been indented to fit the page).

@DEL <Source IP Address> <Destination IP Address> <Context ID> <SAID> <USE_MY_SAID> <USE_ESP_SAID>

The `DEL` command should be preceded by an @ sign and can take any of the three forms of `DEL`, `del`, or `Del`.

You can also call ADMIN.EXE without any parameters and, then, enter the source and destination address of a tunnel to view the parameters from the tunnel context cache. This is essentially the same process as the admin_test utility on the eNetwork Firewall for AIX.

7.5.4 Creating a tunnel between two OS/2 machines

The GUI interface provided can only be used to setup a dynamic tunnel to an eNetwork Firewall on AIX. For setting up a tunnel between two OS/2 machines running TCP/IP 4.1 or later, manual configuration will be needed.

For this, the following components need to be configured:

- Filters to filter IP packets through a tunnel by creating file FWFILTRS.CNF as detailed in Section 7.5.3.1, “Setting up IP filtering” on page 283.
- Tunnel policy to define tunnel endpoint IP addresses and whether to apply encryption and/or authentication and in what order, and tunnel modes by creating the POLICY file as detailed in Section 7.5.3.2, “Setting up IPSec tunnel policies” on page 286.
- Tunnel context containing keys, tunnel ID, and algorithms to be applied by creating file FWMCTX.MAN as detailed in Section 7.5.3.3, “Setting up the IPSec tunnel context cache” on page 287.

Once these files have been created, activate them by running the CFGFILT, FWINSERT, and HAND_K commands as mentioned in these sections of this document.

7.6 Creating a mini-firewall

The IP filtering capabilities of VPN allow an OS/2 system to be set up as a basic firewall. This capability can be used in conjunction with or without IPSec.

For this, the system needs to have tw2o network adapters: One on the secure LAN (intranet) and the other on the insecure LAN (internet). The latter can also be a dial-up connection.

7.6.1 Securing interfaces

To define one of the interfaces as secure, add its IP address to the file \MPTN\ETC\FWSECAD.CNF on a line by itself (multiple secure interfaces can be defined by placing each IP address in the file on a line by itself). Interfaces that do not have their IP address in the file are considered insecure. If the file does not exist, create it.

7.6.2 Configuring filtering

Create a filter configuration file as described in Section 7.5.3.1, “Setting up IP filtering” on page 283.

In order to allow the OS/2 IPSec gateway to forward packets, ipgate must be turned on. To do this enter the following on a command line: `ipgate on`.

The mini-firewall should be active once you have:

1. FWFILTRS.CNF file defined and activated with the `CFGFILT` command
2. IPGATE switched on with the `ipgate on` command
3. your routing tables defined correctly.

You can check whether the firewall is functioning by making use of the `IPTRACE` command to trace the flow of IP packets and formatting the trace entries with the `IPFORMAT` command to verify the correct flow or denial of the packets.

7.7 Network file system

OS/2 Warp Server for e-business includes a new 32 bit NFS client and server that can optionally be installed during the installation of TCP/IP.

This server is there primarily for servicing NetStations, and the client for attaching UNIX NFS shares to your server and shares them further using NETBIOS.

7.7.1 Introduction to NFS

The Network File System (NFS) protocol provides transparent remote access to shared files across networks. The NFS protocol is designed to be portable across different machines, operating systems, network architectures and transport protocols. This portability is achieved through the use of Remote Procedure Call (RPC) primitives built on top of the eXternal Data Representation (XDR). NFS is built on the client-server model. The server is the host holding the actual file and the client is running the application on another host that needs to access it.

The mount protocol allows the server to hand out remote access privileges to a restricted set of clients. It performs the operating system-specific functions that make it possible, for example, to attach remote directory trees to some local filesystem.

7.7.1.1 UNIX filesystem security

UNIX systems use a system of access controls in the filesystem where every file and directory has a number of attributes. These attributes determine if a user or member of a group is allowed to read/write or execute the file.

Each user name and group name is represented by a number. The UID (UserID) and the GID (GroupID).

The best example is to show a directory listing on a UNIX box (in this case, AIX 4.1 with bash as the shell using the `ls -la` command).

```
bash$ ls -la
total 40160
drwxr-sr-x    2 john     staff        512 16 Dec 17:40 .
drwxr-sr-x   63 sys      sys         1536 22 Nov 17:15 ..
-rw-r--r--    1 john     staff    20537463 16 Dec 17:37 aixldap_0920gm.tar.gz
-rw-r--r--    1 john     staff       1096 15 Dec 21:48 .bash_history
drwxr-sr-x   63 john     staff        512 15 Dec 19:59 data
-rwxr-----    1 john     staff        254 08 Sep 1997  .profile
bash$ _
```

The first column on the left is the access controls for the file or directory listed on the right hand side. The third column is the userid of the owner of the file, and the fourth is the name of the group.

In this sample, the file `aixldap_0920gm.tar.gz` is owned by user `john` and the group that has access over the file is the group `staff`.

The sequence is as follows:

Table 58. UNIX file system access controls

bit		Description
1	-	Regular file
	d	Directory
	b	Block special file
	c	Character special file
	p	Pipe special file
	l	Symbolic link
	s	Socket
2	-	The owner cannot read the file.
	s	The owner can read the file.
3	-	The owner cannot write to the file.
	w	The owner can write to the file.
4	-	The owner cannot execute the file (or CD into it for directories).
	x	The owner can execute the file (or CD into it for directories).
	s	Set UID; the file will be executed under the user ID of the owner of the file.
5	-	Members of the group cannot read the file.
	r	Members of the group can read the file.
6	-	Members of the group cannot write to the file.
	w	Members of the group can write to the file.

bit		Description
7	-	Members of the group cannot execute the file (or CD into it for directories).
	x	Members of the group can execute the file (or CD into it for directories).
	s	Set GID; files created in a directory that has this bit set will be owned by the owner of the directory.
8	-	Other users cannot read the file.
	r	Other users can read the file.
9	-	Other users cannot write to the file.
	w	Other users can write to the file.
10	-	Other users cannot execute the file (or CD into it for directories).
	x	Other users can execute the file (or CD into it for directories).
	t	Tacky bit; only the owner can link or unlink the file or directory.

A way to resolve a user name or group name into a UID or GID is to use the `lsname` and `lsgroup` commands on UNIX.

For example, issue lsuser john to find out what the UID for user *john* is. The same can be done for groups with `lsgroup`; issue `lsgroup staff` to find out what the GID of the group `staff` is.

7.7.1.2 UID and GID

When using the OS/2 NFS client to connect to an NFS server that does not support PCNFSD, a UID and GID must be specified. Access to files and directories will first be determined by whether the NFS share is set to read/write or read only and, then, depending on the UID and GID supplied.

When connecting to the OS/2 NFS server from a client that does not support PCNFSD, a UID and GID must be specified. When the user creates a file, the OS/2 NFS server will record the UID and GID with the filename in a file called PERMS.LOG. When someone tries to access the file over NFS, the entry in the logfile is checked for read/write/execute permissions. If there is a clash between the permissions specified in the TCP/IP Configuration Notebook on the NFS tab (EXPORTS file), the most restrictive permission is applied.

7.7.1.3 PCNFSD

The PCNFSD protocol allows non-UNIX hosts to acquire UNIX style authentication for NFS requests. It provides translation between the UNIX style UID and GID to a UserID and password.

The OS/2 NFS client and server supports PCNFSD.

7.7.1.4 File locking

To avoid multiple clients trying to modify the same file on a server, some kind of record locking mechanism needs to be present supporting the NFS protocol implementation. This requires the server to maintain the current state about which the client has locks on that record. Since record locking is not part of NFS, it is provided by additional programs: the Network Lock Manager (NLM) and the Network Status Manager (NSM).

7.7.1.5 WebNFS

The WebNFS server implements extensions to the NFS protocol to support a lightweight binding mechanism for conventional or web-browsing clients that need to communicate with NFS servers accessing the Internet.

For more information on WebNFS, see RFC 2055.

7.7.2 Mounting an NFS share

To be able to mount NFS filesystems, the PORTMAP service and the NFSSTART service must first be started. Then, the `MOUNT` command can be used to mount the filesystem as described in Section 7.7.3.1, “MOUNT command” on page 299.

7.7.2.1 PORTMAP service

Any client or server that uses RPC must have the PORTMAP service started; NFS and NFSD are such services. The PORTMAP service can be started on the Autostart tab of the TCP/IP Configuration Notebook or started manually by issuing PORTMAP from an OS/2 Command Prompt.

7.7.2.2 NFSSTART service

This service does not stay resident in memory but, rather, starts the NFS Client Control program NFSCTL.EXE and mounts any NFS shares defined in FSTAB.INI. This service can be autostarted in the TCP/IP Configuration Notebook on the Autostart tab, or started from an OS/2 Command Prompt by typing `NFSSTART`.

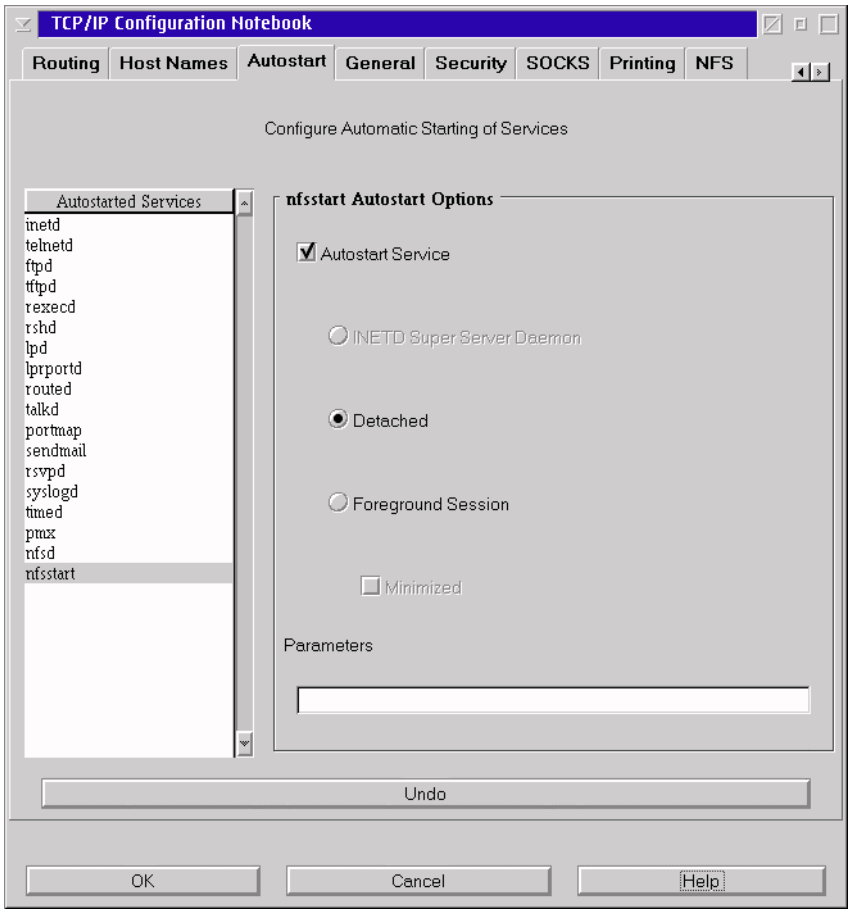

Figure 133. TCP/IP 4.21 configuration notebook: Autostart, nfsstart

NFSSTART has several optional parameters listed in Table 59 that can be entered in the Parameters field on the Autostart tab.

Table 59. NFSSTART parameters

Options	Default	Description
etc_dir	ETC variable	Specifies the location of FSTAB.INI.
-bn	8192	Sets the transfer buffer size for read and write requests. This may not exceed 8192.
-rn	5	Sets the number of RPC retries that the OS/2 NFS client sends to the server before ending the access attempt.

Options	Default	Description
-tn	1	Sets the time-out value for RPC requests in seconds.
-sn	4	Defines the number of Biods that are to be started by NFSCTL, the OS/2 NFS client Control Program. The Biods are a tool to allow parallel reads and writes to a file.
-p		Requests that the OS/2 Client Control Program uses the Biods for both Reading and Writing.
-w		Specifies that the OS/2 Client Control Program uses the Biods for reading only. This is needed with servers that do not support parallel writes to a file.
-c		Respect case when doing filename comparisons. If this is not specified and the first query fails, the OS/2 NFS client will try the name in uppercase. If this also fails, lowercase is tried.
-z		Respect filename case; if this is not specified, files and directories will be created using lowercase.
-i		Specifies that the OS/2 NFS client does serial read/writes instead of parallel read/writes. This is equivalent to the -s0 parameter.

7.7.3 NFS client utilities

This chapter describes the command line utilities available.

7.7.3.1 MOUNT command

To mount an NFS filesystem, use the `MOUNT` command from the OS/2 Command Prompt.

The syntax for `MOUNT` is as follows:

```
MOUNT <options> <drive> <hostname>:<mountpoint>
```

Table 60. Mount options

Options	Default	Description
<drive>		The local drive to mount the NFS volume on
<hostname>		The hostname for the machine the NFS volume is on

Options		Default	Description
<mountpoint>			Directory to attach and any server specific options
<options>	-u[GID]		Set the UNIX user ID (UID)
	-g[GID]		Set the UNIX group ID (GID)
	-l[LOGIN-ID]		Set the server logon ID
	-p[PASSWORD]		Set the password
	-v[PASSWORD]		Set the VM-style password
	-ac[S|M|L]	M	Attributes Cache size ([S]mall, [M]edium or [L]arge
	-acto[TIMEOUT]	15	Attribute Cache Time-out in seconds
	-dc[S|M|L]	M	Default Cache size ([S]mall, [M]edium or [L]arge
	-dcto	15	Default Cache Time-out in seconds
	-c		Enable CR/LF translation
	-crlf		Enable CR/LF translation
	-cs		Enable Case-Sensitive filename comparison
	-cl		Check Lowercase filenames only
	-rc		Respect Case when creating files
	-b[VALUE]	8192	Set transfer buffer size in bytes
	-s		Use Network Lock manager
	-rt[VALUE]	5	Set RPC retry count
	-t[VALUE]	1	Set RPC Time-out in seconds
	-r[VALUE	4	Set number of Read Biods (0 - 8)
	-w[VALUE]	0	Set number of Write Biods (0 - 8)
	-a		Set Archive option
	-um[VALUE]	600	Set umask in octal (000 - 777)

Options		Default	Description
	-fcbits[VALUE]	700	Set directory create permission bits in octal (000 - 777)
	-dcbits[VALUE]	700	Set directory create permission bits in octal (000 - 777)
	-f		Set default attributes for migrated MVS datasets
	-tr		Enable Tracing
	-?		Display command Syntax

Examples:

```
MOUNT Z: catch22:/home/catch22

   Will mount /home/catch22 on server catch22 on Z:

MOUNT -b512 k slipserv:d:\karl

   Will mount d:\karl on server slipserv on K: with a transfer
   buffer size of 512 bytes

MOUNT -u312 -g1 k aix03:/home/wayne

   Will mount /home/wayne on server aix03 on K: with a UID of 312
   and a GID of 1

MOUNT -v v vm1:myid,191,ro,u=myid

   Will mount myid on server vm1 on V: and pass the values 191,
   ro and u=myid to the server

MOUNT -tr -c -lbruce -pmypasswd m: nfsserv:/usrs/bruce

   Will mount /usrs/bruce on server nfsserv on M: with tracing
   enabled, CR/LF translation enabled and use userid bruce with
   password mypasswd
```

7.7.3.2 UMOUNT command

To detach or unmount an NFS share, use the `UMOUNT` command. The syntax is as follows:

```
UMOUNT <drives>
```

Examples:

```
UMOUNT D:

   Will detach D:

UMOUNT DEFG
```

```
   Will detach D:, E:, F: and G:

UMOUNT *
   Will detach all NFS shares
```

7.7.3.3 SHOWEXP command

The SHOWEXP command can be used to see the EXPORTS file from a certain host. The EXPORTS file is essentially a file that describes the shares available and who has access to it.

The syntax is as follows:

```
SHOWEXP [-m] <hostname>
```

-m causes SHOWEXP to pause if the output does not fit on a single screen.

An example of the SHOWEXP command follows:

```
[C:\]showexp w3.rs6000.ibm.com

IBM NFS for OS/2
SHOWEXP Version 3.99
Release:  m27

Export list for w3.rs6000.ibm.com:
/web/mktmatbp       everyone
/web                everyone
/web/mktmat         everyone
/web/logs           everyone
/web/mktmatilnk     everyone

[C:\]_
```

The output is in the format <share> <users>

If the user entry is *everyone*, there is public access to that NFS share.

7.7.3.4 SHOWMOUN command

The SHOWMOUN command can be used to see which hosts have mounted NFS shares on a given server.

The syntax is as follows:

```
SHOWMOUN <hostname>
```

An example of the SHOWMOUN command follows:

```
[C:\]showmoun w3.rs6000.ibm.com
george.austin.ibm.com : /web/mktmatbp
cs1.austin.ibm.com : /web/mktmatbp
cs1.austin.ibm.com : /web/mktmatilnk
cs1.austin.ibm.com : /web/logs
cs1.austin.ibm.com : /web/mktmat
cs1.austin.ibm.com : /web

[C:\]_
```

The output is in the format `<hostname> : <share>`

7.7.3.5 RPCINFO command

The `RPCINFO` command is not part of NFS itself but is very useful to it. It allows you to query the RPC services registered at the PORTMAP service running on a remote machine, and NFS is such a service.

Each RPC service has a program number, name and a version number. In the case of NFS, the program number will be 100003, the name nfs and the version 2 or 3.

Query the status of an RPC service on a host

This queries the status of a given RPC service and optional version number on a given host. The syntax is as follows:

```
Using UDP packets:
RPCINFO [-n portnum] -u host prognum [versnum]
Using TCP packets:
RPCINFO [-n portnum] -t host prognum [versnum]
```

An example from an OS/2 machine using UDP packets follows:

```
[C:\]rpcinfo -u localhost nfs
program 100003 version 2 ready and waiting
program 100003 version 3 ready and waiting

[C:\]_
```

Query all RPC services on a host

The syntax for querying all RPC services running on a host is as follows:

```
RPCINFO -p [host]
```

The hostname is optional; if it is not given, the localhost will be queried.

An example from an OS/2 machine follows:

```
[C:\]rpcinfo -p

   program vers proto   port
    100000    4   tcp    111  portmapper
    100000    3   tcp    111  portmapper
    100000    2   tcp    111  portmapper
    100000    4   udp    111  portmapper
    100000    3   udp    111  portmapper
    100000    2   udp    111  portmapper
    100001    1   udp   1034  rstatd
    100001    2   udp   1034  rstatd
    100001    3   udp   1034  rstatd
    100002    1   udp   1035  rusersd
    100002    2   udp   1035  rusersd
    100008    1   udp   1036  walld
    100012    1   udp   1037  sprayd
    150001    1   udp   1038  pcnfsd
    150001    2   udp   1038  pcnfsd
    100083    1   tcp   1025  ttdbserver
    100068    2   udp   1039  cmsd
    100068    3   udp   1039  cmsd
    100068    4   udp   1039  cmsd
    100068    5   udp   1039  cmsd
    100003    2   udp   2049  nfs
    100003    3   udp   2049  nfs
    100003    2   tcp   2049  nfs
    100003    3   tcp   2049  nfs
    200006    1   udp   2049
    200006    1   tcp   2049
    100005    1   udp   1062  mountd
    100005    2   udp   1062  mountd
    100005    3   udp   1062  mountd
    100005    1   tcp   1047  mountd
    100005    2   tcp   1047  mountd
    100005    3   tcp   1047  mountd
    100024    1   udp    755  status
    100024    1   tcp    755  status
    200001    1   udp    756
    200001    1   tcp    756
    200001    2   tcp    756
    100021    1   udp   1102  nlockmgr
    100021    2   udp   1102  nlockmgr
    100021    3   udp   1102  nlockmgr
    100021    4   udp   1102  nlockmgr
    100021    1   tcp   1049  nlockmgr
    100021    2   tcp   1049  nlockmgr
    100021    3   tcp   1049  nlockmgr
    100021    4   tcp   1049  nlockmgr
1342177279    3   tcp   1053
1342177279    1   tcp   1053

[C:\]_
```

Broadcast an RPC query

The syntax for broadcasting an RPC query on the local network and asking for all hosts running a certain RPC service and version number is as follows:

```
RPCINFO -b prognum versnum
```

Example:

```
RPCINFO -b nfs 3
```

They will return a reply from the PORTMAP service of all machines on the local network running an NFS version 3-compatible RPC service in the format:

```
<ip address> <hostname>
```

This is useful to locate all servers on the network that have a certain RPC service, such as NFS that you might want to use.

7.7.4 Automatically mounting NFS shares on start-up

To automatically mount NFS shares on bootup, the shares to be mounted need to be placed in the FSTAB.INI file located in \MPTN\ETC. This file has the format of a batch file, and the commands in it are run by the NFSSTART service when this service is started.

Each line in FSTAB.INI can contain one mount command. The lines are exactly the same as for mounting an NFS share from a command line as described in Section 7.7.3.1, “MOUNT command” on page 299.

Example FSTAB.INI entries follow:

```
MOUNT Z: catch22:/home/catch22   #mount /home/catch22 on catch22
MOUNT -b512 k slipserv:d:\karl   #mount d:\karl on slipserv
```

Everything behind the hashmark (#) will be considered a comment and discarded when processing the file.

7.7.5 Sharing an NFS share using LAN Server

Once an NFS share has been mounted onto a local drive, this drive can be shared using LAN Server.

The procedure is simply to mount an NFS share as described in Section 7.7.3.1, “MOUNT command” on page 299. Then, create a LAN Server share using either NET SHARE or the *LAN Server Administration* program.

7.7.6 Creating your own NFS shares

First, the PORTMAP service must be started as described in Section 7.7.2.1, “PORTMAP service” on page 297. Then shares can then be created using the TCP/IP Configuration notebook as described in Section 7.3.11, “NFS tab” on page 276. These settings are saved in the EXPORTS file in the \MPTN\ETC directory and processed by NFSD when started.

Making changes to this requires that the NFSD service is restarted. If the NFS service was started as a normal process, simply terminating it from the tasklist and starting NFSD from an OS/2 command line will be sufficient. If NFSD was started detached, a separate process killer will have to be used to terminate NFSD or the computer will need to be restarted.

7.8 TIMED

TCP/IP 4.2 and 4.21 include a Time Server service (TIMED). The time server implements the time protocol as described in RFC 868. The TimeD Protocol returns the time in seconds since midnight Jan 1, 1990 to any client that sends a datagram.

The function of TIMED in TCP/IP for OS/2 is primarily to enable NetStations to retrieve the current date and time. No client service is provided with OS/2.

It can be started from an OS/2 command line by issuing TIMED, or autostarted on the Autostart tab of the TCP/IP Configuration Notebook.

The syntax is as follows:

```
TIMED [-s <portnumber>] [-le] [-lt] [-t <trace file path]
```

Table 61. TIMED parameters

Option	Description
`-s <portnumber>`	Alternative port to listen on for time requests (default is 525)
`-le`	Specifies the network byte order as zero
`-lt`	Adjusts the time to the local time of the server
`-t <trace file path>`	Generates trace records at the location specified

7.9 FTPD improvements

This chapter will only describe the new features in FTPD; knowledge of FTPD is, therefore, assumed.

FTPD has been improved by making it multithreaded and adding support for restarting broken connections.

7.9.1 FTPD multithreading

FTPD has been moved from a multiprocess model to a new, enhanced, multithreaded model that provides faster connection response time and less memory overhead as requests are processed concurrently. Instead of starting a separate instance of FTPDC.EXE for each client, servicing each client with a separate thread improves performance.

For this, a new command line parameter has been added to FTPD.

```
-m[maxthreads]
```

[maxthreads] would be the maximum number of threads that can be used by FTPD. The default is 200, and the minimum value is 10.

7.9.2 Restarting broken connections

The new FTPD and FTP-PM support restarting broken data transfers (reget). The implementation is compatible with the implementations commonly found on Linux and Sun.

As an example, when using the command line version of FTP and the connection to the server is broken during a download, re-establishing the connection and issuing a reget of the file will continue the download where it left of.

Note

To successfully re-establishing a broken connection, the function needs to be supported by both the client and server.

The OS/2 command line FTP does not, for example, support this function while the new version of FTP-PM does.

7.10 Line printer improvements

This chapter will only describe the new features in LPD and LPRPORTD, knowledge of LPD and LPRPORTD is therefore assumed.

Streaming support has been put into LPD and LPRPORTD to allow net stations to print to an OS/2 print server and to allow an OS/2 client to print to a NetStation.

Also, security has been added to LPD to prevent unauthorized hosts from printing.

7.10.1 Streaming LPD

The LPD on an OS/2 print server will automatically detect that streaming is used. No additional parameters or configuration are available or needed.

7.10.2 Streaming LPRPORTD

For LPRPORTD on an OS/2 client, it is necessary to indicate that streaming needs to be used when printing to a NetStation. For this, a *Data Transfer Mode* setting has been added to the configuration page for the LPR port.

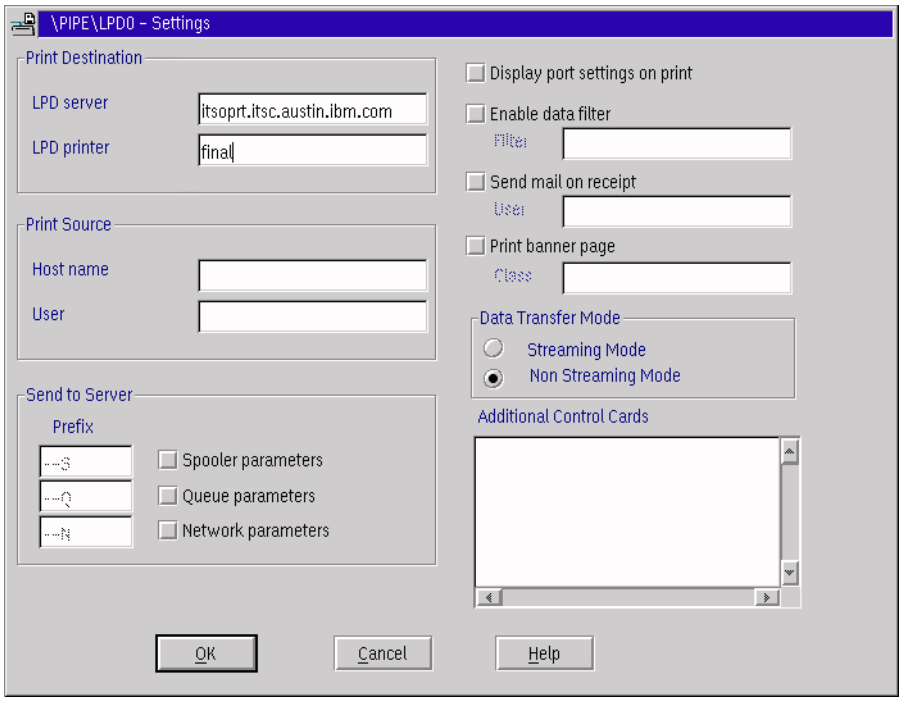

Figure 134. \PIPE\LPD0 settings

Note

Only LPRPORTD for clients has been enabled for streaming mode. Streaming mode is not supported with LPR or LPRMON.

7.10.3 LPD security

Security has been implemented in LPD to allow or disallow certain hosts.

If this is required, a file, HOSTS.LPD, needs to be created in the ETC directory referenced in the SET ETC= environment variable in CONFIG.SYS (usually \MPTN\ETC). If this file is not present, all hosts will have access to all queues.

The format of entries in the HOSTS.LPD file is as follows:

\<queuename> [-]<hostname1>;[-]<hostname2>

<queuename> is the name of the printer queue on the server. A star (*) can be used to represent all printer queues on the server.

<hostname> is a host that is allowed, or, if there is a minus in front of the hostname, disallowed. Multiple hostname entries can be made by using semicolons between them. A star (*) can be used to allow or disallow certain domains.

Examples of entries in the HOSTS.LPD file follow:

```
\final k123.austin.ibm.com;-k125.austin.ibm.com
   This will allow k123.austin.ibm.com to print to printerqueue
   final, but will disallow k125.austin.ibm.com.
\draft *.austin.ibm.com
   This will allow all hosts that are in the austin.ibm.com domain
   access to printerqueue draft.
\* *.itsc.austin.ibm.com;-*.raleigh.ibm.com
   This will allow all hosts that are in the itsc.austin.ibm.com
   domain access to all printerqueues on the server, but will
   disallow anyone from raleigh.ibm.com
```

Note

If security for a queue has been enabled in HOSTS.LPD, only the user that created a print job will be allowed to cancel it. This is checked based on the hostname and username submitted for creating the printjob.

7.11 TFTPD improvements

This chapter will only describe the new features in TFTPD; knowledge of TFTPD is, therefore, assumed.

To better support WorkSpace On-Demand and NetStations with BOOTP, the TFTPD server has been made multithreaded and now supports block sizes up to 8KB.

These changes have the effect that TFTPD can no longer be started from INETD.

Also, a level of security has been implemented to only allow read or write access to certain directories and from certain hosts and to only allow incoming connections on a certain TCP/IP interface.

7.11.1 TFTPD multithreading

For this, two new parameters have been added to TFTPD.

```
-mn[threads]
```

This indicates the minimum number of threads in the pool. The default is 10, minimum is 1 and maximum is the number of threads defined with `-mx`.

```
-mx[threads]
```

This indicates the maximum number of threads that can be used by `TFTPD`. The default is 100, and the minimum cannot be less then the `-mn` value.

7.11.2 TFTPD blocksize

The old block size used to be 512 bytes. This value is negotiated between the client and server, and, for this, a new parameter has been added to TFTPD to indicate the maximum blocksize that should be supported.

```
-b[bytes]
```

The range is 512 to 8192 with a default of 8192 bytes.

7.11.3 TFTPD security

A TFTPD subtab has been added to the Security tab in the TCP/IP Configuration Notebook as shown in Section 7.3.8.3, “TFTPD sub-tab” on page 264. This allows the specification of directories and hosts that should be allowed to access them.

TFTPD also has a new command line parameter that can be used to limit access to the TFTPD server from a certain TCP/IP interface.

```
-a[interface]
```

[interface] would be the IP address of the local interface from where incoming TFTP connections should be accepted. Default is to allow access via all interfaces.

7.12 IP aliasing

IP Aliasing has been available in TCP/IP for OS/2 since version 4.0 but has never been discussed before.

IP Aliasing allows to assign multiple alias IP addresses to an interface. This can be useful when there are multiple IP networks on the same physical LAN or when Multi-Home is needed for a Webserver.

Multi-Home is the technique to have multiple homepages on the same physical server.

For example, you rent out Web space and have been contacted by 2 companies.

- Company A wants a Web site called `www.company-a.com`.
- Company B wants a Web site called `www.company-b.com`.

To save money on hardware and administration, you want to place both Web sites on the same physical server with just one LAN interface. However, customers looking for the Web site of Company A should not get the Web site of Company B, and the reverse is true.

To solve this problem, multiple IP addresses can be assigned to an interface using the ALIAS parameter.

As an example, we are first going to configure LAN0 with an ip address:

IFCONFIG LAN0 10.1.2.3 NETMASK 255.255.255.0

Then, we assign other IP addresses to the same interface using the ALIAS parameter:

IFCONFIG LAN0 ALIAS 10.1.2.4

This command can be repeated for assigning more IP addresses to the interface.

To remove an ALIAS from an adapter, use the -ALIAS parameter.

IFCONFIG LAN0 -ALIAS 10.1.2.4

7.13 X Windows

The PMX product is not included with OS/2 Warp Server for e-business, and the old PMX 2.0 product is no longer available or supported with TCP/IP 4.21.

However, if migrating an existing OS/2 installation that already has PMX installed, it will be migrated.

This chapter will not go into great depth about PMX. If more information is needed, consult the redbook *X Window System Server Guide*, SC31-7070.

7.13.1 Installing PMX over TCP/IP 4.21

If you decide to install PMX after you have installed TCP/IP 4.21, you will have to manually transfer the file TCPCFG.EXE from the \TCPIP\BIN\SAMPLES directory to \TCPIP\BIN and optionally create an icon for it.

Make sure, however, that you do not install any other components from TCP/IP 2.0. Then, install PMX and, if needed, the IBM Library Reader, and apply the latest fixes to it from:

`ftp://ftp.software.ibm.com/ps/products/tcpip/fixes/v2.0os2/pmxfixup/`

Note

When doing a clean install of PMX, or, in some cases, an upgrade from TCP/IP 4.0 or up, you might find that the on-line help for PMX no longer works.

This is due to the fact that COMBO.DLL is no longer part of TCP/IP since version 4.0. The solution is to install the latest PMX fixes to 2.04g since this, again, includes a COMBO.DLL.

If problems are still encountered even after installing this fix, make sure that the directory where the PMX help files (PMX.HLP and CNBPMX.HLP) are located is in your SET HELP= statement in CONFIG.SYS.

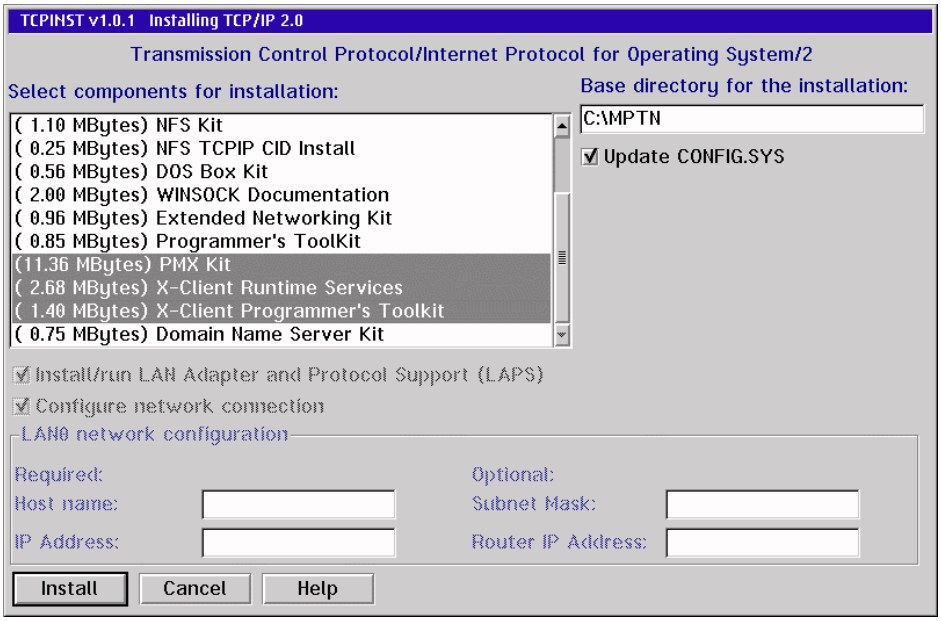

Figure 135. Installing PMX from TCP/IP 2.0

7.13.2 Configuring PMX

There are three ways to configure PMX.

7.13.2.1 Old-style Configuration Notebook

Using TCPCFG.EXE

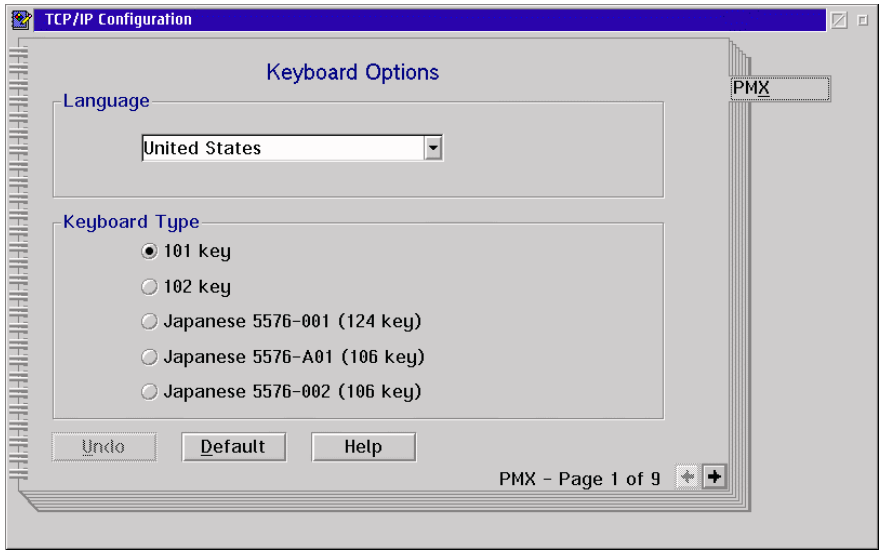

Figure 136. Old-style configuration notebook: PMX tab

Note

If the Configuration Notebook comes up completely empty, make sure you are running at least version 2.04 of PMX.

If you have an older version, install the updates available from:
`ftp://ftp.software.ibm.com/ps/products/tcpip/fixes/v2.0os2/pmxfixup/`

7.13.2.2 From the X-Server

When the X-Server is running, select the **Commands** menu; from there, select **Configuration** and, then, select either **Initial Settings** or **Current Settings**.

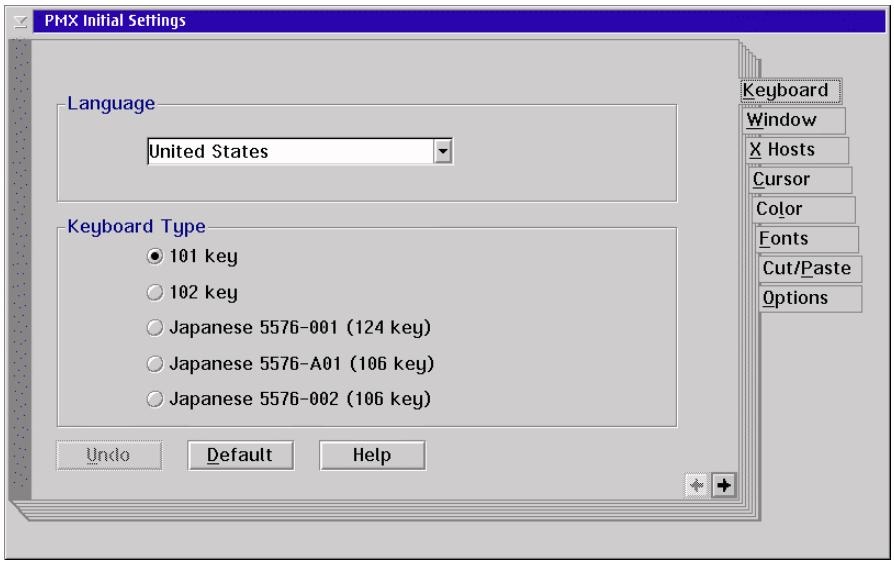

Figure 137. PMX configuration: Initial settings, keyboard

7.13.2.3 Manually editing the configuration files

Manually edit the configuration files by manually editing the TRUSERS file.

Note

If manually editing the TRUSERS file, it will no longer be in sync with TCPNBK.LST, and any changes made using the Configuration Notebook will overwrite the changes you made manually.

7.13.3 Alternatives to PMX

There are two alternatives to PMX

- Hummingbird Exceed for OS/2**

  ```
  This is a commercial product from Hummingbird Communications LTD.
  More information can be found at http://www.hummingbird.com/
  ```

- XFree86/OS2

  ```
  More information on XFree86/OS2 can be found at:
  http://borneo.gmd.de/~veit/os2/xf86os2.html
  Or more general information on XFree86 can be found at:
  http://www.xfree86.org/
  ```

7.14 TCP/IP development toolkit

The toolkit for OS/2 Warp Server for e-business is available as an update to the OS/2 Warp 4.0 toolkit. As such, the OS/2 Warp 4.0 toolkit needs to be installed before the OS/2 Warp Server for e-business toolkit can be installed.

These toolkits are available from DevCon (Developer Connection).

More information on DevCon can be found at the following Web site:

`http://www.developer.ibm.com/devcon/`

7.15 Performance improvements

Many of the enhancements introduced into OS/2 Warp Server for e-business have been for performance and reliability. TCP/IP and the related applications have also been enhanced to more quickly and reliably serve in this e-business transformation. In certain cases, applications taking advantage of the new enhancements will show up to a 40 percent improvement in performance over previous versions of TCP/IP.

The sockets drivers now have two parameters to improve performance. For example:

DEVICE=x:\mptn\protocol\sockets.sys /mem:# /gdt:#

Where

/mem:# is the number of 4KB clusters allocated at initialization time. The default is 75 and the range is 30 to 32766.

and

/gdt:# is the maximum number of 64KB blocks that the stack can allocate. The default is 80.

7.15.1 32-bit stack enhancements

As described earlier, a new ring 0 library enables applications running at ring 0 to make direct calls to the protocol stack. This will dramatically improve the performance of those applications since, now, there will be no ring transitions.

7.15.2 inetcfg

The `inetcfg` command configures or retrieves a current value of a TCP/IP parameter or restores a parameter to its default value.

To simplify the process you can run the command:

```
inetcfg-get all
```

This will retrieve the current parameters from the system and place them in the x:\MPTN\ETC\inetcfg.ini file. If you then wish to change a few of the parameters settings, you can change them in this file and issue the command

```
inetcfg-set all
```

`set all` takes parameters from the ETC\inetcfg.ini file and sets them.

7.15.3 SYN cookies

This is a feature that is, by default, set **on**. It is an enhancement to TCP that is designed to make the protocol stack safe from incoming SYN attacks. When a connection is initiated, that is, on the reception of a SYN frame, memory resources are allocated in the protocol stack memory. A SYN attack occurs when a malicious person simply floods the stack with SYNs resulting in the stack running out of memory.

By using a SYN cookie approach, the information received from the client is stored in an MD5 message digest and a SYN-ACK is generated with the Initial Sequence Number as this digest. If the client is legitimate, there will be a reply, and the sequence number can be verified and resources allocated. Or else, there are no stack level resources wasted.

To turn off this feature, use the command `inetcfg -set syncookie 0`. To turn on this feature, use the command `inetcfg -set syncookie 1`.

7.15.4 Reuse timewait state enhancements

Reuse Timewait tries to reuse the stack resources that are in a Timewait state whenever possible. This resource utilization is based on a matching resource found in the timewait state. The match is done based on local address, local port, and foreign address.

This performance enhancement provides optimal memory usage and faster access to connection resources. The Reuse Timewait deviates from RFC specifications of TCP, but it does not affect the normal operation of TCP.

It can be turned off by using the command `inetcfg -set reusetw 0`. It can be turned on by using the command `inetcfg -set reusetw 1`. The default is on.

7.15.5 HTTP fast path performance

A new feature called the Fastpath for HTTP connections has been added. This improves the performance of HTTP (WEB) servers that bind to port 80. What the server does is when the HTTP server starts up and binds itself, a set of pre-fabricated data structures will be allocated for the future connections to the server. When a new connection is established to the server, it is checked. If this connection is intended for port 80 it uses the *fast path*, that is, using or reusing the already fabricated data structures rather than by going by the usual way of allocating resources only after the requests arrive.

Performance gains are realized by the server pre-fabricating or reusing existing data structures.

It can be turned off by using the `inetcfg -set perfhttp 0` command. It can be turned on by using the `inetcfg -set perfhttp 1` command. The default is on.

If the HTTP server binds to port 80 with the Fast Path HTTP flag ON, then the flag cannot be turned OFF during the period the server remains bound to the specified port. So if you issue the command **inetcfg -set perfhttp 0,** it will only go into effect after the server is restarted.

7.15.6 SMP exploitation

The current OS/2 TCP/IP stack does not scale well on an SMP kernel. This is due to the fact that the OS/2 kernel provides serialization for legacy device drivers on SMP since they are not SMP enabled. The OS/2 SMP kernel provides the serialization that protects the device drivers's data structures from data corruption (concurrent access by multiple threads on SMP) by using a single system-wide spin lock. Because a single spin lock was employed, only one processor can be executing code in a device driver at any one point in time. This leads to poor SMP scaling when there are multiple threads needing to execute in the device driver. The other threads must wait (spin) until the owning processor releases ownership of the SUBSYS spin lock.

To correct the scaling problem, the stack device drivers will use the new 32-bit SMP-enabled Kernel Execution Environment (KEE). KEE provides a set of system service API's that will enable the TCP/IP stack to protect its own data structures with more granular spin locks instead of using the system wide spin lock. This improves SMP scalability of the stack. Due to the nature

of the KEE APIs, performance is improved on both UNI and SMP processor-based machines.

7.15.7 New API calls

New send_file() and accept_and_recv() APIs have been implemented to drastically improve performance. Webservers and other TCP/IP applications can use these APIs to gain significant performance enhancements.

7.15.7.1 send_file()

The main aim of this API is to improve the performance of applications, such as Web Servers that send file data over a connection associated with a socket.

Typically, an application will read the file data from the file system cache to its application buffer (copy no.1) and, then, make a socket API call to send the data. The socket's device driver will copy the file data from the application buffer to a network buffer (copy no. 2).

This new API allows an application to request the sockets device driver to send file data directly from the file system cache to the network resulting in zero buffer copies.

Zero buffer copies will only be achieved on file systems that implement the new IFS entry points. At this point in time, JFS is the only file system that will API results in one buffer copy (read data from file system cache to network buffer) when called for files on legacy file systems, such as FAT, HPFS, and so on.

For this reason, when it is best to hold most of the TCP/IP application data on volumes formatted for JFS, such as web pages, web graphics, and so on.

7.15.7.2 accept_and_recv()

Most socket applications, such as Web Servers, follow a common scenario. They wait for a connection, accept it, then get the first message block from the client and take action according to the received data. This involves several socket API's and kernel transitions which reduces the performance of applications where they need to accept a connection and, then, the first block of a message.

Accepting a connection is usually done by the accept() API. The application then waits for the first block of data and receives it by the recv() API. The accept_and_recv() API combines the socket functions accept() and recv() into a single API/kernel transition. This function accepts a new connection,

receives the first message block from the client and returns the local as well as remote addresses to the application. All this happens in exactly the same time as an accept() call. The thread sleeping on accept_and _recv() wakes up only after it gets the first data block from the client. This not only increases the performance of this thread but also the performance of the system.

7.15.8 Variable cluster sizes

The data that travels from a process to the network interface and from a network interface to a process is held in memory buffers called mbufs. The size of an mbuf is fixed at 256 bytes. To enable the handling of large amounts of data, an external buffer called a cluster is used. The performance of the networking protocols is directly related to the memory management scheme used within the kernel.

Within this version of TCP/IP, the clusters can be made flexible to provide varying sizes of 512, 1024, 2048, and 4096 bytes. Based on the amount of data to be transferred, the appropriate cluster is automatically selected.

This aids in better memory utilization and performance.

Chapter 8. Lotus Domino Go Web Server and WebSphere

The objective of this chapter is to introduce you to the functions and facilities provided in Lotus Domino Go Server and the WebSphere Application Server. In this chapter, we will attempt to indicate the ease of implementing these two products, covering fastpath installation, discussions, and creating sample applications that could perform useful functions for most installations. At the same time, we attempt to illustrate what little effort it takes to enter the arena of processing on the Web.

The explosive interest in the Web, the Internet, intranets, and so on is driven by customer demand to use the internet as a sales, communication, and delivery channel in the marketing of goods and service. It is clear that Web servers must evolve beyond simple HTTP protocol servers into Web application servers that let you use the Web to access data residing on various systems. Existing built-in connectors and a runtime Java environment make it easy to connect to almost all of your existing systems and add new ones.

8.1 Lotus Domino Go Webserver

The Domino Go Webserver provides the foundation for building your company's presence on the Internet and smooths the transition into the 21st century with reliable support for the year 2000 and beyond. Your business can use the Domino Go Webserver on the World Wide Web to reach customers and suppliers around the world. Or, you can use the Domino Go Webserver within your business to communicate with employees.

There are three editions of Domino Go Webserver. The editions differ in the level of security provided. The North American Edition contains the strongest security features and may not be exported outside of the United States and Canada. For information on supported key lengths and encryption modes for each edition, see the *Webmaster's Guide,* one of the documents which gets installed with the on-line documentation as part of the installation process.

8.2 WebSphere application server

WebSphere Application Server supports a wide variety of popular platforms by adhering to the most prevalent open standards, such as HTTP, HTML, JSP, JNDI, and IIOP. WebSphere Application Server is based on a Java-based servlet engine that turns your existing Web servers, such as the Apache Server, Microsoft IIS, Netscape Enterprise Server, and Lotus Domino

Go Webserver into Java Web application servers. As a core element of the IBM e-business initiative, WebSphere Application Server allows you to extend the runtine environment with Java classes from IBM and other sources.

WebSphere Application Server includes everything you need to serve up a Web site: simple installation, graphical user interface (for easier servlet management), Web-based remote administration and security features. It supports standard Java servlets that can exploit services, such as session state, user profiles, high-performance database access with connection management, as well as access to other sources, such as CICS and IMS using the IBM connectors. Release 1.1 includes support for the industry-standard Enterprise Java Beans (EJB) programming model. CORBA support includes code generators that allow developers to start from either IDL or Java interfaces, Java implementation of the CORBA Naming Server, object creation and factory service, and the ability to run object servers as a servlet in the Web server.

8.3 Fastpath Install - Lotus Domino Go Webserver

There are two editions of Lotus Domino Go Webserver. They are:

- The North American Edition referred to as the U.S./Canadian version
- The International Edition referred to as the International export version

The editions differ in the level of security provided. The North American Edition contains the strongest security features and may not be exported outside of the United States and Canada. For information on supported key lengths and encryption modes for each edition, see the *Webmaster's Guide*.

The Lotus Domino Go Webserver installation software CD is included with OS/2 Warp Server for e-Business. Insert this CD in your CD drive, and open an OS/2 window or full screen session. Change to directory lotusgo. Run the `install` command to initiate the installation process. A window, as shown in Table 138 on page 323, is displayed on your terminal.

Figure 138. Domino Go Webserver, Installation Panel with instructions

You can scroll down the Instructions for Installation window which contains updates or changes to the existing installation documentation. Once you have read this information, click on **Continue** to proceed with the installation, or click on **Cancel** to abort the installation.

When you click on **Continue**, you are presented with a window on which you can select or deselect the option of having the server's config.sys updated during the install or not.

You also have the choice of cancelling the installation, obtaining help information, or clicking on **OK** to continue with the installation. Once you click

on **OK**, you are presented with a window in which you select which components you wish to install, as in Table 139 on page 324.

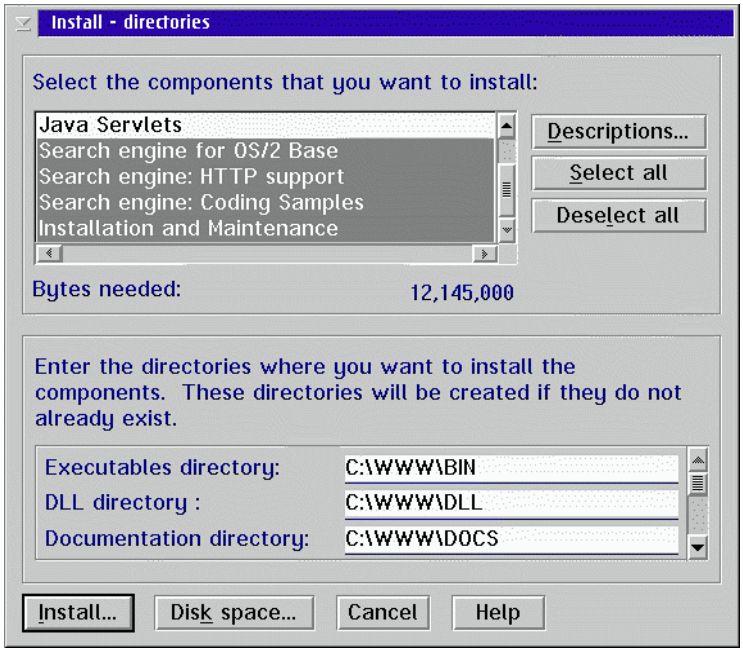

Figure 139. Installation screen, selecting the components and their locations

You can, at this stage, also change the default drive and destination directories where the software will be installed. You have a select box to provide you with a description of the component you can select to install to assist you in your decision if you want or need to install the particular

component. Other select boxes provide facilities to select all components or to deselect all components.

Table 62. Domino Go Webserver, selectable components

Select this package...	To install...
Lotus Domino Go Webserver	The base server files.
Security Files	Files for the secure server.
Java Servlets	Configuration, class and library files for Java servlets.
Search Engine for OS/2 and Search Engine: HTTP support	The optional search engine. You must install both components.
Search Engine: Coding Samples	Code samples for the search engine. The code samples are optional and are useful only if you intend to modify the search engine base functions.
Installation and maintenance	The server installation utility.

Note

If you are going to install WebSphere Application Server, you must deselect Java Servlets from the list of components. The Java servlet support in DGW was the prototype for the current WebSphere App Server. Both are Java servlet support plug-ins to the Go webserver that cannot be loaded and work at the same time because they do the same thing, namely, http requests get passed to the Java servlet engine via the WAS (or DGW Java servlet support) plug-in according to the pass rules specified in the DGW httpd.cnf file. In addition, both plug-ins run in-process (run in the same process as Go), and both create an instance of the JVM. Two instances of the JVM cannot be created in the same process.

If you clicked on **Cance**l, the installation is terminated. If you click on **Disk-space** after making your selections, the window shown Table 140 on page 326 is displayed. The Disk-space select box also indicates how many bytes are needed to install the components you selected. It is recommended that you install on an HPFS volume. Many of the distribution files have long filenames.

Refer to Chapter 3 of the *Lotus Domino Go Webserver: Quick Beginnings* guide for a description of the contents of each of the installation directories.

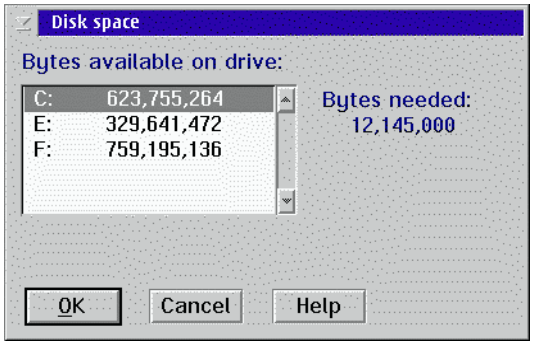

Figure 140. Domino Go Webserver, available disk space.

If you click on **Install**, the window shown in Table 141 on page 327 is displayed. Here, you are presented with installation default values, information required by the installation process, such as the hostname of your server, the key ring file name, HTTP port number, and the SSL port number to be used by the WebServer code. On the same window, you have to supply the administrator ID and password. Take note of these two items. The configuration values are listed in Table 63 on page 327.

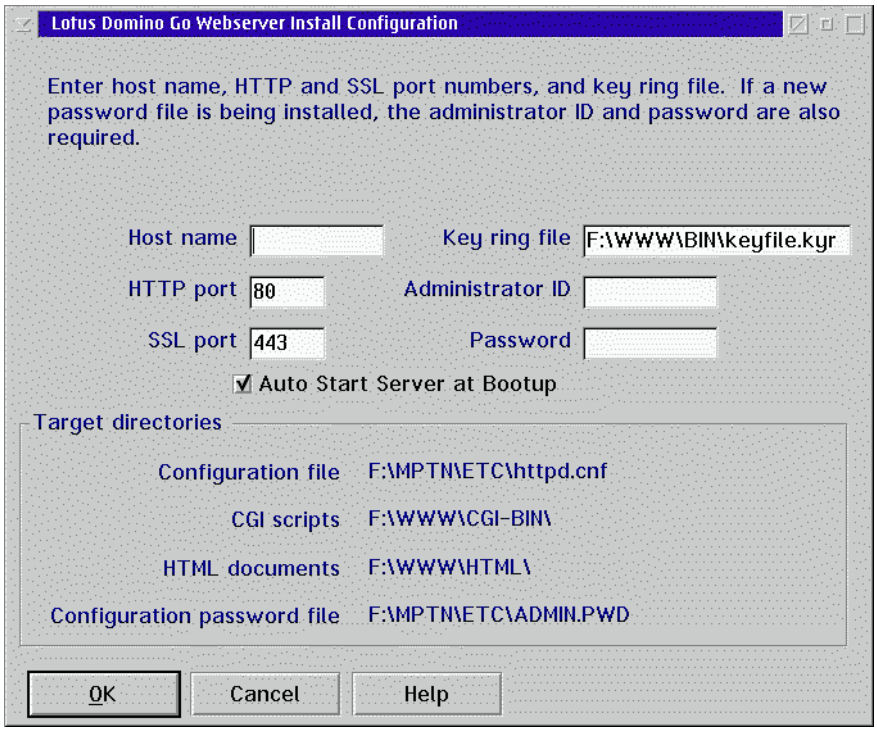

Figure 141. Domino Go Webserver, configuration panel

Table 141 on page 327 and Table 63 below show the configuration panel and corresponding descriptions of each of the configurable parameters.

Table 63. Domino Go Webserver, configuration values

Configuration Value	Description
Host Name	The default value is the host name defined in your CONFIG.SYS file. If you want to use an alias, you can change this field to a fully-qualified host name that is defined in your domain name server.
HTTP Port	The default value of 80 is the well-known port number for Hypertext Transfer Protocol(HTTP). Other port numbers less than 1024 are reserved for other TCP/IP applications. Port numbers 8080 and 8008 are commonly used for testing servers.

Configuration Value	Description
SSL Port	This is the port you want your server to listen to for requests for documents protected by the Secure Sockets Layer (SSL) protocol. The default is 443.
Key ring file	This is the name of the file where you want to store public-private key pairs that the server can use for secure communications. The key ring file that is displayed is taken from your previous HTTP.CNF file if you had one. The default directory is `\WWW\BIN\keyfile.kyr`
Administrator ID	This is the ID of your server administrator. Anyone attempting to use the server Configuration and Administration forms will be prompted to enter this ID. There is no default value. Unless you are using an admin.pwd file from a previous installation, you must specify an Administrator ID
Password	This is the password you use to protect access to the Configuration and Administrations forms.

Once the required information has been entered and you click on **OK**, the install progress window is displayed as in Table 142 on page 328.

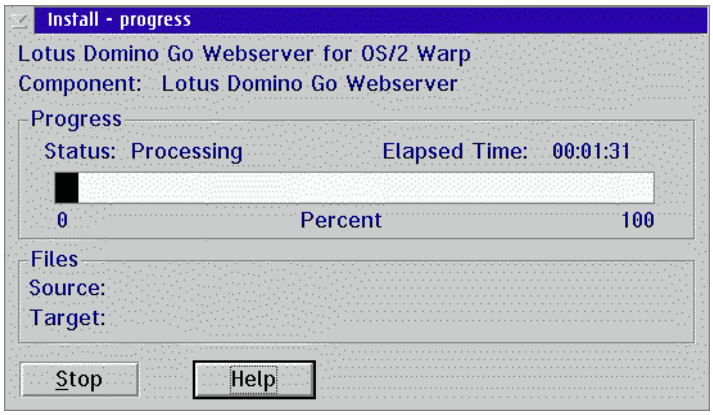

Figure 142. Domino Go Webserver, installation status.

Finally, the window is displayed as in Table 143 on page 329. Click on **OK** to finish the installation.

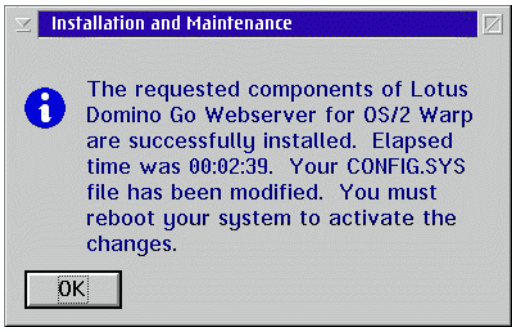

Figure 143. Domino Go Webserver, installation status

8.4 Web server uninstall.

To uninstall Lotus Go Domino Webserver, ensure that the Webserver is not running. If it is, close the application down first. Open the Lotus Go Domino Webserver folder and double click on the installation utility icon as shown in Table 144 on page 329.

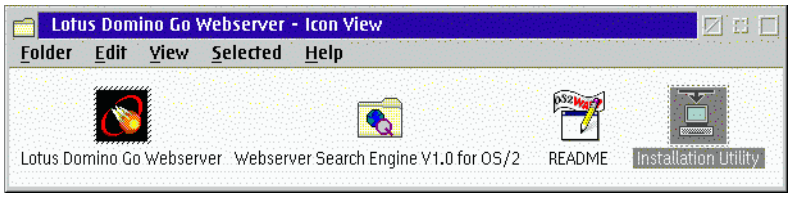

Figure 144. Domino Go Webserver, folder

The uninstall utility will start running and display the window as shown in Table 145 on page 330.

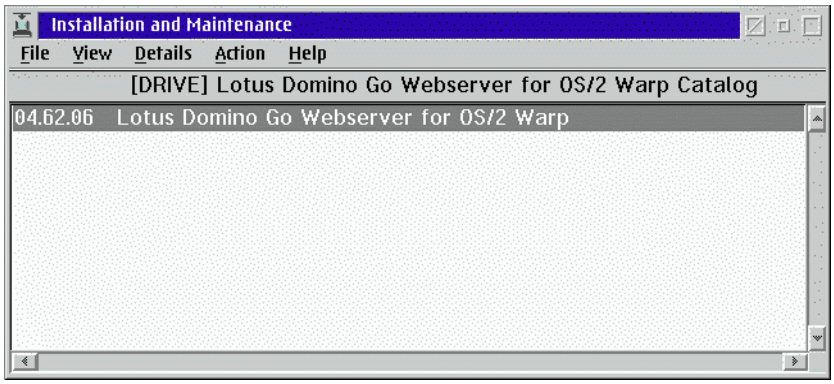

Figure 145. Domino Go Webserver, removing components

Select the **Action** option from the menu bar and then press **Delete** to remove the product. The resulting dialog is shown in Table 146 on page 330, from where you can select the individual components or all the components you want to uninstall. This will enable the Delete action box.

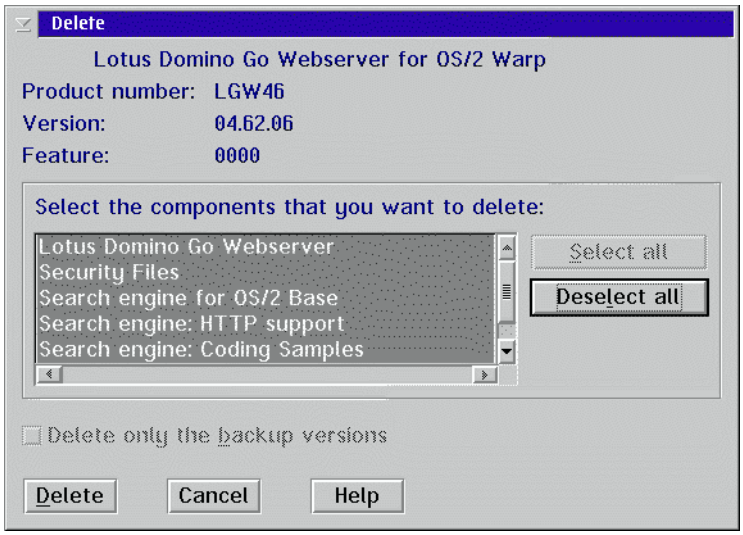

Figure 146. Domino Go Webserver, selecting components for removal

Click on the Delete action box and another window will be displayed, as shown in Table 147 on page 331, with a progress indicator for the selected components as they are deleted.

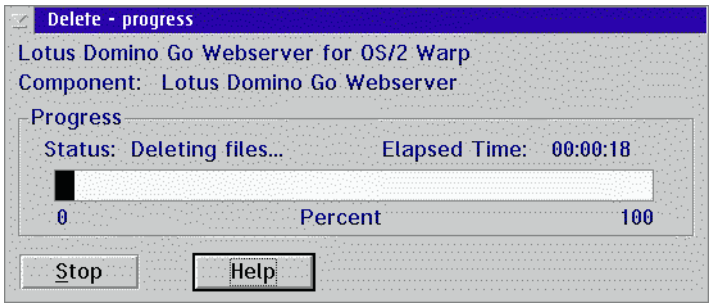

Figure 147. Domino Go Webserver, removal progress

Once all the applicable files have been deleted, a final dialog box is displayed as shown in Table 148 on page 331.

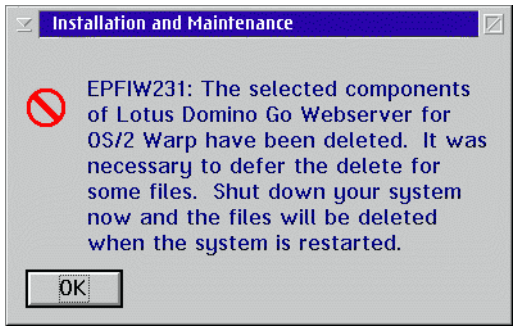

Figure 148. Domino Go Webserver, status message

Click on **OK**, press the **F3** key to exit the installation utility. Final clean-up will be performed once you re-boot the system.

8.5 Functional components

Once the Web server is installed, you may want to tailor the installation to suit your specific requirements. You may want to implement or test some of the functions provided with the Web server. This section briefly discusses these functions or directs you where to get ample information regarding these functions.

The functions we will be discussing here are:

- Tailoring your Webserver to suit your needs
- Using CGI programs, GWAPI programs, and Java servlets with the server

- Managing your server with Simple Network Management protocol
- Restricting access to your server
- Mapping resources on your server
- Logging requests and errors occurring on your server
- Running the server as a caching proxy
- Using server-side includes
- Customizing the server's error messages
- Adding a search engine to your Web site
- Using proxy authentication

The following is just a brief discussion of each of these functions. Once the product is installed, there is comprehensive documentation available on-line, which provide you with ample details in the use, tailoring, implementation and management of the Web server. For more details of these functions, refer to the on-line manuals, *Web Programming Guide* and the *Webmaster's Guide.*

8.5.1 Tailoring the Web server

After you install the server, you can start it using the configuration values you specified during installation. Default values are supplied for all configuration values except your administrator ID and password. Using a Web browser, you can then connect to your server. Key in your machine's host ID in the Web browser, for example, `http://your.server.name/` (This is the hostname as defined in the TCP/IP configuration notebook). The Web server will display its front page, which it uses as its home page. If, however, you have defined your own home page and are using the correct naming conventions (use welcome.html as the filename for your home page), your home page will be displayed instead.

The (default) Front Page provides a valuable assortment of tools and information, including links to:

- Configuration and Administration Forms.
  - A tool that lets you configure your server by interacting with a set of forms. The forms are coded with the Hypertext Markup Language (HTML). Another way to configure/tailor your server to suit your specific requirements is to use your favorite editor to edit the HTTPD.CNF file in the \MPTN\ETC directory. Comprehensive configuration detail can be found in the *Webmaster's Guide*.
- Domino Go Webserver Web site.
  - From this Web site, you can read the latest news about the server and download beta versions of software under development.
- Domino Go Webserver on-line documentation.
  - An on-line HTML version of the following:

- Quick Beginnings
- Webmaster's Guide
- Web Programming Guide
- Tuning your Web Server for Better Performance
- Webserver search engine documentation

If you want to go to the Front Page after you create your own home page, simply include the Front Page file name on the URL. The Front Page file name is Frntpage.html; so, you would go to the following Web site to see it:

```
http://your.server.name/Frntpage.html
```

8.5.2 Using CGI programs, GWAPI programs, and Java servlets

The server supports the Common Gateway Interface (CGI), which allows you to create external programs (CGI programs) that interface with your server and perform tasks such as searching and forwarding e-mail messages.

To assist you in writing CGI programs, the server includes utilities for extracting forms data, writing document headers, and processing image maps.

The server also supports the Go Webserver Application Programming Interface (GWAPI). This interface is designed specifically for the server's threaded processing and allows you to easily extend the server's base processing, such as publishing customized pages based on client's code level, problems, or alert you about serious conditions.

Domino Go Webserver also supports Java servlets. In many cases, Java servlets can be used instead of CGI programs and GWAPI programs. We recommend that you write Java servlets whenever possible. Java servlets can provide better performance that other programming alternatives. Moreover, Java servlets are platform-independent and server-independent. Future programming enhancements to Domino Go Webserver will be directed at the Java servlet support, not GWAPI.

Along with Domino Go Webserver, IBM supplies the WebSphere Application Server, which is totally based on Java servlets. If you want to make use of this product, do not install the Java servlets component of Domino Go Webserver.

8.5.3 Managing your Web server with SNMP protocol

The Simple Network Management Protocol (SNMP) subagent built into Domino Go Webserver maintains server information and performance data in an SNMP Management Information Base (MIB). The MIB data describes the

server being managed, reflects current and recent server status, and provides server performance data. From any SNMP-capable network manager, you can display, monitor, and adjust thresholds for your server's performance to proactively tune or fix server problems before they become server outages.

8.5.4 Restricting access

Most likely, you will not want everyone to be able to access all the information on your server. For example, you probably would not want everyone to be able to access CGI programs.

You can restrict access based on user name and password or the address of the requester. Access authorization is controlled by using the configuration file and, possibly, one or more of the following:

- A protection setup, which defines the protection being used
- A password file, which allows you to define user names and passwords
- A group file, which allows you to define groups of user names
- An Access Control List (ACL) file, which allows you to define access for individual files or groups of files on a protected directory

8.5.5 Mapping resources

You can create a virtual hierarchy of Web resources. As part of the server's configuration, you can specify resource mapping rules, which associates a request template with the actual path to a document or resource. Each request that comes to the server is checked against these rules to determine whether the request should be accepted and where the requested resource is actually located.

8.5.6 Logging requests and errors

To help you determine whether or not your Internet message is reaching the intended audience, you can keep access logs that show who is accessing your server and when. To see internal server errors, you can check your server's error logs. If your server is a caching proxy, you can keep logs of requests for cached files in a cache access log.

In addition, your server creates agent logs, referer logs, and CGI error logs. Agent logs indicate which Web browser was used to access a Web page. Referer logs indicate the page that linked to or referred to the page. CGI error logs contain standard error output (stderr) from CGI programs.

You can control what gets logged by filtering out entries that match a certain particular pattern. Your server automatically creates a report for each log; so, you can view the contents of the log. You can modify these default report templates to include and/or exclude log entries contained in the report.

Because the logs are written in a format that is common to most Web browsers, you can use any of several generic statistical programs to analyze the log contents.

You can compress log data, archive reports, and include old log data in reports.

8.5.7 Running the server as a caching proxy

When you configure your server as a caching proxy server, you can improve performance as well as allow users of your internal network to access documents on the Internet.

You can specify many configuration options for a caching proxy including:

- Which files you want (and do not want) to store in cache
- The maximum amount of space allotted to cache storage
- Automatic reclamation of cache storage space
- Routing of requests to other proxy servers

8.5.8 Running server with multiple IP addresses or virtual hosts

You can configure your server to serve different files based on the IP address of the network connection a request comes in on, or the host for whom a request is made. This is particularly valuable to Internet service providers who want to use one server to provide Web sites for multiple customers. For example, you might want to change one of the following based on the IP address:

- Change Welcome Pages to determine how the server responds to requests that do not contain a file name.
- Change Mapping Rules to set map requests to physical files and to determine whether the server processes a request
- Change Access Control to activate different protection rules for requests depending on which address the request comes in or which host name is specified in a URL

8.5.9 Using server-side includes

Server-side includes enable the server to do some processing of Web pages before the page is sent to the client. The current date, the size of the file, and the last change of a file are examples of the kind of information that can be included in Web pages that are sent to the client.

8.5.10 Customizing the server's error messages

The Domino Go Webserver design is based on the work of CERN. You can customize basic CERN messages that your server sends back to the client when error conditions are encountered. For example, you can change a message to include more information about the cause of the problem and suggest possible solutions to fix it. For internal networks, you might provide a contact person for your users to call.

8.5.11 Adding a search engine to your Web site

The Domino Go Webserver includes an optional search engine that enables you to build and maintain searchable indexes to the information on your Web site. The search engine includes a command interface and HTML forms that aid Web site navigation and administration.

8.5.12 Using proxy authentication

The server distinguishes between authorized user IDs and passwords that are used by the end-point server and the proxy server. This allows you to require authentication for proxy requests.

8.6 Fastpath Install - WebSphere

The installation software for WebSphere Application Server version 1.1 is included with the OS/2 Warp Server for e-Business software. To install WebSphere, insert this CD in the CD-drive and open an OS/2 window or full screen session. Change to the drive letter used to access the CD-rom drive and change to the webspher directory. Run the `install` command. A window as shown in Table 149 on page 337 is displayed on your desktop. If you click on **Cancel**, the installation is terminated. If you want to proceed with the installation, click on **Next**.

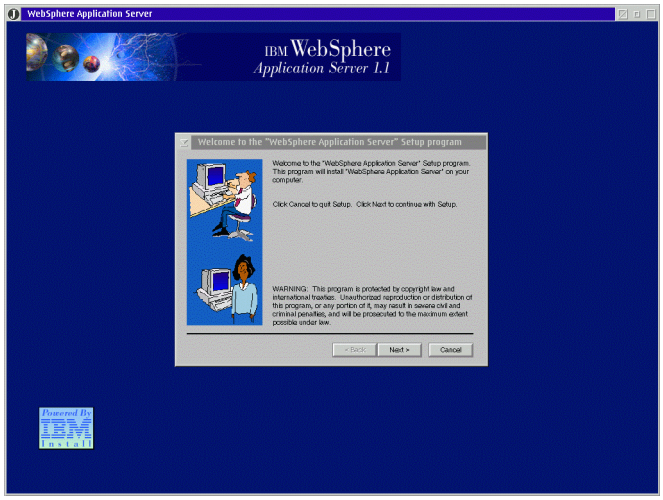

Figure 149. WebSphere, installation

The next window displayed provides the opportunity to choose the destination drive for the installation. You can change the destination drive by clicking on the **Browse** button. A window similar to Table 150 on page 337 is then displayed, which assists in selecting an other than default destination directory.

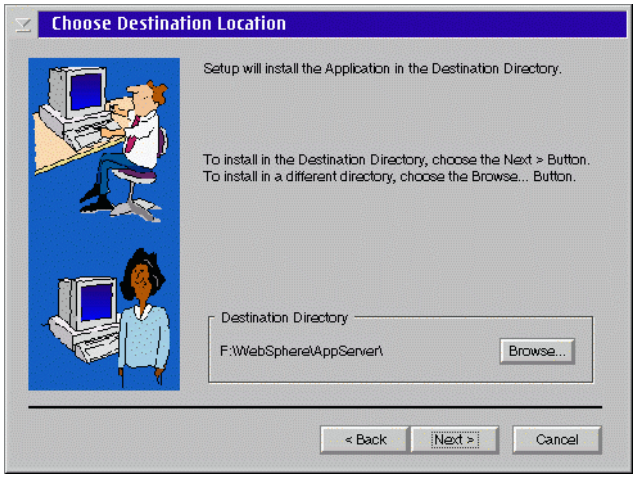

Figure 150. WebSphere, destination

If you click the **Cancel** button, the installation of the software is aborted. You have the option of clicking on **Back** to go back to the previous window.

Once you click on the **Next** button, a window like the one shown in Table 151 on page 338 is displayed. You can select/deselect which components of WebSphere Application Server you want to install.

Table 64. WebSphere, understanding the components

Component	Description
Application Server Base function	Base software
Application Server Administrator	Graphical interface to manage servlets
Documentation	On-line documentation
Samples	Sample applications to demonstrate the basic classes and the extensions
Java Server Pages	Support for a new technology for dynamic page content called JavaServer Pages (JSP)
CORBA	An object request broker (ORB) and a set of services that are compliant with the Common Object Request Broker Architecture (CORBA).

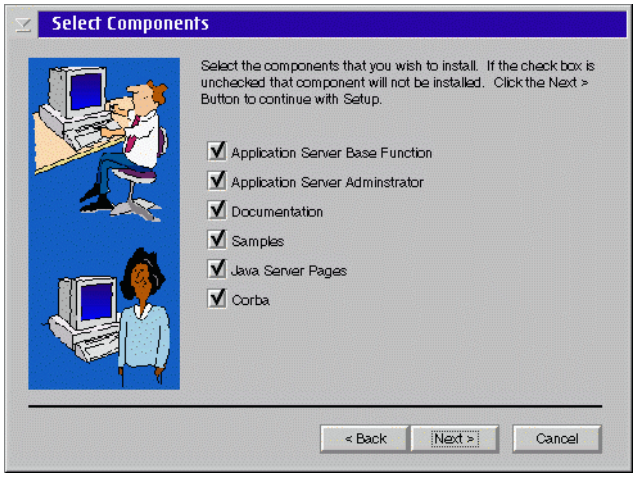

Figure 151. WebSphere, selectable components

WebSphere Application Server supports several Web servers. They are:

- Domino Go Webserver
- Netscape Enterprise Server
- Netscape FastTrack Server

- Microsoft Internet Information Server
- Apache Server

The next screen indicates that WebSphere plug-ins will be installed on Lotus Go Domino Webserver, the Web server installed on our system as shown in Table 152 on page 339. Click on **Next** to continue.

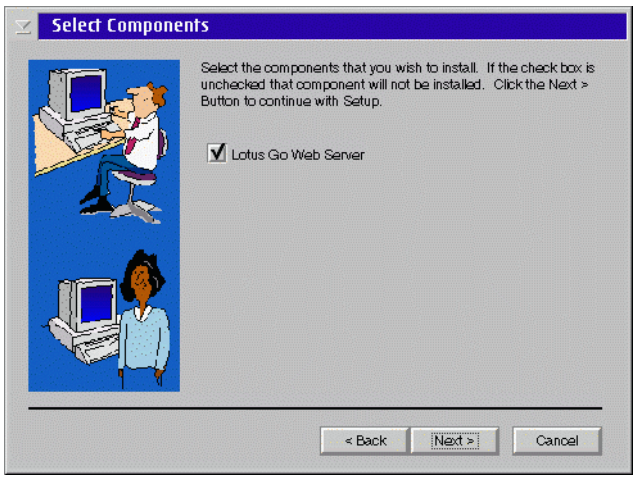

Figure 152. .WebSphere, selecting the Web server

The next window displayed is shown in Table 153 on page 340 and provides the name of the default destination folder where the files will be copied. Additional folder names are also displayed in a scroll down listbox from which you can choose a destination folder for the files to be copied. Once you are happy with the destination folder, click on **Next** to continue.

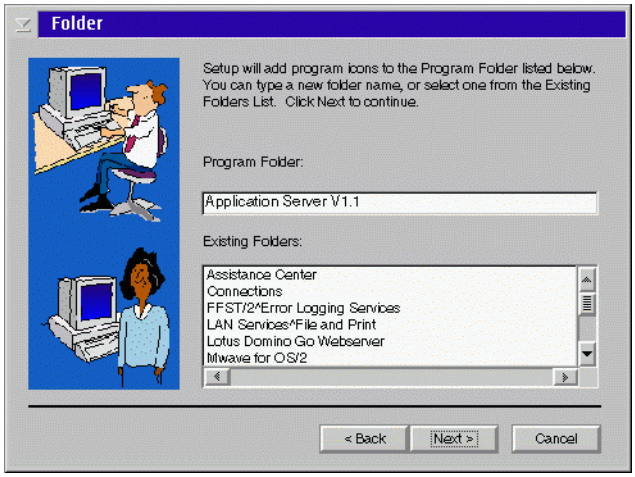

Figure 153. WebSphere, application folder selection

The following window, as shown in Table 154 on page 340, shows the selected destination folder where the files will be copied. If you are not satisfied with your previous choice, click on **Back** to return to the previous window, otherwise, click on **Next** to continue.

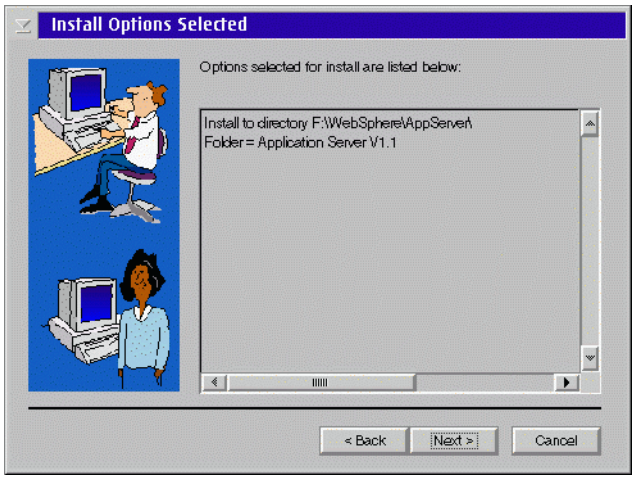

Figure 154. WebSphere, confirm options

The window displayed prompts you if you want to proceed with copying. A progress indication window is displayed showing the full source and destination path names and a progress indicator bar of the files being copied.

A final window is displayed providing a select button asking you if you want to view the readme file at this stage. If you select the Yes, I want to view the ReadMe File, this file is displayed as shown in Table 155 on page 341. When you click on the **Finish** button, the installation is ended.

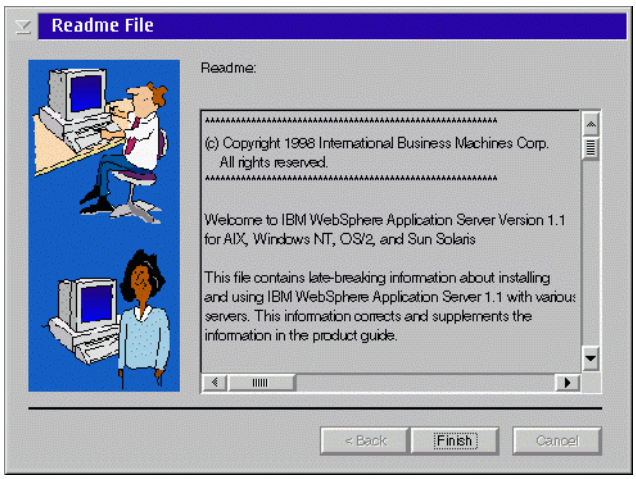

Figure 155. WebSphere, Readme file

8.6.1 Uninstall for WebSphere application server

If you double-click on the WebSphere Application Server folder, it contains two folders. One is the Readme file and the other is the Uninstall executable.

If you double-click on the **Uninstall** folder, a window is displayed on your desktop, as shown in Table 156 on page 342.

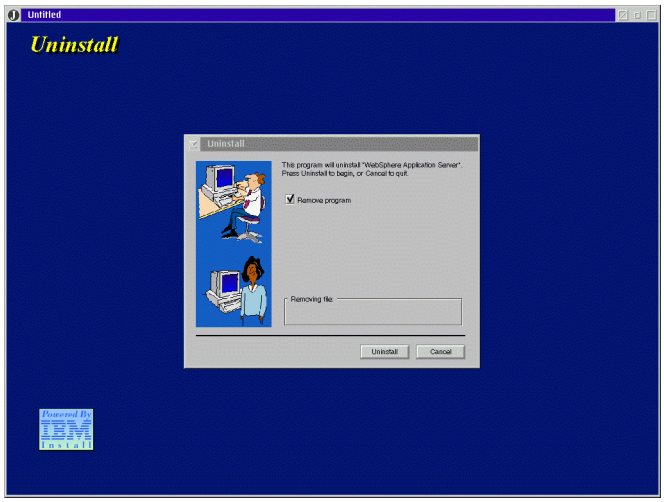

Figure 156. WebSphere, removal

Click on the **Remove Program** button, and the uninstall process will start where the full path filenames are displayed as they are removed from the system. Once this process is finished, the window is closed automatically.

8.7 WebSphere functional components

IBM WebSphere is a set of software products that help customers deploy and manage high performance Web sites. It helps ease the transition from simple Web publishing to advanced e-business Web applications. Because WebSphere is based on open standards, customers benefit from cross-platform portability. With this portability comes scalability. As business needs grow, an application can progressively move to more powerful platforms. As performance demands increase, the same Web application could move from OS/2 Warp Server for e-business to an AS/400 or Solaris and, ultimately, to an OS/390. The application can also be run from various platforms already installed in the company achieving a horizontal scaling by utilizing these existing assets.

Web developers can benefit from the structured and reliable Java servlet environment while encapsulation enables them to leverage their existing CGI- and PEARL-based applications. Compared to CGI programs, Java servlets are more secure and provide better performance and scalability – features that make administration easier. Being Java, the servlets can exploit the beneficial capabilities of Java including security, platform independence, and reusability. These servlets can provide safe and secure access to back-tier

relational databases, transaction-based systems, and applications while generating dynamic content for the Web client. This programming model eliminates the need for access to existing systems from outside a firewall. Businesses can rapidly and securely extend their existing software assets to an e-business model.

Network Computing Online reviewed WebSphere in an article that can be read at the following Web site:

`http://www.networkcomputing.com/913/913sp2.html`

Among the *thunderous capabilities* they reported were that WebSphere:

- Improves CPU utilization and workload balancing for Java servlets
- Replaces old C and PERL CGI scripts
- Is built entirely on open standards like CORBA, JDBC, and Java
- Lets you implement advanced applications on the server using Java
- Simplifies implementing and running servlets

WebSphere Application Server provides the foundation for deploying Web applications using the Java servlet programming model. With this technology, corporations can design and develop server-side e-business applications using a common programming model and placing business logic where it makes the most sense for each application.

Since Java servlets are more secure, have better performance, provide easier administration, and greater scalability than Common Gateway Interface (CGI) programs, WebSphere can help rapidly evolve into dynamic, personalized Web applications that are an integral part of the enterprise processes.

8.7.1 Accessing the product documentation

To view copies of the documentation installed on your Application Server, use one of the following methods:

To access the documentation via your Web server, use your browser to open:

`http://your.server.name/appserver/doc/index.html`

To access the documentation as local files on the Web server, open:

`x:\\WebSphere\Appserver\doc\index.html`

where `x:` is the drive where you installed the Web server.

8.7.2 Starting the application server manager

The Application Server Manager provides a graphical user interface for configuring and managing servlets running on your Web Server. Most of the changes you make to configuration parameters from within the Application Server Manager take effect immediately and do not require you to restart the Web server.

Restart the Web server, however, if you:

- Change the port number for the Application Server Manager
- Change the parameters on the Basic Setup page
- Manually edit the jvm.properties file to control Native DLL or Java standard out logging

To start the Application Server Manager, access the following Web site:

```
http://your.server.name:9090/
```

where `your.server.name` is the fully qualified name of your host. If you start Application Server Manager in a browser on the same machine on which you installed the Application Server, use `http://localhost:9090/` for better performance. The Application Server Manager applet starts and displays the login page as shown in Table 157 on page 345.

To log in to the Application Server Manager for the first time, enter `admin` as the login user ID and password and click **OK**. For security reasons, you should change the login password.

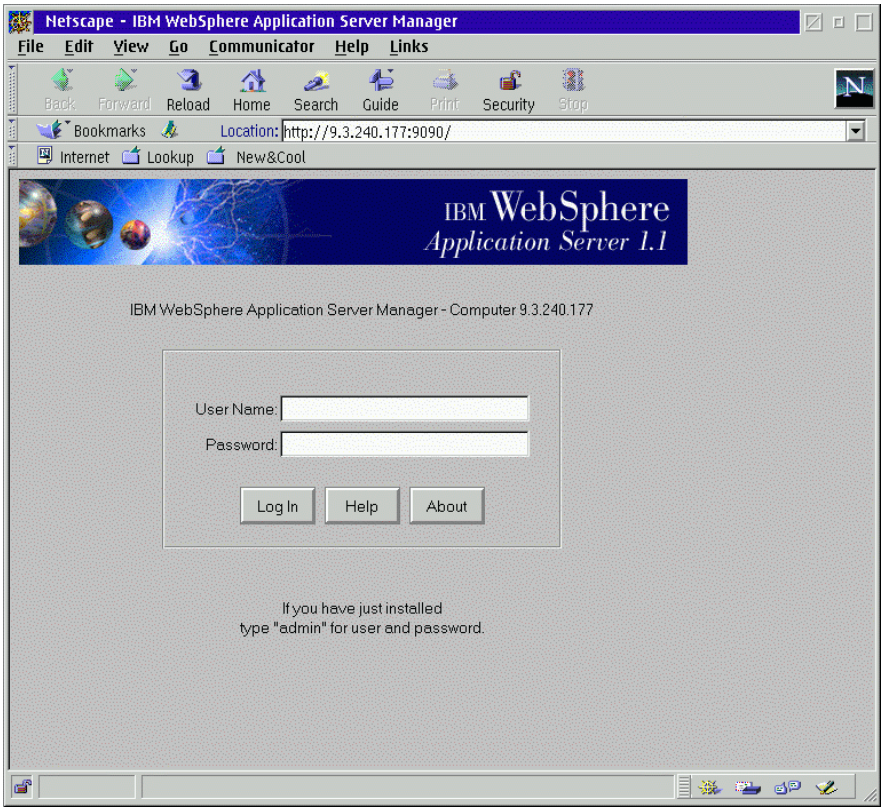

Figure 157. WebSphere, Log-In window

8.7.3 Using the application server manager

After you log in, the Application Server Manager displays the Services page, which shows the current state of the Web server and lists the services that are installed and running on your machine. The service provided is the Application Server servlet administration service, which you use to configure and manage servlets. This service listens on port 9090 unless you change the port using the Properties page.

Current status and summary information for each service is displayed on the Services page. You can stop, restart, and shut down services from this page by selecting the buttons at the bottom of the page.

Configure and manage servlets by highlighting the Application Server servlet administration service, and then clicking **Manage**. A new window opens

displaying the Application Server Manager configuration interface, as shown in Table 158 on page 346 below.

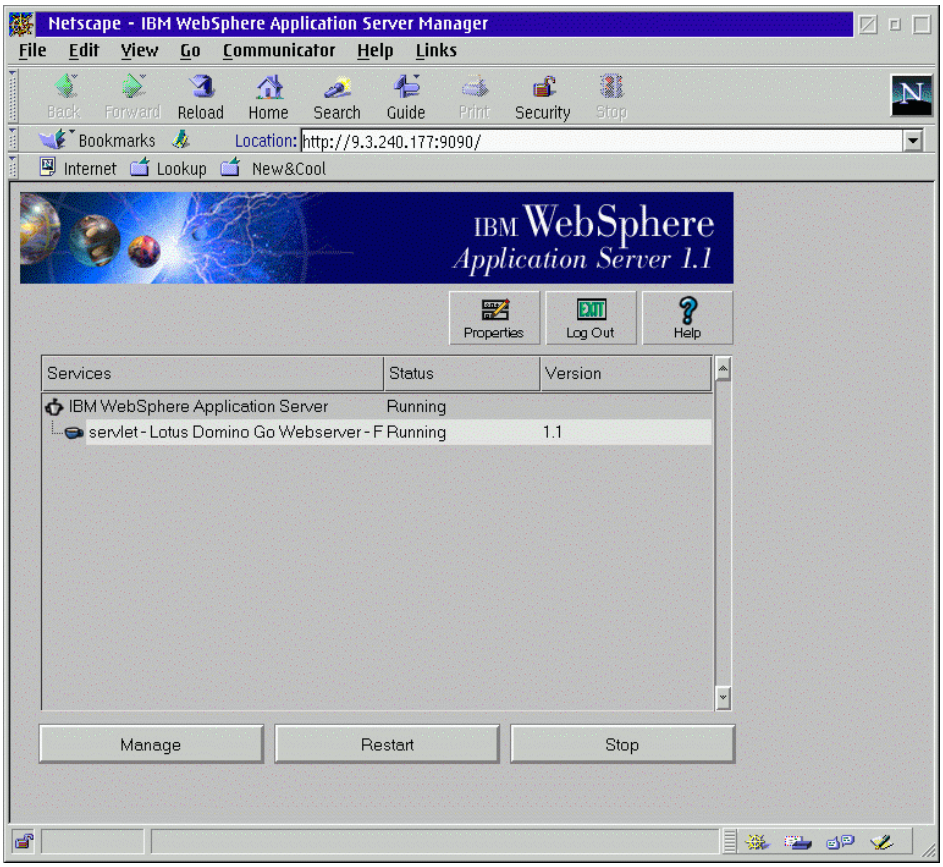

Figure 158. WebSphere, application server manager

To set parameters and view information about the servlet activity, use the navigation buttons at the top of the Application Server Manager interface to display the available configuration and administration tasks. Select a task from the tree-view to display parameters and information for that task. Refer to Table 159 on page 347 for an example of the window displayed.

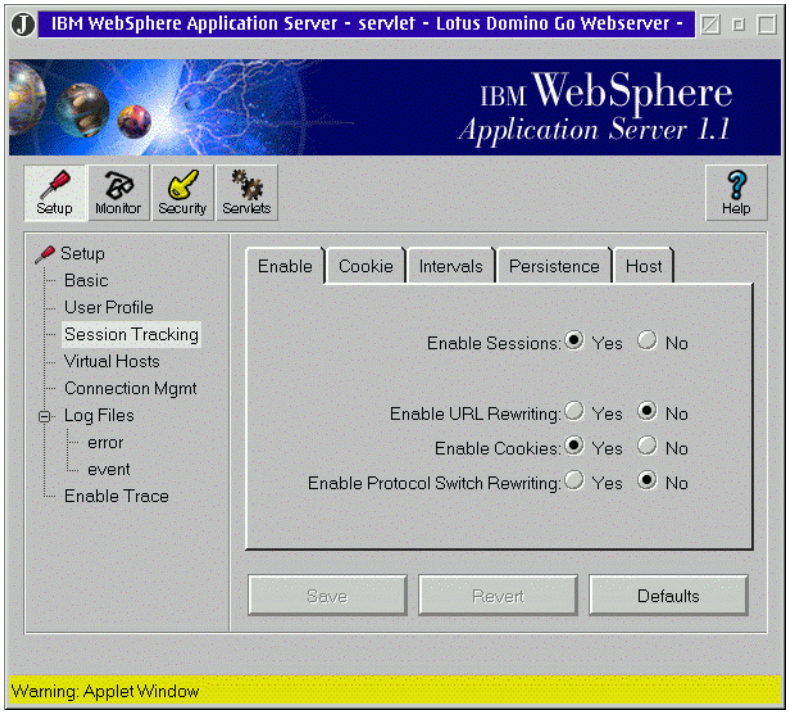

Figure 159. WebSphere, set up Window for an applet

If you modify parameters for a task, select **Save** or **Revert** before viewing another task. If you select Save, Application Server saves the changes you made to the Application Server properties file. If you select Revert, changes you made to parameters for the current task without saving are replaced with the previously saved values.

Refer to the *Application Server Guide*, which is part of the product documentation, for general information about configuring and managing servlets. For more detailed task information and descriptions of configuration parameters, click **Help** from the Application Server Manager configuration interface.

8.7.4 Monitoring servlet activity

Monitor servlet activity by viewing the output of various log files, viewing the status of loaded servlets, and viewing, in real time, how resources are being used. Click the **Monitor** button from the top of the Application Server Manager interface to display the monitor tasks.

Table 65. Servlets Monitor, tasks and their uses

Use this task...	To monitor...
Log output	Information collected in the log files.
Resource usage	How service resources are being used. The monitored resources are memory, a pool of handler objects, requests to the service, and the service response time.
Active sessions	Information about the user sessions that are currently active on the Web server including information about individual sessions and summary information for all active sessions.
Loaded servlets	Status and statistics for individual servlets.

8.7.5 Establishing and maintaining security

By default, there is no restriction imposed on anybody when accessing your server unless you establish a secure environment. Establish and maintain security by defining users, groups, resources, and access control lists. By assigning specific access settings to each user, group, and resource, you can control precisely how the resources of a service are used and by whom. Click the security button from the top of the Application Server Manager interface to display the Security tasks.

Table 66. Servlet security, tasks

Use this task	To specify
Users	Who can access Web pages served by the Application Server and other resources, such as servlets?
Groups	Named lists of users.
Access control lists	Specific access permissions for users and groups.
Resources	Security parameters for specific directories, files, and servlets on the Application Server.

8.7.6 Managing servlets

Servlets placed in the applicationserver_root\servlets directory are automatically loaded and reloaded (if updated) when requested. You can also use the Application Server Manager to manage servlets more directly by defining initialization parameters and creating servlet aliases and filters. To manage servlets, click the Servlets button from the top of the Application Server Manager interface to display the Servlets tasks.

Table 67. Servlet management, tasks

Use this task	To
Add servlets	Add a servlet that you want to manage
Servlet aliases	Specify a path-mapping rule so users can use a shortcut URL to invoke a specific servlet.
Filtering	Associate servlets with MIME-types so that each time a response with a specific MIME-type is generated, a particular servlet is invoked.
Configure servlets	Define configuration information and initialization parameters for individual servlets, such as the associated class file for the servlet, whether the servlet loads at start-up, and whether the Web server loads the servlet from a remote location.

The online document, *Application Server Guide*, offers comprehensive information regarding the installation, management, and use of the Application Server.

8.8 Developing and implementing servlets

This section of this document describes the implementation of two servlets. The purpose of this section is to illustrate the ease of implementing servlets using WebSphere Application Server V1.1. Although the examples used here are very limited in their function, they still serve their purpose in illustrating the little effort required to build servlets that can enable your company to participate in the e-business arena.

The first servlet dynamically builds a Web page and displays it on the client's machine when invoked from a browser. If you are familiar with the Hypertext Markup Language (HTML) or with the Java language, you will easily grasp that the servlet builds the HTML text dynamically and returns these text

statements to the client. You do not have to be an expert in any of these two languages, since the context of the codes (listed below) are very close to English-like statements.

Note

When we tried to compile a Java servlet, the compilation failed because the Java SDK class-file JSDK.JAR could not be found. To correct this problem, we had to manually add the full path specification for this file in the SET CLASSPATH statement in the CONFIG.SYS file. This file resides in the \websphere\AppServer\lib directory.

The samples used here can be found in the Samples directory on the CD included with this document. You could also use your favorite editor and create the following source file calling it FirstServlet.java:

```
import javax.servlet.*;
import javax.servlet.http.*;
import java.io.*;

public class FirstServlet extends HttpServlet {
public void doGet (HttpServletRequest request,
                                    HttpServletResponse response)
                                  throws ServletException, IOException {
          response.setContentType("text/html");   // sending HTML
          ServletOutputStream out = response.getOutputStream();
          out.println("<html>");
          out.println("<head><title>First  Servlet </title></head>");
          out.println("<body>");
          out.println("<center><h1>First Servlet Output</h1></center>");
          out.println("<center>Thank you for visiting this
site!</center>");
          out.println("</body></html>");
          }
        }
```

Create this file in the \\WebSphere\AppServer\servlets directory on your system. To compile this program key in: javac FirstServlet.java. This will create the FirstServlet.class servlet. To run this servlet, a client can start a browser on his machine and go to a Web site patterned after the following:

```
http://your.server.name/servlet/FirstServlet
```

The browser will display the window as shown in Table 160 on page 351.

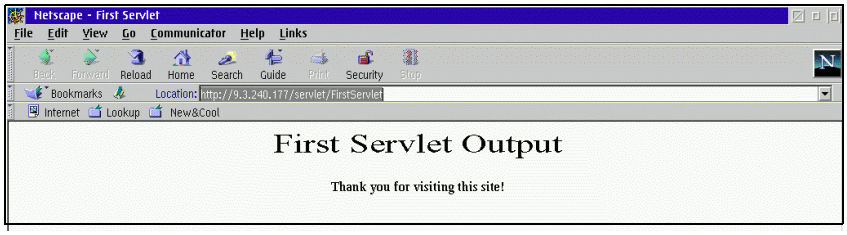

Figure 160. WebSphere, first servlet output

In the second example, there are two files involved. The first is a *hard-coded* HTML file that displays a window collecting some survey information. The window displayed by a browser is shown in Table 161 on page 353. Sample information requested by this HTML file is already filled in on the webpage. This HTML file should reside in the HTML home directory as specified in the httpd.cnf file, which is the configuration file for the Web Server. The default directory where the HTML files are stored is www\HTML.

The second file is the servlet that builds a Web page dynamically and displays the responses entered by the client on the HTML page. Certainly not a complicated application, but, hopefully, illustrating that the servlet could store the information received from the client in a database or a flat file for inclusion in a spread-sheet program or any other program to perform processing on the supplied data as required.

The source text for the HTML page that collects the information can be found in the Samples directory on the CD accompanying this manual under the WebSphere subdirectory with the filename of survey1.html. The contents of this file are listed below:

```
<html>
<head>
<title>Software  Survey</title>
</head>
<body>
<h1><center>Software  Survey</center></h1>
<hr><br>
<form method=POST action="http://9.3.240.177/servlet/survey2">
<table border=0>
 <tr>
 <td align=right>Name:</td>
 <td colspan=2 align=left><input type=text name=name size=40></td>
 </tr>
 <tr>
 <td align=right>Email Address:</td>
```

```
<td colspan=2 align=left><input type=text name=email size=40></td>
</tr>
<tr valign=top>
 <td align=right>Age:</td>
 <td align=left>
 <input type=radio name=age value="<18">Less than 18 <br>
 <input type=radio name=age value="18-25">18 - 25
 </td>
 <td align=left>
 <input type=radio name=age value="26-40">26-40<br>
 <input type=radio name=age value=">40">Over 40
 </td>
 </tr>
 <tr valign=top>
   <td align=right>Version of OS/2 Installed:</td>
   <td align=left>
   <select name=os size=5 multiple>
   <option>OS/2 V3.0 without Win-OS/2
   <option>OS/2 V3.0 including Win-OS/2
   <option>OS/2 V3.0 Connect
   <option>OS/2 V4.0 (Merlin)
   <option>OS/2 Server for e-business
   </select>
   </td>
   </tr>
   <tr>
   <td></td>
   <td><input type=checkbox name=more value="yes">
         Send me more information
   </td>
   </tr>
   <tr>
     <td align=right>Comments:</td>
     <td colspan=2 align=left>
     <textarea name=comments cols=40 rows=4>
     </textarea>
     </td>
     </tr>
     <tr>
     <td></td>
   <td>
     <input type=reset value="Clear Info">
     <input type=submit value="Submit Info">
    </td>
      </tr>
     </table>
     </form>
```

```
</body>
</html>
```

Pay special attention to the eighth statement, (starting with *form*). This invokes the servlet we will discuss next. You will need to modify this line to contain your server's host name or IP address in place of 9.3.240.177.

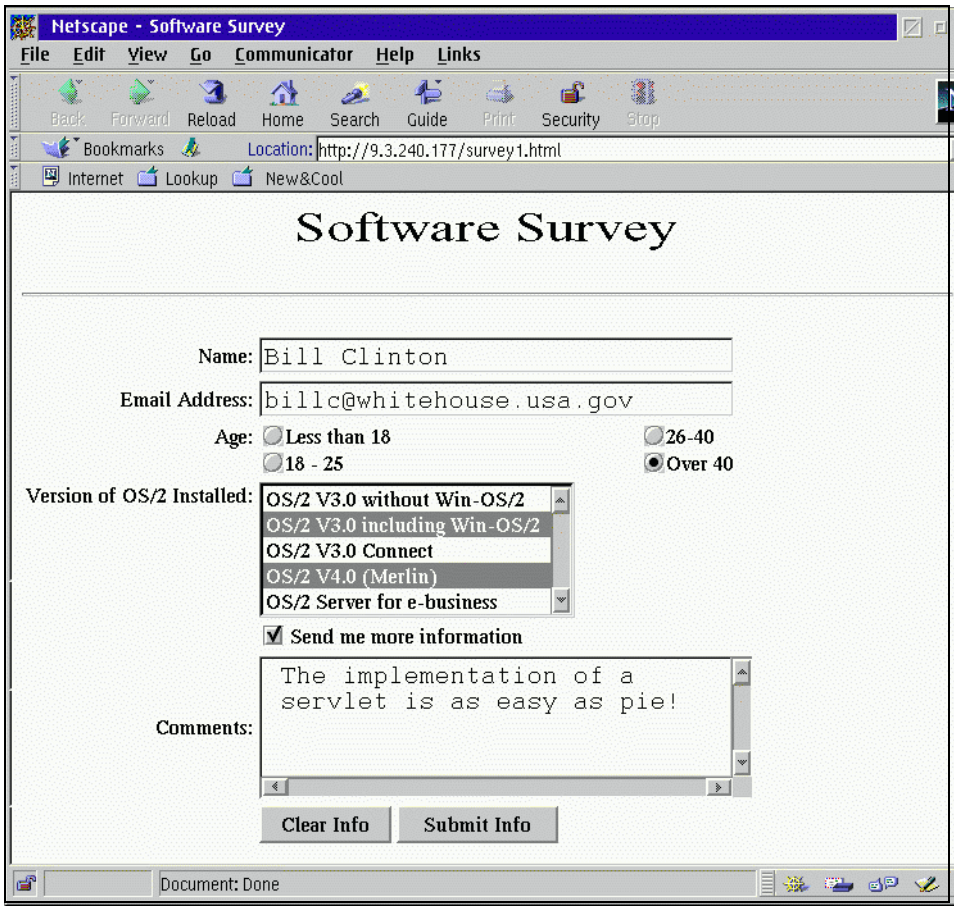

Figure 161. Servlet, sample application

Once the client filled in the requested information and clicked on the submit button, the servlet receives and displays the information as shown in Table 162 on page 357. The source code for the servlet, survey2.java, can also be found in the samples directory mentioned earlier. The source code is very similar to the FirstServlet.java file except that it extracts the values of the fields defined in the survley1.html file, and builds the necessary HTML

statements dynamically to display the entered information on the client screen. The reason for this is just for illustration purposes. In a real life application, one would most probably store the information in a database or a flat file for importation into a spread sheet.

The source code for this servlet (survey2.java) is listed below:

```
import javax.servlet.* ;
   import javax.servlet.http.*;
   import java.io.*;
   public class survey2 extends HttpServlet
     {
        /**
          * <p> Performs the HTTP POST operation
          *
          * @param req The request from the client
          * @param resp The response from the servlet
          */
       public void doPost(HttpServletRequest req,
                                  HttpServletResponse resp)
                throws ServletException, java.io.IOException
                {
                   // Set the content type of the response
                   resp.setContentType("text/html");

                   // Create a PrintWriter to write the response
                   java.io.PrintWriter out = new
java.io.PrintWriter(resp.getOutputStream());

                  // Print a standard header
                  out.println("<html>");
                  out.println("<head>");
                  out.println("<title>Survey Details</title>");
                  out.println("</head>");
                  out.println("<body>");
                  out.println("<h1><center>Your survey has been
processed!");
                  out.println("</center></h1><hr><br>");
                  out.println("You selected the following details:");
                  out.println("<dir>");

                  String values[];

                  // Get the name
                  String name = " ";
                  values = req.getParameterValues("name");
                  if (values != null ) {
                     name = values[0];
```

```
 }
 out.println("Name=" + name + "<br>");

// Get the email address
String email = " " ;
values = req.getParameterValues("email");
if (values != null) {
email = values[0];
}
out.println("Email=" + email + "<br>" );

// Get the age
String age = " ";
values = req.getParameterValues("age");
if (values !=null) {
  age = values[0];
}
out.println("Age=" + age + "<br>");

// Get the operating system. There could be more than one
value
 values = req.getParameterValues("os");
 out.print("Operating System version=");
 if (values != null) {
  for (int i = 0; i < values.length; i++) {
       if (i > 0 ) out.print(", ");
          out.print(values[i]);
                 }
    }
 out.println("<br>");

  // Get the 'more information' flag
  String more = " ";
  values = req.getParameterValues("more");
  if (values !=null) {
  more = values[0];
  }
  out.println("More information=" + more + "<br>");

 // Get the comments
    String comments = " ";
    values = req.getParameterValues("comments");
    if (values != null) {
      comments = values[0];
      }
     out.println("Comments:<br>");
     out.println("<dir>");
```

```
                    // Comment lines are separated by a carriage return/line
feed
                      // pair - convert them to an HTML line break <br>
                    //  out.println(toHTML(comments));
                       out.println(comments);
                      out.println("</dir>");

                     out.println("</dir>");

                    // wrap up
                    out.println("</body>");
                    out.println("</html>");
                    out.flush();
        }
    }
```

To compile this servlet (which should be stored in directory \\WebSphere\AppServer\servlets) from a command prompt after changing to this directory, key in `javac survey2.java`, which will create the survey2.class file. This servlet should be compiled prior to running the survey1.html file from the browser.

Earlier, we stated that the output can be easily changed to, for example, a flat file instead of returning the captured information to the client's browser for display as a Web page. To change the listed servlet to write the captured info to a flat file instead, do the following:

- Replace the statements following '//. Set the content type of the response ' and '// Create a PrintWriter to write the response' with:

```
String filename = "survey2.out";
java.io.FileOutputStream fos = new java.io.FileOutputStream(filename);
java.io.PrintWriter outfile = new java.io.PrintWriter(fos);
```

- Change all subsequent 'out.println' with 'outfile.println'. Also change all 'out.print' with 'outfile.print'Replace the 'out.flush();' statement with 'outfile.close();
- Recompile the program.

After these changes, the servlet should create the file survey2.out in the www\bin directory.

We hope that the very simple example used here illustrates the ease with which you can get started with servlets on the Web allowing you to reap in the benefits provided by the e-business world.

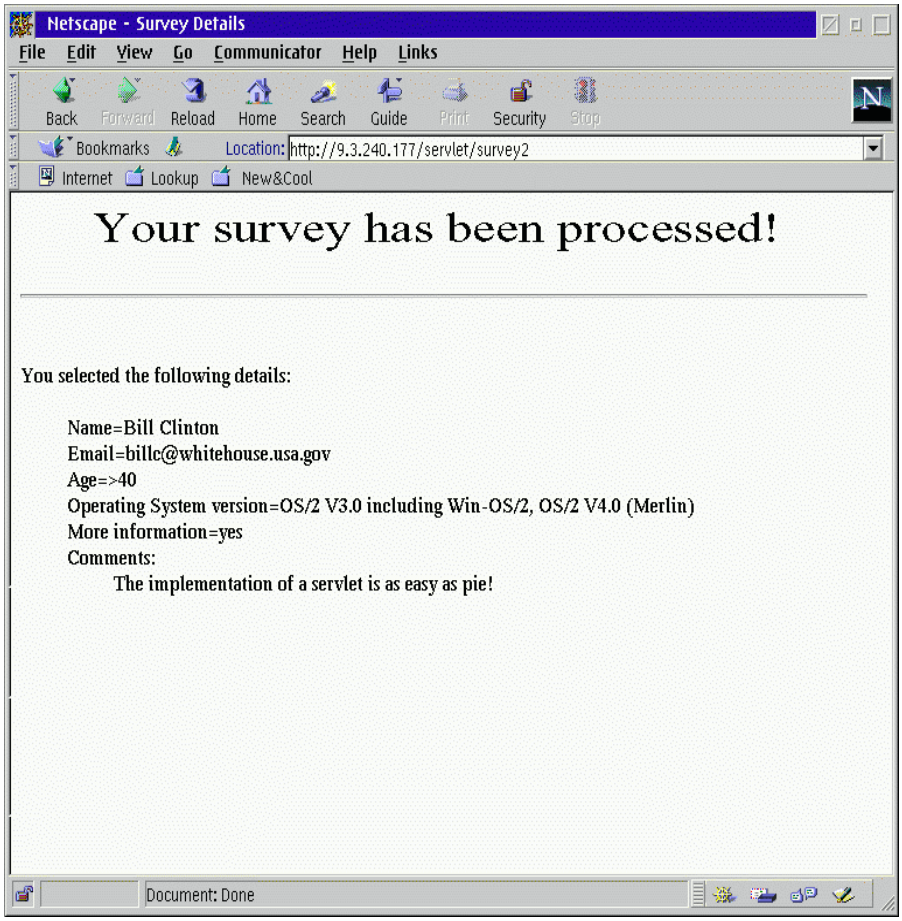

Figure 162. Servlet, application processing

Chapter 9. IBM remote access services

This chapter describes the IBM remote access services Version 5.11 (formally known as LAN Distance) that is shipped with OS/2 Warp Server for e-business. This has been covered in great depth in previous redbooks, which are listed at the end of this chapter.

In this chapter, we will focus on new or changed functions. We will describe how the PPP component is installed, configured, and used and how to set up a PPP-Client. Other functions will be described using a simple scenario. We also describe the configuration and use of various client machines, such as Windows 95 and OS/2.

9.1 Overview

This section discusses the various environments in which IBM remote access connection server for OS/2 Warp Server for e-business can be configured. We also discuss client support including support for PPP Clients.

9.1.1 IBM remote access services environments

The IBM remote access connection server for OS/2 Warp Server for e-business product supports four different types of remote LAN access:

Remote-to-remote: Two remote systems can establish a WAN connection. A typical application of this method is where a traveling employee needs to access their office workstation from a remote location. A remote client can have up to two simultaneous, incoming, remote connections.

LAN-to-LAN: You can establish a connection between two Connection Servers. In this situation, the two servers form a bridge between the two LANs.

LAN-to-remote: LAN-attached workstations can request the Connection Server to establish a connection with a remote workstation.

Remote-to-LAN: This scenario is probably the most common use of IBM remote access services. This solution allows users to access LAN resources from remote locations, such as their home, while visiting customers or while traveling. This method also allows the IBM remote access connection server for OS/2 Warp Server for e-business to be installed as a stand-alone server supporting up to 128 remote workstations (a maximum of 64 PPP connections). No LAN hardware is necessary on the Connection Server, only WAN adapters. The remote workstations can access all the resources at the

server. This may be suitable to allow file sharing for a small, remote workgroup.

Note

Callback is not supported on PPP clients. For more information about client restrictions, see section “Remote client support restrictions” on page 364.

Figure Figure 163 on page 360 shows a typical Remote-to-LAN scenarios. The Connection Server provides an interface for the OS/2 Remote Client, IBM 8235 DIALs, Windows 95/98, and Windows NT Remote Client to the LAN. The clients can access printers and other devices on the LAN. Notice that a Wide Area Network (WAN) is formed by the connection of the LAN-connected Connection Server and Remote Client. The LAN can either be token-ring or Ethernet.

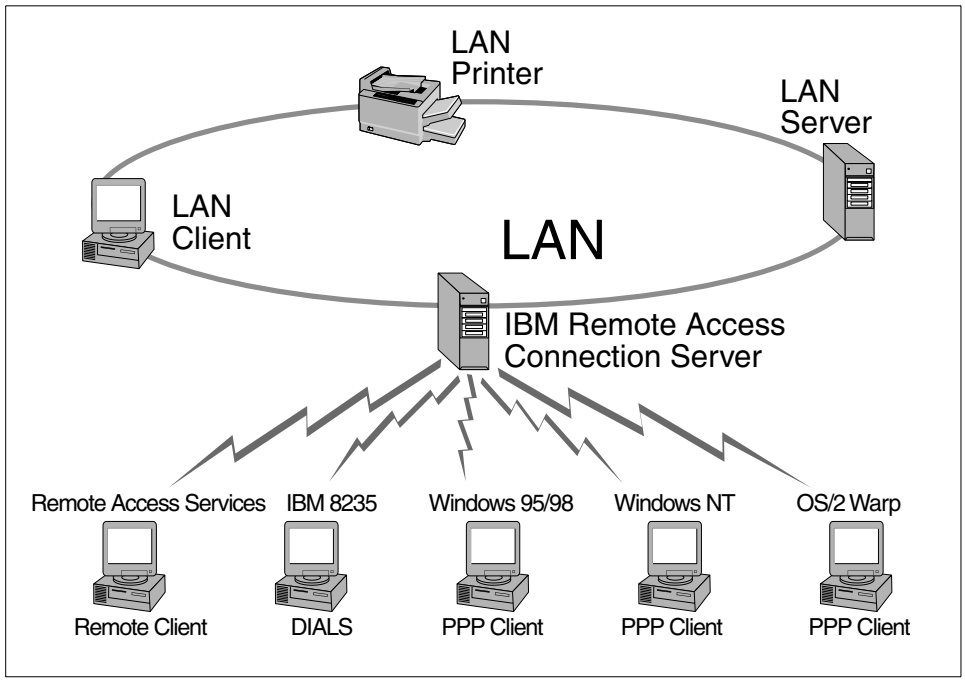

Figure 163. Remote to LAN scenario

9.2 PPP support

IBM remote access connection server for OS/2 Warp Server for e-business adds Point-to-Point Protocol (PPP) support. PPP is currently the best solution for dial-up connections.

IBM remote access connection server for OS/2 Warp Server for e-business accepts calls from PPP clients whose LAN applications use the TCP/IP protocol. This is in addition to being able to accepting calls from LAN Distance/ Remote Access Services clients.

Support for Password Authentication Protocol (PAP) and Challenge Handshake Authentication Protocol (CHAP) has been added. The Password Authentication Protocol is a simple two-way mechanism that is done when the link between the remote client and server is initially established. Challenge Handshake Authentication Protocol (CHAP) is used to authenticate all PPP clients using an integrated database. Previous releases of LAN Distance and Remote Access Services only supported the Two Party Authentication Protocol (TPAP). TPAP will continue to be used to authenticate existing LAN Distance and Remote Access Services clients. IBM remote access connection server for OS/2 Warp Server supports the TCP/IP protocol over PPP connections as well as IEEE 802.2, TCP/IP, IPX, and NETBIOS over LAN Distance and Remote Access Services client connections. For more information about PAP and CHAP, refer to the section “IBM remote access services internal architecture” on page 403.

This provides the capability for an IBM remote access connection server to accept calls from, connect, to and interoperate with any workstation that supports Point-to-Point Protocol and requires the TCP/IP protocol.

9.3 Client support

The IBM remote access connection server for OS/2 Warp Server for e-business supports multiple concurrent sessions:

Up to 64 concurrent PPP connections

- Up to 128 concurrent Remote Access Services or LAN Distance Remote Clients

The following clients have been tested with the IBM remote access services for OS/2 Warp Server for e-business:

- Microsoft Windows 95/98

- Microsoft Windows NT Version 4
- IBM 8235 DIALs Connect for Windows Version 4.1
- IBM LAN Distance Remote for Windows Version 5.0 (from OS/2 Warp Server)
- IBM Dial-Up for TCP/IP Version 2.0 (Rev.1.19)
- IBM 8235 DIALs for OS/2 Version 4.5.2
- IBM 8235 DIALs for OS/2 Version 4.0.3 + patch 3
- IBM LAN Distance Remote for OS/2 Version 5.0 (from OS/2 Warp Server)

The PPP connection support included with IBM remote access services for OS/2 Warp Server for e-business ensures long-term compatibility with other Internet-based products on a variety of computer platforms.

9.4 System requirements

The following sections describe:

- Hardware and Software requirements for IBM remote access connection server for OS/2 Warp Server for e-business
- Hardware and Software requirements for Remote Clients
- Support restrictions for Remote Clients

9.4.1 IBM remote access connection server requirements

The following prerequisite software and the WAN hardware should be installed before you install IBM remote access services:

- OS/2 Warp Server for e-business satisfies all the TCP/IP and MPTS requirements.
- For earlier versions of IBM OS/2 Warp Server:
  - The MPTS Requirements are:
    - For OS/2 Warp Server: MPTS Version 5.11 or later.
    - For OS/2 Warp Server SMP: MPTS Version 5.2 with SMP MPTS APAR IC15968 or FixPak WRx8502.
  - To support PPP clients, TCP/IP Version 3.1 or later is required. (OS/2 Warp Server for e-business comes with TCP/IP V.4.21.)
- The minimum amount of memory recommended for OS/2, the IBM remote access connection server component, and one LAN application is 12MB. The actual requirements of your Connection Server will depend on your

LAN applications, data and response time requirements, and your workstation's processor speed.

- IBM remote access services requires 5.0 MB of fixed-disk storage space.
- A supported LAN adapter.
- A supported WAN adapter.

> **Note**
>
> This adapter will not be available for use by other applications simultaneously. For example, if you are using an ISDN adapter and you have installed Communications Manager/2 (CM/2) on your workstation, you must set up CM/2 so that it is not configured for ISDN.

- A modem to connect to the WAN adapter(s).

9.4.2 Remote client system requirements

All Remote Clients will require the following plus additional requirements depending on the platform:

- A modem and/or adapter to connect to the server.
- For asynchronous communications, it is recommended that the COM port of the Remote Client workstation be FIFO-buffered. This function is provided by a 16550 or 16550A UART chip in the workstation. This ensures you will be able to run your COM ports at up to 115200 baud without buffer over-run problems. However, some non-FIFO workstations with a processor faster than 25 MHz can support a reliable COM port speed of 38400 baud with a 14400 bps modem (or better).

  To verify that your workstation has FIFO buffering, type `MODE COM1` at an OS/2 command prompt. If the response includes `BUFFER=ON`, then, your workstation has FIFO buffering.
- Non-switched (leased) or switched telephone lines to establish an asynchronous, synchronous or ISDN connection.

9.4.3 OS/2 RAS remote client system requirements

- IBM OS/2 Version 3.0 or later

> **Note**
>
> The Windows 3.1 client is no longer supported and is not shipped with OS/2 Warp Server for e-business. If you are using an earlier version of OS/2 Warp Server, see the product documentation for information regarding the Windows 3.1 client.

- The minimum amount of memory recommended for OS/2, the Remote Access Services Remote Client, and one LAN application is 12 MB. The actual requirements of your Connection Server will depend on your LAN applications, data and response time requirements, and your workstation's processor speed.
- The OS/2 Remote Access Services Remote Client requires 5.0 MB of fixed-disk storage space.

9.4.4 PPP client system requirements

One of the following operating systems:

- IBM OS/2 Warp 3.0 or later and one of the following:
  - IBM OS/2 Internet Dialer supplied with the OS/2 Warp BonusPak or IBM TCP/IP Version 3.0 or later
  - IBM 8235 DIALs for OS/2 Version 4.5.2 or Version 4.03 and patch 4
- Windows Version 3.1 or later, running on DOS Version 5.0 or later, and:
  - IBM 8235 DIALs for Windows Version 4.1
- Windows 95 Version 4.0 PPP client (the Dial-Up network component of Network Neighborhood)
- Windows NT Version 4.0 Client

9.4.5 Remote client support restrictions

The Remote Client supports a passphrase, which is case-sensitive, can include spaces, and can be up to 32 characters in length. If you are dialing in with a non-IBM PPP client, be aware of any password length or case-sensitivity restrictions with that client. The following restrictions apply to IBM remote access services Remote Client:

- An OS/2 Remote Client can have up to two concurrent connections.
- A Windows 3.1 or 3.11 Remote Client can have only one asynchronous connection using either a switched or non-switched line.

Native X.25 support is provided by third parties, such as WAN Services for OS/2, announced by Eicon Technology.

We recommend that you run LAN applications on the Remote Clients that do not transfer large amounts of code or data over the WAN connection. For example, if an application loads 5MB of code during startup, it may take 5 to 10 seconds to load on a high-speed LAN. However, over a WAN connection using an asynchronous modem running at 14400 bps, this may take from 5 to 10 minutes. This is clearly not an acceptable response time.

The following restrictions apply to PPP:

- Callback is not supported for PPP clients.
- An IBM remote access connection server does not support the ability to use LAN Distance logical adapter network addresses with a PPP user account.
- You cannot change passphrases from the client or perform security administration functions (see “Allowing PPP-clients to change their PPP passphrase” on page 398 for a workaround).
- Maximum passphrase age and maximum logon attempt policy options are not enforced for PPP clients.
- PPP clients cannot TCP/IP ping LAN Distance and Remote Access Services clients. LAN Distance and Remote Access Services clients cannot TCP/IP ping PPP clients.
- PPP clients can run only LAN applications and networking software that use TCP/IP.
- PPP clients do not require LAN hardware to use a LAN Distance connection in order to access LAN resources. WAN communications hardware, such as a modem and COM port or an adapter, is required for the type of communications connection you want to support.

9.5 Installing IBM remote access services

During the OS/2 Warp Server for e-business installation, you can select **Remote Access Services** from the **Setup and Installation** screen. If you have already installed your OS/2 Warp Server for e-business, you can make up for the Installation by selecting **Selective Install for Networking** in the Install/Remove folder. Both ways will bring you up to the screen as shown in Figure 164 on page 366.

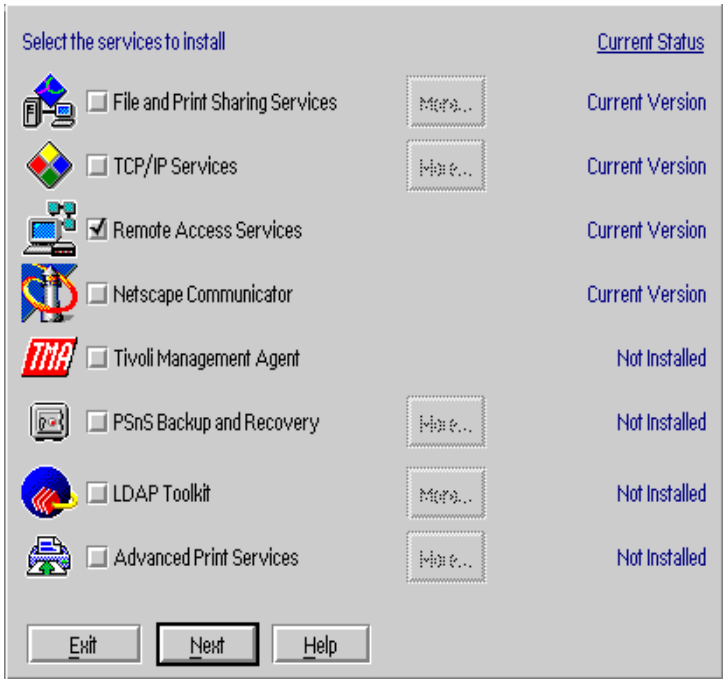

Figure 164. Setup and installation

After selecting **OK**, the following screen will be shown:

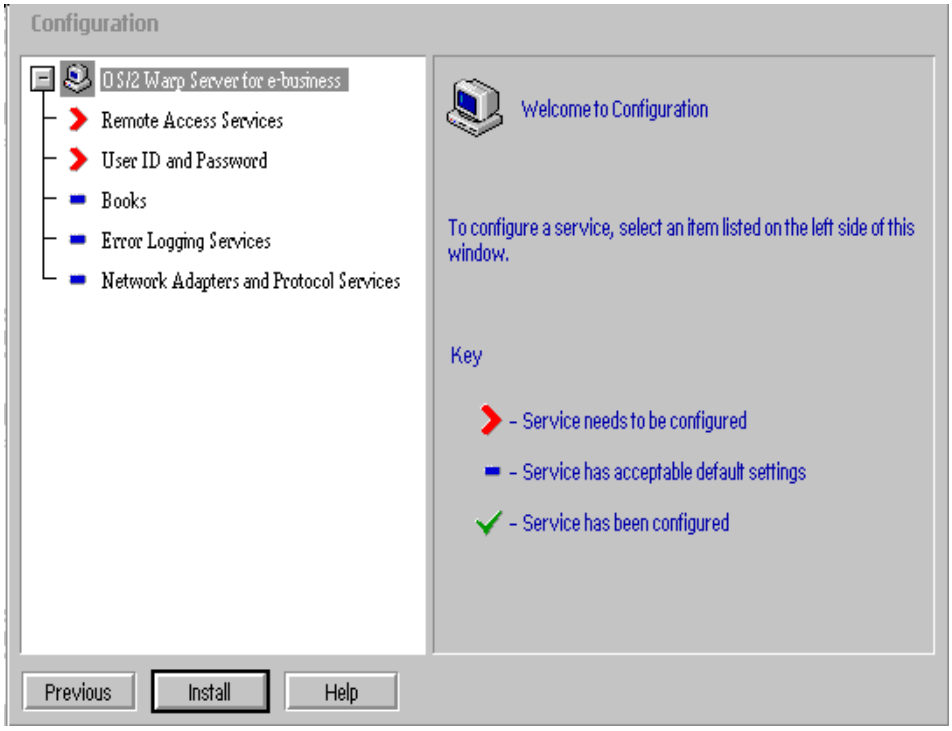

Figure 165. IBM remote access services configuration

Select Remote Access Services by clicking on the red arrow shown in Figure 165 on page 367. In Figure 166 on page 368, we configured the IBM remote access connection server for one communication port (COM2) and an IBM 7852-013 28.8 Modem.

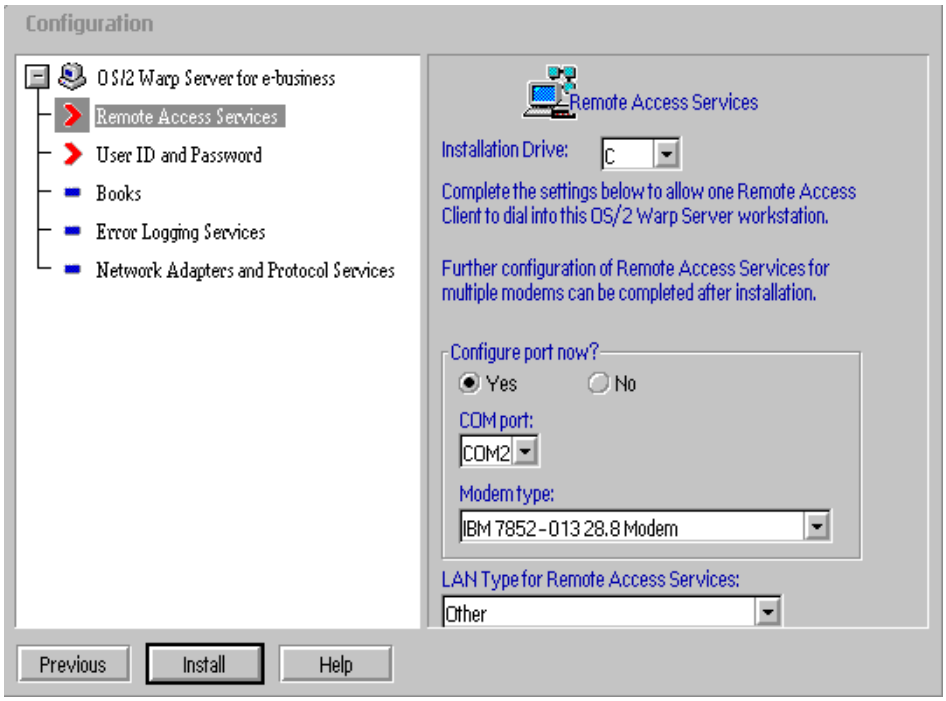

Figure 166. IBM remote access services configuration

Note

If you are installing a WAN adapter, select **No** in the Configure port now? field.

If your modem is not listed in the Modem type field, leave this set to none. You must run the CFMODEM utility after installation has completed.

The two options for LAN type are Ethernet and Other. If you have a token-ring LAN, select **Other**.

Select **User ID and Password** from the left-hand side of the Configuration window in Figure 167 on page 369. You must enter the administrator's user ID and password. Take note of the user ID and password that you enter, since they will be required later when you log on to Remote Access Services.

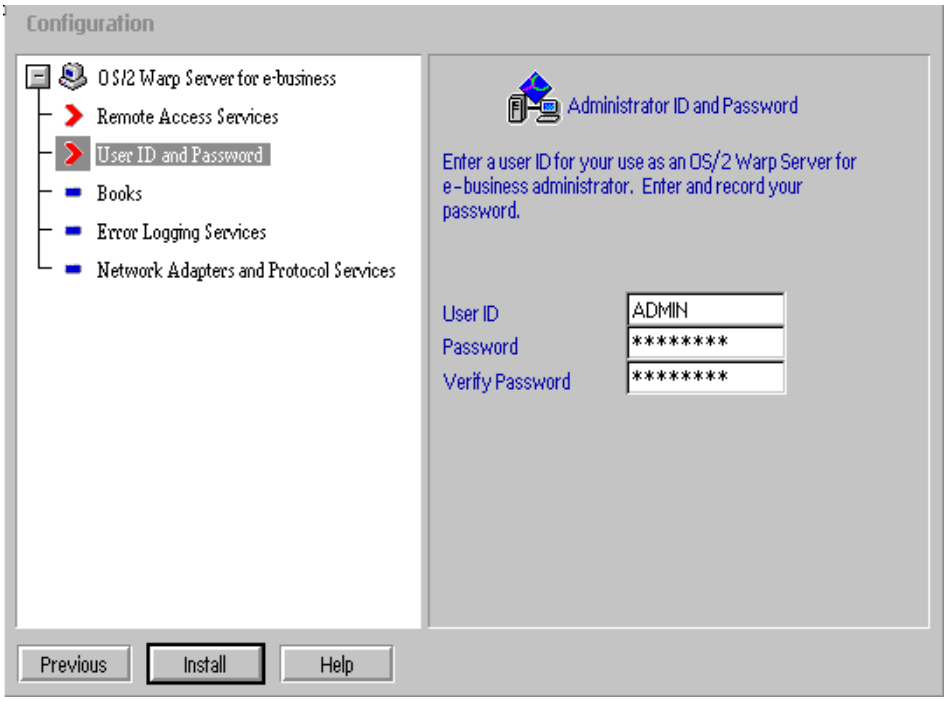

Figure 167. Configure user ID and password

Note

In previous releases of LAN Distance, a default administrator user ID called SECADMIN with a default passphrase of SECADMIN was created automatically. This user ID is no longer created automatically. Instead, the default user ID and passphrase are those that you entered during installation.

If you forget the administrator password and would like to re-initialize the security database, you can use the procedure described in the advanced users guide. If you do that, the default user ID/password is SECADMIN/SECADMIN.

Now, complete the installation by selecting **Install**. After the installation has completed, you have to reboot the server.

Table 68 on page 370 lists the changes that are made to the CONFIG.SYS and PROTOCOL.INI configuration files during installation of IBM remote

access services. Backups of these system configuration files are saved before the files are changed.

Table 68. Modification of CONIFIG.SYS and PROTOCOL.INI

File	Backup	Changes
\CONFIG.SYS	\CONFIG.WAL	- The WAL directory is added to your path specifications for LIBPATH, DPATH, and PATH. - The IBM remote access services helps are added to the HELP specifications. - The specifications for the IBM remote access services device driver are added. - The device drivers for Adapters and protocol services are added. - If FFST/2 is installed, appropriate statements for it are added. Statement for the locked file device driver are added temporarily to the top of your CONFIG.SYS file. The statements are removed the next time you start your workstation.
\IBMCOM\ PROTOCOL.INI	\WAL\ PROTOCOL.WAL	- If NetBIOS is installed during IBM remote access services installation, a section for NETBEUI_NIF is added. - Your NetBIOS timers are adjusted to minimize remote-access timeouts. - The number of NetBIOS NCB's, names and sessions are increased, if not already done by the OS/2 Warp Server for e-business Tuning Assistant. - A section is added for VLAN_kernel. - A section for PDFH_NIF is added.

9.6 Configuring IBM remote access services

After you have installed IBM remote access services and restarted your workstation, a Remote Access Services folder appears on your desktop. Within this folder, you will find an IBM remote access services icon. Double-click this icon to start the IBM remote access services.

Figure 168. IBM remote access services logon option

After IBM Enhanced Remote Access has started as shown in Figure 168 on page 371, select the **Select** pull-down menu; then, **Open As,** and then **Settings**. Alternatively, you can right-mouse-button select **MyWorkstation** (in this case, LABSRV99) and select **Settings**. The configuration of IBM remote access connection server for OS/2 Warp Server for e-business, IBM Enhanced Remote Access Remote Client for OS/2, and IBM Enhanced Remote Access Remote Client for Windows is the same as that in previous releases of Remote Access Services and LAN Distance products, except for configuring support for PPP clients on the Connection Server. In the following sections, we discuss configuring the Connection Server for PPP support and configuring PPP clients on different platforms.

For information on configuring other parts of IBM Enhanced Remote Access, see the on-line documentation by issuing VIEW x:\BOOKS\A3T12MST.INF at an OS/2 command line. Also, refer to Chapter 7.4 of the IBM redbook *Inside OS/2 Warp Server, Volume 1: Exploring the Core Components*, SG24-4602.

9.6.1 Configuring PPP support on the connection server

There are several tasks required to configure the Connection Server to support PPP clients. These include:

- Reviewing IP address considerations for PPP clients.
- Configuring a TCP/IP protocol router, which includes:
  - Bind the IBM TCP/IP protocol to certain adapters.

- Set TCP/IP configuration information.

- Updating the \WAL\WCLLOCAL.INI file and adding the appropriate PPP operating parameters.
- Administering IP addresses for PPP connections. There are three ways of accomplishing this:

  1. Local IP address list
  2. DHCP server services
  3. Client IP address configuration

- Creating user IDs for PPP clients.

9.6.2 Reviewing IP address considerations for PPP clients

IBM Enhanced Remote Access Connection Server supports two ways of routing frames between the LAN and the WAN for PPP clients:

1. LAN and WAN are different networks - the Connection Server workstation has two interfaces configured, each with an IP address from a different subnet. PPP clients must be assigned an IP address from the same subnet as the WAN interface.
2. LAN and WAN are on the same network - the IBM TCP/IP proxy Address Resolution Protocol (ARP) feature allows two interfaces to be configured with IP addresses from the same network. For example, the LAN interface is using a Class B network number of 172.16.x.x, which does not understand subnets. The IP address 172.16.1.26 is assigned to the LAN interface, and 172.16.2.8 is assigned to the WAN interface. Using a router subnet mask of 255.255.255.0 caused frames to be routed between the LAN and the WAN. When a PPP client connects and is assigned an IP address of 172.16.2.x, the Connection Server adds ARP and route entries, as required, for the proxy ARP configuration. The TCP/IP router responds to ARP requests from the LAN to a PPP client with its own hardware address. When the Connection Server receives a frame that is intended for the PPP client, the TCP/IP router forwards the frame to the correct client based on the ARP and Route table entries.

Note

An alternative way to configure the two TCP/IP interfaces on the Connection Server is to use two different subnet masks. For example, all LAN workstations use a subnet mask of 255.255.255.0. Use 255.255.255.0 for the LAN interface and 255.255.255.128 for the WAN interface.

9.6.3 Configuring a TCP/IP protocol router

The IBM remote access connection server bridges all data frames from IBM remote access services Client between the LAN and the WAN. However, for PPP clients, data frames are sent to the TCP/IP stack that is bound to the LAN Distance logical adapter. Therefore, you must configure TCP/IP as a router between the LAN and the WAN. To configure TCP/IP for routing, you must bind the IBM TCP/IP protocol to the:

- LAN adapter
- LAN Distance logical adapter

Note

When the IBM remote access connection server is configured for transparent bridging and the TCP/IP router is configured, IBM remote access services clients will not be able to TCP/IP ping the LAN or WAN TCP/IP interfaces on the Connection Server.

9.6.4 Binding the IBM TCP/IP protocol to the adapters

To bind the IBM TCP/IP protocol to the adapters, use the Adapters and Protocol Support application. Enter `MPTS` from an OS/2 command prompt.

When the Configure window is displayed, select the **LAN adapters and protocols** radio button. Click on the **Configure** pushbutton.

Use the Adapter and Protocol Configuration window to bind the IBM TCP/IP protocol to the adapters.

IBM Enhanced Remote Access assumes the LAN adapter that is used for routing is the first adapter that is the same LAN type as the connection server (Ethernet or token-ring) and is configured for TCP/IP.

In the Current Configuration list box, highlight the LAN adapter you will use for routing. Highlight the **IBM TCP/IP protocol** in the Protocols list box. Click on the **Add** pushbutton directly under the Protocols list box, as shown in Figure 169 on page 374.

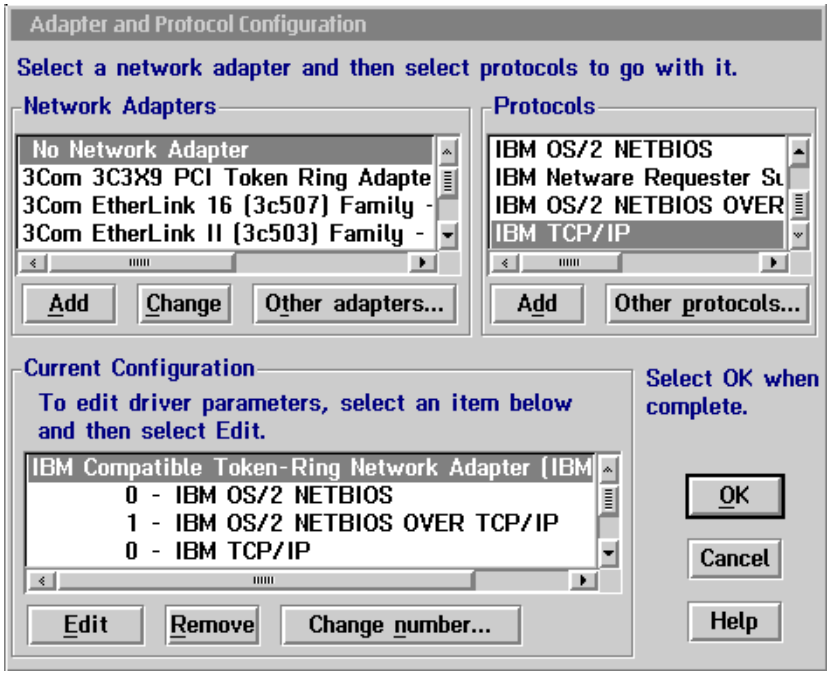

Figure 169. Adapters and protocol services - Binding TCP/IP to LAN adapter

The window refreshes and shows IBM TCP/IP below the LAN adapter you selected.

In the Current Configuration list box, highlight the **LAN Distance Logical Adapter**. Highlight the **IBM TCP/IP** protocol in the Protocols list box. Click on the **Add** pushbutton directly under the Protocols list box. The window refreshes and shows IBM TCP/IP below the LAN Distance logical adapter you selected.

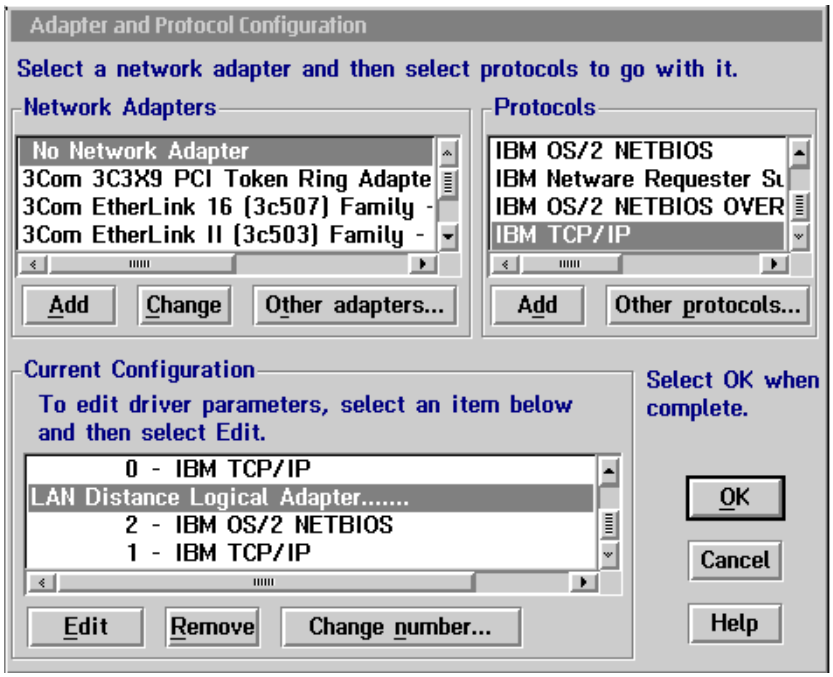

Figure 170. Adapters and protocol services - Binding TCP/IP to LAN adapter

Select the appropriate pushbuttons to close the Adapters and Protocol Support application.

The application updates the CONFIG.SYS file. You are asked to select the appropriate drives for the file and the Update CONFIG.SYS checkbox.

Select the appropriate pushbuttons to close the application.

After you close the application, you must restart OS/2 Warp Server for e-business for the changes to take effect.

9.6.5 Setting TCP/IP configuration for PPP

After you have bound the IBM TCP/IP protocol to the adapters, you must set TCP/IP configuration information for PPP. You must configure a TCP/IP LAN interface for the LAN adapter and, also, for the LAN Distance logical adapter.

1. Start TCP/IP configuration and follow the procedure that best describes your configuration:

- If TCP/IP is not installed, use Multiprotocol Transport Services (MPTS). Type `MPTS` at an OS/2 command prompt or double-click on the **Adapters and protocol services** icon in the System Setup folder.
- If TCP/IP is installed, use the TCP/IP Configuration application, which can be starting by entering `TCPCFG2` from an OS/2 command prompt or by selecting the **TCP/IP Configuration icon** in the TCP/IP folder.

2. With either of the above configurations, you must configure an interface for the LAN adapter. Then, configure an interface for the LAN Distance logical adapter. In this example we use two static IP Addresses (9.3.240.200 and 9.3.240.201) for the IBM remote access connection server.

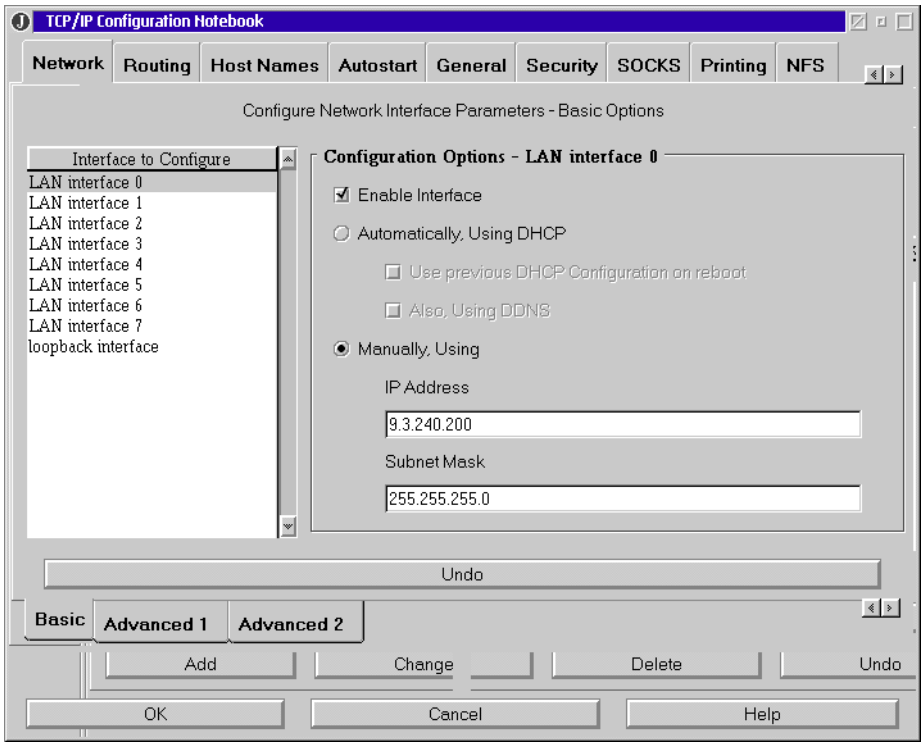

Figure 171. Network section of the TCP/IP configuration notebook

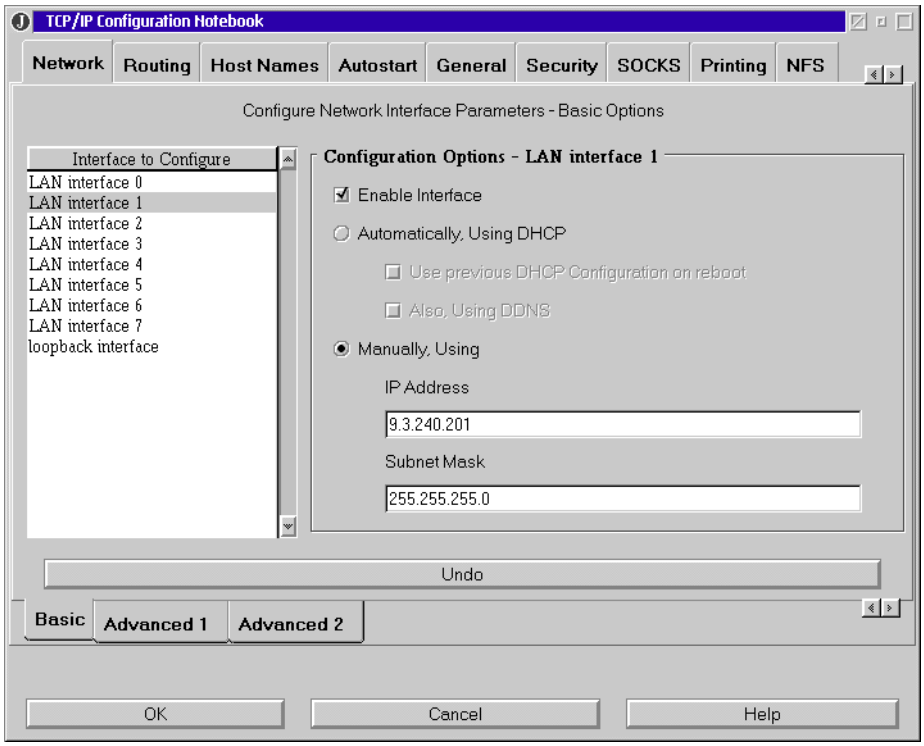

Figure 172. Network section of the TCP/IP configuration notebook

Note

Do not use DHCP to obtain the LAN and WAN interface IP addresses for the IBM remote access connection server. Servers should use preassigned addresses for both LAN interfaces.

3. Next, create a default route.

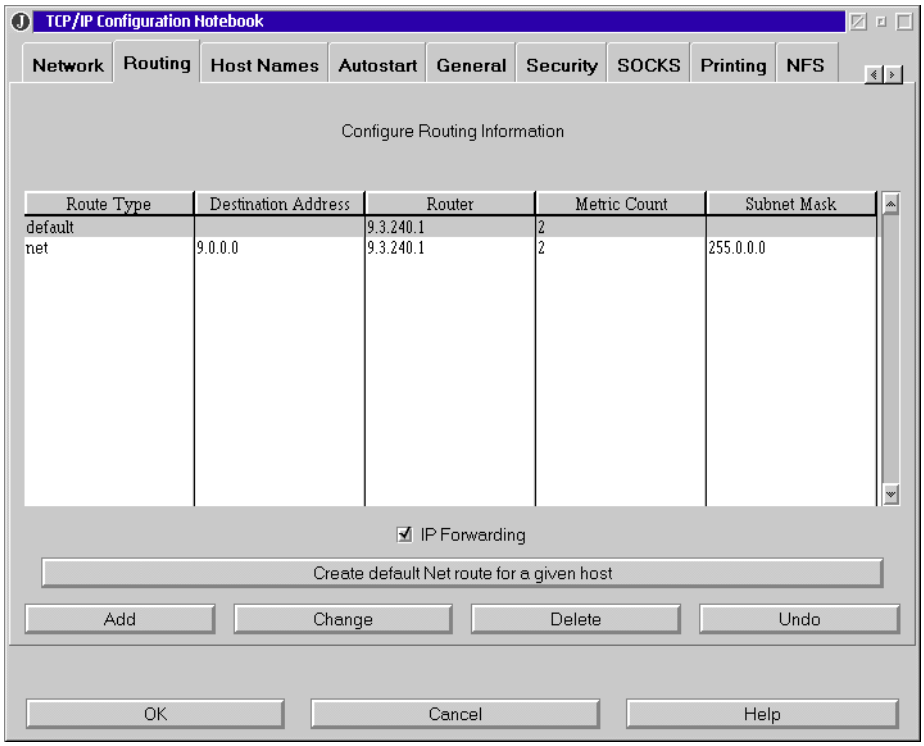

Figure 173. Routing section of the TCP/IP configuration notebook

4. Some networks require dynamic routing. If your network requires dynamic routing, then start ROUTED. TCP/IP must be installed in order to start ROUTED.
5. After you have completed these steps, close the application.
6. You must shut down your workstation and then restart it for the changes to become effective.

9.6.6 Specifying PPP parameters in \WAL\WCLLOCAL.INI

The WCLLOCAL.INI file contains new parameters that indicate if the Connection Server should configure an IP address for the Remote Client, where to get the IP address, and whether a host name should also be configured.

When the LAN Distance connection server starts, it gets specific operating information from the \WAL\WCLLOCAL.INI file. You need to include operating parameters for point-to-point protocol (PPP) in this file.

To configure the \WAL\WCLLOCAL.INI file, edit it using an ASCII editor. Update a PPP section to the file as follows.

```
[PPP]
ObtainIPaddr={LIST, DHCP, USERSPEC}
DHCPMaxWait=seconds
DDNS={YES, NO}
ClientsDomainNm=domain name
pppSecurity={PAP, CHAP}
```

Figure 174. PPP section of the \WAL\WCLLOCAL.INI

The following table lists some additional information about the parameters:

Table 69. PPP section of the WCLLOCAL.INI

Parameter	Value
ObtainIPaddr=	Indicates where PPP gets the IP addresses for a PPP session. Select one or more of the following: – LIST = Get IP addresses from a list in the \WAL\WCLIPADR.INI file. – DHCP = Get IP addresses from a DHCP server that is available on the LAN. – USERSPEC = The client workstation or workstations have configured their own IP addresses. LIST, DHCP and USERSPEC can be used together. The server tries the options in the order specified.
DDNS=	Indicates whether PPP will enable DDNS for PPP sessions. – YES = PPP will enable DDNS. – NO = PPP will not enable DDNS. Either YES or NO must be selected but not both. The DHCP server uses the LAN Distance user ID to update the DNS server when using this method if no name is defined on the Remote Client.

Parameter	Value
pppSecurity	Indicates whether PPP uses PAP security or CHAP security for user authentication. – PAP = Use PAP security for user authentication. – CHAP = Use CHAP security for user authentication. PAP and CHAP can be used together. The server tries the authentication options in the order specified. We recommend that you specify CHAP before PAP.
ClientsDomainNm=	Indicates the domain name of the client. If you do not specify a value, the domain name is obtained from \MPTN\ETC\RESOLV2 on the IBM Enhanced Remote Access system.
DHCPMaxWait=	Indicates the interval, in seconds, that IBM remote access services waits for a response from any DHCP server. You can specify a value from 2 to 40 seconds. The default timeout value is 3 seconds. IBM remote access connection server will retry one time. Since the server will always retry one time, the client must be able to wait for twice the number of seconds specified by the parameter. For example, if the DHCPMaxWait parameter is set to 5, the client must be able to wait for a connection response for at least 10 seconds. Some PPP clients will time out in as little as 8 seconds.

In our example, the WCLLOCAL.INI looks like this:

```
[PPP]
ObtainIPaddr= LIST, USERSPEC, DHCP
DHCPMaxWait= 3
DDNS= NO
ClientsDomainNm= DOMLAB01
pppSecurity= CHAP, PAP
```

9.6.7 IP addresses in WCLIPADR.INI

The WCLIPADR.INI file contains the list of IP addresses maintained by the IBM remote access connection server.

```
[IPADDRESSES]
x.x.x.x
y.y.y.y-z.z.z.z
```

Figure 175. \WAL\WCLIPADR.INI file

Where:

- x.x.x.x represents a single IP address that can be assigned for a PPP connection.
- y.y.y.y-z.z.z.z represents a range of IP addresses that can be assigned for a PPP connection.

In this example, the WCLLOCAL.INI looks like this:

```
[IPADDRESSES]
9.3.240.202-9.3.240.209
```

9.6.8 Using DHCP servers

If you will be using Dynamic Host Configuration Protocol (DHCP) to assign IP addresses to your PPP clients and have the client's LAN Distance user ID associated with the IP address on the DDNS server, your DHCP server must have a version of DHCP that contains the DNS proxy A record.

This feature is shipped with OS/2 Warp Server for e-business. If you are installing this service on an earlier version of OS/2 Warp Server, you will need the following APARS:

- OS/2 Warp Server APAR IC15366
- OS/2 Warp Server SMP APAR IC16980

DHCP servers provide the IP addresses and configuration information to DHCP and Bootstrap Protocol (BootP) clients on the network. DHCP servers contain information about the host operational parameters as specified by the network administrator.

By using a DHCP server, it is possible to assign an IP address to a client for a limited amount of time, and it also offers a way to supply all necessary configuration parameters with no end-user configuration required.

A PPP client that uses TCP/IP may or may not have a configured IP address. Clients that have not configured an IP address must be provided with one during PPP negotiation. This is provided by one of the following:

- DHCP server on the LAN
- IBM remote access connection server with IP addresses contained in the WCLIPADR.INI

9.7 Remote client

The LAN Distance Remote Client is unchanged from previous versions. For information on configuring a Remote Client, see the redbook *Inside OS/2 Warp Server, Volume 1: Exploring the Core Components*, SG24-4602.

9.8 Windows 95 PPP clients

PPP support is included with Windows 95 Version 4.0 through the Dial-Up Networking service. There are three components of Windows 95 that must be configured before this will work. They are:

1. Modem installation and configuration, which enables Windows 95 to operate your hardware.

2. Network configuration, which makes Dial-Up Networking and TCP/IP available.

3. Dial-Up Networking configuration, which allows configuration of the phone number and modem settings. The following sections show an example of configuring Windows 95 PPP on an IBM ThinkPad 760ED with an IBM International 28.8 Kbps data/fax PCMCIA modem.

9.8.1 Modem installation and configuration

You must install and configure your modem in Windows 95 by following the instructions provided by the manufacturer. If you are using a desktop machine, the COM ports and modems will be configured when you are installing Windows 95 or after you have purchased a new modem.

In this example, an IBM International 28.8 Kbps data/fax PCMCIA modem is installed after Windows 95 is installed and running. The Hardware Manager

detected that the card was installed in the PCMCIA slot of the ThinkPad and prompted installation of the device drivers by selecting the **Have Disk** option.

After installation of the modem device drivers, Windows 95 prompts you to restart Windows 95. When this has been completed, Windows 95 will be able to operate your modem.

9.8.2 Windows 95 network configuration

If you have not already installed any networking components for Windows 95, there will be no Network Neighborhood icon on the desktop. The following method of configuring network components will work whether or not you have a Network Neighborhood icon:

1. Select **Start** from the Task Bar.
2. Select **Settings**.
3. Select **Control Panel**.
4. You must select the **Network** icon.
5. You must select **Add** from the Configuration panel.
6. You must highlight **Adapter** and push the **Add** button.
7. Highlight **Microsoft** from the Manufacturers field and the Dial-Up Adapter will appear in the Network Adapters field. Now select **OK**. Notice that the TCP/IP protocol is not included in the default configuration of the Dial-Up Adapter. You must add TCP/IP by selecting **Add**.
8. Notice that IPX/SPX-compatible and NetBEUI protocols were added by default. You can remove the IPX protocol from the Dial-Up Adapter if you are not accessing NetWare servers from this workstation. This protocol can be removed by highlighting **IPX/SPX-compatible Protocol** and selecting **Remove**.

> **Note**
>
> By leaving NetBEUI enabled on the Dial-Up Adapter, you will receive an error message when you connect to the IBM remote access connection server since only TCP/IP is passing over the link between the Remote Client and the Connection Server. Just ignore the error message and continue with the connection so that TCPBEUI will be able to flow over the link.

9. Highlight **Protocol** and select **Add**. Highlight **Microsoft** from the Manufacturers field and highlight **TCP/IP** from the Network Protocols field, then select **OK**.

10. Select **OK**, and you will be prompted to restart Windows 95 for the changes to take effect.
11. Select **YES** to shut down and restart Windows 95. The initial configuration is now complete. The next step is to set up the actual usage of the Dial-Up Adapter.

9.8.3 Dial-up networking configuration

After Windows 95 has restarted, you must configure the Dial-Up Adapter.

1. Select **My Computer** and then select **Dial-Up Networking**.
2. Double-click on the **Make new Connection** icon and enter a name for the computer you are dialing. This will become the name of the icon in the Dial-Up Networking folder you will select to make PPP connections.
3. Select the type of modem you are using.
4. Once the modem has been selected, you must configure it. Select **Configure.**
5. The Maximum speed parameter you select must match your modem's capabilities as well as the Connection Server modem capability. For example, if your modem is capable of 33600 baud, but the Connection Server modem is only capable of 14400 baud, the duration of handshaking at connection time will be longer than necessary due to negotiation of the modems over a compatible speed. Once you have selected a suitable modem speed, select **OK**. Select **Next** to continue.
6. Enter the phone number of the Connection Server modem. This should be provided to you by the IBM remote access services administrator. You must also enter your country. After entering the Telephone number and Country code, select **Next**.
7. Complete the entry by selecting **Finish** when prompted. Once this is done, you will return to the Dial-Up Networking window. Your new configuration is represented by a new icon displayed in the window.
8. Configure the server communications parameters in Dial-Up Networking. To do this, right-mouse-button click on the new icon and select **Properties**.
9. Click on **Server Type**.
10. In the Type of Dial-Up Server field, select **PPP:Windows 95, Windows NT 3.5, Internet**.
11. Select **TCP/IP** in the Allowed network protocols field.
12. Select **TCP/IP Settings**.

13. You must configure your IP address, as requested by your Remote Access Services Administrator. For example, if your IBM remote access services has configured the IBM remote access services workstation to issue IP addresses using the ObtainIPaddr=LIST or ObtainIPaddr=DHCP parameters in \WAL\WCLLOCAL.INI, you must select **Server assigned IP address**. However, if you were advised by your Administrator to use your static IP address, select **Specify an IP address** and then enter the IP address in the IP address field.
14. Configure the domain name server to be used. If the parameter DDNS=YES is specified in \WAL\WCLLOCAL.INI, you should select **Server assigned name server address**. However, if DDNS=NO in \WAL\WCLLOCAL.INI, you should select **Specify name server addresses** and enter the Primary and Secondary DNS IP addresses as specified by your TCP/IP administrator. Leave the Primary WINS and Secondary WINS fields set to 0.0.0.0 since these are not required.
15. Select **OK** to return to the previous window.
16. Select **OK** in the Server Types window to complete the TCP/IP configuration of Dial-Up Networking.
17. Try to connect to the IBM Enhanced Remote Access Connection Server. Select the icon you just modified in the Dial-Up Networking window.
18. The Connect To window is displayed. Enter your PPP user ID (allocated to you by your IBM remote access services Administrator) in the User name field.
19. Enter your password and select **Connect**. The modem will dial the configured number and start a PPP connection. After the modems have completed the negotiation process, the Windows 95 PPP client will be authenticated by the Connection Server.
20. After the authentication has completed, the Connected to window will be displayed using the name of the current remote workstation configuration (Remote Access PPP, in this case). You can check that TCP/IP protocol is active by selecting **Details**.
21. You can now use TCP/IP services, such as Lotus Notes, FTP, Web browsing, and so on. If you have the NetBEUI protocol loaded on your Dial-Up Adapter, you can also logon to LANs using the IBM Networks Client for Windows 95 (a free item available from the IBM Software Choice Website). You must first select the **Logoff** button, and, after logoff has completed, select the **Logon** button and enter your LAN user ID, password and domain in the Logon panel. When you have finished working, you can terminate the connection by selecting **Disconnect** on the Connected.

9.9 Windows NT Version 4 PPP clients

Install and configure Dial-Up Networking. The process used is the same as for Window 95.

For full description of installation and configuration of a Windows NT PPP client, refer to the IBM redbook *Network Clients for OS/2 Warp Server: OS/2 Warp 4, DOS/Windows, Windows 95/NT, and Apple Macintosh*, SG24-2009.

9.10 OS/2 Warp PPP clients

Two types of OS/2 Warp PPP clients will be installed and configured on OS/2 Warp Version 4 in the following sections:

1. IBM Dial-Up for TCP/IP Version 2.0

2. IBM 8235 DIALs Client Version 4.52

9.10.1 IBM dial-up for TCP/IP configuration

In this section, we show how to configure an OS/2 Warp Client to connect to an IBM remote access connection server. In our example, we use the following environment:

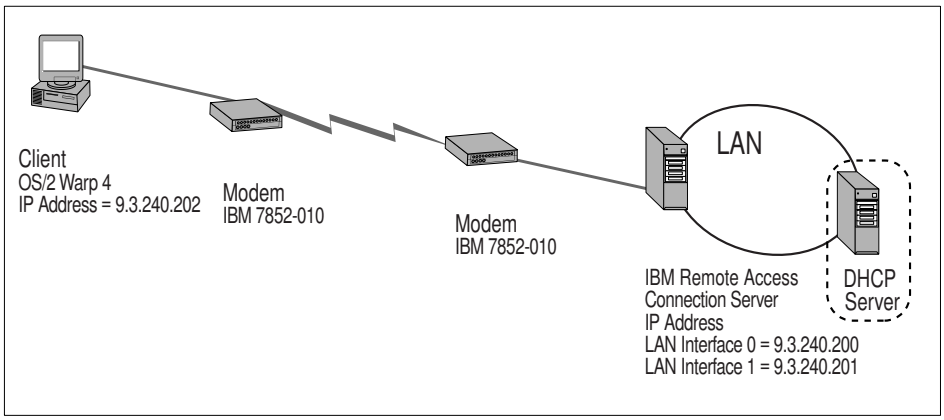

Figure 176. Example environment

Open the **Dial Other Internet Providers** object. You will find the object in the IBM Internet Connection for OS/2 folder under Programs.

Our program example is Revision 1.19, which can be found on the main panel by selecting the **Help** menu option and **Product Information**. You may have Revision 1.16.

1. Select **Add Entry** so that you can define the call and TCP/IP details.
2. Enter the details as follows (Figure 177 on page 388):
   - Name is the name of the call entry, such as PPP.
   - Description is a short description of the call entry, such as PPP ITSO.
   - Login ID is the PPP user ID assigned to you by the Connection Server administrator (in this example, we use DEFUSER).
   - Password is the PPP password assigned to you by the Connection Server administrator.
   - Phone Number is the phone number to dial to reach the IBM remote access connection server.
   - Connection Type should be set to PPP.

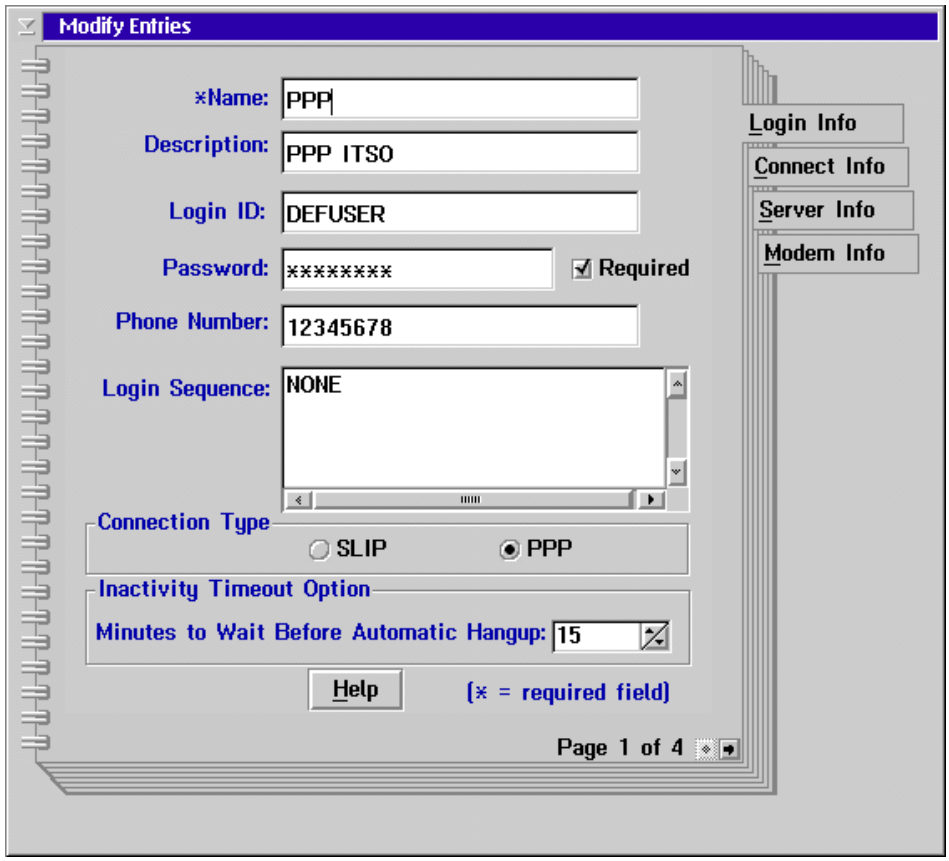

Figure 177. Dial-up configuration - Login info

Required fields are highlighted with an asterisk (*).

3. Select the **Connect Info** tab (Figure 178 on page 389).
4. Fill in details on this panel as provided by your Connection Server administrator. In this example, the IP address field is left blank since it is being allocated by the DHCP function of the Connection Server. The application automatically enters 1500, which is the maximum size for the Maximum Response Unit (MRU, largest data size for PPP transmission) and enables Van Jacobsen (VJ) header compression, which is used on TCP packets. The other two fields that are required to be entered are:
   - **Domain Nameserver** is the 32-bit dotted decimal notation IP address of the server that resolves host names to IP addresses; in our case, it is 9.3.240.2.

- **Your Domain Name** is the name of the TCP/IP domain in which your computer resides; in our case, it is itsc.austin.ibm.com.

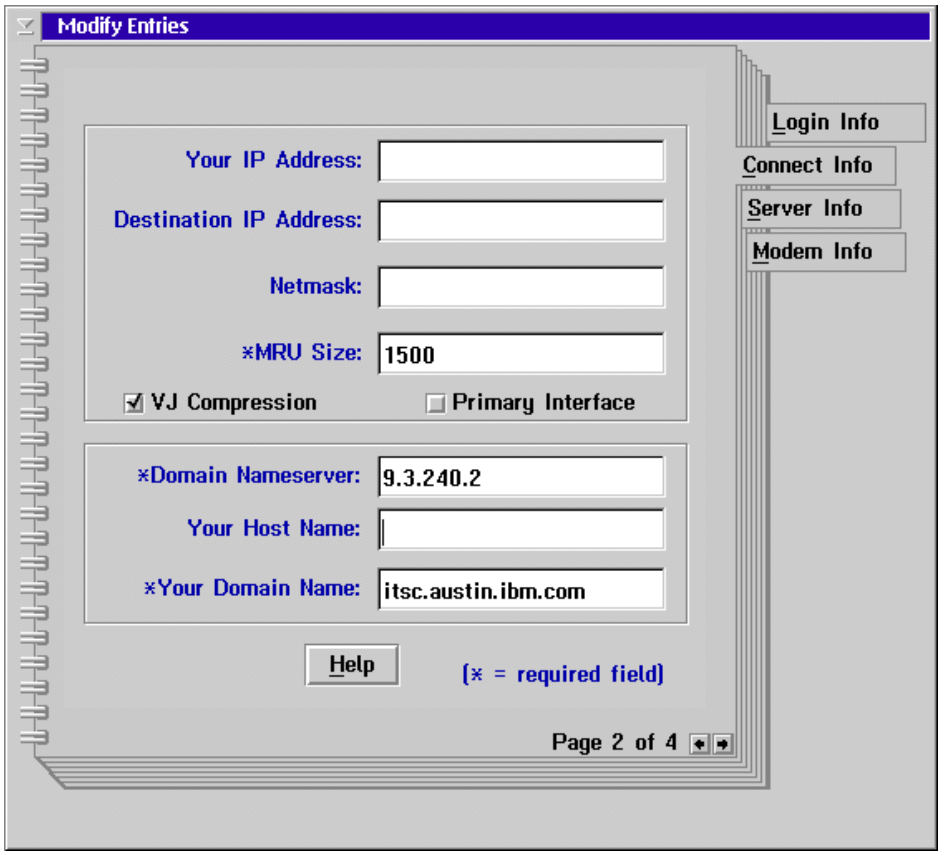

Figure 178. Dial-up configuration - Connect info

5. Select the **Server Info** tab to see the window shown in Figure. Enter the names of optional servers here, such as News Server, Gopher Server, and so on. It is not necessary to enter anything in these fields unless you want to run IBM Web Explorer.
6. Select the **Modem Info** tab to see the next window shown in Figure 179 on page 390.
7. Select your modem from the list in the Modem Type field.
8. Adjust the speed to the maximum your hardware will support.

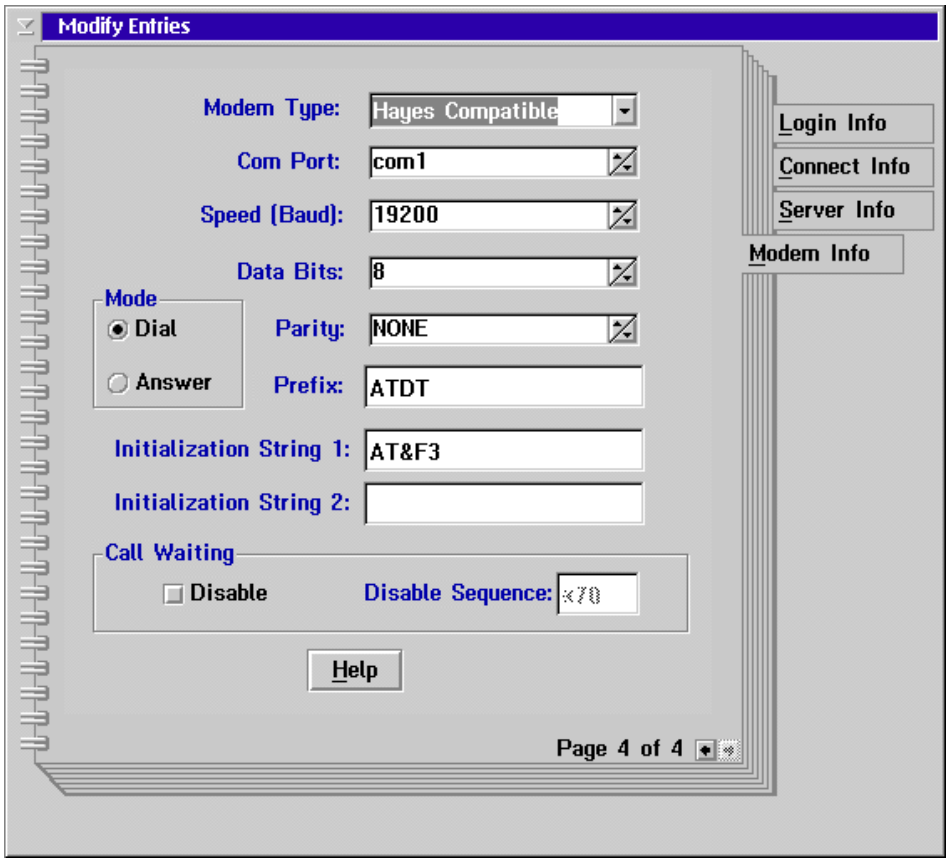

Figure 179. Dial-up configuration - Modem info

9. Close the Add Entries dialog by double-clicking on the title-bar icon, and the following window appears:

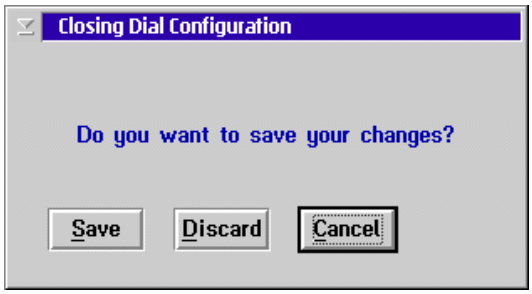

Figure 180. Closing dial configuration

10. Select **Save** to save the dial entry.

11. The IBM Dial-Up for TCP/IP window appears again, but, now, your new entry (PPP) appears.

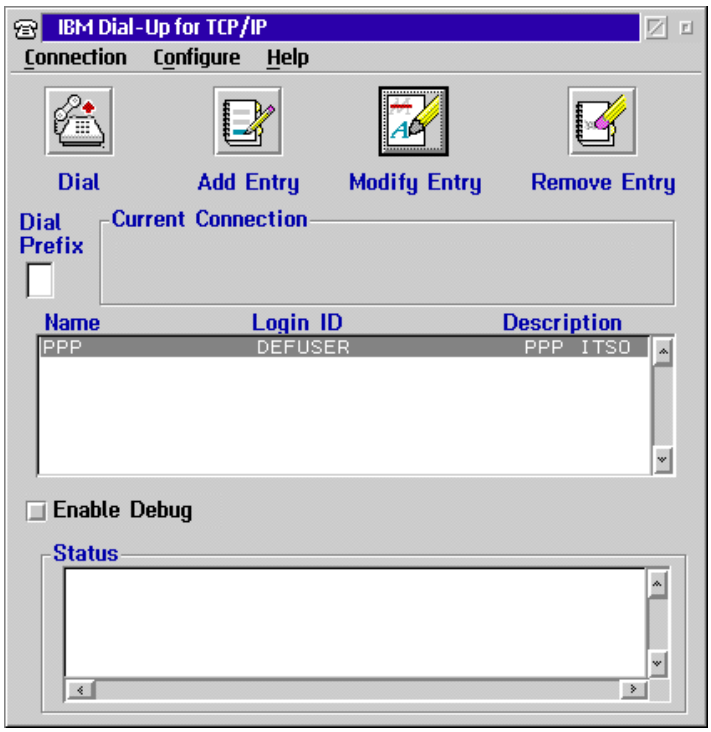

Figure 181. IBM Dial-up for TCP/IP

12. You can dial and connect to the IBM remote access connection server by clicking the **Dial** icon. The connection is made, and the connection messages appear in the Status window as shown in Figure 182 on page 392. It shows the connection information, such as CHAP authentication and the IP Address of the connected IBM remote access connection server.

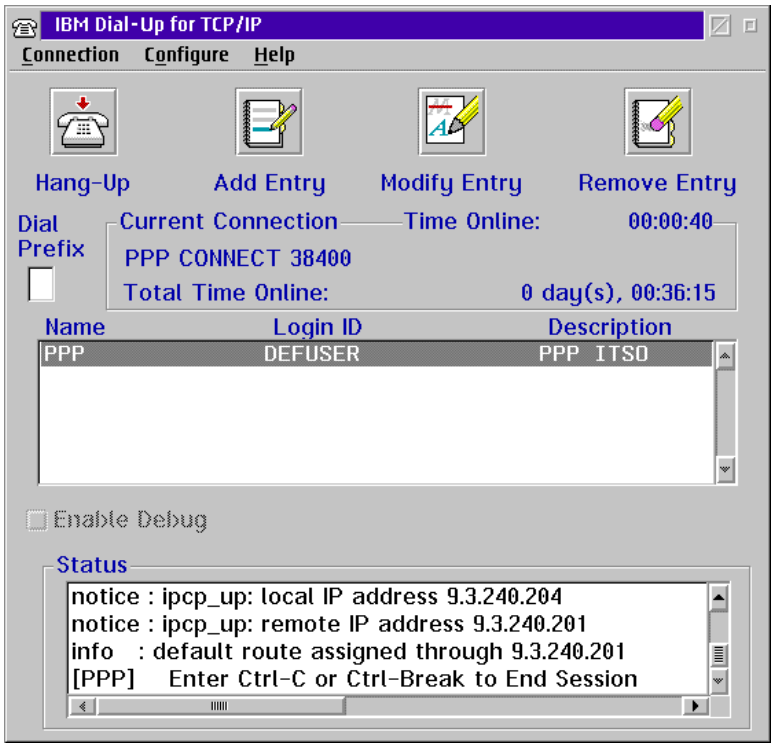

Figure 182. PPP dialer connect info

13. You may now begin using TCP/IP services like Lotus Notes, Netscape Communicator, FTP, and others. When you have finished working, you may end the connection by clicking the **Hang-up** icon.

9.10.2 Configuring IBM 8235 DIALs client for OS/2 Version 4.52

1. Install IBM 8235 DIALs for OS/2 Version 4.52 by inserting the diskette in drive A:, typing the command `A:\SETUP2`, and replying to prompts, such as which directory to install the software in.
2. Select the proper communications port, and, also, select your Modem Name from the drop-down menu list.
3. After the installation program has completed, you will receive an Installation Completed prompt.

Note

Do not run Shuttle to REMOTE yet. You must install the patches for the DIALs client first.

4. You must install the patch for IBM DIALs Client/2 Version 4.5.2. You can get the patch from the World Wide Web (WWW) at the following Web site:
   `http://www.networking.ibm.com/support/8235`
5. Follow the instructions included with the patch in the *.TXT file.
6. Shut down and restart your system.
7. Open an OS/2 window and back up the CONFIG.SYS, IBMLAN.INI and PROTOCOL.INI files. Exit the OS/2 window.
8. Open the **DIALs/2 folder** and select **Shuttle to REMOTE**. The first time Shuttle to REMOTE is run, a number of configuration files are created. A separate configuration for LAN connections are created with the files CONFIG.LAN and PROTOCOL.LAN in a new directory called \SHUTTLE2. There is also a separate configuration for remote connections in the SHUTTLE2 directory with files CONFIG.DIA and PROTOCOL.DIA. Each time you select Shuttle to REMOTE, you are in fact copying \SHUTTLE2\CONFIG.DIA to \CONFIG.SYS and \SHUTTLE2\PROTOCOL.DIA to \IBMCOM\PROTOCOL.INI. Similarly, when you select Shuttle to LOCAL, you are copying \SHUTTLE2\CONFIG.LAN to \CONFIG.SYS and \SHUTTLE2\PROTOCOL.LAN to \IBMCOM\PROTOCOL.INI. Since these configuration files are read at system startup, you must reboot your workstation for the changes to take effect.
9. After the Shuttle to REMOTE has completed, you must shut down and restart OS/2 Warp so that the DIALs Client NDIS driver can load.
10. When the OS/2 Desktop has completed loading, open the **DIALs/2 folder** and select **Connect/2**.
11. Enter a description for the connection in the Description field, such as DIALs to Enhanced Remote Access. Enter your PPP user ID in the Dial-in Name field. Also, enter the phone number to dial in the Phone Number field.
12. Now, select the Options button and enable the IP protocol. Also, enable NetBEUI/LLC if you are going to log on to servers using the TCPBEUI protocol.

> **Note**
>
> You will receive a warning after connecting with NetBEUI/LLC enabled. Ignore this error message because NetBEUI frames are actually sent over the link in IP frames according to RFC 1001/1002.

13. If you must use a static IP address, enter it in the IP Address field. If you wish to be allocated an IP address by either the DHCP server on your LAN or by the IBM remote access connection server, leave the IP Address to the default value - 0.0.0.0.
14. Select **OK** to accept the changes.
15. You can save this configuration by selecting **File** and **Save As**. Enter a suitable file name with an extension of .IR, such as LDPPP.IR. You may also create an icon in the DIALs/2 folder, by selecting **File**, then **Make Icon** and finally enter an appropriate name for the icon in the Icon Name field. Select **OK** to create the icon.
16. You may establish a connection to the IBM remote access connection server by entering your passphrase in the Password field and clicking on the Connect button.
17. After the connection has been established and your user ID has been authenticated, you may begin using IP services.
18. When you have finished work, you may select **Disconnect** to hang up.

9.11 Administration

The following sections discuss various required administration tasks.

9.11.1 PPP security

Since PPP clients authenticate with a PPP server using CHAP or PAP, the MAC address (either Universal Adapter Address (UAA) or Locally Administered Address (LAA)) is not required. This means that a person connecting using PPP can do so from any machine that is correctly configured.

This is different from LAN Distance and Remote Access Services clients, which can be restricted to (up to eight) workstations by including the UAA or LAA of the workstations they use to connect to the Connection Server.

The PPP Server security database has two password fields, one for the LAN Distance client password and one for the PPP client password. If the user ID

is used to dial in from both a LAN Distance client and a PPP client, the two passwords will be different because the first time that the user dials in from the LAN Distance client, he will be prompted to change his password. This is to make sure that only one person (the user) knows his password. So, in this case, the user has to remember two passwords, one for each type of dial in client.

An administrator upgrading an OS/2 Warp Server or OS/2 Warp Server SMP machine to use the new PPP server has to reset all the passwords in existing LAN Distance user IDs or create new user IDs. This is so that both password fields are initialized with the new password.

PPP client passphrases cannot be changed by the PPP client (except when you are using the process described in the section “Allowing PPP-clients to change their PPP passphrase” on page 398). This means that when the Connection Server administrator is requested to reset a passphrase or is entering a new user ID, the administrator must:

1. Logon as an administrator with the User ID and passphrase entered during the installation (Figure 167 on page 369) and right-mouse-button select **User Administration**.

   Note

   If the user ID or passphrase for the administrator is lost or forgotten, refer to Section 3.6.2 "Recovering the Administrator Passphrase" from the ITSO redbook *OS/2 Warp Server Functional Enhancements: Part 1*, SG24-2008.

2. Select the user ID which required a passphrase change (or select **Add** to enter a new user ID).
3. Select the **Passphrase** tab.
4. Enter a passphrase in the Passphrase field and the Verify Passphrase field.
5. Save the changes by double-clicking on the title-bar icon.
6. Logoff the administrator user ID.
7. Logon locally with the new user ID and use the initial passphrase.
8. After being prompted, enter a new passphrase and verify the new passphrase. This is the passphrase to be communicated to the user.

9.11.2 Address configuration

PPP IP addresses can be administered in any of three ways:

- They can be client-specified: If your site uses static IP addressing where a fixed IP address is allocated to a user, the Connection Server does not need to be configured to issue IP addresses.
- They can be listed on the IBM remote access connection server.
- They can be allocated by Dynamic Host Configuration Protocol (DHCP) services on the LAN.

The WCLIPADR.INI and WCLLOCAL.INI files must be customized to suit your site requirements. This means the IBM remote access services administrator will need to communicate with the IP administrator to decide the most suitable method for your site.

9.11.3 Base functions of the open as menu

After starting IBM remote access connection server on your OS/2 Warp Server for e-business, you will see a window like the following (where LABSRV99 is replaced by the name of your server):

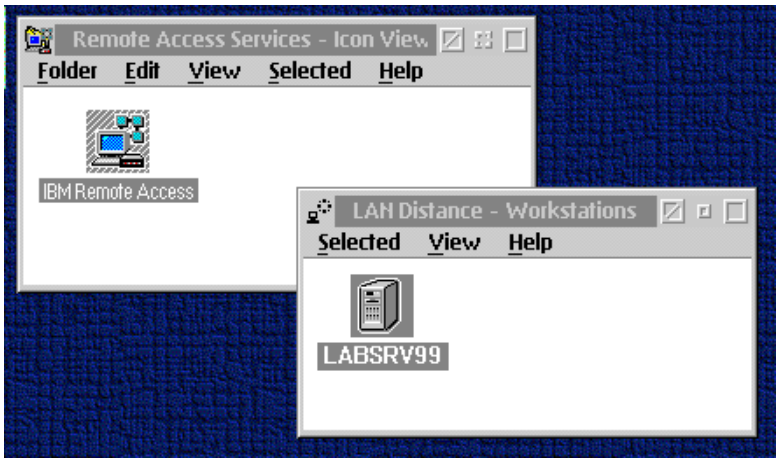

Figure 183. IBM remote access services folders

By clicking on **Selected** and then **Open as** you will see the Base Functions for this server.

- **Phone Book:** The phone book is normally used on a remote workstation to provide a list of LAN Distance Connection Servers that a remote workstation can contact. The phone book is also used on the connection

server to allow the connection server to call back users or to call other remote workstations or connection servers.

- **Call and port management:** Select this action to display information about the ports and modem types. If your user type is system administrator, you can stop, start, and restart port managers from this window.
- **Logged-on users:** Select this action to display a list of the users logged on at your workstation. (This list is displayed only if your IBM remote access connection server has Security enabled.)
- **Settings:** Use this action to configure IBM remote access connection server. This notebook can have up to 12 tabs. The tabs available to you depend on your workstation security and user type. Those not available to you are grayed-out when you open the Settings notebook. The tabs in the Settings notebook are:
  - Information
  - Phone Book
  - Answer
  - Dial
  - Ports
  - Modems
  - Bridge
  - Address/LAN
  - Workstation
  - LAPS
  - Timers
  - Security
- **Error log:** Select this action to start and view the OS/2 system error log.
- **Message log:** Select this action to start and view the message log from FFST/2.
- **Tracking:** Select this action to display and gather problem determination information (including traces, dumps, and audit trails). The functions available to you in this notebook depend on your user type. Those not available to you are grayed-out when you open the Tracking notebook.

9.11.4 Advanced functions of the open as menu

If your are logged on as an administrator (Selected -> Logon), you will find two additional functions in the Open As menu:

- **Personal account information:** This information is available to all users. Users can dial in, log on remotely, and change their passwords on this page. Use the first page of this notebook to view passphrase status information and to define or change your passphrase. Use the second page to view general information about your passphrase and to update the Comment field for your account. (This list is displayed only if your IBM remote access connection server has Security enabled).
- **User Account Management:** Use this notebook to add, change, delete, and view user accounts. (This list is displayed only if your IBM remote access connection server has Security enabled).

9.11.5 Allowing PPP-clients to change their PPP passphrase

The IBM remote access connection server for OS/2 Warp Server for e-business did not provide the same set of security features for its PPP clients that it provided for its LAN Distance remote clients.

This lack of security features for the PPP clients was the result of a requirement to have the PPP server work with existing PPP clients, which had not implemented such features due to a lack of direction in the industry.

One feature that was not provided was the capability to allow PPP clients to change their own PPP server access passphrases.

A workaround to this limitation that was developed was to have the IBM remote access connection server security administrator periodically hand out new passphrases based on his organization's security policy on a passphrase's maximum age.

This workaround was of concern not only to security administrators who had to spend extra resources to manage the passphrases but to many companies and organizations with security policies that required that an individual's access passphrase be under the individual's control and known only to that individual.

This section describes a solution that addresses those concerns.

9.11.5.1 Overview

An Internet solution has been developed to address the requirements that no changes be required to existing PPP clients and that the solution be platform-independent.

After a PPP client connection is established, the user will use a Web browser to access a *change passphrase* Web page that is provided by the IBM remote access connection server security administrator.

The URL resides on an IBM Internet Connection Server that has been installed on the IBM remote access connection server machine. The URL will contain the HTML for the change passphrase Web page and a CGI program to process input from the HTML form.

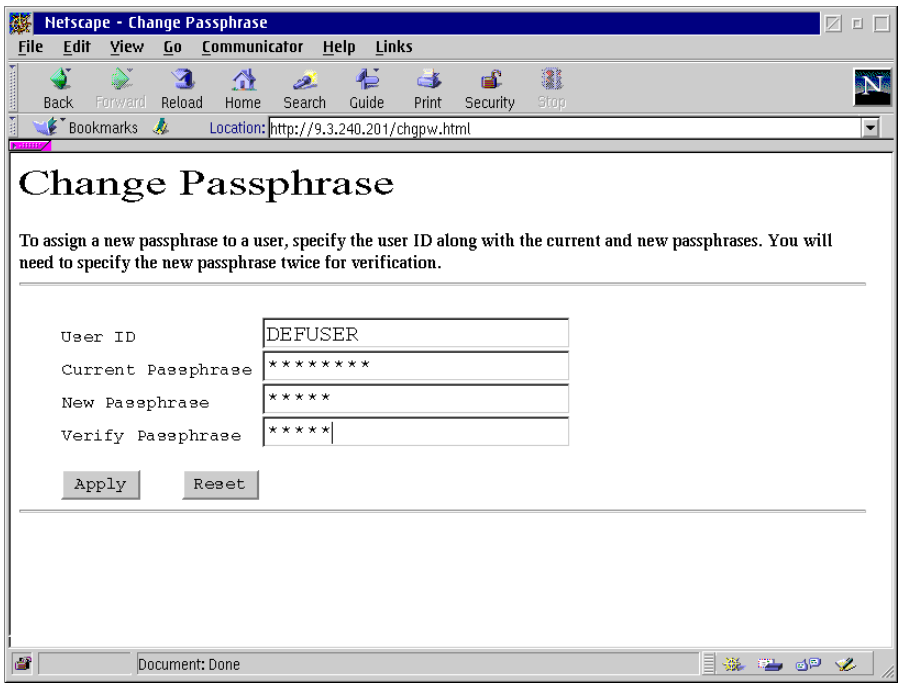

Figure 184. Change passphrase page

The CGI program will invoke a new IBM remote access connection server security API to update the server security database.

9.11.5.2 Prerequisites

The change passphrase Internet solution requires that you have the following software installed on your machines:

PPP server:

- IBM remote access connection server Version 5.11 for OS/2 Warp Server for e-business and all its prerequisites
- An Internet server (such as the Lotus Domino Go Webserver) and all its prerequisites
- A Web Browser (such as Netscape Communicator for OS/2)

PPP client:

- PPP client software (such as the IBM 8235 DIALs for OS/2 version 4.5.2 client)
- A Web Browser, such as Netscape Communicator for OS/2

9.11.5.3 Installation

This section describes the steps you should take to install the change password Internet solution on your IBM remote access connection server.

Step 1. Install the IBM remote access connection server security component.

Step 2. Install the HTML form and CGI program.

This step is dependent on the Internet server software that you installed on the IBM remote access connection server and on your firewall configuration.

For example, if you installed the Lotus Domino Go Webserver, a simple installation would be to copy the following files to the following directories:

1. copy \wal\chgpw.cmd to \www\cgi-bin\chgpw.cmd
2. copy \wal\chgpw.htm \www\html\chgpw.html

> **Note**
>
> If you do not place chgpw.cmd in the \cgi-bin directory, you will have to edit chgpw.html and change the *action* parameter on the `form` command to point to the appropriate directory.

9.11.5.4 Security issues

The following security items should be considered and addressed by the PPP server security administrator:

- Encryption of the passphrase
- Access to the change passphrase HTML form and CGI program

9.11.5.5 Password encryption

The change passphrase Web programs provide a very simple encoding method while transmitting the passphrases over the network.

You can enhance your security of the passphrases by using a SECURE version of the Internet Connection Server, which has the Secure Sockets Layer (SSL) enabled.

SSL uses a security handshake to initiate the TCP/IP connection between the client and the server. During the handshake, the client and server agree on the security level they will use, and the client authenticates the server. After that, SSL is used to encrypt and decrypt the information in both the client requests and server responses.

If you are interested in this level of security, see the Lotus Domino Go Webserver *WebMaster's Guide (Part 3. Security)* for more information.

9.11.5.6 Access to the change passphrase programs

The PPP server security API verifies that an authorized user is attempting to change the access passphrase by checking that the user ID is logged on from a PPP client, that the user ID is a *user* user ID, and that the IP address that the client used to log on with matches the IP address of the machine that is being used to change the passphrase.

It is also possible to prevent access to the change passphrase URL by configuring the Internet Connection Server to protect the directory that the change passphrase HTML form is located in. Access can be allowed based on ICS user names and passwords and/or through a list of authorized IP addresses.

For example, if you wish to limit access to the change passphrase form to machines with the IP addresses = 9.3.240.200-9.3.240.209, you would set up your Internet Connection Server as follows:

Copy C:\WAL\CHGPW.HTML into C:\WWW\HTML\PPP-SRV.

Copy C:\WAL\CHGPW.CMD into in C:\WWW\CGI-BIN.

Edit the ICS configuration file HTTPD.CNF and add the following statements in the protection section:

```
Protect /PPP-SRV/* {

Mask        @9.3.240.20*
```

}

Pass /PPP-SRV/* C:\WWW\HTML\PPP-SRV*

These statements protect access to the \www\html\ppp-srv directory. Client requests from (f.e.) 9.3.240.202 are allowed to access the directory. The following lines are extracted from the C:\MPTN\ETC\HTTPD.CNF file to reflect what the entry should look like:

```
# ================================================================== #
#
#       User authentication and document protection
#
# ================================================================== #


#       Within the configuration file, there are three directives that
#       define file access protection:
#         Protect, DefProt, and Protection.
#
# ....
#
#       Protection setup by hosts; you can use both domain name
#       templates and IP number templates. Inna i.l.D. F.v.S.
#
# Protection PROT-SETUP-HOSTS {
#         ServerId        YourServersFancyName
#         AuthType        Basic
#         PasswdFile      /where/ever/passwd
#         GroupFile       /where/ever/group
#         GET-Mask        @(*.cern.ch, 128.141.*.*, *.ncsa.uiuc.edu)
# }
# DefProt /very/secret/URL/*
# Protect /very/secret/URL/*      PROT-SETUP-USERS
# Protect /another/secret/URL/*   PROT-SETUP-HOSTS
#
#================================================#
# The following lines where added for the PPP-Server
#
Protect /ppp-srv/* {
        Mask    @9.3.240.20*
}
Pass /ppp-srv/* c:\www\html\ppp-srv\*
#
#================================================#
#

Protection  PROT-ADMIN  {
      PasswdFile   C:\MPTN\ETC\ADMIN.PWD
      Mask         All@(*)
      PostMask     All@(*)
      PutMask      All@(*)
      GetMask      All@(*)
      AuthType     Basic
      ServerID     Private_Authorization
}

Protect  /admin-bin/*  PROT-ADMIN
```

```
Protect /Usage*      PROT-ADMIN
Protect /servlet/extConfigServlet PROT-ADMIN


# ================================================================== #
```

See the Lotus Domino Go Webserver documentation for more information about the different options that are available for securing the Internet Connection Server.

9.12 IBM remote access services internal architecture

This section describes the architecture of IBM remote access connection server for OS/2 Warp Server for e-business, which allows to support PPP, LAN Distance and Remote Access Services Remote Client workstations. The PPP implementation is based on the current IBM TCP/IP implementation of the protocol.

When the Connection Server is configured for *flows* mode and an answer to a call is accepted, the server always waits for the client to respond in order to determine the type of client that has just connected. A LAN Distance or Remote Access Services client always sends an XID frame, and a PPP client always sends a PPP frame. The server will respond with the appropriate response based on the client type.

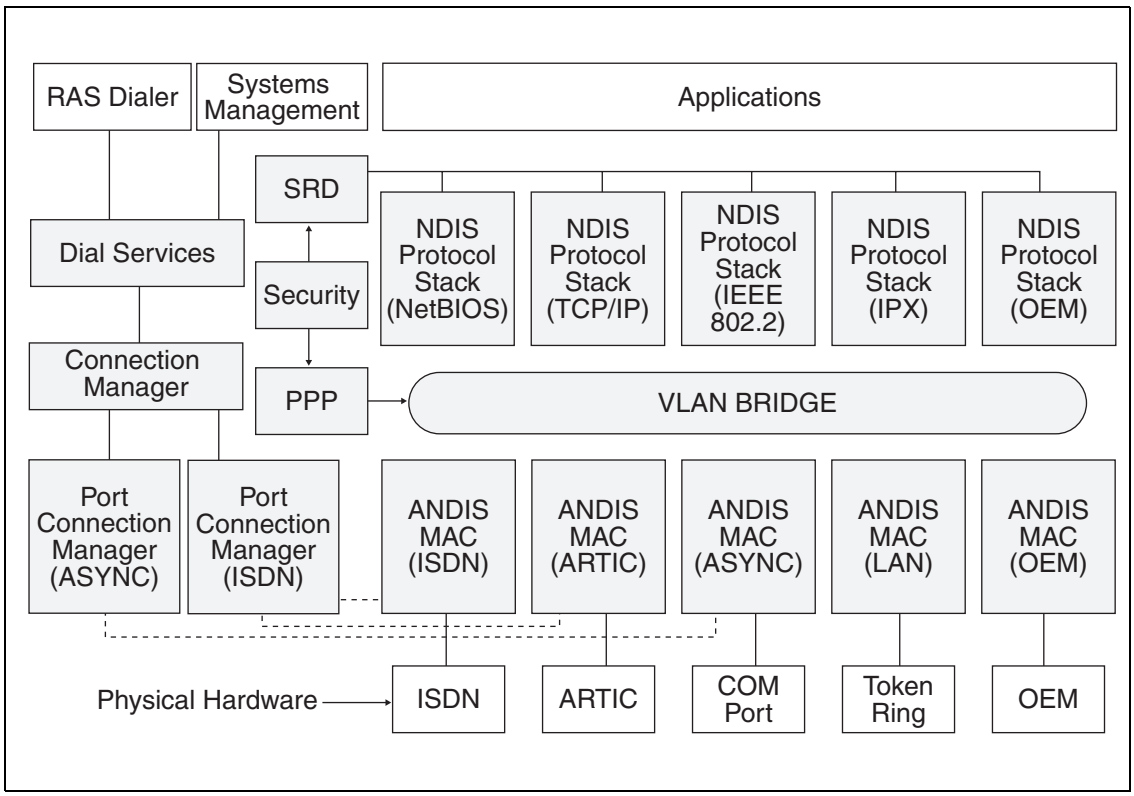

Figure 185. RAS, protocol architecture

The security component authenticates the client based on the client response at connection time. PPP clients use CHAP or PAP authentication while LAN Distance and Remote Access Services clients continue to use TPAP authentication. The differences between PAP and CHAP are now described:

- **PAP**: Password Authentication Protocol is a simple two-way mechanism which is done when the link between the remote client and server is initially established. Although RFC 1334 specifies that the default is to have passwords sent in clear text format, vendor implementations can encrypt this password before sending it over the link to be processed by the server.
- **CHAP**: Challenge Handshake Authentication Protocol (CHAP) is a stronger authentication mechanism (although defined by PPP extension RFC 1334). The server sends a CHAP Challenge packet to the remote client, which must calculate the proper value using a one-way hash function (MD5) and return it to the server. This is done at link

establishment and can be repeated with a new Challenge value anytime after the link establishment. This protects against intruders using a *record and playback* scheme to gain access to servers.

A LAN Distance or Remote Access Services client uses bridgeable LAN frames as the frame type to transport data between the client and the server and, for asynchronous links, uses ISO 3309 to encapsulate the LAN frame.

All PPP clients send raw protocol frames encapsulated in a PPP frame. Raw protocol frames can only be processed by the appropriate protocol stack. This requirement mandates that a routing function exists in the IBM remote access connection server. The Network Device Interface Specification (NDIS) IP router must be configured on the Connection Server to support PPP clients. The IP router function is provided by the MPTS component of OS/2 Warp Server.

A router requires that each side of the route be a separate IP subnet. IBM remote access services uses the Proxy ARP function provided by MPTS Version 5.11 which allows the LAN and WAN to be the same subnet.

PPP clients may or may not have an IP address configured. Clients that have not configured an IP address must be provided with one as part of PPP negotiation. The IBM remote access connection server provides the PPP client with an IP address if the \WAL\WCLLOCAL.INI file has the ObtainIPaddr=LIST parameter. The first IP address allocated is the first IP address in the file \WAL\WCLIPADR.INI, and subsequent IP addresses allocated are the next number from this file.

If the \WAL\WCLLOCAL.INI file parameter is ObtainIPaddr=DHCP, the IP address is assigned by a DHCP server on the LAN.

When an IP address is allocated by a DHCP server or the IBM remote access connection server, the PPP user ID is used as the TCP/IP host name.

The major component of IBM remote access connection server for OS/2 Warp Server for e-business is a component called PPP, which is contained in a DLL module. By necessity (for performance reasons), PPP is tightly coupled with the Virtual LAN (VLAN) component.

The VLAN has been modified to implement PPP framing and to interoperate with the PPP component. This design structure allows additional non-LAN Distance clients to be added in the future with minimal changes.

The VLAN tasks are divided into two major states:

- Logical connection phase, which includes authentication
- Runtime phase

The objective of this design is to minimize the changes to the VLAN component and the tasks it performs especially while in runtime phase. To accomplish this, most of VLAN's PPP frame handling is placed in the lowest layer VLAN subcomponent called MACFH.

During the connection phase, MACFH gives the incoming PPP frame to the PPP component for processing. When in the runtime phase, MACFH handles PPP encapsulation, PPP decapsulation, and adding a MAC header to the frame. The intent of this design is to limit the changes to the rest of the VLAN subcomponents by making the frame look like a bridgeable LAN frame that is destined for the protocol stack in the server.

The flows between the Connection Server and the LAN Distance or Remote Access Services client remain unchanged from previous releases.

When the Connection Server is configured for *no flows* mode, all frames pass through the VLAN unmodified as they do for previous releases of Connection Server code.

9.13 Remove IBM remote access services

Use LDREMOVE.EXE to uninstall IBM remote access services from your server.

Appendix A. More on Windows NT administration

This appendix contains more information on Windows NT administration. The aim is to try to consolidate the administration platforms. This section contains information on tools and utilities that are available.

A.1 Administration tools and utilities

Administration of a domain consisting of NT and Warp Servers requires the use of some special tools and utilities. These tools will help with the following NT Server administration tasks:

- Managing Security Access Control Management of NT Resources
- Starting and Stopping NT Services
- Viewing NT Server Event Logs

The network administration tools and utilities discussed will include:

- Netfinity Network Management
- Web Administration of Microsoft Windows NT Servers
- Utilities from the NT Server Resource Kit
- Utilities from the Internet

These will augment the regular administration functions provided with:

- Warp Server LAN Administration GUI

  and

- NT Server Manager

A.1.1 Server management

User and Group management within a domain can be handled seamlessly by the Warp Server LAN Administration. However, resource management for NT servers can only be partially handled through the Warp Server LAN Administration. Specifically, access permissions for NT 4.0 resources must be handled independently.

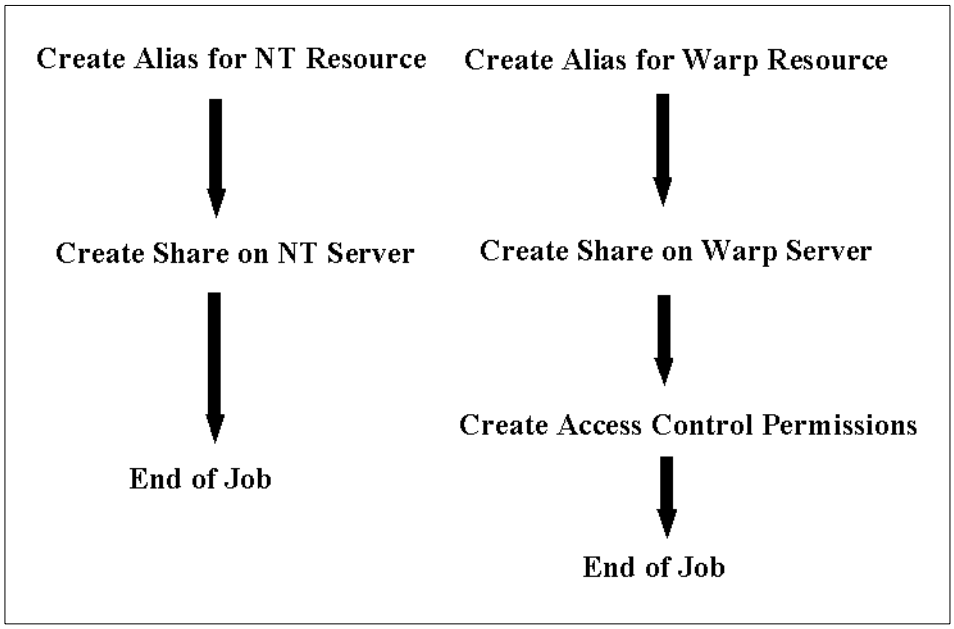

Figure 186. Warp Server alias creation flow

As you can see in Figure 186, the resource management for an NT resource leaves out the access permission step. There are a number of possible solutions to help you manage these resources. Of the tools and utilities just mentioned, we will look very closely at a couple. Namely, *Netfinity* and the *Web Administration tool from Microsoft.*

A.1.2 Netfinity client for NT 4.0

The following is an example of using the Netfinity Client for NT 4.0 as a management tool for the administration of NT 4.0 resource servers within an OS/2 Warp Server for e-business Domain. This topic will not cover the installation or setup of the Netfinity product in any great detail. For more information, view the product documentation. This topic will discuss the use of Netfinity as a Network Administration tool for the purpose of administration and maintenance of NT Servers within a Warp Server Domain.

After the Netfinity Client services are installed on the NT Server, you should immediately change the security access. This can be done by clicking on **Start**, then through Programs, Netfinity, to **Netfinity Service Manager**. Figure 187 shows the Netfinity Service Manager.

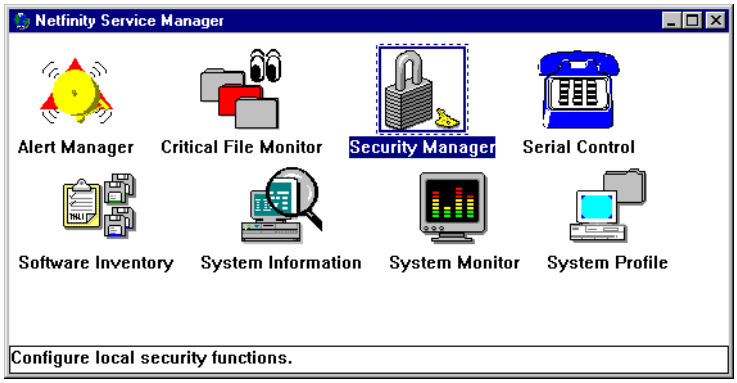

Figure 187. Netfinity service manager

The Security Manager should be selected to allow modification of the security access of this workstation. Figure 188 shows the post installation default, which allows public access to all services on this workstation.

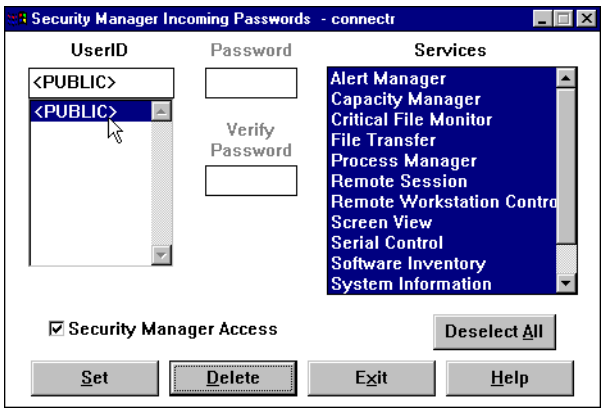

Figure 188. Security manager

Your first step should be to create an administrator ID that can be used for the remote NT Server administration. In our case, we created an ID called NTADMIN and assigned all services to that ID as shown in Figure 189.

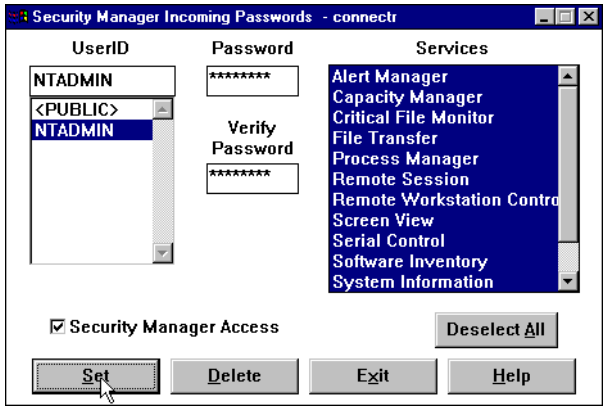

Figure 189. Security manager: Add user ID and enable services

Next, you should select the **PUBLIC UserID** and revoke all its services including the Security Manager Access. This is shown in Figure 190.

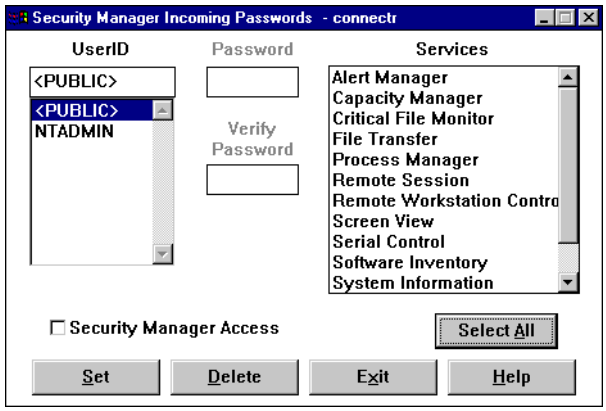

Figure 190. Security manager: Revoke services from public

Now that you have set up your Netfinity client services on the NT Servers within your domain, you can easily do remote control of these servers, which gives you full remote control of these servers and allows you to do NT graphical administration on these servers.

A.1.2.1 Example of setting up an NT file resource

Lets look at an example of setting up an NT Server file resource in the domain, then adding the necessary access permission on the NT Server for the resource. This is a two-step process:

Step 1 Set up an Alias for the resource in the domain

This is done on the Warp Server Domain from a Warp client. Start the LAN Server Administration GUI -> Domain Object -> Directory Resource Definition -> Create a directory from a Directory Template. Then, finish the process by filling in the *Directory Create* notebook. The example below shows just the final step (the notebook).

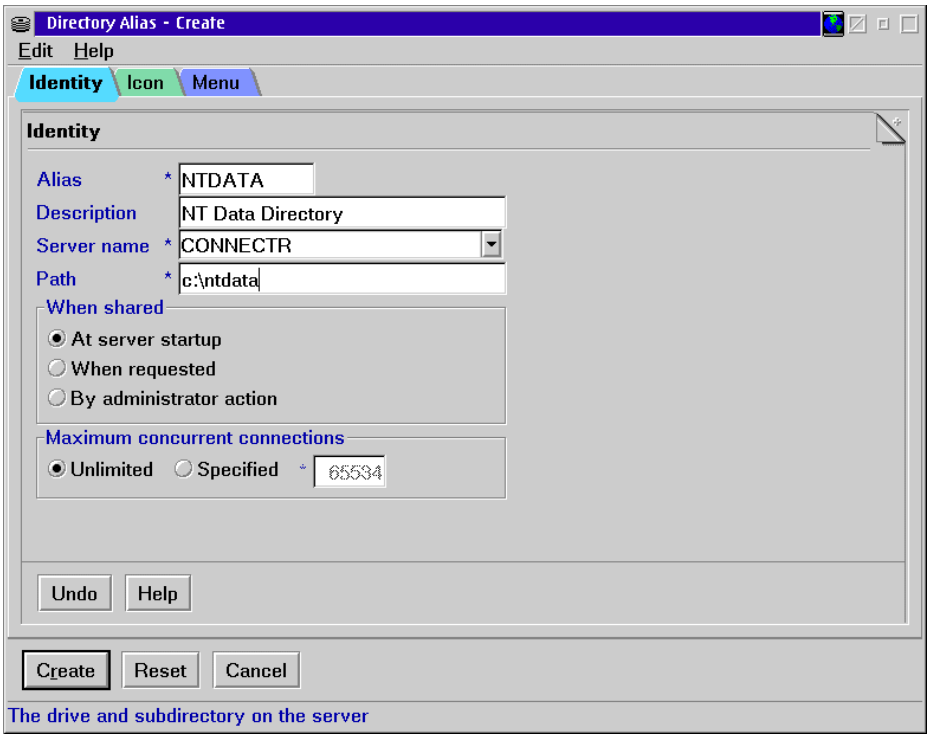

Figure 191. Directory alias - Create

Defining the alias will create the resource on the NT Server. From this point, you can use Netfinity administration to create the access controls for this resource on the NT Server.

Step 2 Create the ACL for the resource at the NT Server (Using Netfinity).

This is a matter of resource administration on the NT Server. You can use the NT GUI to set up share permissions and NTFS security access. The following is an example of using the NT GUI to check the share access permissions of the resource we just created via Warp Server administration. This NT administration can be done from any client with Netfinity manager to remotely access the NT Server. No need to run around and you have your choice of graphical or command line.

Using a Netfinity console, you would select the NT Server from the group of managed clients. Then you would select **Remote Workstation Control**.

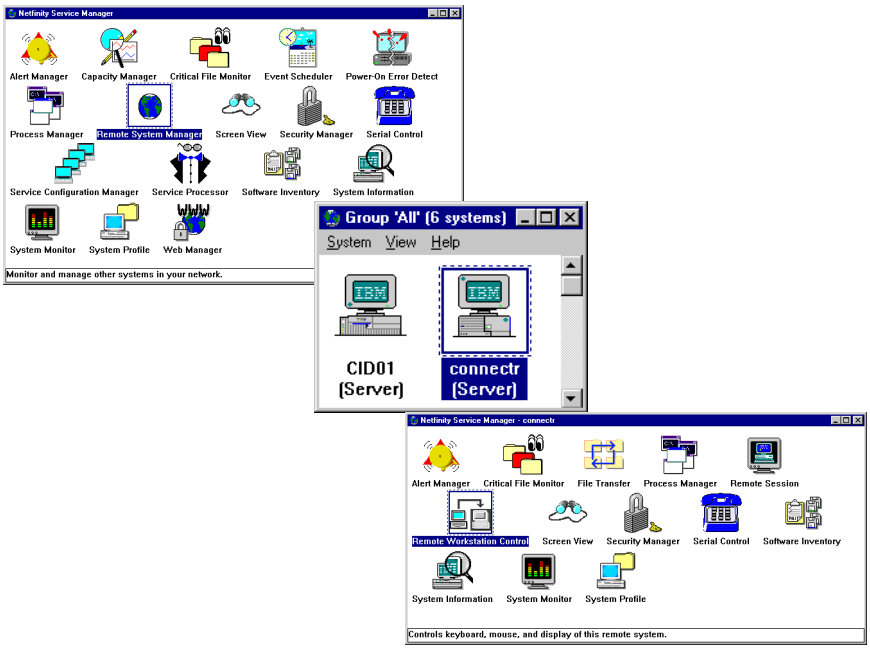

Figure 192. Select the NT server for remote administration

Use the NT GUI to check the share access permissions of the resource.

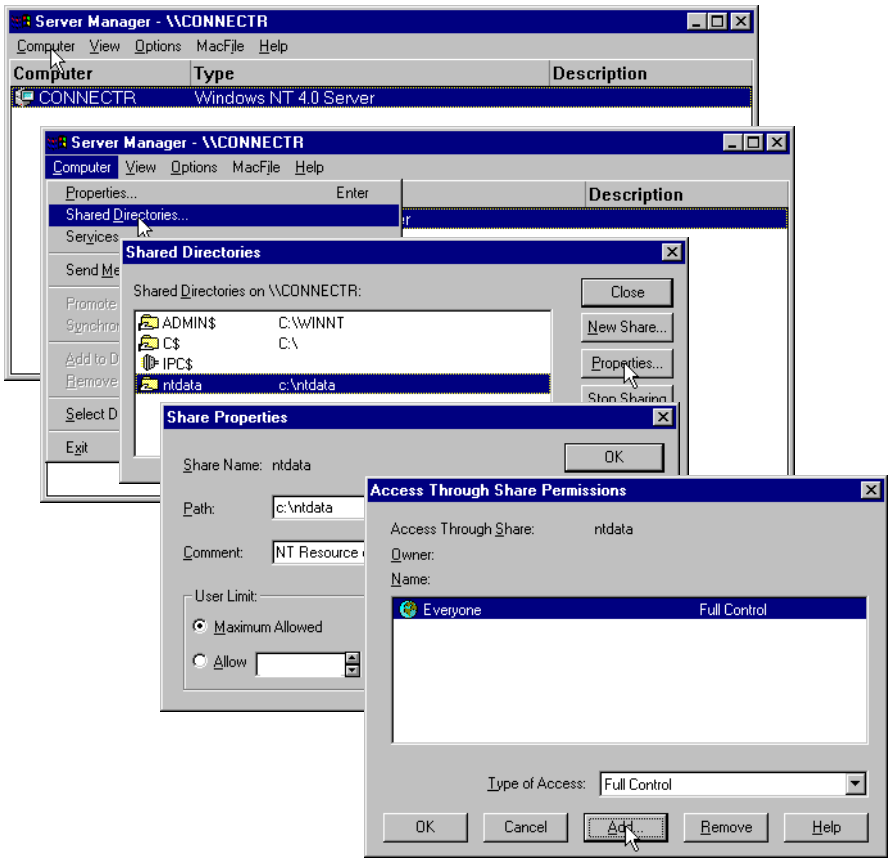

Figure 193. NT server manager - Share management

Figure 193 shows the use of the NT Server Manager to manage the share resources on the server. The resource and its associated permissions can be managed with this interface. Figure 194 on page 414 shows NTFS security resource management. Again, all this can be done from a central site on a single workstation using the Netfinity network manager, which is already a component of the OS/2 Warp Server for e-business product.

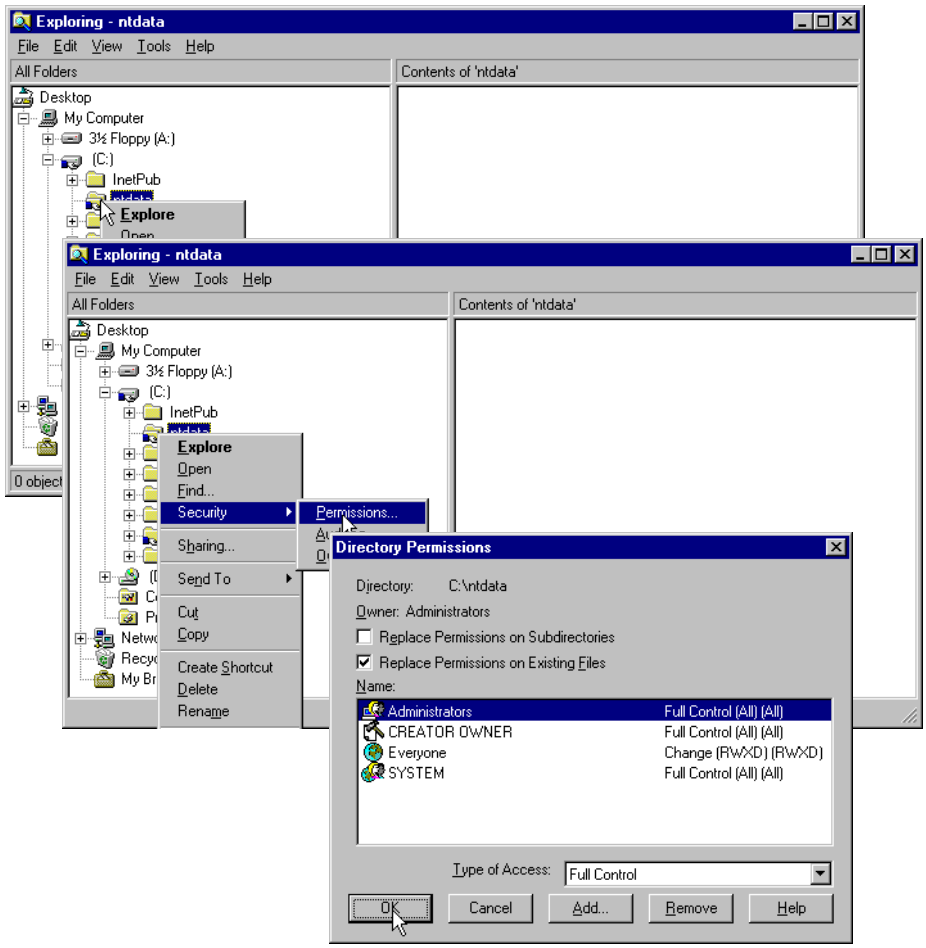

Figure 194. NTFS security access management

Netfinity provides a client for most workstations in the industry. This includes Netware and Warp in addition to Microsoft NT. With Netfinity Client installed on all your servers, your network of servers becomes very manageable. This makes Netfinity the most useful network management tool for administration of the servers within your Warp Sever for e-business domains.

If Netfinity Client is not an option for your NT servers, other options are discussed in the following sections.

A.1.3 Web administration of Microsoft Windows NT servers

The following is an example of installing and using the Microsoft Web Administration for Microsoft Windows NT Server Utility.

A.1.3.1 What is it?

Web Administration for Microsoft Windows NT Server enables you to remotely administer Microsoft Windows NT Server using existing HTML browsers running on Microsoft Windows, Macintosh, UNIX, and OS/2 platforms. Web Administration is not designed to replace existing administrative tools for Windows NT Server; instead, it is to enable you to perform limited administrative tasks when you are roaming away from your usual workstation and without access to traditional tools. Web Administration is a tool that is implemented to work in conjunction with Microsoft Internet Information Server.

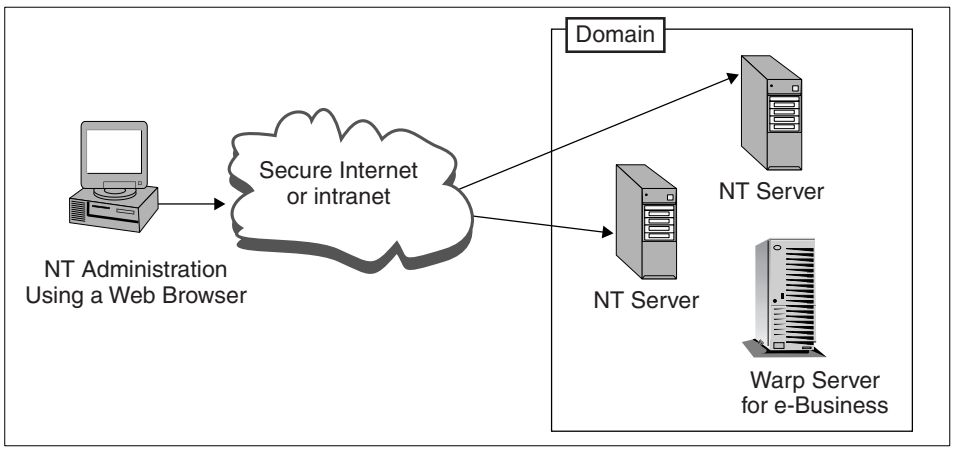

Figure 195. Microsoft Web administration

The Web Administration tool is intended for existing Windows NT Server administrators who have performed tasks with the regular administrative tools on Windows NT 4.0.

A.1.3.2 How does the software work?

You can install the Web Administration software on any server that runs Windows NT Server 4.0 and Microsoft Internet Information Server (IIS). Installing the Web Administration software on the server causes the server to publish web pages that include forms you can use to administer that particular server.

You can then use any web browser that supports either Basic or Windows NT Challenge Response authentication. You simply type `http://<your_server_name>/ntadmin/ntadmin.htm` in the address section of your Browser and begin administering the server.

A.1.3.3 Administrative tasks you can perform

The tasks most commonly performed by roaming administrators are supported by Web Administration. Our interest in this tool focuses on the administration of file and printer resources and the access control associated with these NT resources.

The actual interface is a series of HTML pages that the Administrator navigates through using a web browser. The interface is intended for administrators familiar with existing administrative tools (such as User Manager, Control Panel, Performance Monitor, and so on). Tasks to be performed can be very detailed, and the pages include some Wizard-like explanations to assist the user.

Web Administration supports several modes of security. Each server you administer must support Basic authentication, Windows NT Challenge Response security, or both. In addition, Secure Sockets Layer (SSL) can be used with either or both of these modes for encryption of your sessions data.

A.1.3.4 Security

As mentioned earlier, this tool supports two types of authentication security. Basic and Windows Challenge Response as well as session data encryption through the use of Secure Sockets Layer (SSL).

Basic authentication simply prompts the user for a name and password when the administrator accesses the server. The name supplied is checked against the members of the Administrators group on the server. Passwords are transmitted in clear text.

Windows NT Challenge Response is more sophisticated, and passwords are not transmitted over the wire. With this security, the administrator must be logged on to his or her computer with a user name that is a member of the Administrator group on the machine they want to administer.

In addition to these, you can also configure Web Administration to use the Secure Sockets Layer protocol. As mentioned earlier, SSL is complementary to user authentication. To use SSL, in addition to setting up the server to use SSL, you must obtain a certificate from a certificate authority such as VeriSign.

If your browser supports only Basic authentication, it is recommended that you also use SSL. You may also want to use SSL even if you use Windows NT Challenge Response because SSL encrypts all data in the session.

When you choose between Basic and Windows NT Challenge Response, you must take into account what is supported by the web browser you will use to administer the server. In our case, because we were using the Netscape browser that comes with OS/2 Warp Server for e-business, we had to choose Basic authentication.

As mentioned previously, you can set up a server to require the use of SSL to administer it using the Web Administration tools. To do so, after installing Web Administration on the server, use a web browser to connect to the server over the web to administer it. Click **Maintenance**, click **Web Admin Preferences**, then select the **Ensure use of SSL secure channel** check box and click **OK**. This sets the registry entry SSLRequired to 1. SSLRequired is in the HKEY_LOCAL_MACHINE\Software\Microsoft\Inetsrv_NTAdmin key.

When it comes to security, use common sense. Use the most secure method possible and do not leave your workstation while logged on to an administrative account or during an administrative session.

A.1.3.5 Example of Using the Web Administration Tool

The following is an example of using this tool to manage the resources of the NT 4.0 servers over TCP/IP. The tool is distributed on the Windows NT Resource Kit. The readme document details the installation on the NT Server; so, we will not discuss that here.

After Web Administration is installed on the NT Server and the Internet Information Server (IIS) service is started, you must make some security choices. On the NT Server go to the Internet Service Manager by clicking on **Start -> Programs -> Microsoft Internet Server (Common)** and click on **Internet Service Manager**. The following window shows the IIS services that are currently installed and their status.

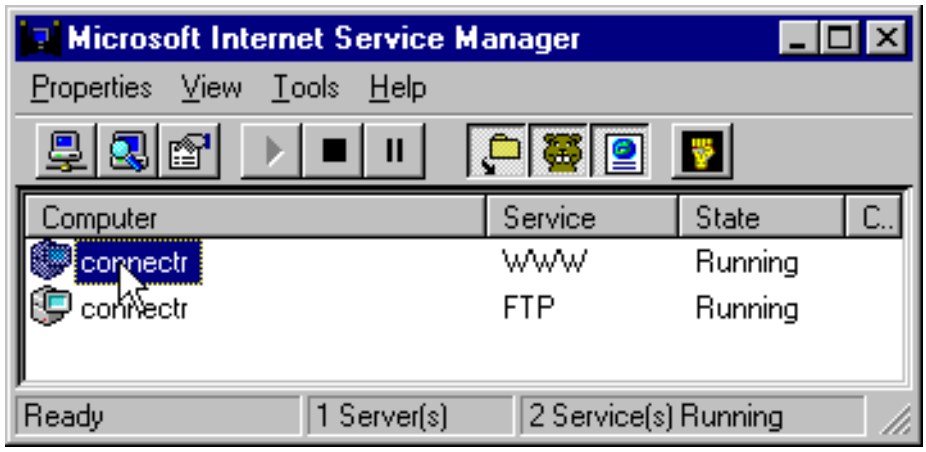

Figure 196. Microsoft Internet service manager

Bring up the properties of the WWW service by double clicking on **the WWW service line**. The properties page, as shown in Figure 197, will allow you to configure the type of security access to this WWW server. You should choose the highest type of security that your browser will support. As mentioned above, you should also enable SSL. In our case, we chose Basic authentication because we used the Netscape browser, which does not support Microsoft Challenge/Response.

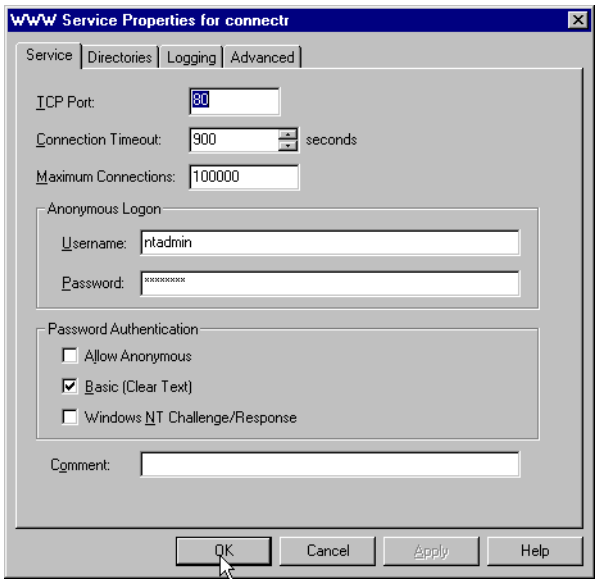

Figure 197. WWW service properties

With the WWW daemon running on the NT Server, a Web browser can be used from almost anywhere to administer the NT Server.

The format of the Web address is:
`http://your_server_name/ntadmin/ntadmin.htm`

This is the only way to start the Web Administration tool. You cannot start it by double-clicking any file in the file system or using localhost in place of the real server name. In our example, we typed:

`http://connectr/ntadmin/ntadmin.htm`

as the address.

Note

You may need to add an entry to the hosts file on the client in order to locate the NT Servers. This depends on the availability and status of DNS servers in your network.

If the user ID you are currently logging on with does *not* belong to the Administrators group on the NT Server, you will be prompted for a user ID and password. The user ID entered must belong to the administrators group of the NT Server being accessed.

Note

You many need to use the full name *domain_name\user_name.*

Provided the user ID entered (or the user ID you are logged on with) has administrative rights, you will see the introduction window as follows:

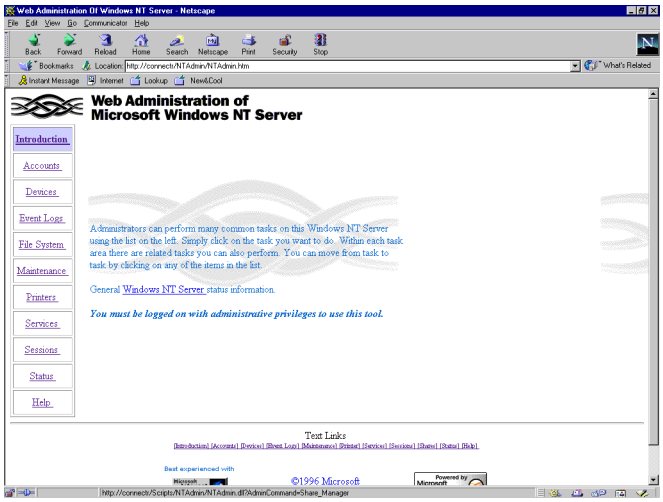

Figure 198. Web administration of NT server

As you can see from the following figures, you can do remote NT Server administration including security access control of file resources.

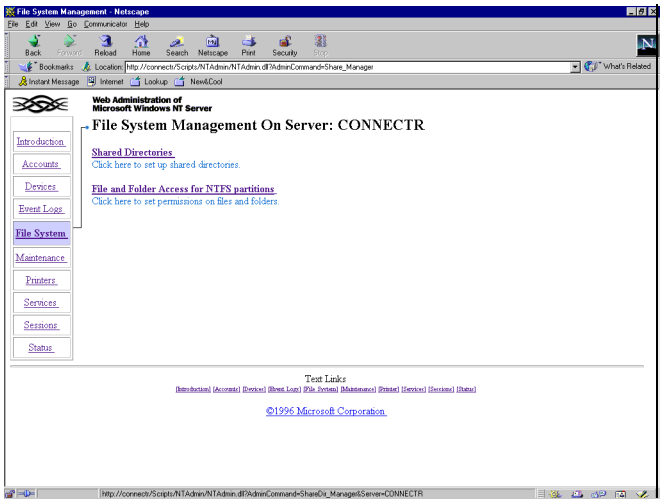

Figure 199. Web administration - File system management

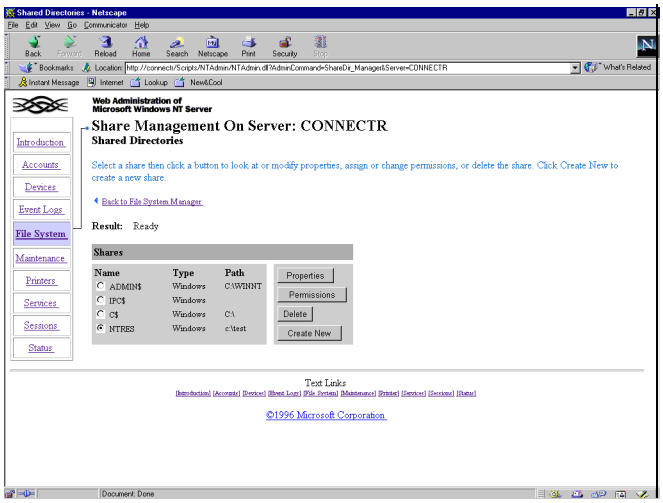

Figure 200. Web Administration - Share management

Figure 201. Web Administration - NTFS management

In addition to resource management, this tool can also be used to start and stop services including the IBM Network User Account Manager service on NT Servers. In fact, almost all NT Server administration can be done with this tool. One of the most useful is viewing the event log.

Note

The Viewing the Event Log tool gives you access to the event log on the NT Servers in your domain. This can be extremely useful for problem

The biggest concern with using this tool will probably be security. If you have a good firewall in your organization and there are no concerns from within your firewall, this tool could be very useful.

A.1.4 Other administration utilities from NT Resource Kit

Some command line utilities that require a Microsoft Client (usually an NT Client) to execute are:

- Remote Command Services

  The Remote Command Service (RCMD.EXE) provides a secure, way to remotely administer and run command-line programs.

- Remote Command Line

  The REMOTE utility allows you to run command-line programs on remote computers.

- Remote Console

  Remote Console is a client/server application that you can use to run a remote command-line session. In this command-line session, you can start any other application remotely. A Remote Console session resembles a Telnet session under UNIX.

- Remote Share

  RMTSHARE is a command-line utility that allows you to set up or delete shares remotely.

- Remote Service List

  SCLIST is a command-line tool that can show currently running services, stopped services, or all services on a local or remote computer.

- Perms

  PERMS.EXE is used to display users' access permissions for a specified file or set of files.

- XCACLS (Improved CACLS)

  XCACLS.EXE enables you to set from the command line all file-system security options that are accessible in File Manager or Windows Explorer. Before XCACLS, the standard way of setting file and directory access restrictions from the Windows NT command line was by using CACLS. With CACLS, however, it was not possible to specify access rights in the detail possible with File Manager (in Windows NT 3.51 or higher) or Windows Explorer (on the Windows NT 4.0 platform). Now, with XCACLS, you can view and modify the access control lists (ACLs) of files from the command line. XCACLS comes on the NT Server Resource Kit Supplement.

If you have a Microsoft NT Client available from your central administration facility, these utilities can be very useful for the administration of the NT resource servers.

Using these utilities in combination would be most useful for accomplishing remote administration. For example, Remote Console could be used to allow execution of XCACLS in order to update the access controls of NTFS file resources on the NT Servers in the domain. Batch utilities could also be created to facilitate routine maintenance or logging. These batch utilities could also utilize Rexx since the Regina Rexx interpreter comes on the NT Server Resource Kit.

In addition to command line utilities, there are GUI utilities distributed on the Resource Kit. Some of these may be useful in monitoring, maintaining, and administering NT servers such as:

- ShareUI
- Service Monitor

A.1.4.1 Tools from the Internet

If you want more function than the CACLS or XCACLS tools can provide, a very good ACL tool is Super CACLS and is available from:

`http://www.trustedsystems.com`

This tool allows you to capture ACLs into files using a number of filters or masks. You can then modify or update the captured files, which are batch files. The batch files can then be used to replace the ACLs. This gives the administrator the ability to back up, modify, and restore the access controls of the NT Server file resources. At the time of writing, this tool cost $200 for a 3 server 5 client license.

There are a number of other tools and utilities to help you manage NT file resources available from the internet. A Web site that lists some of these is:

`http://www.sysinternals.com`

Internet sites are moving targets; so, you may need to search for these and other tools of which there are many.

A.1.4.2 Technet subscription

The single most valuable source of technical information for all Microsoft products is the Technet database. The subscription fee is around $300 per year, but for most Microsoft related technical questions, Technet has the answers.

In addition to the technical database, there are utilities, Service Packs and other information, which include seminars, evaluation software, and Option Packages.

The Web site for subscription information is:

`http://www.microsoft.com/technet/`

A.1.5 Summary of administration tools and utilities

This topic has discussed some of the tools and utilities from which you have to choose. There is no shortage of tools and utilities. In many cases, the tool

you choose is a matter of personal preference: Command Line or Graphical User Interface. Editor's choice would be Netfinity in combination with NT Resource Kit Utilities.

Appendix B. CD-ROM contents

OS/2 Warp Server for e-business is a product that contains a significant amount of content. No single document could possibly contain enough information to cover all of it. A CD-ROM accompanies this redbook and includes some links to sites containing information on each of the chapters. All the sites were live at the time of printing this book.

The CD-ROM associated with this redbook is also availiable in softcopy format on the Internet from the redbooks Web site. Point your Web browser to:

`ftp://www.redbooks.ibm.com/redbooks/sg245393`

Alternatively, you can go to the redbooks Web site at:

`http://www.redbooks.ibm.com`

Select **Additional Materials** and open the file that corresponds with the redbook form number.

In most cases, for comfortable viewing, it is best to open links in their own windows. You can do this by selecting the second mouse button and selecting the option.

Video content has also been included to give you a quick overview of the various topics. In order to view the video content, you will need the Bamba player for either OS/2 or Windows depending on which platform you are using. Go first to the downloads section and download and install the Bamba plugin.

Some simple sample scripts are also provided. Unless otherwise stated, these products are supplied AS-IS, and IBM does not provide support for them.

Appendix C. Special notices

This publication is intended to help systems administrators and support personnel to understand the new features incorporated into OS/2 Warp Server for e-business. The information in this publication is not intended as the specification of any programming interfaces that are provided by OS/2 Warp Server for e-business. See the PUBLICATIONS section of the IBM Programming Announcement for the OS/2 Warp Server for e-business product 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.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact IBM Corporation, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee

that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Any pointers in this publication to external Web sites are provided for convenience only and do not in any manner serve as an endorsement of these Web sites.

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

IBM	Netfinity
Netfinity Manager	Network Station
Operating System/2	OS/2
OS/390	Presentation Manager
Print Services Facility	RS/6000
S/390	SP
Streamer	System/390
SystemView‘	ThinkPad
VoiceType	WebExplorer
WebSphere	XT
400	

The following terms are trademarks of other companies:

C-bus is a trademark of Corollary, Inc.

Java and HotJava are trademarks of Sun Microsystems, Incorporated.

Microsoft, Windows, Windows NT, and the Windows 95 logo are trademarks or registered trademarks of Microsoft Corporation.

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

MMX, Pentium and ProShare are trademarks of Intel Corporation in the U.S., other countries, or both.

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

SET and the SET Logo are trademarks owned by SET Secure Electronic Transaction LLC. (For further information, see http://www.setco.org/aboutmark.html)

Other company, product, and service names may be trademarks or service marks of others.

Appendix D. Related publications

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

D.1 International Technical Support Organization publications

For information on ordering these ITSO publications, see “How to Get ITSO Redbooks” on page 437.

- *TCP/IP Tutorial and Technical Overview*, GG24-3376
- *IBM LAN Distance Version 1.1 Configuration Customization Guide*, GG24-4158
- *Inside OS/2 Warp Server, Volume 1: Exploring the Core Components*, SG24-4602
- *Inside OS/2 Warp Server, Volume 2: System Management, Backup/Recovery and Advanced Print Services*, SG24-4702
- *OS/2 Warp Server Functional Enhancements: Part 1*, SG24-2008
- *Workgroup Management Using SystemView for OS/2*, SG24-2596
- *Inside OS/2 LAN Server 4.0*, SG24-4428
- *Network Clients for OS/2 Warp Server: OS/2 Warp 4, DOS/Windows, Windows 95/NT, and Apple Macintosh*, SG24-2009
- *WorkSpace On-Demand Handbook*, SG24-2028
- *OS/2 Warp Server, Windows NT, and NetWare: A Network Operating System Study, SG24-4786*
- *Understanding LDAP*, SG24-4986
- *TCP/IP Implementation in an OS/2 Warp Environment*, SG24-4730
- *Beyond DHCP - Work Your TCP/IP Internetwork with Dynamic IP*, SG24-5280
- *A Comprehensive Guide to Virtual Private Networks, Volume 1: IBM Firewall, Server and Client Solutions*, SG24-5201
- *NetFinity V5.0 Database Support*, SG24-4808
- *NetFinity V5.0 Command Line and LMU Support*, SG24-4925
- *Netfinity Server Management*, SG24-5208

D.2 Redbooks on CD-ROMs

Redbooks are also available on the following CD-ROMs. Click the CD-ROMs button at `http://www.redbooks.ibm.com/` for information about all the CD-ROMs offered, updates and formats.

CD-ROM Title	Collection Kit Number
System/390 Redbooks Collection	SK2T-2177
Networking and Systems Management Redbooks Collection	SK2T-6022
Transaction Processing and Data Management Redbooks Collection	SK2T-8038
Lotus Redbooks Collection	SK2T-8039
Tivoli Redbooks Collection	SK2T-8044
AS/400 Redbooks Collection	SK2T-2849
Netfinity Hardware and Software Redbooks Collection	SK2T-8046
RS/6000 Redbooks Collection (BkMgr Format)	SK2T-8040
RS/6000 Redbooks Collection (PDF Format)	SK2T-8043
Application Development Redbooks Collection	SK2T-8037

D.3 Product documentation

The following publications are product documentation:

- *Lotus Domino Go Programming Guide*
- *Lotus Domino Go Quick Beginnings*
- *Lotus Domino Go Webserver: Quick Beginnings*
- *Lotus Domino Go Webserver: WebMaster's Guide*
- *Microsoft Windows NT Server Books*
- *Network Administrator Tasks*
- *OS/2 Warp Server for e-business, Network Administrator Tasks*
- *Quick Beginnings: Installing OS/2 Warp Server for e-business*
- *Web Programming Guide*
- *Webmasters Guide*
- *WebSphere Application Server Guide*

D.4 Other publications

These publications are also relevant as further information sources:

- *IBM TCP/IP for OS/2: Network File System Guide*, SC31-7069

- *IBM TCP/IP for OS/2: X Window System Server Guide*, SC31-7070
- *IBM Firewall for AIX, Reference Guide*, SC31-8418
- *Mastering Windows NT Server 4*, *Third Edition, Sybex Press,* ISBN:0-7821-1920-4
- *Microsoft Windows NT Server*, *Networking Guide*, ISBN:1-57231-344-7
- *NetBIOS Name Servers - Extending Client-Server Applications over Routed IP Networks (Technical Paper), NTS Network Telesystems, Sunnyvale, CA*. Their Web site is located at: `www.nts.com`
- *RFC 1001*
- *RFC 1002*
- *RFC 2055*

How to Get ITSO Redbooks

This section explains how both customers and IBM employees can find out about ITSO redbooks, redpieces, and CD-ROMs. A form for ordering books and CD-ROMs by fax or e-mail is also provided.

- **Redbooks Web Site** http://www.redbooks.ibm.com/

  Search for, view, download, or order hardcopy/CD-ROM redbooks from the redbooks Web site. Also read redpieces and download additional materials (code samples or diskette/CD-ROM images) from this redbooks site.

  Redpieces are redbooks in progress; not all redbooks become redpieces and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

- **E-mail Orders**

  Send orders by e-mail including information from the redbooks fax order form to:

	e-mail address
In United States	usib6fpl@ibmmail.com
Outside North America	Contact information is in the “How to Order” section at this site: http://www.elink.ibmlink.ibm.com/pbl/pbl/

- **Telephone Orders**

United States (toll free)	1-800-879-2755
Canada (toll free)	1-800-IBM-4YOU
Outside North America	Country coordinator phone number is in the “How to Order” section at this site: http://www.elink.ibmlink.ibm.com/pbl/pbl/

- **Fax Orders**

United States (toll free)	1-800-445-9269
Canada	1-403-267-4455
Outside North America	Fax phone number is in the “How to Order” section at this site: http://www.elink.ibmlink.ibm.com/pbl/pbl/

This information was current at the time of publication, but is continually subject to change. The latest information may be found at the redbooks Web site.

IBM Intranet for Employees

IBM employees may register for information on workshops, residencies, and redbooks by accessing the IBM Intranet Web site at http://w3.itso.ibm.com/ and clicking the ITSO Mailing List button. Look in the Materials repository for workshops, presentations, papers, and Web pages developed and written by the ITSO technical professionals; click the Additional Materials button. Employees may access MyNews at http://w3.ibm.com/ for redbook, residency, and workshop announcements.

IBM Redbook Fax Order Form

Please send me the following:

Title	Order Number	Quantity

First name	Last name	
Company		
Address		
City	Postal code	Country
Telephone number	Telefax number	VAT number

☐ Invoice to customer number _____

☐ Credit card number _____

Credit card expiration date	Card issued to	Signature

We accept American Express, Diners, Eurocard, Master Card, and Visa. Payment by credit card not available in all countries. Signature mandatory for credit card payment.

List of Abbreviations

ACL	Access Control Lists
API	Application Programming Interface
CID	Configuration, Installation, and Distribution
DDNS	Dynamic DNS
DHCP	Dynamic Host Configuration Protocol
DLS	DOS LAN Services
DNS	Domain Name Server
DOS	Disk Operating System
EMEA	Europe, Middle East, and Africa
FAT	File Allocation Table
FTP	File Transfer Protocol
FTPD	FTP Daemon
HPFS	High Performance File System
HTTP	Hypertext Transfer Protocol
HTTPD	HTTP Daemon (Web server)
IBM	International Business Machines Corporation
IP	Internet Protocol
ITSO	International Technical Support Organization
JFS	Journaled File System
LAN	Local Area Network
LDAP	Lightweight Directory Access Protocol
LP	Line Printer
LPD	LP Daemon
LVM	Logical Volume Management
MPTS	Multiple Protocol Transport Services
NFS	Network File System
NFSD	NFS Daemon
OS/2	Operating System/2
PMX	Presentation Manager Xwindows
PS/2	Personal System/2
PSnS	Personal Safe and Sound
RIPL	Remote Initial Program Load
RSH	Remote Shell
RSHD	RSH Daemon
SHTTP	Secure HTTP
SMP	Symmetric Multi Processing
SSL	Secure Sockets Layer
TCP	Transfer Control Protocol
TFTP	Trivial FTP
TFTPD	TFTP Daemon
VPN	Virtual Private Networking
WSOD	WorkSpace On-Demand

Index

Numerics
386 HPFS 158
8235 DIALs Client 392

A
abbreviations 439
accept_and_recv() 319
acronyms 439
AIC78U2.ADD 163
Architecture
  Server 8

B
Backup and Recovery 12
Boot Manager 118
Bootable CD 22

C
Capacity Enhancements 166
  keepdossearch 170
  maxconnections 167
  maxopens 168
  maxsearches 169
  maxshares 171
  summary 172
CHAP 361
CHKDSK 140
Client/Server 3
Computing trends 1
Configuring PPP 371

D
DDNS 279
Deep computing 6
DHCP 279
Discontinued Components 20

E
e-business 6

F
Fault Tolerance 158
FDISK 99
File and Print 10
  administration 156
  Architecture 148
  folder 155
  installation 150
  Services 146
File System
  Comparison 93
File Systems 13
filtering 293

H
Host Computing 1
HPFS386 92
  Fault Tolerance 93

I
IBM Networks Coordinated Logon Client 165
IFSM 100
Installable File System Manager
  See also IFSM
IPSec Filters 282

J
Java Application Server 14
JFS 121
  aggregate structure 133
  Cache 121
  CHKDSK 137
  DEFRAGFS 138
  Disk Layout 123
  EXTENDFS 139
  Format 137
  Structure 123
  System Structure 124
  Utilities 136

K
keepdossearch 170

L
Lightweight Directory Access Protocol (LDAP) 14
Local 160
Local Security 160
Logical Volume Manager

See also LVM
Lotus Domino Go 321
Authentication 336
Caching Proxy 335
CGI 333
Configuration 327
errors 334
Fastpath Install 322
Functional Components 331
GWAPI 333
Java Servlets 333
Restricting Access 334
Search Engine 336
SNMP 333
Uninstall 329
Virtual Hosts 335
LVM
Bad Block Relocation 112
Benefits 120
comparison to FDISK 99
expanding a volume 97
Fixed Disk Utility 103
LVMGUI.CMD 95
Operation 112

M

mini-firewall 293
Multiple Server Names 172

N

Neighborhood Browser 163
NET USE 178
Network Computing 4
Network File System 294
See also NFS
Networks Primary Logon Client 165
NFS 294
NFS Utilities 299

P

PAP 361
Performance Considerations 142
Personal Computing 2
Pervasive computing 6
PPP
Internal Architecture 403
OS/2 Warp 386
Parameters 378
passphrase 398
Security 394
TCP/IP Configuration 375
Windows 95 382
Windows 95 Configuration 383
Windows NT 386
PPP Configuration 371
PPP Support 361

R

Remote Access Services 359
Client Support 361
Client System Requirements 363
Configuring 370
Installing 365
IP Considerations 372
LAN-to-LAN 359
LAN-to-remote 359
Remote-to-LAN 359
Remote-to-remote 359
Remove 406
Scenarios 360
System Requirements 362
TCP/IP Protocol Router 373
Using DHCP 381

S

SECADMIN 369
send_file() 319
Server
Adapter and Protocol Services 9
Backup and Recovery 12
File and Print 10
File Systems 13
Hard Space Disk 19
Hardware Requirements 18
Java Application Server 14
Licencing 21
Packaging 21
Remote Access Services 11
System requirements 17
Systems Management 12
TCP/IP Services 10
Versatile 15
Server Packaging 21
Streaming LPD 308
SYN cookies 317

System requirements 17
Systems Management 12

T

TCP/IP 4.21
- Autostart 250
- changes 227
- General 252
- Host Names 246
- Installation 229
- Network Tab 238
- New API calls 319
- NFS 276
- Performance Improvements 316
- Printing 274
- Routing tab 244
- Security 254
- SMP Exploitation 318
- SOCKS 268

TCP/IP Services 10
TCP/IP Stack Structure 226
TCP/IP Toolkits 228
TFTPD improvements 309
Time Server 306
Time Wait State 317
TPAP 361
Tunnels 282

V

Variable Cluster-Sizes 320
Vinca StandbyServer 177
Virtual Memory Support 23
Virtual Private Networks 279

W

Web Server 14
WebSphere 321
- Application Server Manager 345
- Components 338
- Developing Servlets 349
- Fastpath Install 336
- Implementing Servlets 349
- Security 348
- Servlet Activity 347

X

X Windows 312

ITSO Redbook Evaluation

Inside OS/2 Warp Server for e-business
SG24-5393-00

Your feedback is very important to help us maintain the quality of ITSO redbooks. **Please complete this questionnaire and return it using one of the following methods:**

- Use the online evaluation form found at http://www.redbooks.ibm.com
- Fax this form to: USA International Access Code + 1 914 432 8264
- Send your comments in an Internet note to redbook@us.ibm.com

Which of the following best describes you?
_ Customer _ Business Partner _ Solution Developer _ IBM employee
_ None of the above

Please rate your overall satisfaction with this book using the scale:
(1 = very good, 2 = good, 3 = average, 4 = poor, 5 = very poor)

Overall Satisfaction _____

Please answer the following questions:

Was this redbook published in time for your needs? Yes___ No___

If no, please explain:

_____

_____

_____

_____

What other redbooks would you like to see published?

_____

_____

_____

Comments/Suggestions: (THANK YOU FOR YOUR FEEDBACK!)

_____

_____

_____

_____

SG24-5393-00
Printed in the U.S.A.

Inside OS/2 Warp Server for e-business

SG24-5393-00